

Air Force Institute of Technology

AFIT Scholar

Theses and Dissertations

Student Graduate Works

3-2022

Deep Learning Approach to Multi-phenomenological Nuclear Fuel Cycle Signals for Nonproliferation Applications

Preston J. Dicks

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Nuclear Engineering Commons](#)

Recommended Citation

Dicks, Preston J., "Deep Learning Approach to Multi-phenomenological Nuclear Fuel Cycle Signals for Nonproliferation Applications" (2022). *Theses and Dissertations*. 5460.
<https://scholar.afit.edu/etd/5460>

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact AFIT.ENWL.Repository@us.af.mil.



**Deep Learning Approach to
Multi-phenomenological Nuclear Fuel Cycle
Signals for Nonproliferation Applications**

THESIS

Preston J. Dicks, Second Lieutenant, USAF
AFIT-ENP-MS-22-M-087

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENP-MS-22-M-087

DEEP LEARNING APPROACH TO MULTI-PHENOMENOLOGICAL
NUCLEAR FUEL CYCLE SIGNALS FOR NONPROLIFERATION
APPLICATIONS

THESIS

Presented to the Faculty
Department of Engineering Physics
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Nuclear Engineering

Preston J. Dicks, B.S.
Second Lieutenant, USAF

March 2022

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

DEEP LEARNING APPROACH TO MULTI-PHENOMENOLOGICAL
NUCLEAR FUEL CYCLE SIGNALS FOR NONPROLIFERATION
APPLICATIONS

THESIS

Preston J. Dicks, B.S.
Second Lieutenant, USAF

Committee Membership:

Abigail Bickley, Ph.D.
Chair

Maj James Bevins, Ph.D.
Member

Brett Borghetti, Ph.D.
Member

Anthony Franz, Ph.D.
Member

Tony Kelly, Ph.D.
Member

Abstract

The challenges facing the United States in detecting, identifying, and characterizing evidence of nuclear proliferation have expanded in the past decade. In order to reduce the time required for data analysis and decision-making relevant to nuclear proliferation detection, Artificial Intelligence (AI) techniques are applied to multi-phenomenological signals emitted from nuclear fuel cycle facilities to identify non-human readable characteristic signatures of operations for use in detecting proliferation activities. The operational state classification ability of nuclear facility operations using traditional Machine Learning (ML) classification tools such as Random Forest, K-Means, K-Nearest Neighbors, and DBSCAN to name a few is compared to customized Deep Learning (DL) neural networks. In this research, the signals exploited include seismic and magnetic emanations collected in the vicinity of the High Flux Isotope Reactor (HFIR) and the McClellan Nuclear Research Center (MNRC). A DL Bi-Headed 1D-Convolutional Neural Network with shared weights between branches is designed to test the viability of transfer learning between nuclear reactor facilities. It is found that the network produces an 84.1% accuracy for predicting the operational state of the MNRC reactor when trained on the operational state of the HFIR reactor. The accuracy improves to 99.4% when data collected during transient state time periods is removed from the training and testing. In comparison, the best performing traditional ML single-phenomenology classification tool, K-Means, produces a 67.8% accuracy for predicting the operational state of the MNRC with a 80.5% accuracy when transient state time periods are excluded. The applicability of these techniques to alternate sites and fuel cycle operations is also explored.

AFIT-ENP-MS-22-M-087

*Dedicated to all the individuals who know their curiosity will never be satisfied, but
yet still seek answers to unending questions...*

Acknowledgements

I would like to thank all of the family, friends, faculty, sponsors, and collaborators. Both directly and indirectly, each and every one of these people have provided a system of support and direct feedback that have both kept me focused and motivated.

I would like to thank my advisor and committee for their overall care, support, and overall contribution to my academic, personal, and professional development. Through all the sleepless nights, everyone on the committee devoted their personal time and professional passion in supporting me within this project.

I would like to thank Dr. Bethany Goldblum and the University of California - Berkeley Complexity group for allowing me to utilize their datasets and their collaboration.

I would like to thank Ms. Hilary Pederson for her 24/7 support in limiting technical, procedural, and administrative setbacks. Her care and promptness has allowed me to fully focus on conducting research.

Finally, I would like to thank the United States Air Force (USAF) for providing me with the opportunity of serving in the world's greatest military and the Defense Threat Reduction Agency (DTRA) for sponsoring this project.

Preston J. Dicks

Table of Contents

	Page
Abstract	iv
Dedication	v
Acknowledgements	vi
List of Figures	ix
List of Tables	xiii
I. Introduction	1
1.1 Overview	1
1.1.1 Dual-Use Technology	1
1.1.2 Nuclear Fuel Cycle (NFC)	2
1.2 Motivation	3
1.3 Research Hypothesis	5
1.4 Research Assumptions & Limitations	6
II. Theory	8
2.1 Background	8
2.1.1 Magnetic Phenomenology	8
2.1.2 Acoustic Phenomenology	11
2.1.3 Seismic Phenomenology	14
2.1.4 Nuclear Fuel Cycle (NFC) Facilities	16
2.2 Nuclear Reactors and Nuclear Fuel Use	24
2.2.1 Flux-Trap Reactor	26
2.2.2 TRIGA Reactor	28
2.3 Machine Learning	30
2.3.1 Task T : Classification	31
2.3.2 Performance Metrics P	32
2.3.3 Data Distribution Metrics	34
2.4 Deep Learning	35
2.4.1 Convolutional Neural Networks (CNN)	37
2.5 Multimodal Approaches	39
III. Data	41
3.1 Description	41
3.1.1 Data Collection Locations	41
3.1.2 Merlyn Multimodal Sensor Platform	43
3.2 Preliminary Work	44

	Page
3.2.1 Pre-Processing	44
3.2.2 Data Distribution	48
IV. Methodology & Results	53
4.1 Machine Learning (ML): Single-Phenomenology (SP) Trials	53
4.1.1 ML Methodology	53
4.1.2 ML-SP Results: Steady & Transient State	56
4.1.3 ML-SP Results: Steady State	58
4.2 Deep Learning (DL): Single-Phenomenology (SP) Trials	64
4.2.1 DL Methodology: Traditional Configuration.....	64
4.2.2 DL-SP Results: Steady & Transient State	70
4.2.3 DL-SP Results: Steady State.....	74
4.3 Deep Learning (DL): Multi-Phenomenology (MP) Trials	76
4.3.1 DL Methodology: Novel Configuration	76
4.3.2 DL-MP Results: Steady & Transient State.....	82
4.3.3 DL-MP Results: Steady State	87
V. Conclusion	91
5.1 Research Conclusions	91
5.2 Future Work	94
Bibliography	96

List of Figures

Figure		Page
1	An illustration of the Nuclear Fuel Cycle with uranium and plutonium fuels	3
2	Exline’s initial assessment of the practicality in using machine learning for weapons of mass destruction applications	4
3	An illustration of the application of a Black Box Model to nuclear fuel cycle facilities.	8
4	The different types of seismic waves propagating through the Earth.....	15
5	The different pressure gradients illustrated within a gaseous diffusion stage.	21
6	Several cascades illustrated from the K-25 Gaseous Diffusion Facility	22
7	The different parts of a Zippe-Type centrifuge stage	24
8	The world’s first nuclear reactor: Chicago Pile I	25
9	The chain reaction of the nuclear fission process	26
10	The system of a nuclear reactor.....	27
11	The reactor core components of a flux-trap type of nuclear reactor	28
12	The UZrH fuel rods used in the TRIGA type of nuclear reactors	29
13	An example of a confusion matrix.....	34
14	Multiple perceptrons create and artificial neural network	36
15	An example of a 2D CNN classification task	37
16	The detection platforms deployed around the HFIR and MNRC	42
17	HFIR Node 9 Cooling Tower	43

Figure		Page
18	The Merlyn detection platform	44
19	The process of stratification within sampling	46
20	The process of stratifying during systematic sampling.....	47
21	The On/Off binary operational distribution of the HFIR and MNRC data sets	50
22	Principle Components Analysis performed on the MNRC data set	52
23	ML-SP algorithms performed on the HFIR & MNRC data sets separately	55
24	The predicted labels resulting from the ML-SP trial with the inclusion of transient states	56
25	The confusion matrix resulting from the ML-SP trial with the inclusion of transient states	57
26	The predicted labels resulting from the ML-SP trial with the exclusion of transient states	59
27	The confusion matrix resulting from the ML-SP trial with the exclusion of transient states	60
28	The NearestNeighbor algorithm performed on the MNRC data set used within the ML-SP trial, excluding transient states	61
29	The DBSCAN algorithm performed on the standardized MNRC data set used within the ML-SP trial, excluding transient states	62
30	The DBSCAN algorithm performed on the normalized MNRC data set used within the ML-SP trial, excluding transient states	63
31	The data organization for the <i>Traditional</i> Input Configuration	65
32	The Generic Traditional Bi-Headed 1DCNN for the DL-SP Trial	66
33	The DL-SP network in the configuration used for training	68

Figure		Page
34	DL-SP network in the configuration used for testing	69
35	The DL-SP network outputted by Keras	70
36	The resulting Manhattan distances from the DL-SP network	71
37	The DL-SP binary classification results with the inclusion of transient states	72
38	The resulting DL-SP confusion matrix with the inclusion of transient states	73
39	The DL-SP binary classification results with the exclusion of transient states	75
40	The resulting DL-SP confusion matrix with the exclusion of transient states	76
41	The data organization for the <i>Novel</i> configuration	77
42	The Generic Novel Bi-Headed 1DCNN for the DL-MP Trial	79
43	The DL-MP network in the configuration used for training	80
44	The DL-MP network in the configuration used for testing	81
45	The DL-MP network outputted by Keras	82
46	The resulting Manhattan distances from the DL-MP network with the inclusion of transient states	83
47	The DL-MP binary classification results with the inclusion of transient states	84
48	The resulting DL-MP confusion matrix with the inclusion of transient states	86
49	The resulting Manhattan distances from the DL-MP network, excluding transient states	87
50	The DL-MP binary classification results with the exclusion of transient states	88

Figure		Page
51	The resulting DL-SP confusion matrix with the exclusion of transient states	90

List of Tables

Table		Page
1	The worldwide distribution of uranium resources	18
2	The silhouette scores resulting from using the un-scaled MNRC data set	49
3	The three dominant features resulting from PCA	52
4	The Same/Not-Same Equivalence Gate	65

DEEP LEARNING APPROACH TO MULTI-PHENOMENOLOGICAL NUCLEAR FUEL CYCLE SIGNALS FOR NONPROLIFERATION APPLICATIONS

I. Introduction

1.1 Overview

1.1.1 Dual-Use Technology

Nuclear technology provides a multitude of positive contributions to daily life through applications that support a wide range of fields [1]. However, these benefits come with the risk of weaponized proliferation from technologies applicable to the development of nuclear weapons. As such, the demarcation lines that signify the threshold between peace and conflict are constantly adjusted for dual-use technologies relative to fuel cycle facilities [2]. In an attempt to help solidify these lines, international nuclear treaties, e.g. the Treaty on the Non-Proliferation of Nuclear Weapons (NPT), attempt to satisfy this balance between peace and security with large-scale agreements against the spread of nuclear weapons [3]. International organizations, such as the International Atomic Energy Agency (IAEA), implement safeguards to monitor treaty compliance. The analysis of data collected via IAEA inspections and facility monitoring relies upon techniques and technical advances from the field of nuclear forensics to continuously improve monitoring fidelity [4, 5]. The pace of technological innovation has led to an increase in the quantity and types of data to be analyzed from monitoring activities, producing slow processing times with the use of conventional analytical methods [4, 6]. Furthermore, as the monitoring domain evolves, the

growing influx of information can lead to delays in decision-making and hinder steps towards identifying and preventing nuclear proliferation. One of the domains of interest within the nuclear proliferation detection field includes the production, handling, transportation, and use of nuclear material termed the *nuclear fuel cycle*.

1.1.2 Nuclear Fuel Cycle (NFC)

The nuclear fuel cycle (NFC) consists of physical processes like the mining and milling of nuclear fuel and technological processes like enrichment and nuclear facility-use, e.g. nuclear reactors [7, 8, 9]. These are dual-use processes as actors may use parts of the nuclear fuel cycle for peaceful applications, e.g. nuclear energy, allowing for the potential concealment of non-peaceful, clandestine, nuclear proliferant activities. Figure 1 illustrates an example of the nuclear fuel cycle steps with uranium and plutonium used for nuclear fuel.

Although compressed in Figure 1, the nuclear fuel cycle physically takes place over a multitude of locations within a variety of different facilities. For example, uranium mining may take place in Australia, Canada, South Africa, or the United States to name a few [9]. Furthermore, there are significant variations in both the types of nuclear facilities and their purposes. Representative diverse examples of fuel cycle facilities include a gaseous diffusion plant in Paducah, Kentucky USA for uranium enrichment, pressurized water reactors (PWR) in France, and boiling water reactors (BWR) in the United States [9]. According to the Department of Defense (DoD), traditional nuclear proliferation indicators have included fissile material production, acquiring a workable nuclear weapon design, weapon manufacturing, weapon delivery system, and the most difficult indicator to detect being the harvesting of fissile material [10]. Therefore, in order to advance lead times for detecting proliferation, the front-end of the NFC will be the focus of this project.

Nuclear fuel cycle

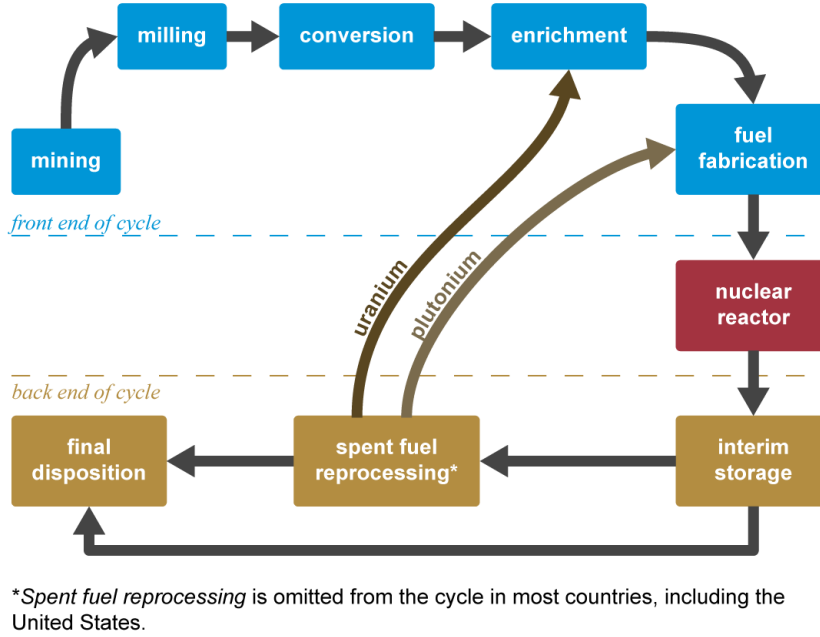


Figure 1. The front end of the nuclear fuel cycle, shown in blue, is this project’s focus for early proliferation indicators. Reproduced from Pennsylvania State University Radiation Science and Engineering Center (public domain) [7].

1.2 Motivation

As the spheres of interest involved in the development of nuclear technology are similar to the vastness and growing complexity of the NFC’s domain size and heterogeneous modality, an exploration of the application of Machine Learning (ML) in support of passive nuclear monitoring is warranted. With recent developments in computational power, ML has been applied to solving problems with large probability spaces, Figure 2, given sufficient amounts of data in real-world observations [10]. These problems typically have well-defined datasets and input feature importance from well-characterized domain knowledge. In nuclear proliferation detection applications, a need exists for efficient tools that adjust to processing different data sizes and account for the non-availability/non-practicality of *truth-labeled data*, data

where observation states are known. This makes typical supervised ML methods for categorizing similar data, termed *classification*, not adequate. Examples of using machine learning as a tool for technical proliferation detection have often included truth-labeled methods: e.g. using potential textual indicators [11, 12].

	Current			5-10 Years			10-20 Years		
	Data Set	Uncertainty	Value	Data Set	Uncertainty	Value	Data Set	Uncertainty	Value
Raw material acquisition									
Excess production of low-enrichment U-235									
Excess enrichment levels of U-235									
PU-238 short fuel cycle enrichment									
Power generation excess									
Excavations / hidden facilities									
Refusal of international inspections									
Nuclear weapons scientists									
Nuclear weapon design acquisition									
Nuclear Suppliers Group controlled list acquisitions									
Chemical explosives expertise									
Complex electronic expertise									
Delivery system development									
Nuclear testing									
State decision to proliferate									
State decision to employ									
Fully Capable/Feasible									
Mostly Capable/Feasible									
Partially Capable/Feasible									
Not Capable/Feasible									
Not Applicable									

Figure 2. Exline’s assessment shows that the majority of weapons of mass destruction machine learning applications are typically limited by the specific environment and amount of human intervention during implementation. Reproduced with permission from Springer Nature [10].

As truth labels are often inaccessible or impractical, the focus of this research is on

predicting a nuclear fuel cycle facility’s binary (On/Off) operational labels through signatures identified by machine learning (ML) models. Deep learning, a subset of ML, has historically been shown to be successful in this application space by Brinker *et al.*, who identified operational conditions at a nuclear fuel cycle facility with higher classification fidelity than on/off status by training a 1D Convolutional Neural Network (1DCNN) using only magnetic field data [13]. In this project, the approach is broadened to include the consideration of acceleration/seismic data to test whether an increased classification fidelity through utilizing fundamental physical dependencies between phenomenologies can be obtained. In addition to the testing of different ML algorithms, a bi-headed twin neural network is constructed and trained using multi-phenomenological data collected at a single facility prior to being tested on data collected at a separate NFC facility using the pre-trained weights and biases. This approach allows the feasibility of *transfer-learning* to be tested within the NFC domain with the goal of identifying signatures emitted during the course of operations that may be used to distinguish specific equipment operations from balance of plant (BOP) operations.

