EVALUATING SYSTEM CONFIDENCE OF NEAR-FIELD ANTINEUTRINO-BASED NUCLEAR REACTOR SAFEGUARDS

A Thesis
Presented to
The Academic Faculty

By

Matthew Dunbrack

In Partial Fulfillment
of the Requirements for the Degree
Masters of Science in
Nuclear Engineering

Georgia Institute of Technology

December 2022

EVALUATING SYSTEM CONFIDENCE OF NEAR-FIELD ANTINEUTRINO-BASED NUCLEAR REACTOR SAFEGUARDS

Thesis committee:

Dr. Anna Erickson
Nuclear & Radiological Engineering Program
Georgia Institute of Technology

Dr. Steven Biegalski
Nuclear & Radiological Engineering Program
Georgia Institute of Technology

Dr. Nathaniel Bowden
Nuclear and Chemical Sciences Division
Lawrence Livermore National Laboratory

Date approved: June 10, 2022

ACKNOWLEDGMENTS

I would like to thank all of my professors, national laboratory collaborators, and committee members for their continuous support and guidance throughout my graduate school experience. In particular, Dr. Kyle Hartig opened my eyes to the world of nuclear research. By joining his research group as an undergraduate research assistant, I was given many opportunities to connect with national laboratories and learn about state-of-the-art research. Dr. Rachel Whitlark provided me with context to the world of nuclear policy. After taking the Politics, Tech & Proliferation course, I began to consider the role of nuclear in our society beyond just technological limitations. Dr. Meghan McGarry reached out and introduced me to the world of national laboratories. Through her initiative, I have become involved in projects I never imagined. Dr. Nathaniel Bowden guided me as I joined the world of antineutrino research. With his help, I was able to gain perspective to the importance of my graduate research on a larger scale. Dr. Steven Biegalski always supported me as I navigated the world of balancing meetings, conferences, and graduate school. It is always a pleasure hearing stories, both old and new, within the evolving nuclear field. Dr. Anna Erickson made endless opportunities for me to grow both in the world of research and academia. Prior to my enrollment at Georgia Tech, I thought graduate school was divided into two, distinct portions: being a student and being a researcher. With her help these last two years, I realized that graduate school is really about learning how to fit the two portions together. Lastly, this work would not have been possible without the ongoing support of the Consortium for Monitoring, Technology, and Verification (MTV). Through this consortium, I have become involved in diverse and pioneering research in the nuclear nonproliferation community.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

The International Atomic Energy Agency (IAEA) relies heavily on surveying facilities and verifying inventories to ensure that special nuclear material (SNM) pathways are correct and complete. This process, conducted through on-site inspections, draws a significant amount of the limited resources from the IAEA. Through implementing near-field antineutrino detection systems, changes in reactor core composition can be continuously monitored without the need of any expensive and invasive inspection. Our confidence in such a system, however, needs to be carefully considered for the IAEA to implement antineutrino detection systems for nuclear reactor safeguards. In this work, system confidence, or the certainty of the predicted antineutrino spectra, is evaluated to outline current antineutrino-based safeguard capabilities as well as to highlight the leading causes of uncertainty. The proposed system under evaluation is the Reactor Evaluation Through Inspection of Near-field Antineutrinos (RETINA) system, which utilizes high-fidelity modeling to predict the antineutrino spectra emitted from a simulated reactor. Certain deviations in real-time antineutrino spectra would indicate a shift in fissile inventory and a possible diversion of SNM from the reactor core. To fully analyze the role of reactor designs and diversion scenarios in the system evaluation, the antineutrino spectra was simulated for various next generation reactor designs as well as processed for possible diversion scenarios the IAEA would aim to detect. The results indicate that larger reactors with more common fissile inventories lead to lower system uncertainty. While some simulated diversion scenarios were consistently detected, the overlapping spectra led to low confidence of diversion following IAEA standards. Future work will go into modeling new reactor-detector systems as well as applying modern machine learning methods for confidence improvement.