1.3 Research Hypothesis

Although previous deep learning techniques have yielded high classification accuracy for binary operational states (On/Off) with single-phenomenology nuclear fuel cycle signals [13], to the author’s knowledge, the exploration of deep learning techniques across different facilities within a given step of the nuclear fuel cycle using multiple phenomenologies has not been conducted. With the large domain-space considered, direct deep learning implementations run the risk of indirectly over-fitting to specified scenarios, limiting transferability to similar facilities within the same step of the nuclear fuel cycle. As a result, this possibility limits the power derived from machine

learning over traditional analytical techniques. Therefore, in order to increase signature identification, analysis, and operational classification accuracy within nuclear fuel cycle steps, this project investigates the following objectives:

- Design a deep learning algorithm to explore the use of multiple phenomenologies and their effect on NFC operational classification accuracy as compared to single-phenomenology analysis.
- Analyze multi-phenomenology signals collected by sensors positioned at varying positions relative to the target facility of interest to allow the distance dependence of the operational state classification accuracy to be determined.

It is hypothesized that the use of multiple phenomenologies will decrease information loss and increase analytical fidelity relative to single-phenomenology analysis, producing greater relative classification accuracy further from sources/facilities of interest.

1.4 Research Assumptions & Limitations

This project comes with a variety of assumptions and limitations present in the datasets and model implementation.

Assumptions:

- Specific equipment correlates with nuclear facility power level.
- Inter-dependencies are present between the phenomenologies of interest.
- Similarities exist between nuclear facilities such that trained models are transferable.

Limitations:

- The 16 Hz sampling rate used may limit the ability to distinguish operational signatures.
- The time-series based analysis may limit signature attribution to specific equipment.
- The signal-to-noise ratio (SNR) may be low due to an unfiltered background present within the data.

II. Theory

2.1 Background

This project considers each scenario of interest as a hypothetical black box as there are many possible facilities within the nuclear fuel cycle (Figure 1), as well as limitations on available representative data for each facility [10]. This results in the inputs, facilities, and outputs as the three fundamental components considered; e.g. input power, nuclear facility operation, and outputted phenomenology, respectively (Figure 3). The output phenomenologies are assumed to have physics-based interdependencies, signatures, and correlations with respect to different facility operations. However, without the knowledge of real-time domain parameters, accurately modeling the dynamics between these phenomenologies without information loss becomes computationally challenging. Thus, it is important to first macroscopically consider the fundamental physics of each phenomenology prior to exploring modeling methods.




Figure 3. An illustration of the application of a Black Box Model to nuclear fuel cycle facilities.

2.1.1 Magnetic Phenomenology

Magnetic fields are vectorized fields arising from the movement of charged particles [14]. Created from either an intrinsic movement of charge within a material or

a transference of energy, magnetic fields are fundamentally related to other interactive fields through field changes arising from charged particles and kinetic energy, e.g. charge-generated electric fields. When particles interact with one another within a field, energy is transferred between particles through *forces* or impulses through the fields. Conceptualizing that forces represent the interaction between particles, the electric field signifies the range of influence exerted by a charged particle and the magnetic field subsequently signifies how this influence changes over time. The sources that generate magnetic fields are assumed to generate bi-directional fields as no natural magnetic monopoles have been observed in nature. This directional dependence manifests in a net summation with stronger fields overshadowing weaker fields, allowing strong fields to principally drive the net field direction detected. The relationship between magnetic fields, directional dependence, electric fields, and the movement of energy results in a complex, time-dependent system.

The magnetic field of this system can be described using a few equations, where bold terms represent vector components [14]. The Lorentz Force (Eq. 1) describes the fundamental relationship between the force \mathbf{F} [N], electric field \mathbf{E} [N/C], and magnetic field \mathbf{B} [T]. Additionally, q [C] represents the electric charge and \mathbf{v} [m/s] represents the source particle velocity. This relation states that a traveling charged particle interacts with the electric fields and magnetic fields from other sources, where the cross-product between both fields represents the induced force vector.

$$\mathbf{F} = q\mathbf{E} + q\mathbf{v} \times \mathbf{B} \quad (1)$$

As the Lorentz force equation illustrates the fundamental relations between the vector fields of interest, the following equations are utilized for the magnetic field within more complex systems:

$$\mathbf{B} = \hat{\mathbf{z}} \frac{\mu_0 I}{2\pi r} \quad (2)$$

$$d\mathbf{B} = \frac{\mu_0 I}{4\pi} \frac{d\mathbf{l} \times \hat{\mathbf{r}}}{r^2} \quad (3)$$

Equation 2 states that a constant current I [A] scaled by the vacuum permeability constant $\mu_0 = 4\pi \times 10^{-7}$ [N/A²] divided by 2π and flowing in the unit-vector direction $\hat{\mathbf{z}}$ induces a magnetic field with a strength inversely proportional to the radial distance r [m] from the source. Equation 3, termed the ***Biot-Savart Law*** is the foundation for calculating the instantaneous strength of the induced magnetic field $d\mathbf{B}$ resulting from the cross-product of a slowly-changing or steady current I [A] within an infinitesimal wire element $d\mathbf{l}$ and the radial unit vector $\hat{\mathbf{r}}$ from the wire, scaled by the vacuum permeability constant μ_0 [N/A²] divided by 4π . Both of these equations allow the magnetic field to be visualized and interpreted through black-box scenarios within the nuclear fuel cycle. For example, the current flowing through the power lines to a facility is travelling through a wire of infinitesimal length $d\mathbf{l}$. The projection of the magnetic field onto an XYZ reference frame originating at $d\mathbf{l}$ (typically a detector) shows the orthogonality resulting from the cross product between $d\mathbf{l}$ and $\hat{\mathbf{r}}$ within Equation 3.

Equation 4, termed ***Faraday's Law***, represents the induced electromotive-force ϵ or *voltage* resulting from the integration of the electric field \mathbf{E} passing through a surface area \mathbf{s} along the line C , the integration of the total magnetic field \mathbf{B} over the area \mathbf{a} of Gaussian surface \mathbf{S} changing through time $\frac{d}{dt}$, or a changing magnetic flux $-\frac{d\Phi_B}{dt}$ over time. From this relation, the magnetic field arising from an operational fuel cycle facility is detectable through the changing magnetic flux resulting in an induced voltage within a detector system.

$$\epsilon = \int_C \mathbf{E} \cdot d\mathbf{s} = \frac{d}{dt} \int_S \mathbf{B} \cdot d\mathbf{a} = -\frac{d\Phi_B}{dt} \quad (4)$$

Within the context of nuclear fuel cycle facilities, principle drivers of localized mag-

netic fields can be attributed to two fundamental systems: environmental magnetism and facility-based magnetism. Furthermore, these systems cause successive induced magnetic fields that add to the overall system complexity. Environmental magnetism includes the earth’s magnetic field, naturally magnetic materials, induced magnetism through ferromagnetic materials, and magnetic anomalies [15, 16]. However, under the assumption that environmental magnetism has weaker fields than those arising from nuclear fuel cycle facilities and that these fields are consistent throughout detection efforts, environmental magnetism is treated as a characterizable and, therefore, subtractable background. Facility-based magnetism can be attributed to a variety of facility operations that, within this project, are separated into balance-of-plant (BOP) and specific-equipment (SE) magnetism, where BOP includes any support equipment that is not of interest and successively induced magnetic fields as in the case of facility lights. Magnetic anomalies within facility-based magnetism occur during nonstandard events such as within an earthquake or localized, directionally-inverted magnetic fields resulting from magnetic-shielding material such as mu-metal [17, 18].

2.1.2 Acoustic Phenomenology

The acoustic phenomenology represents the signals obtained from detecting and processing sound waves. Using basic linear theory, disturbances in homogeneous, idealized media produce waves that interact with one another additively through constructive and destructive *superposition* [19]. Although scientists have historically attempted to formulate a variety of theories to fully describe the physics of acoustic waves, real-world applications have shown a high degree of modeling complexity and unpredictability [20, 21, 22]. Furthermore, as real-world application domains become increasingly complex, linear continuum mechanical theories have oversimplified the interdependence between terms during modeling as the linear theories treat acous-

tic waves as being independent from one another albeit though additively interfering during detection [19, 23]. In the case of physical and nonlinear acoustics, further modeling has indicated an underlying physical-relationship between acoustic waves, often leading to in-determinism and chaos-theory modeling [19, 23].

Although the physics of the acoustic phenomenology is highly complex and still developing, the following relations allow the general nature of acoustic waves to be understood for ML applications [19]:

$$\nabla\left(\frac{1}{\rho}\nabla p\right) - \frac{1}{\rho c^2}\frac{\partial^2 p}{\partial t^2} = 0 \quad (5)$$

$$\nabla^2 p - \frac{1}{c^2}\frac{\partial^2 p}{\partial t^2} = 0 \quad (6)$$

$$p(x, t) = f(x - ct) + g(x + ct) \quad (7)$$

As acoustic detection arises from changes in the pressure p of the travelled medium, Equation 5 applies the *wave equation of linear acoustics* to represent the pressure change within a medium of density ρ over time $\frac{\partial^2}{\partial t^2}$ with respect to the speed of light c . However, if the ambient density is independent of position, such as in idealized scenarios, then Equation 5 reduces to the Eulerian-based Wave Equation of Linear Acoustics. Often, as shown in Equation 6, this relation is written as a Laplacian equation using the d'Alembertian operator. The Laplacian indicates homogeneity within the traveled medium as the equation shows a dependence only on the pressure of the adiabatic medium. Equation 7 is a generic solution to Equation 6 with $f(x - ct)$ and $g(x + ct)$ as generic functions with changing positional inputs. This equation will be used to represent the initial wave emitted from the source prior to generalizing for perturbations such as energy dissipation.

Equation 8 represents a distance x perturbation that decays exponentially $e^{-\alpha x}$. This perturbation is dependent on the traversed medium's transmission property α

applied to the initial pressure of the plane-wave $|\hat{p}(0)|$ arising from energy absorption and thermal conduction through the transported medium or bulk material. Assuming spherical diffraction, the spread of energy in all directions can be modeled using the Helmholtz Equation leading to an inverse dependence on r , or the radial distance away from the source within Equation 9, where k is the wave number and A is a solution coefficient. Deriving the time-averaged intensity \mathbf{I}_{av} of the diffracting spherical acoustic wave, the real solution $Re(\hat{p}\hat{v}_r^*)$ is found to be the inverse-squared dependence on the radial distance $\frac{1}{r^2}$ and density of the traversed medium ρ scaled by the speed of light c .

$$|\hat{p}(x)| = |\hat{p}(0)|e^{-\alpha x} \quad (8)$$

$$p = A \frac{e^{ikr}}{r} \quad (9)$$

$$\mathbf{I}_{av} = \frac{1}{2} Re(\hat{p}\hat{v}_r^*) = \frac{1}{2} \frac{1}{\rho c} \frac{|A|^2}{r^2} \quad (10)$$

As the nuclear fuel cycle consists of a variety of facilities at different locations, there are many possible sources for acoustic waves. These sources include environmental noise from passing cars, people, and weather processes, as well as internal acoustic waves generated through the use of process-specific equipment. However, simplifying this problem domain to consider only linear-physical acoustic waves generated from a black-box facility, the detection of acoustic waves is expected to be inversely dependent on the squared radial distance of the sensor from the source. Linear theory allows the interaction between different acoustic waves to be additive; therefore, this project's detection efforts focus on the overall signal received as well as the conversion of acoustic phenomenologies into seismic phenomenologies.

2.1.3 Seismic Phenomenology

Similar to acoustic phenomenology, seismic phenomenology is modeled using wave continuum theory through bulk material [24, 25, 26]. This indicates the need to consider any inter-dependencies between the seismic and acoustic phenomenologies. Within this project, seismic waves are detected through the Earth's crust with accelerometers. However, as the Earth's crust is inhomogenous with nonlinear parameters, a first-order approximation is often used. This assumes the Earth's crust is an isotropic material that can be modeled with the Theory of Elasticity through the following equations:

$$\nabla^2 \phi - \frac{1}{r^2} \frac{\partial^2}{\partial r} (r^2 \frac{\partial \phi}{\partial r}) = \frac{1}{\alpha^2} \frac{\partial^2 \phi}{\partial t^2} \quad (11)$$

$$\phi = \frac{1}{r} f(t - \frac{r}{\alpha}) \quad (12)$$

$$\mathbf{u} = \nabla \phi \quad (13)$$

$$u_r = -\frac{1}{\alpha r} f'(t - \frac{r}{\alpha}) - \frac{1}{r^2} f(t - \frac{r}{\alpha}) \quad (14)$$

Equation 11 is the generic wave equation represented radially with a Laplacian ∇^2 in accordance with wave continuum theory, where ϕ represents the energy potential within a given medium with parameter α . The solution for this potential is expressed radially, r , within the time, t , dependent Equation 12 with a generic function $f(t - \frac{r}{\alpha})$, with α determined by the propagation material. Equation 13 is the resulting displacement \mathbf{u} of the disturbed ∇ energy potential ϕ derived from Equation 12. The radial component of this displacement is represented in Equation 14. This suggests that seismic wave propagation (Figure 4) can be categorized in two different forms: *Body Waves* and *Surface Waves*. Body waves include longitudinal primary waves (P-Waves) acting in a compressional manner and transverse secondary waves which shear through the material. Surface waves include Love Waves and Raleigh Waves,

categorized by either the linear or circular motion of constituent particles, respectively.




Figure 4. The type of seismic wave produced influences how the particulates present within the medium move, influencing detected signatures. Reproduced from *VectorMine/Shutterstock.com* through a Standard Image License.

Within the NFC, the detection of processes that drive seismic propagation is dependent on the distance away from the driver. The derived radial component of displacement away from the driver (Equation 14) allows for two detection distances: near-field ($\frac{1}{r^2}$) and far-field ($\frac{1}{r}$). By deriving the velocity from the radial displacement component, it follows that the arrival time of primary waves is less than that of secondary waves, the arrival time of secondary waves are less than that of Love waves, and the arrival time of Love waves are less than that of Rayleigh Waves. Therefore, the Near-Field regime is driven by the detection of body waves and the Far-Field regime is driven by the detection of surface waves.

2.1.4 Nuclear Fuel Cycle (NFC) Facilities

As shown in Figure 1, the NFC consists of all facilities involved in handling, distributing, and using nuclear fuel, often categorized in terms of the “front end,” “middle,” or “back end” of the cycle. However, the cycle may look different within different environments depending on resources, scale, and intent. For example, the *Once-through Fuel Cycle* does not have a back end that is similar to the *Reprocessing Fuel Cycle* as the former lacks the nuclear fuel transfer and storage steps post-use [9]. Therefore, in scenarios such as these, certain steps like the interim storage of spent fuel can be considered either “back end” post-use or “front end” pre-use, depending on context. As such, generalizing further, the NFC categories can also be correspondingly considered either “*preparatory*”, “*in-use*”, or “*conclusive*” with respect to how the nuclear fuel is handled [9, 27]. Within this writing, references to the different steps in the NFC will be in correspondence with the general steps shown within Figure 1 and described within the upcoming sections. Furthermore, to align with this project’s focus, this paper’s discussions will be limited to the steps of the NFC which allow for nuclear fuel use in reactors.

2.1.4.1 NFC Front-End “Preparatory”

The front end of the NFC includes facilities that “prepare” or source, configure, and purify nuclear material for use. The first two steps of the cycle, *mining and milling*, are when nuclear fuel is sourced and unwanted contaminants are removed. Typically, fuels are naturally occurring (e.g. uranium); therefore, most mining and milling facilities involve open mining or in-situ leaching (ISL) to extract ore to be fabricated into nuclear fuel from deposits in the Earth’s crust [9, 27]. The next three steps: *conversion*, *enrichment*, and *fuel fabrication* are the processes for converting the nuclear fuel into products for use, e.g. within a nuclear reactor. This writing will

use uranium as the primary element to describe the front-end processes. Uranium is composed of several *isotopes* or atoms that are distinguished through different neutron numbers with the most naturally abundant of uranium’s isotopes being ^{238}U and ^{235}U with atom percent abundances of 99.27% and 0.72%, respectively [9, 28]. ^{235}U is a *fissile* material as it is able to *fission* (Section 2.2), while ^{238}U is a *fertile* isotope as it can be converted into a fissile material [27]. Because these materials are atomically dense, both ^{235}U and ^{238}U are sought as large amounts of power are able to be derived from interactions with these materials during fission, resulting in technology such as nuclear reactors and atomic weapons.

Within the field of *nuclear forensics*, the analysis and attribution of the processes governing the use and misuse of nuclear material relies heavily on a variety of physical, chemical, visual, and isotopic signatures [4]. Furthermore, depending on the type of materials present, the IAEA typically recommends recording eight types of characterization data ranging from physical characteristics to trace element concentrations for materials characterization analysis [29, 30]. While some of the practices within the Front-End of the NFC usually create large environmental signatures due to processes such as large fuel transportation vehicles, patterned human activity, drilling, and chemical composition changes, depending on the context, environment, and resources, these processes may come to fruition differently, making the characterization of signatures difficult. Therefore, in an attempt to widen the scope of interest, as well as to reduce processing times, the possibility of passive monitoring assisted by machine learning algorithms has garnered much interest from organizations like the IAEA [31].

Mining and milling nuclear fuel is an expensive and tedious process. For example, in cases where uranium is used as fuel, resource availability, environmental concerns, financial onus, and detect-ability are all taken into account. Uranium ore concen-

trations vary globally, but concentrations in the Earth’s crust typically range from 1-500,000 parts-per-million (ppm) with average estimates often quoted as around 3ppm [6, 9, 32]. In the 2020 Nuclear Energy Agency (NEA) report, the global distribution of *identified conventional resources* or uranium resources that are assured or inferred across 16 countries is identified with the 7 highest distributions shown in Table 1 [33].

Table 1. Worldwide distribution of identified conventional uranium resources across the top seven countries [33].

Country	Distribution Percentage %
Australia	28%
Kazakhstan	15%
Canada	9%
Russia	8%
Namibia	7%
Brazil	5%
China	4%

In addition to the regional resource deposits often facilitating the relocation of mining and milling, environmental byproducts and radiation safety are also of concern. For example, miner’s *Bergkrankheit* or “mountain sickness” has been attributed to high uranium concentrations within rocks [9]. Furthermore, as the radionuclides within the uranium ores decay, both the uranium and thorium decay series pose health challenges to miners as well as challenges to atmospheric pollution. Often times, mining and milling operations take place away from urban districts where both health and environmental measures can be taken to prevent excessive pollution. Mining and milling typically takes place side-by-side to reduce overall costs as the transportation

of mined uranium ore from remote locations to more-urban milling locations becomes very costly [32]. However, mining in remote locations also leaves room for clandestine proliferant activities through difficulties in accountability and monitoring.

Several initiatives have been undertaken to implement ML for supporting the characterization of nuclear proliferation signatures within mining and milling. However, as the act of mining and milling itself is not the greatest indicator of proliferant activity, the application of ML techniques has shifted the focus from answering “what” to answering “how.” For example, through using open-source commercial materials like Google Earth, both Sundaresan *et. al.* and Kim were able to identify the “what” through ML image processing techniques [34, 35]. While Kim identified and supported the viability of using ML to identify facilities of interest within satellite imagery, Sundaresan *et. al.* attempted to quantify the performance of the ML identification process through comparing the target facility of interest to process-based observables needed for known mining, milling, and production techniques [34].

Diab et. al. used language to discriminate potential nuclear proliferant mining and milling activities from traditional non-nuclear mining and milling activities [36]. As Sundaresan *et. al.* and Kim’s analyses focused on the “what,” this implied the already known knowledge of a proliferant action taking place as satellite imagery is capable of high resolution images but not necessarily in knowing where to observe. However, *Diab et. al.*’s analysis attempts to limit the overall problem space by focusing on relevant language indicators. The differences in technique between Sundaresan *et. al.*, Kim, and *Diab et. al.* highlight the need for explainability and interpretation when developing models for nuclear fuel cycle-based facilities of interest.

Conversion, enrichment, and fuel fabrication are the processes that physically and chemically prepare the attained nuclear fuel for use in technologies ranging from peaceful applications (e.g. nuclear reactors) to conflict-oriented preparation (e.g. nu-

clear weapons). Within NFC's including uranium enrichment, uranium ore (U_3O_8) is converted into (UF_6) during *conversion* as the temperature dependence of the compound's phase allows UF_6 to be in the gaseous form at slightly elevated temperatures as required for many enrichment methods to exploit the mass-based density gradient needed for isotopic separation [9, 6]. However, as *conversion* is a chemical process, resources are needed to not only perform the conversion, but to also store the highly corrosive UF_6 in specially designed containers [37].

This gas is then eligible used in several different types of *enrichment* processes including [9]:

- Gaseous Diffusion
- Centrifugal Separation
- Aerodynamic Processes
- Electromagnetic Separation
- Laser Enrichment

Historically, although all five processes listed above have been explored with variable success rates at industrial levels, gaseous diffusion and centrifugal separation are the most common and will therefore be explained within this section.

Gaseous diffusion is a process dependent on Fick's law as the gaseous fuel is passed through a permeable membrane with pressure gradients for isotopic separation. Fick's law, shown in Equation 15, states that the diffusion flux, J , is dependent on how the concentration ϕ of the gaseous particles change with respect to position x and is scaled by the probability of interaction within the membrane termed the *diffusion coefficient* D . In the case of uranium fuel, as UF_6 is passed through a *gaseous diffusion stage* as shown in Figure 5, the difference in mass between the ^{235}U and ^{238}U allows separation to occur as the ^{238}U passes more slowly through the semi-permeable membrane.

$$J = -D \frac{d\phi}{dx} \quad (15)$$




Figure 5. The gaseous diffusion stage creates a pressure gradient resulting in an enriched uranium stream and depleted uranium stream. Reproduced from the U.S. Nuclear Regulatory Commission (public domain) [38].

Within one stage, separation allows a greater ratio of ^{235}U particles to pass into the *head* or enriched stream of the stage onto the next stage, while the greater ratio of ^{238}U particles pass into the *tails* or depleted stream to be recycled into the previous stage's head or to be removed as waste. In addition to gaseous diffusion being highly inefficient, the superposition of micro-inefficiencies within the system leads to difficulties in calculating the number of stages needed to achieve a specified target enrichment level resulting in large facilities with several stages or *cascades* (Figure 6). Facilities such as these often require a large amount of energy, incurring significant cost in money and resources. As a rule of thumb, as many as 1400 stages are needed to achieve an enrichment level of 3% [6]. Depending on the target enrichment level, gaseous diffusion is often replaced by more efficient methods that take less space and time.



Figure 6. The K-25 gaseous diffusion facility operated with several cascades over a large span of space due to the high separative inefficiencies resulting from gaseous diffusion. Reproduced from the Department of Energy (public domain) [39].

An enrichment method that is more time and energy efficient is *gaseous centrifugal enrichment*. Centrifuge systems separate material primarily by the centrifugal force arising from rotational motion. In simplified terms, with *gaseous centrifugal enrichment* using uranium fuel, rotating cylinders (Figure 7) leverage the mass difference between the uranium isotopes to separate $^{235}\text{UF}_6$ from $^{238}\text{UF}_6$ in accordance with the centrifugal force equation (Equation 16) arising from pseudo forces from a non-inertial reference frame that are modeled through Isaac Newton’s Second Law of Motion ($F = ma$):

$$F = m\omega^2 r = m\left(\frac{v}{r}\right)^2 r = m\frac{v^2}{r} = ma. \quad (16)$$

As the cylinder’s rotor revolves at a high speed, the centrifugal force F is scaled proportionally to the mass m , angular velocity ω , and radius r from the center of the cylinder. The angular velocity arises from the linear tangential velocity v through the equation: $\omega = \frac{v}{r}$. In accordance with momentum conservation laws, the gaseous UF_6 maintains the momentum resulting from the inertia-generated centrifugal force

by scaling the radius of the two primary isotopic sub-components ($^{235}\text{UF}_6$, $^{238}\text{UF}_6$) to account for the difference in mass. Therefore, as the heavier isotope particles have more inertia as a byproduct of having a greater mass, similar to the gaseous diffusion process, this results in an overall lower angular velocity through the scaling of the linear velocity $v = \frac{\delta x}{\delta t}$. Consequently, in order to maintain the same centrifugal force applied on both isotopes, Equation 16 is scaled through the heavier isotope (^{238}U) concentrating at a greater radius away from the cylinder’s center. These cylinders are then grouped together to form a system of cascades in which the ^{238}U is taken from the cylinder’s outer radius to be recycled into the previous cylinder, while the ^{235}U is taken from the cylinder’s inner radius to be sent into the next cylinder for further enrichment. After the process is completed, the enriched product is transferred for fuel fabrication.

As *conversion*, *enrichment*, and *fuel fabrication* are chemical and physical processes, detectable signals for proliferation vary between processes. Historically, there have been initiatives to explore chemical indicators of proliferation through the use of ML algorithms [40, 41, 42, 43]. The majority of these initiatives are similar to *Gum*’s analysis with respect to pursuing the characterization of particle morphology or spectroscopic indicators in a field of study named *Nuclear Proliferomics* by *Schwerdt et. al.* As shown by *Pastoor et. al.*, this field has utilized the analysis of uranium-based chemistry to aid in NFC forensics. However, Nuclear Proliferomics techniques traditionally involve on-site presence and destructive analysis (DA) to obtain high-confidence measurements with regards to nuclear monitoring.

Other initiatives have attempted to explore the physical indicators of proliferation with the use of ML algorithms [31, 44, 45, 46]. These physical indicators include flow rate, satellite images, gamma counts, and non-destructive assays (NDA). However, the overall uncertainty of the physical indicators has led to the dominance of chemical




Figure 7. The Zippe-Type of centrifuge stage is more efficient when compared to gaseous diffusion and the system is spatially contained within one stage. Reproduced from the U.S. Nuclear Regulatory Commission (public domain) [38]

indicators and the need for the fusion of multi-source physical indicators as described by Cui [45].

2.2 Nuclear Reactors and Nuclear Fuel Use

Nuclear Reactors are facilities that generate interactions on a sub-atomic scale. Preceding the development of the atomic bomb and the utilization of sub-atomic interactions during its deployment, the concept of *nuclear energy* was explored when Enrico Fermi built the world's first nuclear reactor in 1942 named "Chicago Pile-1" [47] shown

in Figure 8.




Figure 8. The world’s first nuclear reactor was a complex project that took much planning and effort. Reproduced from *Argonne National Laboratory* (CC BY-NC-SA 2.0) [47]

This nuclear reactor relied on utilizing neutron particles and uranium atoms for its interactions in a process called *fission*. As shown in Figure 9, the neutron bombardment of the target nucleus (^{235}U) results in the expenditure of energy in the form of additional free neutrons and the split of the target nucleus into fission product nuclei. This technique became the underlying concept behind how fission-based atomic weapons function.

However, in nuclear reactors, the chain-reaction process is utilized by the balance of *moderators* and *control rods* [48]. As shown in Figure 10, uranium is typically fabricated into nuclear reactor cores through the steps described in Section 2.1.4.1. These cores are bombarded with neutrons to expend heat within the process of fission. As the energy is released as free neutrons and fission products, more fission reactions take place within the uranium core and are slowed by the moderator to sustain the reaction. Interactions with the moderating medium, uranium rods, control rods, and structure produce heat, typically used to heat the moderating medium, such as in the boiling water reactor (BWR) shown in Figure 10 [48]. The steam resulting from this




Figure 9. The nuclear fission reaction is a probabilistic reaction that can be sustained to form a chain. Reproduced from *Karla Beatty/Shutterstock.com* through a Standard Image License.

heat is then converted from kinetic energy into electrical energy through a generator. Although, these reactor processes are traditionally used as sources for nuclear energy, as shown in Figure 1, reactors can also be used to produce another dual-use element: *plutonium* [9].

2.2.1 Flux-Trap Reactor

Reactors with cores that consist of moderating material surrounded by nuclear fuel are termed *Flux-Trap* reactors as fast neutrons are converted to thermal neutrons that emanate from the moderating material prior to being “trapped” by the surrounding *reflector* material [49]. Several research reactors [50, 51, 52] use the flux-trap principle since the convenience of trapping thermal neutrons allows for multiple experiments to be performed concurrently through the *beam tube* and multiple *target rods* as shown in Figure 11.

As described in Section 3.1, one of the datasets used in this research was collected

Components of Nuclear Reactor




Figure 10. The heat produced within a nuclear reactor is converted into electricity for derivative applications. Reproduced from *DKN0049/Shutterstock.com* through a Standard Image License.

at a flux-trap type of research reactor called the High Flux Isotope Reactor (HFIR). Originally built for transuranium isotope production, typical HFIR cycles operate at 85 MW and use highly enriched uranium (HEU) as fuel, beryllium as a reflector, light water as a moderator, and, typically, fabricated transplutonium isotopes for target rods [49]. As uranium's natural isotopic abundance has low amounts of ^{235}U , in addition to the technical challenges posed by the front end of the NFC, an enrichment method was required to create HEU, Section 2.1.4.1. The HFIR's standard highly enriched ^{235}U fuel allows nuclear fission to take place more efficiently, Section 2.2, producing fission products and fast free neutrons to be thermalized by the surrounded light water moderator. Encased with the beryllium reflector, the cycle's primary interactions then take place between the trapped thermal neutron flux and target isotopes. This flux-trap within the HFIR, regulated by the control plates shown in Figure 11, allows an adequate source of thermal neutrons for experiments in which the




Figure 11. The core of the HFIR has target rods surrounded by moderating material and nuclear fuel in accordance with the neutron *flux-trap* principle in order to allow for multiple neutron-based experiments to be concurrently performed. Courtesy of Oak Ridge National Laboratory, U.S. Dept. of Energy. (public domain)

flux is controlled in accordance with user needs. Through the required NFC actions, technological capabilities, 30 target positions within the flux trap, and several other experiment facilities, the HFIR's size and exceptionally high neutron flux creates a large imprint as it is the western world's only supplier of ^{252}Cf [49, 51].

2.2.2 TRIGA Reactor

The Training, Research, Isotope production, General Atomic (TRIGA) reactor was originally commissioned to provide research capabilities with an *inherent safety* mechanism implemented in case of failures within the physically engineered mechanisms (e.g. reactor control rods) [53]. Historically, nuclear reactor accidents have been

attributed to failures in implementing mechanical safety mechanisms, causing catastrophic secondary and tertiary effects such as fuel melting caused by accidental uncontrolled fission (Section 2.2) from the removal of control rods or a buildup of overpressure caused by reactor core interactions with the surrounding environment (e.g. zirconium cladding interacting with steam) [54, 55]. Therefore, TRIGA reactors are designed with a neutron source irradiating fuel rods made of light-enriched ($< 20\%$) uranium, zirconium, and hydrogen (UZrH) for intrinsic neutron regulation with neutrons within the hydrogen of the fuel rod being warmer than the neutrons within the surrounding water, causing less fission in the event of bypassed mechanical safety mechanisms [53, 56]. This property of TRIGA reactors is termed a “prompt negative temperature coefficient of reactivity” [53]. As shown in Figure 12, compared to historical parameters [54, 55] and flux-trap reactor requirements (Section 2.2.1), the light-enriched UZrH composition allows less uranium with lower levels of enrichment to be used with nuclear reactions supported through graphite reflectors and aluminium (historical) or stainless-steel (modern) cladding for increased robustness.




Figure 12. UZrH fuel rods are used in TRIGA reactors, allowing an intrinsic fail safe system based on fundamental interactions between the neutrons within the composition’s hydrogen and the surrounding environment. Reproduced from Pungertič *et. al.* [57] through Creative Commons Attribution-NonCommercial-No Derivatives License (CC BY-NC-ND 4.0)

As described in Section 3.1, a second data set used in this research project was collected at a TRIGA research reactor located at the McClellan Nuclear Research Center (MNRC). Originally commissioned by the United States Air Force, the MNRC reac-

tor typically operates at a 2 MW steady-state with a 30 second, 400 MW, pulsing capability [58]. In agreement with historical TRIGA reactors, MNRC operations uses light-enriched UZrH fuel rods with stainless-steel cladding and control rods that have speeds of 24 inches per minute (37.5 seconds for a 15 inch total travel length) [59]. Compared to the flux-trap type HFIR reactor (Section 2.2.1), the MNRC’s steady-state operations are at a lower power level with less uranium and lower ^{235}U enrichment requirement.

2.3 Machine Learning

Machine Learning (ML) is a subset of Artificial Intelligence (AI), which deals with developing tunable algorithms that are trained on data to be tested on a wide range of applications including: clustering, regression, prediction, etc. Often times, the algorithms are separated into supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning depending on the amount of information available about the data [60]. As referenced in *Goodfellow et. al* [61], *Mitchell* [62] describes ML as a program that is learning from experience E in reference to tasks T with a performance measure P . The methodology of ML algorithms traditionally follows implementing **training** to learn experience E on a dataset from the dataset attributes or *features*, **validating** on a separate un-seen dataset or portion of data to measure how well the experience E makes the model generalizable to other datasets with performance metric P , and **testing** on a targeted dataset to complete the task T . Within this paper, *Mitchell’s* description will be used as a baseline to help describe ML and its application to this project.

2.3.1 Task T : Classification

This project's objective of testing the feasibility of implementing ML and deep learning (DL) techniques on data collected from the nuclear fuel cycle environment sets the task T as the classification of binary reactor operations. Therefore, both supervised and unsupervised learning are used for this project's ML models as the experience E that is taken from these types of designations cover the range of primary ML algorithms used traditionally. Supervised learning is defined as the ML algorithm having access to the labels or targeted datum states used during training. Geron represents supervised learning within an example of a ML spam detector algorithm. Geron shows with this example how the data (emails) are fed into the algorithm along with the dataset's corresponding target labels represented with green check marks and red interdictory symbols signifying *spam* and *not spam*, correspondingly. The features learned for experience E may range from the types of words used and their connotation to a generated metric based on the significance for which letters are used within the email. Therefore, within Geron's example, the ML algorithm would organize the data named *Training set* to train and validate on the dataset with each email's given labels prior to testing on an unlabeled datum named *New instance* with the goal of classifying whether the unlabeled datum is *spam* or *not spam*.

For an unsupervised learning scenario, however, the dataset is fed into the network without labels. Geron also illustrates an example of unsupervised learning in an unlabeled plot of individuals. In this scenario, the model trains, validates, and tests without labels, implying a different, unlabeled, task T in learning the fundamental relationships present in the data, e.g. clustering. For a classification task, the model has to be generalizable enough to perform well on unseen data. Therefore, how well the model can fit to a dataset is an attribute called *capacity* or the amount of potential a model has for fitting a dataset ranging from *underfitting* to *overfitting*.

Goodfellow et. al. illustrates the difference between overfitting and underfitting data for a regression task where a fundamentally quadratic set of datapoints is fit by the three following models [61]:

$$\hat{y} = b + wx \quad (17)$$

$$\hat{y} = b + w_1x + w_2x^2 \quad (18)$$

$$\hat{y} = b + \sum_{i=1}^9 w_ix^i \quad (19)$$

Equation 17 represents a linear model fit that is *Underfitting*, Equation 18 represents a quadratic model fit that matches the underlying data and is represented by a labeled of *Appropriate capacity*, and Equation 19 represents a 9-degree polynomial represented by a label of *Overfitting* [61].

This concept is also represented within classification tasks as, analogously, Equation 17- 19 could represent models that create a classification boundary within the dataset such that data on one side of the boundary are labeled differently than those on the other side. The most optimal decision boundary is called the *Bayes decision boundary*. For a classification task, an appropriate capacity model would attempt to train, validate, and test on datasets such that the classification boundary is similar to the Bayes decision boundary.

2.3.2 Performance Metrics P

To determine the performance of a ML model at classification tasks, four primary metrics are used to compare the labels produced by the ML model or the *predicted classes* and the test dataset's actual labels or *true classes*. Within this project, *actual labels*, *true classes*, and *truth labels* are all used interchangeably. These four metrics

are shown in the Equations 20-23 [63, 64]:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (20)$$

$$F1 = \frac{TP}{TP + \frac{FN + FP}{2}} \quad (21)$$

$$Precision = \frac{TP}{TP + FP} \quad (22)$$

$$Recall = \frac{TP}{TP + FN} \quad (23)$$

These equations are written in terms of true positives TP , true negatives TN , false positives FP , and false negatives FN . These terms correspond to a cross-comparison between the actual labels or classes and predicted labels or classes represented within a *confusion matrix* such as in Figure 13. In a binary classification scenario with two classes (Positive, Negative), e.g. Bernoulli trials, the true predictions are represented by the True Positive TP and True Negative TN categories as these occur when both the predicted and actual classes match. The False Positive FP and False Negative FN categories signify a mismatch between the actual and predicted classes and are synonymous with the statistical *Type I* and *Type II Errors*.