CHAPTER 1
INTRODUCTION

1.1 Proliferation Concerns in the Nuclear Power Industry

The world is in need of a reliable and sustainable energy source [1]. While many decision makers look towards solar and wind as the soul solution to climate change, nuclear reactors could substantially decrease the global carbon footprint [2]. With the spread of nuclear reactors, however, there is a growing nuclear proliferation concern [3]. For there to be a nuclear renaissance, the agency inspecting a majority of these reactors, the International Atomic Energy Agency (IAEA), would need to be prepared to safeguard widespread and novel nuclear technologies [4].

Since the implementation of the Treaty on the Non-Proliferation of Nuclear Weapons (NPT) in 1970, nuclear reactors spread across the globe [5]. The NPT has three pillars of focus: nonproliferation, disarmament, and peaceful use. These three components of the NPT are verified through the IAEA. Ensuring nuclear nonproliferation and the peaceful use of nuclear energy, the IAEA has developed a safeguarding framework through comprehensive safeguards agreements to prevent diversion of nuclear energy to nuclear weapons for all non-nuclear-weapon agreement states. This framework, however, might not suffice in effectively safeguarding a shifting global nuclear order [4].

The IAEA is looking towards modernizing their infrastructure to account for the evolving global nuclear power industry [6]. Some aspects of modernization include the Safeguards by Design concept [6], in which reactor designers incorporate safeguards measures prior to development, as well as implementing remote safeguards technologies [7]. While the IAEA can encourage the development of Safeguards by Design reactor designs, the employment of proliferation resistant reactors are beyond the agency's control [6]. Remote

sensing safeguards provide a mechanism for effective and efficient reactor monitoring and treaty verification without the need for invasive inspections. The potential of these remote sensing technologies as nuclear reactor safeguards remain in the authenticity and reliability of the gathered information [7].

1.2 Antineutrino Detection for Reactor Safeguards

Many researchers believe antineutrino detectors can be used to remotely safeguard nuclear reactors [8, 9, 10, 11, 12]. Antineutrinos, due to their low interaction cross section, can be used to gather reliable and continuous real-time information from a nuclear reactor core [8]. With enough antineutrinos collected, these safeguards system can verify a reactor's on-off status, thermal power level, and, to some extent, the fissile inventory [9]. Antineutrino detectors are far from optimised as IAEA safeguards and vary widely in parameters depending on user-needs and limitations [10]. The capabilities of these antineutrino-based safeguards systems continue to be explored across the globe for IAEA application [11]. The Nu Tools Report [12] in 2021 explored potential roles for neutrinos in nuclear energy and security based on the opinions of subject matter experts in neutrino physics, reactor design, and nuclear security. The diverse findings support that they have potential beyond current IAEA safeguards, and can be used for advanced reactor safeguards, future nuclear deals, reactor operations, non-cooperative reactor monitoring, spent nuclear fuel assay, and for post-accident response [12].

CHAPTER 2

BACKGROUND

There are many different antineutrino detector configurations. These detectors can be divided into three classes, near-field ($\sim$10 - 200m), mid-field ($\sim$200 - 1,000 m), and far-field ($\sim$10,000 m) [13]. Near-field antineutrino detectors are placed as close to the reactor as possible to maximize the geometric efficiency and experience the largest antineutrino flux possible from the source. While this is an ideal location for antineutrino collection from a concentrated source, the bounded size of reactor facilities limit the available detector mass that can be deployed. Mid-field detectors, typically located just outside the facility fence, allow for a larger detector but sacrifices detector geometric efficiency. This distance increase might also be necessary depending on the cooperative agreement between all involved parties [14]. The last antineutrino detector class, far-field, is significantly further away from reactor cores and can be scaled to megaton volumes, as is the case with the WATCHMAN detectors [15]. These far-reaching detectors can theoretically detect antineutrinos being emitted from undeclared facilities [13].