Within classification tasks, accuracy (Equation 20) or a direct comparison between predicted and actual classes does not provide a full illustration of the model's performance in the datum being identified and the datum being mis-classified. Therefore, the F1 score (Equation 21) represents a balance between its two sub-components: precision (Equation 22) and recall (Equation 23). Precision is a metric that provides how well the model classifies on the identified targeted data, while recall, synonymous with *sensitivity*, is a metric that provides how well the model identifies targeted data. Therefore, Equations 20-23 illustrate both the accuracy of a model and the parameterization of its predictive capability. This allows for hyper-parameter tuning




Figure 13. This figure is an example of a confusion matrix between the predicted classes and actual classes. Reproduced from *Sidarta Carvalho/Shutterstock.com* through a Standard Image License.

to lower the occurrence of mis-classification.

2.3.3 Data Distribution Metrics

In addition to the statistics used to describe a data set, this project utilizes two additional descriptors: *silhouette coefficient* and *Manhattan Norm*. Sklearn describes the silhouette coefficient as a statistical metric that indicates the overlap of clusters ranged from $[-1, 1]$ with scores near 0 being an indication of overlapping clusters [65]. Geron describes the Manhattan Norm as a metric that "measures the distance between two points in a city if you can only travel along orthogonal city blocks" [63].

Visualizing this through a right angle triangle, the Manhattan Norm is the total distance represented by the sum of the legs, excluding the hypotenuse. Within this project, the silhouette score is used to indicate the amount of separation present between clusters while the Manhattan norm is used to both reduce the power given to outliers and as a distance metric that treats the difference between zero and non-zero elements with importance [61].

2.4 Deep Learning

Deep Learning (DL) is a subset of ML tasked with developing dynamic algorithms that have tunable parameters which mimic the concentration gradients within human-based neurological pathways. Therefore, certain DL models are called *neural networks* with each neuron representing a neural network’s most fundamental unit called a *perceptron*.

Each perceptron functions as a dynamic logic gate that decides whether to let information through by weights, biases, and an activation function. As illustrated in Figure 14, the amount of information that is passed into the output $h_{\mathbf{w}}(\mathbf{x})$ is dependent on leveraging matrix multiplication between the inputs, weights, and step function. Within a network consisting of multiple perceptrons, the transference of information from one *layer* of parallel perceptrons to the next *layer* is called *forward propagation* [61]. The flow of information is then reversed during *back propagation* and combined with an optimization algorithm to compute the adjustments needed to optimize the network, updating a *loss function* to signify the difference between the predictive performance and targeted performance towards task T [61].

For the nuclear fuel cycle environment, several multi-layer perceptron DL models have been applied to classification tasks [66, 67]. *Bae et. al.* directly tested and applied a fully-connected artificial neural network (ANN) or a multi-layer percep-




Figure 14. An artificial neural network is composed of *perceptrons* in which the amount of information passed is determined through a series of weights and biases. Reproduced from *ShadeDesign/Shutterstock.com* through a Standard Image License.

tron network to determine nuclear fuel composition from nuclear reactor parameters like burnup and initial enrichment. *Nabeshima et. al.* applied a fully-connected ANN called an *auto-associative network* to monitor nuclear power plant operational anomalies. Although both *Bae et. al.* and *Nabeshima et. al.* achieved high performance metrics for their respective tasks, each project lacks the ability to be generalizable within different environments, and its use limited to the given parameters and scenarios. This can be partially attributed to the difference between the NFC's large probability space and the perceptron's limited range of functionality. Therefore, within classification tasks that require specific features to be learned in depth to achieve a

more generalizable model, a different DL unit may be needed.

2.4.1 Convolutional Neural Networks (CNN)

Convolutional Neural Networks, called LeNet-5, were first introduced by Yann Lecun in 1998 with the goal of character recognition [68]. To train the DL model to learn the in-depth features present within an image, *LeCun et. al.* used 2D Convolutional Neural Networks (2DCNN) as network units responsible for transforming the input data in parts to highlight any fundamental features present within the image. Therefore, as shown in Figure 15, a *single filter* or *kernel* convolves over the pixels of the *original input* to reproduce a response map that highlights certain features while suppressing others according to the kernel's initialization.

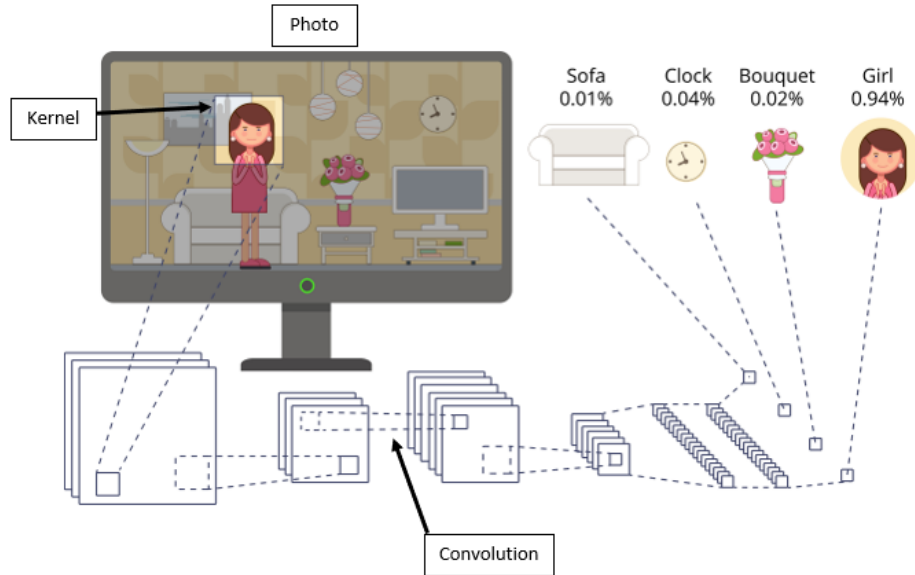


Figure 15. As the kernel slides across the picture, it convolves and maps the results to feature maps in a process that eventually leads to a classification task. Reproduced from *Arkadiy Chumakov/Shutterstock.com* through a Standard Image License.

As 2DCNNs are typically used for image processing [68], 1D Convolutional Neural Networks are used for signal processing where the kernel is instead convolving over the input data in the direction of time. This results in the dot product between the

extracted patch within each window stride and the kernel is taken as a convolution to produce a transformed output. Chollet *et. al.* describes the properties needed to perform a 1D convolution: the kernel length size, stride, kernel initialization, and activation function [60].

For a 1DCNN implementation, the number of kernels, or *filters*, is first established. Conventionally, a temporally-based input dataset is organized with timestamps as the dataset's second axis is the shape = (row = samples, column = time, depth = features) [60]. Using the Keras API, the filter's considered samples and features (rows, depth) matches the shape of the batched input dataset with only the filter's length, or considered time-period (columns), being adjustable as the 1D *kernel_size* parameter [69]. Once the number of filters and time period over which to perform the convolution are chosen, the kernel is initialized with values recommended by Glorot *et. al.* chosen from a uniform distribution over the range [-limit, limit] where *limit* is given by the following relation:

$$limit = \sqrt{\frac{6}{(fan_{in} + fan_{out})}}$$

where the fan_{in} and fan_{out} parameters are equivalent to the number of input and output units, respectively [70, 69]. The 1DCNN then performs a *non-causal* convolution signified by the default *valid* padding parameter and mathematically equivalent to *cross-correlation* shown in Equation 24, where the 1DCNN layer x of length n applies a kernel h of length k with stride s [71].

$$y = \begin{cases} \sum_{i=0}^k x(n+i)h(i), & \text{if } n = 0 \\ \sum_{i=0}^k x(n+i+(s-1))h(i), & \text{otherwise} \end{cases} \quad (24)$$

The non-causal designation signifies the assumption of no temporal dependence

between the output and input data. This assumption is applied as a default within the 1DCNN layer as traditional implementations using the mathematical non-causal correlation have shown no-change to greater performance improvement over the *causal* mathematical convolution. Physically, the 1DCNN is analogous to the filters applied within traditional signals analysis [71] e.g. a band-pass filter. Traditional implementations of 1DCNNs have been used in applications ranging from standard time-series based classifications [72] to sentence classification [73].

2.5 Multimodal Approaches

One of the primary concerns in implementing DL classification algorithms is that incorporating domain fidelity, feature interdependence, and data fusion into models is still in infancy [74, 75]. Because of the complexity within multimodal and multi-sensory data, it becomes difficult to build a network that optimizes both the analysis of individual signal contributors and their dependencies without a significant amount of information loss. This challenge limits the overall generalizability of a network and biases outputs to address only certain input features.

Historical multimodal approaches to classification tasks have ranged from combining data from similar sensors (e.g. imagery analysis) to data from completely different phenomenologies (e.g. magnetic and acceleration). *Hostettler et. al.* determined that the fusion of signals from a magnetometer and accelerometer with different sampling rates for vehicle tracking activity is feasible with stochastic particle filtering [76]. However, the paper also recommends that both phenomenologies are used for vehicle tracking as the fusion of both phenomenologies allows for a lower error (Root-Mean-Squared-Error 'RMSE') throughout the trials. *Piramanayagam et. al.* applied data concatenation prior to processing multi-sensor images with a convolutional neural network and found that data fusion earlier in the network performed

better than or equal to data fusion later in the network [77]. *Anguita et. al.* used a support vector machine (SVM) ML model with multi-sensor accelerometer and gyroscopic data to classify six different types of human activity ranging from walking to laying down with >89% classification accuracy [78] from which *Brownlee* created a multi-headed 1DCNN and achieved similar accuracy [79]. These approaches either filled the unknown relations between signals through stochastic approaches [76], used comparable signals fused in more optimized manners [77], or analyzed the individual features per signal for an overall classification [79]. Furthermore, Koch designed a One-Shot Learning Network with parallel branches which is referenced to construct this project’s multi-phenomenology DL model [80]. The primary recommended model presented within this project (Section 4.3) is developed in accordance with the use of magnetic and acceleration signals recommended by [76], the use of a Convolutional Neural Network by [77], the human activity classification through a multi-headed 1DCNN by [78, 79], and the parallel DL configuration presented by [80].

III. Data

3.1 Description

3.1.1 Data Collection Locations

Data sets 1 and 2 are used for method development and the collection details were reported by Goldblum [81]. These data sets were collected using multi-phenomenological, packaged, 'Merlyn' sensors positioned at the Oak Ridge National Laboratory (ORNL) High Flux Isotope Reactor (HFIR) and the McClellan Nuclear Research Center (MNRC) Training, Research, and Isotope Production General Atomics (TRIGA) Reactor. As such, this document refers to the first data set as the *HFIR data set* and the second data set as the *MNRC data set*. A subset of the available data was used in this analysis. Specifically, the HFIR Cycle-484 [82] collection from 8 January 2019 to 1 February 2019 was selected to be representative of HFIR operations, and a continuous 24-hour period collected on 17 August 2020 was selected as representative of MNRC operations. There are twelve Merlyn nodes stationed around the HFIR and five Merlyn nodes around the MNRC; both sensor layouts are illustrated in Figure 16. Further discussion of the Merlyn sensor is provided in Section 3.1.2.

For both the HFIR and MNRC, detailed operational logs exist through the Multi-Informatics for Nuclear Operations Scenarios (MINOS) venture [84]. As these two reactors are very different facilities, the operational logs are used to determine operational classification labels for machine learning pursuits. For binary operation (On/Off) determination, the labels are *one-hot encoded*, or classified with respect to the overall power level of the NFC facility. The HFIR Off:0 label corresponds to a power $\leq 7\%$ and On:1 label corresponds to a power $> 7\%$; the MNRC Off:0 label corresponds to a power $\leq 2.5\%$ and On:1 label corresponds to a power $> 2.5\%$ [81]. The HFIR research reactor operates at a 85 MW steady-state (~ 5.95 MW On/Off




Figure 16. Locations of the (a) 12 multi-sensory platforms around the HFIR at Oak Ridge National Laboratory and (b) 5 multi-sensory platforms around the MNRC in Davis, CA [81, 83]

classification boundary) with a different surrounding environment than the MNRC TRIGA research reactor operating at a 2 MW steady-state (~ 0.05 MW On/Off classification boundary) with high-power pulsing capability. While it is anticipated that there may be some degree of similarity between the operational signatures within the data sets, notable differences exist due to the different environments, reactor parameters, and power-level classification boundaries. This research explores whether

these differences may be overcome to produce generalized ML and DL models. The existence of facility operational logs allows the testing accuracy of the applied models to be assessed, providing a quantitative performance metric for model transferability between facilities within similar steps of the nuclear fuel cycle.

Previous analysis has shown that the performance of ML models is dependent on sensor location when fitting to the HFIR operations [83]. Therefore, the HFIR Node 9 location, shown in Figure 16, was chosen for this research. Furthermore, as it is near the HFIR’s cooling towers (Figure 17), the dominating detected phenomenologies are expected to be generated from the same source and correlated to reactor operations. The MNRC Node 4 location was chosen as the most comparable point for comparison because this location historically resulted in higher ML performance metrics [85].

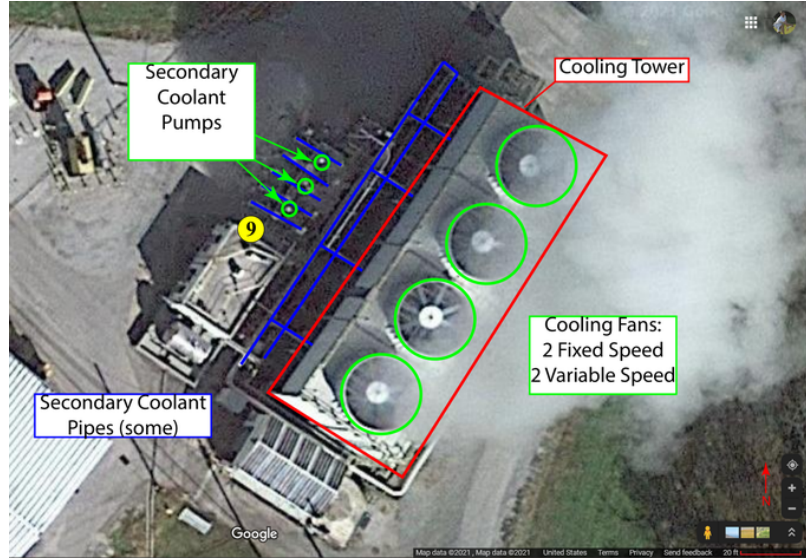


Figure 17. HFIR cooling tower located at Node-9 [85].

3.1.2 Merlyn Multimodal Sensor Platform

The Merlyn multi-phenomenological detector system (Figure 18) was designed by Special Technologies Laboratory and deployed by the University of Berkeley, California Complexity Group [83]. Each detector records GPS, magnetic field (3-axis),

acceleration (3-axis), pressure, temperature, light (red-green-blue), and proximity at a 16 Hz sampling rate with synchronized time stamps. This project explores the magnetic field and acceleration phenomenologies with units of microTesla (μT) and meters per second² ($\frac{m}{s^2}$), respectively.



Figure 18. The Merlyn multimodal detection platform measures multi-phenomenological signals simultaneously at a 16 Hz sampling rate [81].

3.2 Preliminary Work

3.2.1 Pre-Processing

Prior to feeding the data into the ML and DL models, the data were organized with respect to input shape and scaling. The data organization configurations for all the models are named either *Traditional* (Section 4.2.1) or *Novel* (Section 4.3.1) to indicate the difference in data organization when compared to typical DL implementations that are in accordance with DL package expectation [69]. Furthermore, as the models are implemented through the Scikit-learn (Sklearn) python package [86] for ML and the Keras application programming interface (API) [69] for DL, their processing, focus, and classification are dependent on how the data set is organized during input. The sub-parameters of applied models were adjusted accordingly to shift the focus of the network to highlight different parts of the input data. The scal-

ing of the data sets consisted of raw, standardized, and normalized transformations as no assumptions were made on the underlying data distributions. Subsequently, to reduce analysis time and conserve any hidden fundamental relationships between temporally-based data points, different sampling methods were tested. Within this project, the focus on magnetic and acceleration signals results in 6 total features for analysis from the 3-positional components per signal.

3.2.1.1 Sampling

Sampling was used in the ML models (Section 4.1) and multi-phenomenology DL model (Section 4.3). The single-phenomenology DL model (Section 4.2) used the entirety of the HFIR and MNRC data sets. Probability sampling methods are used in conjunction with non-parametric statistical tests for sample size determination. As a baseline for sampling, the *On* operational label is considered a "success" within considering the observations from each data set as Bernoulli trials. Therefore, a conservative sample size based on the binomial test's estimator for the probability of success results in equation 25; where n is the sample size with respect to the *On* operational label success probability, $(z_{\alpha/2})^2$ is the z-score confidence level, and D is the margin of error [87].

$$n = \frac{(z_{\alpha/2})^2}{4D^2} \quad (25)$$

This results in a sample size of ~ 384 for a 95% confidence level ($(z_{\alpha/2})^2 = 1.96$) and 5% margin of error. However, when sampling for a *balanced data set* or a data set where the classes are even, systematic-stratified sampling is used to conserve temporal relationships between sampled data and balance between the two sampled classes (On/Off). This is accomplished by first stratifying the data into two groups in correspondence with operational labels (On, Off) as shown in Figure 19. Treating these two groups as independent samples with corresponding success probabilities, the




Figure 19. The "Off" operational label is colored in red and blue, while the "On" operational label is colored in green. These binary operational labels are first separated into their corresponding operational labels with temporal causality preserved.

Off-group has success probability p_1 corresponding to an observation being labeled *off* and the On-group has success probability p_2 corresponding to an observation being labeled *on* with respect to the entire data set. Within a balanced data set, the assumption is made that these two probabilities should be equivalent $p_1 = p_2$ with equivalent corresponding group sizes $n_1 = n_2$. Therefore, the sample size within a balanced data set is determined from the Pearson's test for two success probabilities between Bernoulli trials shown in equation 26 where the sample size m is equivalent to $n_1 = n_2$ and the other terms have the same definitions as in equation 25 [87].

$$m = \frac{(z_{\alpha/2})^2}{2D^2} \quad (26)$$

This results in a sample size of ~ 768 for a 95% confidence ($(z_{\alpha/2})^2 = 1.96$) and 5% margin of error. Systematic sampling, shown in Figure 20, is performed at an interval determined by the sample size from the aforementioned Pearson’s test for two success probabilities between Bernoulli trials with a 95% confidence interval and 5% margin of error. The resulting sampled data set is then used within a model

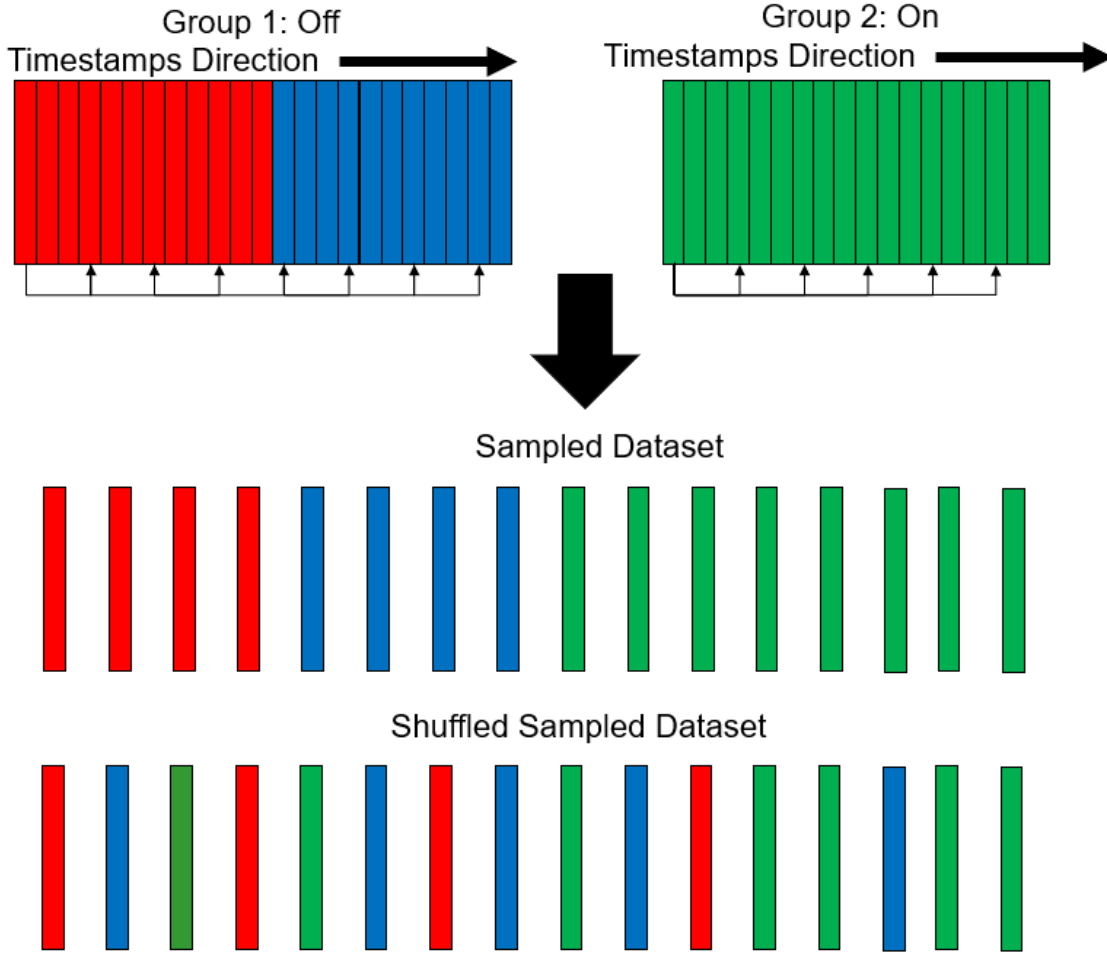


Figure 20. For a balanced data set, a sample size of 768 is sufficient. However, to ensure a non data-starved model, 768 observations were taken sequentially at an interval from the Group 1: Off and Group 2: On operational labels indicated by the combined red and blue observations separated from the green observations, respectively.

or shuffled to prevent learning biases within non-causal models. Within trials that exclude datum from transient states, this systematic-stratified sampling procedure is narrowed from the entire data set to the middle 31.73% of data within each binary

(On/Off) operational state. The transient time-frames included positive, steady, then negative rates of change in reactor power level for the 'On' steady-state and vice-versa for the 'Off' steady-state (Section 3.2.2). Due to the abundance of data and to limit the potential for including datapoints within transient time periods, an adjustable subset of the middle 31.73% of data was chosen.

3.2.1.2 Scaling

Scaling transformations are applied with the Scikit-learn (Sklearn) python package [86] for comparisons between signals on equivalent scales and input processing with no distribution-based assumptions. For all models, the data was processed with raw, standardized, and normalized scaling. Standardized scaling is implemented through Sklearn's *StandardScaler()* function and normalized scaling is implemented through Sklearn's *MinMax()* function [86]. Although raw, standardized, and normalized scaling was tested, this writing describes the scaling that resulted in the best performance within each model.

3.2.2 Data Distribution

To determine the possibility of utilizing transfer learning for a given step of the nuclear fuel cycle, the HFIR and MNRC data sets were chosen as they differ temporally in distribution with respect to binary operations. Within these data sets, *transient* periods are defined as time frames where the operational label changes. As shown in Figure 21a, the HFIR data set was collected over a span of 24 days with 94% of the collection classified in the *On* operational state with only one transient period from the *On* state to the *Off* state. In contrast, Figure 21b illustrates that the MNRC data set was collected over a smaller time period of 24 hours with 25% of the collection classified in the *On* operational state with two transient periods: from the

'Off' state to the 'On' state followed by a period of 'On' steady-state, and then from the 'On' state to the 'Off' state. The data sets were evaluated with and without the inclusion of transient states within the ML models (Sections 4.1.2 & 4.1.3), single-phenomenology DL model (Sections 4.2.2 & 4.2.3), and multi-phenomenology DL model (Sections 4.3.2 & 4.3.3) to garner an accurate performance assessment of the steady-state classification of binary operations. For the accuracy metric, a "coin flip's" random chance (50%) is the fundamental baseline of performance comparison given the binary operations considered.

The contribution of each phenomenological component to the data sets' overall distribution was measured through silhouette scores (described in Section 2.3.3). As this project's primary recommended model utilized raw data to improve model performance (Section 4.3), the silhouette scores shown in Table 2 are calculated on the raw, un-scaled, HFIR and MNRC data sets. Each phenomenology (magnetic, acceleration) has components that correspond to their 3-axis (XYZ) directional feature. Therefore, Table 2 references these components with their abbreviated phenomenologies (Mag \equiv magnetic, Acc \equiv acceleration) and corresponding directional feature (X, Y, or Z).

Table 2. Sampled silhouette scores with raw data.

Feature	HFIR	MNRC
Mag X	0.19	0.40
Mag Y	0.19	0.34
Mag Z	0.22	0.03
Mag 3-Axis	0.20	0.34
Acc X	0.1	0.40
Acc Y	0.03	0.34
Acc Z	0.01	0.03
Acc 3-Axis	0.02	0.22
All Features	0.20	0.34

Because the majority of scores are near zero, these silhouette scores represent a sig-




Figure 21. Binary operational (On/Off) label distributions of the (a) HFIR data set collected over 24 days and (b) MNRC data set collected over 24 hours where the label 1 corresponds to the *On* state and the label 0 corresponds to the *Off* state.

nificant overlap between On/Off binary operational states within each reactor's data set, separately. However, these scores also illustrate that the training data from the HFIR data set has less clustering separation between operations than the testing data from the MNRC data set, as demonstrated by HFIR scores being closer to zero when compared to MNRC scores. This implies that a model should be able to distinguish clustering separability more effectively on the MNRC data set. Furthermore, the

magnetic phenomenology was able to distinguish more separation between states relative to the acceleration phenomenology in both data sets. For the HFIR data set, the consideration of both phenomenologies lowers the maximum silhouette score from 0.22 (Mag Z) to 0.20. Similarly, within the MNRC data set, the consideration of both phenomenologies lowers the maximum silhouette score from 0.40 (Mag X or Acc X) to 0.34. Therefore, due to the silhouette score results and Brinker *et. al.*'s analysis, the magnetic phenomenology is selected for models that use only one phenomenology.

Additionally, Principal Components Analysis (PCA) was performed on the phenomenological components of the test data set from the MNRC to determine which phenomenologies and components dominate. Figure 22 shows the percentage of variance explained within the MNRC test data set through a choice of 3-principal sub-components in a bi-phenomenology data set of six overall components (3-axis magnetic and acceleration). Three principal sub-components were chosen on the assumption that one phenomenology and its corresponding 3-axis directional sub-components would dominate the measurements dependent on scaling transformations used within the model and phenomenological amplitude dependence described in Sections 2.1.1 & 2.1.3. Furthermore, as will be explained in future work (Section 5.2), the choice of 3-principal sub-components allows for models to highlight the assumed inter-dependencies between the phenomenologies (Section 1.4) during different scaling transformations. Within raw data, Figure 22 shows that 3-principal sub-components explain 100% of the variance within the data. The sub-component correlations are determined with respect to the PCA analysis' loading vectors shown in Table 3. The 3-principal components are most correlated with the magnetic phenomenology. Correspondingly, the first principal component (PC) is most explained by the magnetic phenomenology in the X-direction (Mag-X), the second PC is most explained by the magnetic phenomenology in the Z-direction (Mag-Z), and the third PC is most ex-

plained by a negative correlation in the Y-direction (Mag-Y). Due to the overall test data set's correlation with the magnetic phenomenology, in addition to the results founded by Brinker *et. al.* [13], this project's trials principally use the magnetic phenomenology as the primary signal for single-phenomenology ML and DL analysis.




Figure 22. PCA indicates that 100% of the MNRC data is explained by three features out of the six total features considered.

Table 3. In reference to Figure 22, PCA shows that the three dominant features correlate to the magnetic phenomenology within the raw MNRC data set.

Feature	PC-1	PC-2	PC-3
Mag-X	0.92	0.09	-0.37
Mag-Y	-0.37	0.41	-0.83
Mag-Z	0.08	0.91	0.41
Acc-X	-0.001	0.001	0.00
Acc-Y	0.00	0.00	-0.001
Acc-Z	0.003	-0.00	0.001

IV. Methodology & Results

4.1 Machine Learning (ML): Single-Phenomenology (SP) Trials

4.1.1 ML Methodology

To determine a baseline metric against which *transfer-learning* accuracy can be compared, ten ML algorithms were utilized by training on the HFIR data set and testing on the MNRC data set. The data were sampled in accordance with the systematic-stratified sampling described in Section 3.2.1.1 with raw, standardized, and normalized transformations defined in Section 3.2.1.2 and organized in the Novel Configuration as shown in Section 4.3.1 without the use of batches or mini-batches. Sklearn accepts this configuration with features along the columns and corresponding samples along the rows [65]. Furthermore, as described in Section 3.2.2, the silhouette scores were calculated using the binary operational labels per phenomenology and 3-axis vector components to statistically determine intra-label overlap. The following mix of supervised and unsupervised ML algorithms utilized Sklearn's *fit()* and *predict()* functions to train on the HFIR data set and test on the MNRC data set in determining a baseline accuracy for identifying facility operational status with respect to NFC transfer learning applications [86]:

- Random Forest/Random Forest Regressor
- K-Means
- Density-Based Spatial Clustering of Applications with Noise (DBSCAN)
- Logistic Regression (LR)
- Linear Discriminant Analysis (LDA)
- K-Nearest Neighbors (KNN)

- Classification and Regression Trees (CART)
- Naive Bayes
- Support Vector Machine (SVM)
- Principal Components Analysis (PCA)

To approximate a baseline on how ML performs on each facility, eight primary machine learning algorithms were trained and validated separately between facility data sets as shown in Figure 23 [86]. K-Fold Cross-Validation was used with five splits, random shuffling, and a seed of 87 to determine validation accuracy. All of the ML models were implemented using Sklearn’s default parameters [86] except for the SVM kernel which was set to *linear* and the K-Means number of clusters which was set to *two*. All of the algorithms performed significantly above coin-flip accuracy (>50%) across both data sets. Furthermore, in accordance with Section 3.2.2, all the models performed significantly higher (>80%) on the MNRC data set than that of the HFIR data set. Similar results were found with standardized and normalized scaling within both data sets. This shows that use of ML without hyper parameter tuning yields well-fit models and high-performance metrics given a single facility of interest.

Although the ML algorithms have high performance metrics using magnetic phenomenology data limited to individual facilities, this does not address the transferability of these models to separate, like-type facilities (*transfer learning*) through comparable performance metrics. To the author’s knowledge transfer learning has not been accomplished using nuclear reactor operations. Therefore, the same ML algorithms are applied within the next section to address the overall focuses of *transfer learning* and *multi-phenomenology*.




Figure 23. Training and validation performance of eight ML algorithms using the 3-axis magnetic phenomenology on the (a) HFIR dataset and (b) MNRC dataset yielded high-accuracies significantly above a coin-flip ($>50\%$) across both data sets.

4.1.2 ML-SP Results: Steady & Transient State

All ten ML algorithms listed in Section 4.1.1 were examined for transfer learning ability through training on the HFIR data set and testing on the MNRC data set with supervised learning algorithms (Random Forest/Random Regressor, LR, LDA, KNN, CART, Naive Bayes, SVM) or unsupervised learning algorithms (K-Means, DBSCAN). The trials used only the magnetic phenomenology and included transient states present in the sampled datasets. However, none of the algorithms performed significantly above a 'coin-flip' accuracy ($> 50\%$) except for K-Means as shown in Figure 24.




Figure 24. Labels resulting from K-Means using the standardized single-phenomenology (magnetic) MNRC testing dataset. Green represents the correct label classification; red represents incorrect label classification.

Figure 24 illustrates the accuracy of binary operational states determined by the ML K-Means algorithm based on the 3-axis magnetic phenomenology when trained on standardized, magnetic phenomenology samples from the HFIR data set and tested on standardized, magnetic phenomenology samples from the MNRC data set. Using the Sklearn library [86], the K-Means algorithm is fit through the *fit()* function and the experience (Section 2.3) gained from learning the training data set's intra-

clustering distances is used through the *predict()* function to determine the label of each MNRC observation. The binary classification performance metrics are as follows:

- Accuracy: 0.678
- F1-Score: 0.583
- Precision: 0.83
- Recall: 0.45

Despite a significant overlap within the *On* and *Off* operational states (silhouette score of 0.29), the ML algorithm was able to classify 67.8% of the sampled MNRC dataset correctly. Based on the precision metric, the model is able to effectively classify 83% of the identified *On* states. However, according to the recall metric, the model only identifies 45% of the total *On* observations within the data set. The model's misclassification of the binary states is represented in the confusion matrix illustrated in Figure 25.




Figure 25. Using the standardized, sampled magnetic-phenomenology from the MNRC data set, the K-Means algorithm mis-classifies observations with truth labels in the *On* operational state more often than observations in the *Off* operational state.

4.1.3 ML-SP Results: Steady State

In order to assess the potential origin of the misclassification and adjust the classification threshold between *On* and *Off* states, the ML algorithms were applied to datasets where the transient states were excluded in order to enhance the separability between operational states. However, despite this change, none of the algorithms performed significantly above a “coin-flip” accuracy of 50% except for K-Means and DBSCAN. For trials with standardized samples, both K-Means and DBSCAN performed with clustering metrics significantly above 50%, and DBSCAN also performed well within trials with normalized samples.

Figure 26 illustrates the accuracy of binary operational states determined by the ML K-Means algorithm based on the 3-axis magnetic phenomenology through training on standardized, steady-state, magnetic phenomenology samples from the HFIR dataset and testing on standardized, steady-state, magnetic phenomenology samples from the MNRC dataset. The binary classification metrics are as follows:

- Accuracy: 0.805
- F1-Score: 0.762
- Precision: 0.98
- Recall: 0.62




Figure 26. Labels resulting from K-Means using the standardized single-phenomenology (magnetic) MNRC testing dataset. Green represents the correct label classification; red represents incorrect label classification.

Due to the removal of transient states, a decrease in overlap of the operational states was seen (silhouette score of 0.56). Consequently, the ML algorithm was able to classify 80.5% of the sampled MNRC dataset correctly. Although the overlapping labels within Figure 26 show misclassifications concentrated within the *On* state, based on the precision and recall metrics, the model performs better at identifying the majority of *On* states, accurately classifying 98% of the identified *On* states. This indicates a significant improvement in the performance metrics when compared to the metrics in Section 4.1.2. This model’s confusion matrix is represented in Figure 27.




Figure 27. The Confusion Matrix for K-Means Predicted Labels w/standardized sampled steady-state, magnetic-phenomenology MNRC dataset is a significant improvement above the metrics resulting from the inclusion of transient states.

Overall, the model performs better in all metrics when considering only the steady-state operations. As a result, a significant portion of K-Mean’s initial mis-classification can be attributed to transient state mis-classification.

When excluding transient states, DBSCAN also performed substantially above a coin-flip accuracy on both normalized and standardized data scaling with the magnetic phenomenology. Contrary to K-Means, however, DBSCAN is completely unsupervised and therefore did not train on the HFIR dataset. Through the Sklearn *fit_predict()* function [86], DBSCAN was implemented only on the sampled MNRC data set. Furthermore, the performance metrics from DBSCAN must be interpreted differently due to the DBSCAN’s process in determining labels. DBSCAN worked in accordance with the hyper-parameters determined from the resulting K-Nearest Neighbors (KNN) distances from two neighbors as shown in Figure 28.




Figure 28. The NearestNeighbors algorithm with normalized sampled steady-state, magnetic-phenomenology from the MNRC data set determines that $\epsilon = 0.027$ is the optimized hyperparameter for DBSCAN at $n=2$ resulting from the maximum curvature within the graph.

This algorithm allows the distance-based epsilon hyper-parameter to be determined from the place of maximum curvature on the distance graph. After determining $\epsilon = 0.027$ and $min_{points} = 2$, the intra-clustering observation density was then used to determine operational state classification as shown in Figure 29.




Figure 29. DBSCAN implemented with standardized sampled steady-state, 3-axis magnetic-phenomenology from the MNRC dataset determined a number of samples from which the two most dense samples are colored in red and blue. Smaller groupings are colored in green, orange, brown, and yellow. This highlights the separation detailed by the cluster’s silhouette coefficient of 0.56, resulting in an clustering metric of 89.9%.