2.1 Fundamental Properties of Near-Field Antineutrino Safeguards

Most antineutrino detectors for reactor safeguards have been designed for near-field implementation [13]. Near-field antineutrino monitoring capabilities include verification of reactor operational status, thermal power, and fissile content. While these capabilities have been studied to some extent, more research and development is needed for effective system implementation [16].

2.1.1 Production

Within fission reactors, there are a significant portion of neutron-rich radioactive fission-products [17]. These products are most likely to decay rapidly through a series of beta decays, outlined in Equation Equation 2.1 [18]. There are multiple types of antineutrinos, referred to as 'flavors,' and the original flavor of the antineutrino is dictated by the counter particle in formation [19]. Nuclear reactors are a massive source of electron antineutrinos, $\bar{\mu_e}$, as seen with the electron formation of beta decay [10].

$$ {}^{A}_{Z}X \rightarrow {}^{A}_{Z+1}Y + e^{-} + \bar{\nu}_e \tag{2.1} $$

When an atom fissions, fission product libraries are used to predict the various possible paths of fission product formation. These product libraries are dependent on the isotope undergoing fission. From the known fission product, the various decay chains are predicted via well-defined branching ratio libraries [20]. Since a majority of these neutron-rich products will undergo beta decay, the electron antineutrino yields and energies can be deduced [16]. On average, there are 6 antineutrinos generated per fission [19]. From well-understood fission product and branching ratio libraries, the average amount and energies of these antineutrinos can be reasonably simulated [16]. Due to the isotopic variation in fission products, there are slight deviations between the antineutrino yield spectra for the different fissioning isotopes, as seen in Figure 2.1. By effectively detecting this alteration in isotopic antineutrino flux, even diversion scenarios in which fuel is replaced with fissionable material could be monitored [8].

2.1.2 Detection

There are many different mechanisms used for neutrino detection [21]. Neutrinos, lacking electrical charge and significant mass, interact very little with other forms of matter [22]. Their relative properties and interaction mechanisms depend on their leptonic flavor and

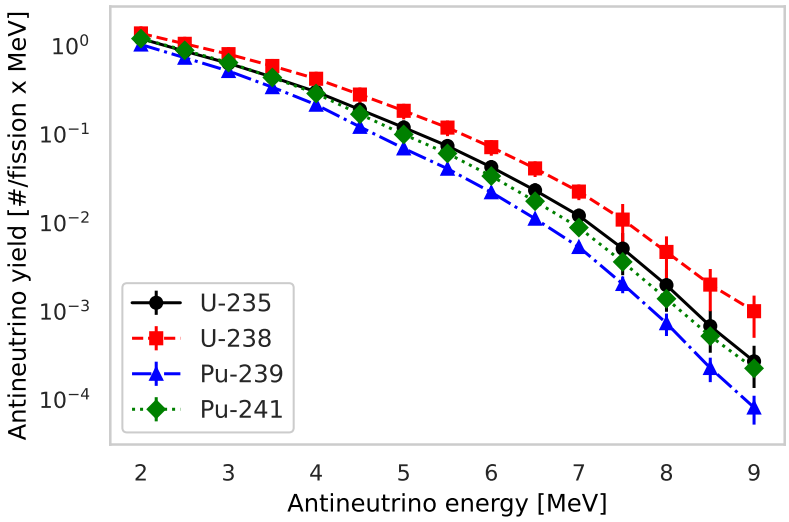

Figure 2.1: Antineutrino yield per fission per MeV for U-235, U-238, Pu-239, and Pu-241

charge. There exists three different active neutrinos: electron, $\bar{\nu_e}$, muon $\bar{\nu_\mu}$, and tau $\bar{\nu_\tau}$. There are also equivalent antiparticles, or antineutrinos, to balance various leptonic charges through existing processes [23].