As DBSCAN is not forced to cluster data into two groups, performance metrics such as precision and recall are challenging to obtain for binary classification. The predicted labels per observation are determined through a direct comparison between true operational labels and DBSCAN labels from which only a pseudo-accuracy metric is reported. A significant separation between the *On* and *Off* states was found as shown by the calculated silhouette coefficient of 0.56. DBSCAN was able to cluster 89.9% of the samples from the standardized MNRC dataset. The 89.9% pseudo-accuracy metric is not a true measure of the classification correctness as the unsupervised nature of DBSCAN permits the reverse interpretation of its predicted labels - e.g. 0:On, 1:Off - thereby requiring a heuristic to label each cluster. Rather, this

metric represents how well DBSCAN clusters states of interest from the MNRC. Similar results are obtain with normalized scaling. A greater separability in operational states, 0.68 Silhouette Score, was observed with a clustering metric of 86.1% shown in Figure 30. Therefore, as both DBSCAN and K-Means are distance-based algorithms, the concept of **distance metrics** is considered within the DL models described in the next sections.




Figure 30. DBSCAN implemented with normalized sampled steady-state, 3-axis magnetic-phenomenology from the MNRC dataset determined a number of samples from which the two most dense samples are colored in red and blue. Smaller groupings are colored in green, orange, brown, and yellow. This highlights the separation detailed by the cluster's silhouette coefficient of 0.68, resulting in an clustering metric of 86.1%.

4.2 Deep Learning (DL): Single-Phenomenology (SP) Trials

4.2.1 DL Methodology: Traditional Configuration

In order to determine if *transfer learning* can be achieved through a direct signature comparison between reactors, a single-phenomenology bi-headed 1D-Convolutional Neural Network (1DCNN) was constructed with the DL TensorFlow package and Keras API [88, 69]. This network was trained and validated on the HFIR data set prior to testing on the MNRC data set. The data were not sampled, but were scaled with raw, standardized, and normalized transformations as described in Section 3.2.1.2.

The input organization for the *Traditional* configuration consisted of organizing the data per phenomenology with 3-axis vectors within rows, components separated in depth, and corresponding observations in columns as shown in Figure 31. As Keras' 1DCNN accepts a 3D configuration, this input shape in (row, column, depth) format is (batch size, timesteps, features) [69]. Historical implementations have used this configuration in 1D Convolutional Neural Network (1DCNN) implementations [78]. When processed by a 1DCNN, Figure 31 illustrates the goal of the input shape is to focus on all features or, in this case, 3-axis vectors (XYZ) simultaneously while convolving over a given time period. As described in Section 2.4.1, the kernel matrices within a 1D-Convolutional Neural Network are initialized randomly. Therefore, as the vector components are stacked depth-wise per timestep, each kernel simultaneously considers each observation as a tensor of 3-axis vectors from which the initializations highlight the magnitude of the X, Y, or Z vectors prior to sliding the window across to the next observation tensor in a loop until the end of the mini-batch, completing a convolution. Each time a mini-batch is processed through the 1DCNN, the weights and biases of the network are updated. The *Traditional* configuration is only used for the DL single-phenomenology model within this project.

As shown in Figure 32, the single-phenomenology DL model has two input *heads*

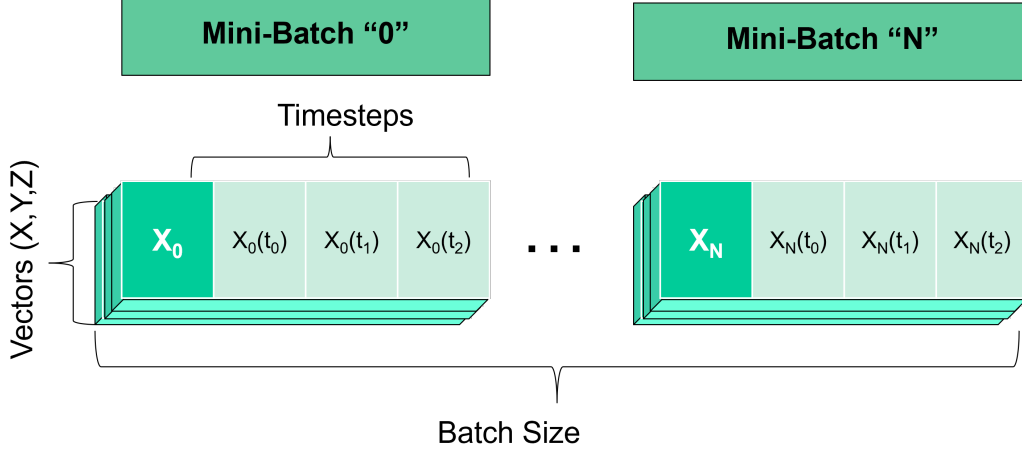


Figure 31. The Traditional Input Configuration stacks vectors depth-wise with each timestep separated in columns, progressing forward temporally with the row size. The entire data set is separated into Mini-Batches to be passed into the network.

or branches that share the same weights and biases. The network’s construction was designed with reference to Koch and Brownlee [80, 79], but the network parameters were changed to accompany this project’s transfer learning goal. Furthermore, the *Traditional* naming convention signifies the input data shape configuration (Section 4.2.1) and how the Same/Not-Same parameter performs as a matching function between branch labels [89]. This Same/Not-Same parameter is in agreement to the functionality of the equivalence gate or Exclusive-Not-Or (XNOR) logic gate shown in Table 4.

Table 4. The Same/Not-Same Equivalence Gate

Input 1	Input 2	Output (1 XNOR 2)
1	1	1
1	0	0
0	1	0
0	0	1




Figure 32. This network is comprised of two branches or *heads* within which the information of the input data is transformed for training the shared weights between each branch or for predicting labels during testing applications.

Per input branch, the data is passed through the first 1DCNN which performs convolutions over the 1D (temporal) space through a kernel size of 4, kernel depth of 3 (3-axis XYZ vectors), stride of 1, 64 filters, and *relu* activation function. The data is then passed through the second 1DCNN with the same parameters as the first. The *Dropout* layer then deactivated 50% of the nodes as a *regularizing parameter* to prevent the model from overfitting. With a pooling size of 2 and stride of 1, the *Max Pooling* layer then takes the maximum value of remaining tensors two at a time. The data is flattened to be sent to 500 fully-connected nodes within the dense layer, each with a *relu* activation function. The Manhattan distance is calculated within the lambda layer between the tensors resulting from the 500 nodes of each branch. This resulting tensor is mapped to a prediction node with a sigmoid activation function. During training, then the corresponding truth labels corresponding to the input data are passed into each branch, resulting in the network's adjusted weights and biases learning characteristic signatures. However, during testing, truth labels are

only passed into the top branch and the model predicts the bottom branch’s binary operational labels from the last prediction node.

Within this project, using only the HFIR Node 9 location and MNRC Node 4 location, the traditional configuration is trained using the 3-axis (XYZ) magnetic phenomenology from the HFIR prior to testing the performance of the network in predicting the binary operational labels of the MNRC data set. Therefore, the network’s training is conducted on the HFIR’s magnetic phenomenology passed into both input branches. The network’s testing is conducted with the HFIR’s magnetic phenomenology passed into the top branch and the MNRC’s magnetic phenomenology passed into the bottom branch. The following sections describe the transfer learning training and testing of the single-phenomenology DL model using raw HFIR and MNRC data. Additional trials were performed with additional input data transformations (standardized, normalized) and resulted in lower performance metrics.

4.2.1.1 Single-Phenomenology Training

To build a model that can recognize the magnetic field signal characteristics of binary reactor operations, the single-phenomenology DL model is trained on the HFIR data set collected at the Node-09 location shown in Figure 16a with only the magnetic phenomenology as shown in Figure 33. All three-axis vectors (XYZ) of the magnetic phenomenology are passed into each branch in the traditional configuration (Section 4.2.1) with each 3D input shape = (batch size = 2,330,957, timesteps = 16, channels = 3) along with the corresponding binary operational state labels. The HFIR data set is sampled at 16 Hz. Therefore, the input shape allows each batch to represent 1 second of data. Due to the 1DCNN non-causality 24, the second branch of the network is shuffled prior to entering the network, providing a balanced representation in both the Same/Not-Same metric and binary operational states. Furthermore, this




Figure 33. HFIR Signals are passed into both branches of the network during training along with the corresponding truth labels. This network then matches and labels the observations between branches using the Same/Not-Same metric resulting in a pre-trained network for testing applications.

randomized shuffling aims to prevent any training biases, strengthening the network’s ability to recognize similar signals with different reactor operational states such as during *transient operations*. The network is trained and validated on all 37,295,312 observations of the collection during HFIR Cycle-484. The validation set uses the latter 25% of the input data through the *validation_split* function from Keras [69] with accuracies approaching 100%. This network was trained with a batch size of 500 for 10 epochs using an ADAM optimizer with a learning rate of 0.001.

4.2.1.2 Single-Phenomenology Testing

As the network’s testing relies on matching test-data signal characteristics to known pre-trained characteristics representative of reactor operations, the testing phase inputs the HFIR data set’s magnetic signal within the top branch to match with the MNRC data set’s magnetic signal within the bottom branch as shown in Figure 34.

Ranging from 0 to 1, the matching metric is determined through a probability per observation resulting from the prediction node's sigmoid function when generated through the Keras *predict()* function [69]. The observations within the probability distribution are attributed to Same/Not-Same labels through a classification boundary at 0.5. When an MNRC observation is found to match (Same:1) the known signal's pre-trained characteristic signature within the simultaneously passed HFIR observation, the MNRC observation is then labeled with the HFIR signal's operational label. Conversely, if the MNRC observation is found to not match (Not-Same:0) the known signal's characteristic signature, then the MNRC observation is then given the opposite operational label with respect to the HFIR signal's operational label. The accuracy, F1-Score, precision, and recall metrics are calculated by comparing the predicted labels to the truth labels of each MNRC observation.




Figure 34. Same/Not-Same testing framework for the traditional bi-headed 1DCNN using the MNRC data set compared to the HFIR data set using single-phenomenology magnetic 3-axis data.

4.2.2 DL-SP Results: Steady & Transient State

Modeled after historical input data organization and parallel network Same/Not-Same metric parameters [78, 79], the bi-headed 1DCNN single-phenomenology DL model establishes a baseline for testing the concept of *transfer learning* through DL algorithms. The network uses magnetic-phenomenology signals batched in $\frac{1}{16}$ second increments. This results in a (row, column, depth) training dataset input shape of (233,0957, 16, 3) and testing dataset input shape of (86,244, 16, 3). Compiling the model using the Keras API [69] results in the model summary shown in Figure 35. This summary matches the network described in Section 4.2.1 and Figure 32 where each layer is contained within the *sequential* layer of the Keras summary.

Model: "model"

Layer (type)	Output Shape	Param #	Connected to
input_1 (InputLayer)	[(None, 16, 3)]	0	
input_2 (InputLayer)	[(None, 16, 3)]	0	
sequential (Sequential)	(None, 500)	177780	input_1[0][0] input_2[0][0]
lambda (Lambda)	(None, 500)	0	sequential[0][0] sequential[1][0]
dense_1 (Dense)	(None, 1)	501	lambda[0][0]
Total params: 178,281			
Trainable params: 178,281			
Non-trainable params: 0			

Figure 35. The single-phenomenology DL model was built through the Keras API [69] with two branches indicated by *input_1* and *input_2*. This information is passed through the model's *sequential* layers prior to the conjoining lambda layer and *dense_1* prediction node.

Probing the output Manhattan distances determined by the network's lambda layer during testing shows a roughly bi-modal distribution between distances within Figure 36. As these distances are projected onto a sigmoid function to determine predicted label states, the separation between clusters suggests the possibility of clustering operational states. However, as the peaks of the separated clusters are multi-

modal, this suggests a level of confusion may be present if classification is conducted based on the Manhattan distances.





Figure 36. The true-density histogram of resultant Manhattan Distances shows a separation between two slightly bi-modal clusters which indicates a possibility for classification.

The Same/Not-Same metric distribution resulting from the Manhattan distance's projection onto the sigmoid function is shown in Figure 37a. As described in Section 4.2.1, this metric represents the probability of the MNRC test datum's operational label matching (Same:1) the HFIR datum's operational label of the other branch. The probability distribution shows a concavity with the center around 12.0 hrs. As this metric compares the HFIR magnetic observations to the MNRC magnetic observation, with the assumption of a classifier boundary at 0.5 applied manually and resulting from the K-Means ML algorithm for comparison (Figure 37a). When compared to Figure 37b, the two largest clusters (0.0-16.0 hrs and 16.0-24.0 hrs) within the probability distribution appear to correspond to the binary operational truth labels of the MNRC. However, the classification boundary implemented at a probability of 0.5 results in a significant mis-classification in binary operational states biased towards the mis-classification of observations with 'On' truth labels and

(a) Same/Not-Same Probability Distribution



(b) MNRC Truth Labels




Figure 37. Predicted Same/Not-Same probability distribution of the (a) MNRC data set with a classification boundary applied at a probability of 0.5 and (b) MNRC truth labels where the label 1 corresponds to the *Same* metric and the label 0 corresponds to the *Not-Same* metric.

the following performance metrics:

- Accuracy: 0.697

- F1-Score: 0.534
- Precision: 0.431
- Recall: 0.70

The single-phenomenology DL model was able to classify 69.7% of the MNRC dataset correctly. Based on the precision metric, the model is able to effectively classify 43.1% of the identified *On* states. However, according to the recall metric, the model identifies 70% of the total *On* observations within the data set. The model's confusion between the binary states is represented in the confusion matrix illustrated in Figure 38.




Figure 38. The model often misclassifies the observations with *Off* truth labels of the MNRC data set more than observations with *On* truth labels.

This mis-classification may result from the combination of a lack of data cleaning or filtering and the inclusion of transient states, as was the case in Section 4.1.3.

4.2.3 DL-SP Results: Steady State

In order to prevent further confusion and adjust the classification threshold between *On* and *Off* states, the single-phenomenology DL model was applied with the same parameters as mentioned in Section 4.2.2; however, the transient states were excluded from the HFIR and MNRC data sets without sampling through an adjustable middle 31.76% data subset in accordance with Section 3.2.1.1 to enhance the separability between operational states.

Referring to the HFIR comparison’s operational status data illustrated in Section 3.2.2, the operational labels are determined through a 0.7 classification boundary as large observation ranges (e.g. inclusion of outliers) skew the probability distribution, resulting in low separation between clusters. Therefore, a classification boundary applied at 0.7 (Figure 39a) resulted in the best performance metrics when compared to the results in applying the boundary elsewhere. These metrics show a significant mis-classification in binary operational states biased towards the classification of observations with *Off* truth labels (Figure 39b). This results in the following performance metrics:

- Accuracy: 0.742
- F1-Score: 0.566
- Precision: 0.402
- Recall: 0.960

Without the inclusion of transient states, the single-phenomenology DL model was able to classify 74.2% of the MNRC dataset correctly. Based on the precision metric, the model is able to effectively classify 40.2% of the identified *On* states. However, according to the recall metric, the model identifies 96% of the total *On* observations




Figure 39. Excluding transient states, the predicted Same/Not-Same probability distribution of the (a) MNRC data set with a classification boundary applied at a probability of 0.7 and (b) MNRC truth labels where the label 1 corresponds to the *Same* metric and the label 0 corresponds to the *Not-Same* metric.

within the data set. The model's confusion between the binary states is represented in the confusion matrix illustrated in Figure 40.

Overall, the model performs better in all metrics, except for a slight performance reduction in *precision*, when considering only the steady-state operations. As a result,




Figure 40. Confusion matrix of predicted MNRC binary operational labels where there is a significant mis-classification in binary operational states biased towards the classification of observations with *Off* truth labels.

a significant portion of the single-phenomenology DL model’s initial mis-classification can be attributed to the confusion presented within transient state classification. However, the vast majority of the mis-classification remains unattributed and may result from a lack of data cleaning or filtering. At this time, background measurements and filtered phenomenologies of binary reactor operational states are unavailable.

4.3 Deep Learning (DL): Multi-Phenomenology (MP) Trials

4.3.1 DL Methodology: Novel Configuration

Taking into account historical multimodal implementations described in Section 2.5, to determine if *transfer learning* can be achieved through the use of multiple phenomenologies between reactors, the multi-phenomenology DL model was constructed with the DL TensorFlow package and Keras API [88, 69]. This network was trained and validated on the HFIR data set prior to testing on the MNRC data set using

the 3-axis magnetic and acceleration phenomenologies from both data sets. The data were sampled in accordance with Section 3.2.1.1 and were scaled with raw, standardized, and normalized transformations as described in Section 3.2.1.2. The best performance metrics arose using the raw transformation and therefore, will be the principal transformation presented.

The input organization for the *Novel* configuration consists of organizing the data per phenomenology with the 3-axis vector components within columns and the corresponding observations in rows as shown in Figure 41. For 3D configuration, this input shape in (row, column, depth) format is (batch size = timesteps, features, 1). Historical implementations primarily use this configuration to process data into ML models without the use of mini-batches [65] as in Section 4.1.




Figure 41. The Novel configuration results in the unravelling of 3-axis vector components when compared to the Traditional configuration. This allows each vector component per phenomenology to be analyzed individually prior to model calculation.

However, when processed into a 1DCNN, Figure 41 illustrates how the input shape is the projection of the XYZ 3D tensor onto a 2D input plane. This projection provides

the option for each vector component to be considered individually during convolution within a given sequence of time. Furthermore, as convolutional neural networks have historically been used for image processing applications [90], this configuration's 3D to 2D projection is analogous to how an image is the projection of a 3D figure with unit-vector dimensions of (\hat{x} = width, \hat{y} = height, \hat{z} = depth) onto 2D pixels with dimensions (width, height). These 2D pixels have traditionally been analyzed with discernment along the red, green, and blue (RGB) color sequence in groups of images, leading to the traditional configuration input shape of (batch size = # of images, row = width, column = height, depth = RGB). Within this configuration, the image's dimensions of width and height are analogously equated to a signal's physical 3-axis vector (XYZ) dimensions with unit-depth magnitude. The image's depth of RGB colors are equated to unit-depth magnitude as the signal does not have a comparable physical discernment. This equates to the Novel configuration's input shape of (batch size = timesteps, row & column = 3-axis vectors (XYZ), depth = 1).

As shown in Figure 42, similar to the single-phenomenology traditional configuration, the multi-phenomenology DL network has two input *heads* that share the same weights and biases. Except for the 1DCNN layers' kernel size parameter set to 1 and the *Max Pooling* parameter set to 2, all other model parameters remain the same. This has the effect of allowing the 1DCNN kernel to convolve over each positional vector individually as opposed to the simultaneous vector convolution shown within the single-phenomenology DL model (Section 4.2.1). However, this model differs both in its data input organization and the Same/Not-Same parameter determination.