Although these subatomic particles are outlined in the Standard Model, evidence supports that neutrinos follow physics beyond our current understanding [24, 25, 26]. Along with leptonic flavors, antineutrinos have associated mass flavors. Due to these flavors having slightly varying wavelengths, the antineutrinos form is predicted by an oscillating term, or a mixing angle. While the antineutrino is born in a specific form, the probability of antineutrino transition alters as a function of neutrino energy and distance between the source and detector. Through experimentation, another neutrino flavors, referred to as sterile neutrinos, are theorized to exist with masses in the eV scale. Their existence would explain discrepancies seen in short-baseline reactor experiments [26].

Depending on the desired neutrino form, different detection mechanisms can be leveraged to design an optimal system. A few neutrino detector technologies include Cherenkov, liquid Argon TPC, liquid scintillator, and sampling detectors [21]. One of the most commonly used antineutrino detector interaction, inverse beta-decay (IBD), typically is detected via scintillation and/or Cherenkov detectors [27]. While there are other measurable

neutrino interactions, such as elastic neutrino-electron scattering and coherent neutrino-nucleus scattering [28], the IBD interaction is antineutrino flavor specific [29] and has a relatively high cross section for low MeV antineutrinos [30]. With a majority of the reactor antineutrinos emitted in the electron flavor, there is a benefit to the flavor-specific nature of IBD detection [29].

In an IBD interaction (Equation Equation 2.2), an antineutrino interacts with a proton, forming a positron and neutron. For this interaction to take place, the initial antineutrino must have a minimum threshold energy of 1.8 MeV. Accounting for this threshold energy, only about 1.92 and 1.45 antineutrinos for U-235 and Pu-239, respectively, are produced per fission in the detectable range [31]. The IBD-generated positron quickly annihilates within the detector, depositing its initial kinetic energy as well as two 511 keV photons. The neutron, after some thermalization, is typically captured by a neutron-absorbing dopant, such as gadolinium or lithium [32], and after some time produces several photons through de-excitation. Through careful positron-neutron detection coincidence timing, IBD events can be reconstructed [33].

$$\bar{\nu}_e + p \rightarrow e^+ + n \; (E_{\bar{\nu}_e} > 1.8 \; MeV) \tag{2.2}$$

2.2 Current Antineutrino-Based Safeguard Capabilities

2.2.1 Production

There are two main approaches to predicting the antineutrino spectra emitted from a reactor core: the summation method and the conversion method [34]. The summation method consists of computing the total antineutrino spectra via fission product libraries and weighted beta decay branches. Through this method, the spectra can deconvolved to analyze each component contributing to the antineutrino spectra. This method is severely limited by available nuclear data. The nuclear beta decay data is especially limiting in the summation method due to the Pandemonium effect, in which the higher energy antineutrino yields are

overestimated [35]. The summation method has been applied as both deterministic, as described in Mueller *et al.* [36], and stochastic calculations, as outlined in the Oklo toolkit [37].

Alternatively, the antineutrino spectra can be calculated via the conversion method. This method utilizes experimental beta particle measurements to deduce possible antineutrino spectra. These spectra values are normalized by a user-defined factor to control the spectra used in calculation. The conversion method is highly dependent on the beta particle measurements, which took place in the 1980s at the Institute Laue-Langevin (ILL) for the U-235, Pu-239, and Pu-241 antineutrino spectra [38, 39, 40]. The antineutrino spectra from fast fissioning U-238 was later added to this method [41].

The Huber-Mueller model [42] is a widely used conversion-method-based library for the antineutrino spectra [43]. The spectra was developed through a combination of virtual beta-decay branches and Monte Carlo simulations to quantify correction terms to match the ILL experimental data. This model led to an overall 2-3% upwards shift in antineutrino yield compared to previous spectra. The library is comprised of experimentally-driven values for U-235, Pu-239, Pu-241, and U-238 as well as theory-driven values for all other fissionable isotopes of interest. They are divided in antineutrino energy bins ranging from 2 MeV to 9 MeV with 0.5 MeV energy bin widths. This library also includes well-evaluated estimated errors for the spectra [42].

2.2.2 Detection

There have been many experiments that have pioneered the path towards near-field antineutrino-based safeguards across the globe, including Rovno [44] in Russia, SONGS [45] in the U.S.A, and Double Chooze [46] in France. These projects utilized Gd-doped liquid scintillator detectors to measure IBD events and deduce characteristics of nearby reactors [11]. Detector characteristics, such as dopant type and scintillation material, can vary depending on the project. For example, compared to the Gd-doped, liquid scintillation experiments

previously mentioned, the PANDA experiment utilizes a plastic scintillation material target [47] and the PROSPECT experiment utilizes a Li-6 dopant [48].

The Precision Reactor Oscillation and Spectrum Experiment (PROSPECT) is of particular interest within this study. The project is designed to take precise antineutrino spectra measurements from the High Flux Isotope Reactor (HFIR) at the Oak Ridge National Laboratory as well as to probe eV-scale sterile neutrinos. The 4-ton, segmented, Li-6 doped detector, antineutrino detector - 1 (AD-I), is placed from 7 m to 13 m away from the 85 MW_{th}, highly enriched uranium fuel loaded HFIR core. PROSPECT has been collecting data since March 2018 with these set parameters [48]. While AD-I, the prospect collaboration is preparing an upgraded detector, antineutrino detector - 2 (AD-2), for system improvements [49].

Background radiation fields play a large role in these antineutrino detector experiments due to the low event rates of inverse-beta decay. The sources of background include naturally occurring radiation within the facility structures, reactor-generated radiation, and cosmogenic background. The majority of background radiation within the facility, namely γ-rays and thermal neutrons, can largely be shielded. The fast neutron background, $>$10 MeV, is the most significant concern since fast neutrons can mimic the signature of IBD events. Although there are fast neutrons born within reactor cores, reactor-correlated fast neutrons are typically moderated to sub-10 MeV energies prior to reaching the detector. Through detailed simulation and experimental verification, background values can be reasonably determined for a reactor-detector system at specific locations [50]. The relative background rates can be scaled according to the reactor-detector system and configuration, as seen in Figure 2.2, in which the radiation background at HFIR has been scaled to match a small reactor, the Advanced Fast Reactor (AFR) and a large reactor, the ultra-long cycle fast reactor (UCFR), as well as an altered detection system, the Reactor Evaluation Through Near-field Antineutrinos system (RETINA) [51].