Because the input data for this configuration is systematically sampled with a subset size in accordance with the non-parametric Pearson's test to preserve processing and analysis time, the input shape is organized to allow the network to perform convolutions over the 3-axis (spatial) vectors (X,Y,Z) within the input signal, pre-




Figure 42. The multi-phenomenology DL model’s network is the same as the single-phenomenology network except for the data input configuration and Same/Not-Same metric functionality.

serving information from the inter-dependency between signal components within the original signal as much as possible. The Same/Not-Same parameter is also changed from the traditional approach of performing as a matching algorithm to directly representing the binary reactor operational states per observation. Rather than the Same/Not-Same parameter being synonymous with *Match/No-Match*, this parameter is configured as Same = On (1) and Not-Same = Off (0). This metric change implies that observation pairs simultaneously input into each head must have the same collection timestamp. Therefore, contrary to the single-phenomenology network, the bi-headed multi-phenomenology network requires a minimum of **two** phenomenologies from the same detector platform at a time. Subsequently, the novel network’s training, validation, and testing is performed using the magnetic and acceleration phenomenologies.

4.3.1.1 Bi-Phenomenology Training

The multi-phenomenology DL model is trained on the magnetic and acceleration phenomenologies from the HFIR data set collected at the Node-09 location shown in Figure 16a. These phenomenologies have simultaneous timestamps and therefore the same operational state and truth labels. As shown in Figure 43, the 3-axes of both phenomenologies are input into the network in the top and bottom branches with the Same/Not-Same metric now configured to pass binary operational truth labels into the network. The goal of the training process is to learn the Manhattan distance between phenomenologies within latent space to determine operational labels in the event that the single-phenomenology traditional DL model mis-classified in transfer learning scenarios due to fundamental label thresholds and differences in reactor power levels, operations, and environments as discussed in Section 1.4.




Figure 43. Using only the HFIR data set, the network is trained using the 3-axis magnetic signal in the upper branch and 3-axis acceleration signal in the lower branch concurrent with the simultaneous binary operation truth labels.

4.3.1.2 Bi-Phenomenology Testing

With pre-trained weights, the testing phase of the novel network can be used to assess the transferability of identified signal characteristics for the magnetic and acceleration phenomenologies, regardless of facility, operational, or environmental differences. Therefore, as the network was trained on the operational labels corresponding to the Manhattan distances within the bi-phenomenology latent-space projection, testing on the MNRC data set consists of processing the magnetic and acceleration phenomenologies sampled only at the MNRC. As shown in Figure 44, the signals are fed into the network with the goal of projecting them onto the same latent space used during training. This projection acts as a comparable medium between the HFIR and MNRC to assess signatures that correspond to binary operations. Thus, the *Lambda* layer calculates the Manhattan distances within the latent space from which predictions on these distances are made to determine facility operations.




Figure 44. Using only the MNRC data set, the pre-trained network is tested using the 3-axis magnetic signal in the upper branch and 3-axis acceleration signal in the lower branch prior to predicting the operational labels of the MNRC.

In an attempt to decrease the amount of information loss and the analysis time while testing the feasibility of *transfer learning*, the Novel configuration differs from the Traditional configuration (Section 4.2.1) through the input dataset organization and the definition of the Same/Not-Same metric (Section 4.3.1). Therefore, as shown in Figure 45, compiling the model with the Keras API [69] results in changing the shapes that are used within the 'input_1' and 'input_2' layers from [(None, 16, 3)] to [(None, 3, 1)], matching the network presented in Figure 42. The network uses non-parametric sampled magnetic and acceleration phenomenology signals with balanced classes batched in 2-second increments (32 observations, 16 Hz sampling rate) for an initial HFIR training input shape of (1538, 3, 1) and MNRC testing input shape of (1539, 3, 1).

Model: "model"

Layer (type)	Output Shape	Param #	Connected to
input_1 (InputLayer)	[(None, 3, 1)]	0	
input_2 (InputLayer)	[(None, 3, 1)]	0	
sequential (Sequential)	(None, 500)	36788	input_1[0][0] input_2[0][0]
lambda (Lambda)	(None, 500)	0	sequential[0][0] sequential[1][0]
dense_1 (Dense)	(None, 1)	501	lambda[0][0]

=====

Total params: 37,289
Trainable params: 37,289
Non-trainable params: 0

Figure 45. The multi-phenomenology DL model was built through the Keras API [69] with two branches indicated by *input_1* and *input_2* with different shapes than the Traditional configuration. This information is passed through the model's *sequential* layers prior to the conjoining lambda layer and *dense_1* prediction node.

4.3.2 DL-MP Results: Steady & Transient State

Probing the output Manhattan distances determined by the network during testing shows a bi-modal distribution between distances in Figure 46. Compared to the

network’s traditional single-phenomenology counterpart shown in Figure 36, the Manhattan distances from the novel multi-phenomenology network have less separation between center peaks, but result in very little spread within each mode. This lack in spread suggests the possibility of greater clustering between operational states when compared to the single-phenomenology outputs.





Figure 46. The output Manhattan distances from the bi-phenomenology testing show a bi-modal distribution with very little spread in each mode.

The Same/Not-Same metric distribution resulting from the multi-phenomenology network’s projections of the Manhattan distance outputs onto the sigmoid function is shown in Figure 47a. This metric does not signify the XNOR equivalence gate as explained in Section 4.2.1. Instead, the metric represents the probability of the MNRC test datum’s label being classified in the ‘On’ state. Similar to the single-phenomenology distribution, the probability distribution resulting from the multi-phenomenology network shows a concavity within its center around 12 hrs with peaks arising early and late in the distribution around 0.0 hrs and 24.0 hrs.

A classification boundary implemented at a probability of 0.5 is shown in Fig-

(a) Predicted Binary Operational State Probability Distribution



(b) Predicted MNRC Labels




Figure 47. The predicted binary operational probability distribution of the (a) MNRC data set with a classification boundary applied at a probability of 0.5 and (b) the accuracy of predicted MNRC truth labels where green represents the correct predictions and red represents the incorrect predictions.

ure 47a. In accordance with the predicted label probability distribution's density, a classification boundary applied at a 0.5 probability results in the majority of data

between 0.0 and 16.0 hrs classified in the *Off* state and the majority of data between 16.0 and 24.0 hrs classified in the *On* state. As shown in Figure 47b, green colored labels represent an accurate prediction of the MNRC datum’s operational state and red colored labels represent an inaccurate prediction of the MNRC datum’s operational state. The majority of inaccuracies or mis-classifications are between 0.0-3.2 hrs, 16.0-18.0 hrs, and 22.4-24.0 hrs. These mis-classifications match the time period distributions identified as *transitory states* in Section 3.2.2.

As a result, the following performance metrics are computed:

- Balanced Accuracy: 0.841
- F1-Score: 0.851
- Precision: 0.802
- Recall: 0.905

The multi-phenomenology DL model was able to classify 84.1% of the MNRC dataset correctly. Based on the precision metric, the model is able to effectively classify 80.2% of the identified *On* states. According to the recall metric, the model identifies 90.5% of the total *On* observations within the data set. The model’s confusion between the binary states is represented in the confusion matrix illustrated in Figure 48.




Figure 48. The confusion matrix shows that the multi-phenomenology DL model misclassifies observations with truth labels in the *Off* operational state.

The confusion matrix shows that the model tends to mis-classify datum with truth labels in the *Off* operational state. With reduced mis-classification errors from the multi-phenomenology model when compared to the single-phenomenology model, the multi-phenomenology model has a higher predictive capability using the MNRC test data set's binary operations. The higher performance metrics of the multi-phenomenology model when compared to the single-phenomenology model suggests that the use of bi-phenomenologies allows the Manhattan distance metric used during classification to preserve inter-phenomenology information for dynamism and noise robustness. Furthermore, a majority of mis-classifications occur in transient operations, suggesting a path forward to improve the model performance.

4.3.3 DL-MP Results: Steady State

To test how much the inclusion of *transient states* as defined in Section 3.2.2 affect the multi-phenomenology model's performance, the sampling was adjusted from the entire dataset to the middle 31.73% of datapoints within each binary state as described in Section 3.2.1.1. The multi-phenomenology model for these trials is the same as summarized in Section 4.3.2 and Figure 45. The Manhattan distances determined by the network during testing shows a bi-modal distribution between distances, shown in Figure 49. Compared to the Manhattan distance distribution output by the multi-




Figure 49. The exclusion of transient states has resulted in a bi-model distribution of the output Manhattan distances.

phenomenology model with the inclusion of transient states, Section 4.2.2, there is a greater separation between center peaks without a significant increase in the spread of each peak. This increased separation indicates a possibility of greater fidelity to determine a classification boundary when compared to the boundary resulting from the inclusion of transient state data.




Figure 50. With the exclusion of transient states, the predicted binary operational probability distribution of the (a) MNRC data set with a classification boundary applied at a probability of 0.6 and (b) the accuracy of predicted MNRC truth labels where green represents the correct predictions and red represents the incorrect predictions.

The binary operational label distribution resulting from the network’s projections of the Manhattan distance outputs onto the sigmoid function is shown in Figure 50a.

Similar to the boundary implemented in Section 4.2.3, the operational labels are determined through a 0.6 classification boundary as large observation ranges (e.g. inclusion of outliers) skew the probability distribution, resulting in a translated separation between clusters. Therefore, a classification boundary applied at 0.6 resulted in the best performance metrics when compared to the results in applying the boundary elsewhere. The classification accuracy of this boundary is represented in Figure 50b. 'Green' colored labels represent an accurate classification when compared to the MNRC truth labels, and 'red' colored labels represent inaccuracy or misclassification when compared to the MNRC truth labels. Subsequently, the following performance metrics are computed:

- Accuracy: 0.994
- F1-Score: 0.993
- Precision: 0.996
- Recall: 0.990

With the exclusion of transient states, the multi-phenomenology DL model was able to classify 99.4% of the MNRC dataset correctly. According to the recall metric, the model identifies 99.0% of the total *On* observations within the data set and, according to the precision metric, effectively classify 99.6% of the identified observations. These metrics show that the vast majority of the test data is accurately classified. The resulting confusion matrix is shown in Figure 51.




Figure 51. The confusion matrix shows that the multi-phenomenology DL model rarely mis-classifies observations.

With respect to the entire sampled test dataset, there is almost a 50-50 split between the model’s performance in the determination of *Off* (49.81%, 0.19% error) and *On* states (49.55%, 0.45% error) with very little error. The accuracy represents a 99.4% multi-phenomenology model prediction capability of the MNRC dataset’s steady-states. Similarly, all associated metrics (F1, Precision, Recall) are have substantially high accuracies, with the majority of inaccuracies being attributed to external factors and limitations described in Section 1.4.

V. Conclusion

5.1 Research Conclusions

Transfer Learning Summary		
Model	With Transient	Without Transient
ML-SP	67.8%	80.5%
DL-SP	69.7%	74.2%
DL-MP	84.1%	99.4%

This research shows that within *transfer learning* scenarios between nuclear fuel cycle facilities the use of multi-phenomenology data with deep learning (DL) models results in higher predictive capabilities, robustness to noise, and a less than or equivalent size data set when compared to the use of single-phenomenology data in either machine learning (ML) models or DL models. The goal of this research was to test the effectiveness of utilizing multi-phenomenology (acceleration, magnetic) data within predicting the binary operations (On/Off) of an unknown test reactor after training the model on an operationally known reactor as described in Objective 1 (Section 1.3). As such, each model was trained and validated on a data set collected at the High Flux Isotope Reactor (HFIR) facility prior to testing the model’s predictive capabilities on a data set collected at the McClellan Nuclear Research Center (MNRC) facility. The binary operational states were *one-hot encoded* with 0:Off and 1:On. Within each data set, *transient states* were also identified to be excluded in latter testing. Three principle configurations are identified and reported in this text as machine learning utilizing the 3-axis magnetic single-phenomenology (ML-SP), deep learning utilizing the 3-axis magnetic single-phenomenology (DL-SP), and deep learning utilizing both the 3-axis magnetic and acceleration multi-phenomenologies (DL-MP). As shown in Table 5.1, the DL-SP trials performed better in predictive accuracy during the trans-

fer learning binary classification for MNRC operations in both the inclusion (84.1%) and exclusion (99.4%) of transient states.

Within the ML-SP implementation, systematic-stratified sampling was utilized for both the HFIR and MNRC data sets to re-balance the binary classes prior to being passed into 10 assorted supervised and unsupervised ML algorithms within python's Sklearn ML package [65]. Each data set was organized with the columns representing features and rows representing observations in accordance with Sklearn's expected inputs [65] (Section 4.3.1). Training and validating ML algorithms on each facility separately resulted in significantly ($>50\%$) high accuracies indicating the ability for ML to adequately fit and predict the embedded signatures that correspond to binary operations within each facility's data set. The implementation of transfer learning with the 3-axis magnetic phenomenology by fitting to the HFIR data set and predicting the MNRC data set resulted in high performances with the K-Means and DBSCAN algorithms. Standardization scaling for K-Means and both standardization and normalization scaling for DBSCAN resulted in the best model performance. As shown in Table 5.1, the highest ML-SP trials resulted in a 67.8% predictive accuracy with the inclusion of transient states and a 80.5% accuracy (+12.7%) with the exclusion of transient states. Therefore, as machine learning techniques are typically used for their ease of implementation, ability to work with smaller data sets, and fast processing times, the ML-SP results support the potential use of ML for transfer learning applications between similar facilities within the nuclear fuel cycle (NFC). However, as both K-Means and DBSCAN are clustering algorithms, a semi-supervised approach is recommended for direct cluster-to-label attribution. Furthermore, due to both the lack of fine-tuning within each algorithm and the performance difference between the inclusion and exclusion of transient states, the ML-SP implementation establishes a proof-of-principle baseline for the use of artificial intelligence (AI) models

within the NFC transfer learning scenario.

The DL-SP model was implemented without sampling through the Keras API [69]. Each data set was organized with the columns representing observations, rows representing temporal space, and depth representing features in accordance with Keras' expected inputs [69] (Section 4.2.1). The implementation of transfer learning with the 3-axis magnetic phenomenology by fitting to the HFIR data set and matching pre-trained HFIR signatures to the MNRC signatures for prediction resulted in substantial ($>50\%$) performances with un-scaled data. As shown in Table 5.1, the highest DL-SP trials resulted in a 69.7% predictive accuracy with the inclusion of transient states and a 74.2% accuracy (+4.5%) with the exclusion of transient states. These metrics support the potential use of DL for transfer learning applications between similar facilities within the nuclear fuel cycle (NFC). However, as deep learning techniques typically require larger data sets, fine tuning, and slower processing times, the ML-SP model is recommended when compared to the DL-SP model. Furthermore, due to both the lack of fine-tuning within the model parameters and the performance difference between the inclusion and exclusion of transient states, the DL-SP implementation establishes a potential for robustness to signatures that are not of interest and a proof-of-principle baseline for the use of DL models within the NFC transfer learning scenario.

In order to address the classification confusion within the single-phenomenology models, the DL-MP model combined the attributes of both the ML-SP and DL-SP models through utilizing the same sampling and same data organization as the ML-SP trial (Section 4.3.1), while also utilizing the same structure within the DL-SP model. However, rather than matching phenomenological signatures as described in the DL-SP model, the implementation of transfer learning with the un-scaled 3-axis magnetic and acceleration phenomenologies were passed separately into each parallel

branch of the network. Furthermore, each observation was labeled with respect to binary reactor operation rather than whether or not observations between branches match. This allows the network to fit the HFIR data set through training on the corresponding binary operation’s latent space projection of the magnetic signal with respect to the acceleration signal. During testing, the pre-trained HFIR latent space projection is used to predict the operations associated with the MNRC projection. As shown in Table 5.1, the highest DL-MP trials resulted in a 84.1% predictive accuracy with the inclusion of transient states and a 99.4% accuracy (+15.3%) with the exclusion of transient states. These metrics support the potential use of multi-phenomenological data coupled with DL for transfer learning applications between similar facilities within the nuclear fuel cycle (NFC). Furthermore, the combination of sampling technique and data organization from the ML-SP trial paired with the DL-SP model allows for faster processing times and the potential for increased robustness. Therefore, the DL-MP model is recommended when compared to either of the single-phenomenology models.

5.2 Future Work

The work presented in this writing accomplishes Objective 1 from Section 1.3 through recommending the use of deep learning coupled with multi-phenomenological data (DL-MP) within NFC transfer learning applications. However, as all models used within this research lacked fine-tuning, specific equipment attribution, and data cleaning, the recommendations presented within this writing are based on the establishment of proof-of-principle concepts and potential applications. Furthermore, as the specific equipment attribution to operational signatures was not established, Section 1.4, the distance-dependence of these signals as described in Objective 2 of Section 1.3 was not established as the source of these signals remain unattributed.

As a result, the future work of this project includes the following:

- The collection of additional data at further distances to better determine the distance-dependence of phenomenologies and accomplish the second objective of this research, Section 1.3.
- The testing of both ML and DL model parameters to enhance model capacity, goodness of fit, and robustness to noise with the inclusion of transient time periods.
- The analysis of data with respect to frequency space for the ability to attribute signatures to specific equipment.
- The implementation of additional phenomenologies to test the robustness and applications of both objectives within this research, Section 1.3.

Bibliography

1. “The many uses of nuclear technology,” May 2021. [Online]. Available: <https://world-nuclear.org/information-library/non-power-nuclear-applications/overview/the-many-uses-of-nuclear-technology.aspx>
2. A. Witze, “How a small nuclear war would transform the entire planet,” *Nature*, vol. 579, no. 7800, p. 485–487, 2020.
3. *Treaty on the Non-Proliferation of Nuclear Weapons*. United Nations, 1968.
4. V. Fedchenko, *The new nuclear forensics: analysis of nuclear materials for security purposes*. Oxford University Press, 2015.
5. *Nuclear Forensics: Beyond the Science*, ser. TECDOC Series. Vienna: International Atomic Energy Agency, 2020, no. 1896. [Online]. Available: <https://www.iaea.org/publications/13648/nuclear-forensics-beyond-the-science>
6. E. C. Morse, *Analytical Methods for Nonproliferation*. Springer International Publishing, 2018.
7. “Nuclear explained: The nuclear fuel cycle,” Jun 2021. [Online]. Available: <https://www.eia.gov/energyexplained/nuclear/the-nuclear-fuel-cycle.php>
8. B. Collum, “1 - nuclear fuel cycle,” in *Nuclear Facilities*, B. Collum, Ed. Woodhead Publishing, 2017, pp. 1–43. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780081019382000015>
9. D. Bodansky, *Nuclear energy: principles, practices, and prospects*. Springer, 2004.