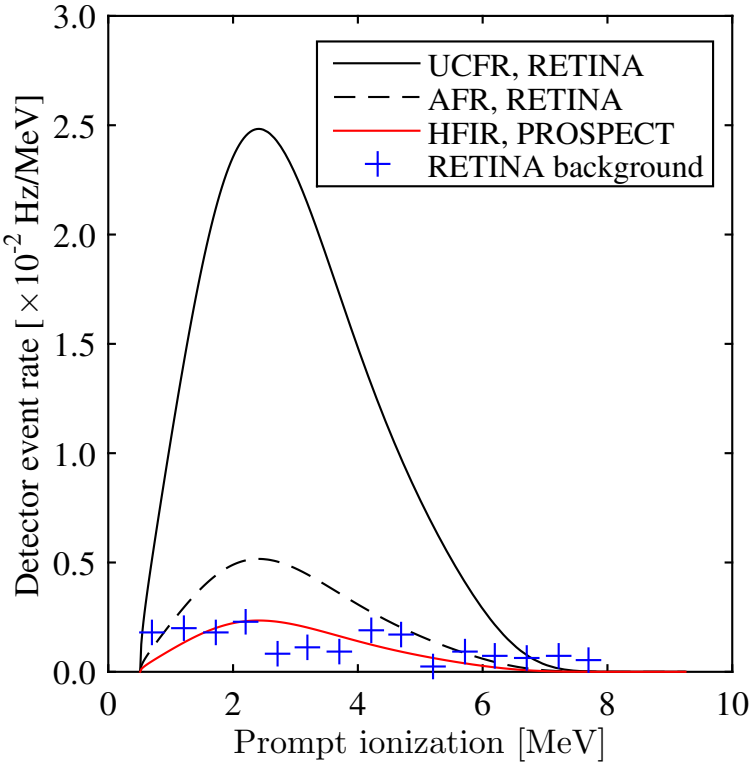

Figure 2.2: PROSPECT AD-I background estimates as described in Ashenfelter *et al.* [52] Figure from Stewart [51]

2.3 The RETINA System

The proposed system of this study is the Reactor Evaluation Through Inspection of Near-field Antineutrinos system, or RETINA system. The system includes a suite of detectors for data acquisition and a sensitivity tool for safeguards implementation. In practice, the suite of antineutrino detectors are placed near any well-understood reactor core to isolate the antineutrino flux emitted from the reactor core. Based on high-fidelity modeling, we can reasonably predict the quantity and energies of antineutrinos collected by our suite of detectors. If the real-time antineutrino spectra deviates from our predict spectra, this could indicate that the fissile inventory of the core has been altered [53]. Sensitivity studies for the RETINA system have been performed in Stewart [51] using statistical methods for diversion confidence analysis.

The detector design of choice for the RETINA system closely resembles the initial antineutrino detector design (AD-I) used in the PROSPECT experiment [52]. These detectors are well understood and the experiments associated with the detector provide a fair baseline for current detection capabilities [48]. The detector suite characteristics are listed in Table 2.1. These parameters will likely change as antineutrino detectors continue to evolve but these initial values include realistic properties for system implementation.

Table 2.1: Properties of the RETINA detection suite

Parameter	
Scintillator	EJ-309
Specific Gravity	0.959
Neutron Capture Dopant	Li-6
Proton Density	5.5 x 10^{28} /m^3
System Mass	$<$10 t
Target Mass	5 t
Efficiency	42 %
Standoff Distance	17 - 25 m

2.4 Reactor Designs of Interest

Three different reactor designs were included in this study: the small fluoride salt-cooled high-temperature reactor (FHR), the advanced fast reactor-100 (AFR-100), and the lead-cooled flexible conversion ratio fast reactor in unity configuration (LFR-FCR1). These varying reactor designs were chosen for their unique isotopic inventories, power levels, and fuel cycle lengths.

2.4.1 FHR

The FHR reactor design is a well-researched generation IV reactor concept utilizing a lithium beryllium fluoride salt (Li_2Be_4F), or FLiBe, as the fuel coolant to reach elevated temperatures with a high volumetric heat capacity. While there are many reactor designs that include this coolant, such as the advanced high-temperature reactor (AHTR) [54] and the small modular advanced high-temperature reactor (SmAHTR) [55], they can also vary in fuel type. The previously mentioned FHR reactor designs both use tri-structural isotropic (TRISO) particle fuel, which is made up of a uranium oxide/carbide kernel. This relatively novel fuel type technology could benefit from an antineutrino-based safeguard [56], but would involve a complex system-preparation scheme to account for the bulk inventory. Rather than redesign the current system, a UO_2 pin-type, graphite-moderated, small FHR developed by Mohamed *et al.* [57] was modeled to closely reflect the other reactor designs outlined in this study.

This reactor design includes many different configurations to match user requirements. User needs include the size of the reactor and the enrichment and longevity of the fuel. To optimize the system to these requirements, system variations include the number of fuel rings around the assembly control rod as well as the channel diameter. The different ring options are illustrated in Figure 2.3. In this study, we model the 4 ring, 6 variation (4R6) assembly configuration. Parameters matching this configuration are listed in Table 2.2 for

a 13 effective full-power month (EFPM) burnup cycle [57].