10. P. R. Exline, *Machine Learning in the Countering Weapons of Mass Destruction Fight*. Cham: Springer International Publishing, 2020, pp. 71–92. [Online]. Available: https://doi.org/10.1007/978-3-030-28342-1_5
11. S. Lee, W. Song, and J.-S. Yang, “Nuclear fuel cycle-related rd classification for implementing the iaecas additional protocol,” *Progress in Nuclear Energy*, vol. 139, p. 103884, 2021.
12. J. A. Pike, T. L. Danielson, T. Whiteside, B. Mayer, N. Muralidhar, N. Self, and P. Butler, “Machine learning using open data sources for detection of nuclear proliferation activities,” Tech. Rep. SRNL-STI-2021-00047. [Online]. Available: <https://www.osti.gov/biblio/1764825>
13. M. Brinker, A. Bickley, B. Borghetti, A. Franz, B. Goldblum, J. Whetzel, and J. Bevins, “Machine Learning-Based Characterization of Nuclear Fuel Cycle Operations,” in *Symposium on Radiation Measurements and Applications*, Virtual Conference, 2021.
14. E. M. Purcell and D. Morin, *Electricity and magnetism*, 3rd ed. Cambridge University Press, 2013.
15. Q. Liu, A. P. Roberts, J. C. Larrasoana, S. K. Banerjee, Y. Guyodo, L. Tauxe, and F. Oldfield, “Environmental magnetism: principles and applications,” *Reviews of Geophysics*, vol. 50, no. 4, 2012.
16. R. Lanza and A. Meloni, *The Earths Magnetism*. Springer-Verlag Berlin Heidelberg, 2006.
17. D. Ma, M. Ding, J. Lu, J. Zhao, K. Yang, X. Fang, K. Wang, N. Zhang, and B. Han, “Magnetic noise calculation of mu-metal shields at extremely low fre-

- quencies for atomic devices,” *Journal of Physics D: Applied Physics*, vol. 54, no. 2, p. 025004, 2020.
18. K. Yamazaki, K. Kato, K. Muramatsu, A. Haga, K. Kobayashi, K. Kamata, K. Fujiwara, and T. Yamaguchi, “Incremental permeability of mu-metal in low magnetic fields for the design of multilayer-type magnetically shielded rooms,” *IEEE Transactions on Magnetics*, vol. 41, no. 10, pp. 4087–4089, 2005.
 19. T. D. Rossing, *Springer handbook of acoustics: with 91 tables*. Springer, 2007.
 20. I. Y. Solodov and B. A. Korshak, “Instability, chaos, and “memory” in acoustic-wave–crack interaction,” *Physical review letters*, vol. 88, no. 1, p. 014303, 2001.
 21. E. V. Vetchanin and A. O. Kazakov, “Bifurcations and chaos in the dynamics of two point vortices in an acoustic wave,” *International Journal of Bifurcation and Chaos*, vol. 26, no. 04, p. 1650063, 2016.
 22. W. Lauterborn and J. Holzfuss, “Acoustic chaos,” *International Journal of bifurcation and Chaos*, vol. 1, no. 01, pp. 13–26, 1991.
 23. P. Filippi, A. Bergassoli, D. Habault, and J. P. Lefebvre, *Acoustics: basic physics, theory, and methods*. Elsevier, 1998.
 24. C. Chapman, *Fundamentals of seismic wave propagation*. Cambridge university press, 2004.
 25. P. Bormann, B. Engdahl, and R. Kind, “Seismic wave propagation and earth models,” in *New manual of seismological observatory practice 2 (NMSOP2)*. Deutsches GeoForschungsZentrum GFZ, 2012, pp. 1–105.
 26. E. S. Krebes, *Seismic wave theory*. Cambridge University Press, 2019.

27. J.-O. Liljenzin, G. R. Choppin, and J. Rydberg, *Radiochemistry and nuclear chemistry (third edition)*. Butterworth-Heinemann, 2002.
28. K. S. Krane, *Introductory nuclear physics*. Wiley, 2016.
29. B. Briyatmoko, “Identification of high confidence nuclear forensics signatures for mining, milling and conversion processes,” Tech. Rep. IAEA-TECDOC-1820, 2017.
30. *Nuclear Forensics in Support of Investigations*, ser. Implementing Guides. Vienna: International Atomic Energy Agency, 2015, no. 2-G (Rev. 1). [Online]. Available: <https://www.iaea.org/publications/10797/nuclear-forensics-in-support-of-investigations>
31. N. Shoman and B. B. Cipiti, “Unsupervised machine learning for nuclear safeguards.” Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), Tech. Rep., 2018.
32. M. Benedict, T. H. Pigford, and H. W. Levi, *Nuclear chemical engineering*. McGraw-Hill, 1981.
33. L. Grancea, M. Mihalasky, M. Fairclough, J. R. Blaise, A. Boytsov, A. Hanly, L. Lopez, J. Marlatt, G. Schneider, and R. Vance, “Uranium resources, production and demand 2020,” Organisation for Economic Co-Operation and Development, Tech. Rep., 2020.
34. L. Sundaresan, S. Chandrashekar, and B. Jasani, “Monitoring uranium mining and milling using commercial observation satellites,” *ESARDA Bulletin*, vol. 53, pp. 73–82, 2015.

35. N. K. Kim, "Object-based land cover classification for pyongsan uranium mine and concentration plant using machine learning based classifier," in *Transactions of the Korean Nuclear Society Autumn Meeting*, 2020, pp. 13–14.
36. J. Diab, P. Burr, and R. Stohr, "Using machine learning and natural language processing to enhance uranium mining and milling safegaurds," Tech. Rep., 2018.
37. F. Hilbert, M. Hennebach, and Y. v. Wijk, "Developing the dn30," Dec 2019. [Online]. Available: <https://www.neimagazine.com/features/featuredeveloping-the-dn30-7548245/>
38. "Nrc infographics — nrc.gov - nrc web," Nov 2021. [Online]. Available: <https://www.nrc.gov/reading-rm/doc-collections/infographics/index.html>
39. U. D. of Energy, "K-25," Jan 2022. [Online]. Available: [https://en.wikipedia.org/wiki/K-25#/media/File:K-25_\(7609929206\).jpg](https://en.wikipedia.org/wiki/K-25#/media/File:K-25_(7609929206).jpg)
40. D. A. Gum, "A machine learning approach to characterizing particle morphology in nuclear forensics," Master's thesis, Air Force Institute of Technology, March 2020. [Online]. Available: <https://scholar.afit.edu/etd/3598>
41. I. J. Schwerdt, A. Brenkmann, S. Martinson, B. D. Albrecht, S. Heffernan, M. R. Klosterman, T. Kirkham, T. Tasdizen, and L. W. McDonald IV, "Nuclear proliferomics: A new field of study to identify signatures of nuclear materials as demonstrated on alpha-uo3," *Talanta*, vol. 186, pp. 433–444, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0039914018304508>
42. K. J. Pastoor, R. S. Kemp, M. P. Jensen, and J. C. Shafer, "Progress in uranium chemistry: Driving advances in front-end nuclear fuel cycle forensics," *Inorganic*

Chemistry, vol. 60, no. 12, pp. 8347–8367, 2021, pMID: 33619961. [Online]. Available: <https://doi.org/10.1021/acs.inorgchem.0c03390>

43. C. Ly, A. M. Olsen, I. J. Schwerdt, R. Porter, K. Sentz, L. W. McDonald, and T. Tasdizen, “A new approach for quantifying morphological features of u3o8 for nuclear forensics using a deep learning model,” *Journal of Nuclear Materials*, vol. 517, pp. 128–137, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0022311518313497>
44. D. R. Farley, M. G. Negus, and R. N. Slaybaugh, “Industrial internet-of-things data analytics for nuclear power safeguards.” Tech. Rep. SAND2018-12807 669746. [Online]. Available: <https://www.osti.gov/biblio/1481947>
45. Y. Cui, “Scoping study of machine learning techniques for visualization and analysis of multi-source data in nuclear safeguards,” Tech. Rep. BNL-203606-2018-FORE. [Online]. Available: <https://www.osti.gov/biblio/1436245>
46. E. Charatsidou, “A machine learning approach to the nuclear fuel fabrication process,” Master’s thesis, Department of Physics, KTH Royal Institute of Technology, June 2020. [Online]. Available: <http://www.diva-portal.org/smash/get/diva2:1441542/FULLTEXT01.pdf>
47. N. E. D. of Argonne National Laboratory, “Reactors designed by argonne national laboratory,” accessed on 3 March 2022. [Online]. Available: <https://www.ne.anl.gov/About/reactors/early-reactors.shtml#:~:text=Chicago%20Pile%201%20was%20the,Chicago's%20Stagg%20Field%20football%20stadium>.

48. H. Sekimoto, *Nuclear reactor theory: Coe-Ines Textbook*. The 21st Century Center of Excellence Program "Innovative Nuclear Energy Systems for Sustainable Development of the World", Tokyo Institute of Technology, 2007.
49. "High flux isotope reactor (hfir) user guide: A guide to in-vessel irradiations and experiments," Nov 2015. [Online]. Available: HighFluxIsotopeReactor(HFIR) USERGUIDE
50. G. J. Ehrhardt, A. R. Ketrin, and L. M. Ayers, "Reactor-produced radionuclides at the university of missouri research reactor," *Applied Radiation and Isotopes*, vol. 49, no. 4, pp. 295–297, 1998. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0969804397000389>
51. R. D. Cheverton and T. M. Sims, "Hfir core nuclear design." [Online]. Available: <https://www.osti.gov/biblio/4008514>
52. G. J. Ehrhardt, A. R. Ketrin, and L. M. Ayers, "Reactor-produced radionuclides at the university of missouri research reactor," *Applied Radiation and Isotopes*, vol. 49, no. 4, pp. 295–297, 1998. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0969804397000389>
53. D. M. Fouquet, J. Razvi, and W. L. Whittemore, "Triga research reactors: A pathway to the peaceful applications of nuclear energy," *Nuclear News*, vol. 46, no. 12, pp. 46–56, 2003.
54. "Chernobyl accident 1986." [Online]. Available: <https://world-nuclear.org/information-library/safety-and-security/safety-of-plants/chernobyl-accident.aspx>

55. “Fukushima daiichi accident.” [Online]. Available: <https://world-nuclear.org/information-library/safety-and-security/safety-of-plants/fukushima-daiichi-accident.aspx>
56. H. Boeck and M. Villa, “Triga reactor characteristics,” Vienna University of Technology, Tech. Rep., 2007.
57. A. Pungerčič, D. Čalič, and L. Snoj, “Computational burnup analysis of the triga mark ii research reactor fuel,” *Progress in Nuclear Energy*, vol. 130, p. 103536, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0149197020302845>
58. S. Moody, “The McClellan Nuclear Research Center - About Us,” Aug 2021, accessed on 03 March 2022. [Online]. Available: <https://mnrc.ucdavis.edu/about-us>
59. “McClellan Nuclear Radiation Center Safety Analysis Report,” Feb 2000. [Online]. Available: <https://www.nrc.gov/docs/ML0532/ML053250012.pdf>
60. C. Francois, *Deep Learning with Python*. Manning Publications Co., 2018.
61. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. The MIT Press, 2017.
62. T. M. Mitchell, *Machine Learning*. MacGraw-Hill, 1997.
63. A. Geron, *Hands-on machine learning with scikit-learn, Keras, and tensorflow: Concepts, tools and techniques to build Intelligent Systems*. O’Reilly, 2019.
64. G. James, D. Witten, T. Hastie, and R. Tibshirani, *An introduction to statistical learning with applications in R*. Springer, 2021.

65. “Scikit learn documentation.” [Online]. Available: <https://scikit-learn.org/stable/modules/clustering.html#clustering>
66. J. W. Bae, A. Rykhlevskii, G. Chee, and K. D. Huff, “Deep learning approach to nuclear fuel transmutation in a fuel cycle simulator,” *Annals of Nuclear Energy*, vol. 139, p. 107230, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306454919307406>
67. K. Nabeshima, T. Suzudo, K. Suzuki, and E. Türkcan, “Real-time nuclear power plant monitoring with neural network,” *Journal of Nuclear Science and Technology*, vol. 35, no. 2, pp. 93–100, 1998. [Online]. Available: <https://doi.org/10.1080/18811248.1998.9733829>
68. Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
69. F. Chollet *et al.*, “Keras,” <https://keras.io>, 2015.
70. X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, Y. W. Teh and M. Titterton, Eds., vol. 9. Chia Laguna Resort, Sardinia, Italy: PMLR, 13–15 May 2010, pp. 249–256. [Online]. Available: <https://proceedings.mlr.press/v9/glorot10a.html>
71. R. S. Srinivasamurthy, “Understanding 1d convolutional neural networks using multiclass time-varying signals,” Ph.D. dissertation, Clemson University, 2018.
72. Z. Cui, W. Chen, and Y. Chen, “Multi-scale convolutional neural networks for time series classification,” *arXiv preprint arXiv:1603.06995*, 2016.

73. Y. Kim, “Convolutional neural networks for sentence classification,” *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
74. J. Gao, P. Li, Z. Chen, and J. Zhang, “A Survey on Deep Learning for Multimodal Data Fusion,” *Neural Computation*, vol. 32, no. 5, pp. 829–864, 05 2020. [Online]. Available: https://doi.org/10.1162/neco_a_01273
75. F. Castanedo, “A review of data fusion techniques,” *The Scientific World Journal*, vol. 2013, p. 1–19, 2013.
76. R. Hostettler and P. M. Djurić, “Vehicle tracking based on fusion of magnetometer and accelerometer sensor measurements with particle filtering,” *IEEE Transactions on Vehicular Technology*, vol. 64, no. 11, pp. 4917–4928, 2015.
77. S. Piramanayagam, E. Saber, W. Schwartzkopf, and F. W. Koehler, “Supervised classification of multisensor remotely sensed images using a deep learning framework,” *Remote Sensing*, vol. 10, no. 9, 2018. [Online]. Available: <https://www.mdpi.com/2072-4292/10/9/1429>
78. D. Anguita, A. Ghio, L. Oneto, X. Parra, J. L. Reyes-Ortiz *et al.*, “A public domain dataset for human activity recognition using smartphones.” in *Esann*, vol. 3, 2013, p. 3.
79. J. Brownlee, “1d convolutional neural network models for human activity recognition,” Aug 2020. [Online]. Available: <https://machinelearningmastery.com/cnn-models-for-human-activity-recognition-time-series-classification/>
80. G. Koch, “Siamese neural networks for one-shot image recognition,” Master’s thesis, Graduate Department of Computer Science, University of Toronto, 2015. [Online]. Available: <http://www.cs.toronto.edu/~gkoch/files/msc-thesis.pdf>

81. B. Goldblum, *NSD Staff Meeting*. Lawrence Berkeley National Laboratory, Oct 2020. [Online]. Available: <https://conferences.lbl.gov/event/463/contributions/3899/attachments/3025/1536/Goldblum-Transferability-Oct2020.pdf>
82. “HFIR operating forecast & planning schedule,” Aug 2018, accessed on 03 March 2022. [Online]. Available: https://neutrons.ornl.gov/sites/default/files/-%20OperatingForecast_2018-08-03.pdf
83. J. Tibbetts, “Explainable classification of nuclear facility operational state using node and region importance for sensor networks,” Master’s thesis, EECS Department, University of California, Berkeley, May 2021. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-95.html>
84. A. Rajadhyaksha, N. Hubley, G. Fairchild, P. Schuster, and E. Casleton, “Multi-Informatics for Nuclear Operations Scenarios ,” in *Center for Verification Technology Workshop*, Ann Arbor, MI, 2018.
85. C. Stewart, “HFIR-109 cooling fan schematic,” personal communication, September 2021.
86. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
87. M. Hollander, D. A. Wolfe, and E. Chicken, *Nonparametric statistical methods*. Wiley-Blackwell, 2014.
88. M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp,

- G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
89. J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. Lecun, C. Moore, E. Sackinger, and R. Shah, “Signature verification using a “siamese” time delay neural network,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 07, no. 04, pp. 669–688, 1993. [Online]. Available: <https://doi.org/10.1142/S0218001493000339>
90. Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel, “Handwritten digit recognition with a back-propagation network,” *Advances in neural information processing systems*, vol. 2, 1989.

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY) 03-24-2022		2. REPORT TYPE Master's Thesis			3. DATES COVERED (From — To) August 2020 — Mar 2022	
4. TITLE AND SUBTITLE Deep Learning Approach to Multi-phenomenological Nuclear Fuel Cycle Signals for Nonproliferation Applications				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Dicks, Preston J, 2d Lt, USAF				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765					8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENP-MS-22-M-087	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Threat Reduction Agency 8725 John J. Kingman Rd., Fort Belvoir, VA					10. SPONSOR/MONITOR'S ACRONYM(S) DTRA	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT In order to reduce the time required for data analysis and decision-making relevant to nuclear proliferation detection, Artificial Intelligence (AI) techniques are applied to multi-phenomenological signals emitted from nuclear fuel cycle facilities to identify non-human readable characteristic signatures of operations for use in detecting proliferation activities. Seismic and magnetic emanations were collected in the vicinity of the High Flux Isotope Reactor (HFIR) and the McClellan Nuclear Research Center (MNRC). A novel bi-phenomenology DL network is designed to test the viability of transfer learning between nuclear reactor facilities. It is found that the network produces an 84.1% accuracy (99.4% without transient states) for predicting the operational state of the MNRC reactor when trained on the operational state of the HFIR reactor. In comparison, the best performing traditional ML single-phenomenology algorithm, K-Means, produces a 67.8% prediction accuracy (80.5% without transient states).						
15. SUBJECT TERMS artificial intelligence, machine learning, deep learning, transfer learning, nuclear fuel cycle, nuclear reactors, nuclear proliferation, 1DCNN						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Abigail Bickley, AFIT/ENP	
U	U	U	UU	122	19b. TELEPHONE NUMBER (include area code) (937) 255-6565; abigail.bickley@afit.edu	