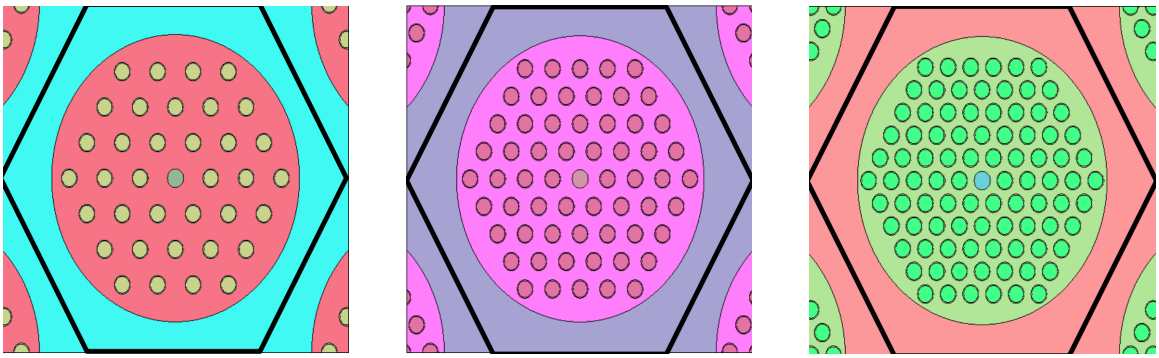

Figure 2.3: FHR configurations for 3 (left), 4 (middle), and 5 (right) rings [58]

2.4.2 AFR-100

The AFR-100 is a small, sodium-cooled fast reactor design developed by Argonne National Laboratory. The main objectives for this core design was to develop a small modular reactor for small local grids that could then run for many years without refueling. To do so, this design relies on an onion-enrichment scheme of U-Zr metallic alloy fuel in which the outer portions of the core have a higher enrichment level than the inner portions of the core. Towards the beginning of the core's operation, the outer enrichment provides a majority of the power. After a few years, the flux moves inwards as the inner content of U-238 breeds into a significant portion of Pu-239. This fuel development allows for the AFR-100 to remain operational without refueling for 30 effective full power years (EFPYs) while still maintaining a reasonable reactivity swing [59]. The enrichment scheme can be seen in Figure 2.4. Other key reactor parameters are listed in Table 2.3.

The breeding-nature and longevity of this reactor design make it desirable for a novel safeguarding method. Using current inspection techniques, the core fissile inventory cannot be properly inspected until after the 30 EFPYs of operation. Considering the core is designed to breed more than a significant quantity of plutonium, this reactor design could lead to a serious proliferation problem if not properly addressed. With antineutrino-based safeguard systems being non-intrusive and long-lasting, they show potential as a safeguard

Table 2.2: Properties of the 4R6 FHR reactor design [58] [57]

Parameter	
Power	125 MW_{th}
Fuel Cycle Length	13 EFPM
Primary Coolant	FLiBe
Number of Fuel Pins per Assembly	60
Number of Fuel Assemblies	60
Pin Outer Radius	1.068 cm
Cladding Thickness	0.057 cm
Fuel Pin Pitch	1.944 cm
Active Core Height	100 cm
Composition [Weight %]	
O	12.4
U	87.6
Enrichment	16 %U-235
Smear Density	10.32 g/cc
Specific Power Density	128.81 MW/t
Initial Heavy Metal Inventory	1.07 t
Average Discharge Burnup	3.04 MWd/kg

for this reactor design.

2.4.3 LFR-FCR1

The LFR is another next-generation reactor design considered in this work. The LFR-FCR can be implemented with either a conversion ratio of 0 or a conversion ratio of 1. The conversion ratio dictates specific parameters of the core to fulfill a certain objective. With a conversion ratio of 0, the initial reactor inventory is optimized to burn TRU with maximum efficiency. With a conversion ratio of 1, or the unity configuration, the core is designed to produce energy in a sustained cycle [60].

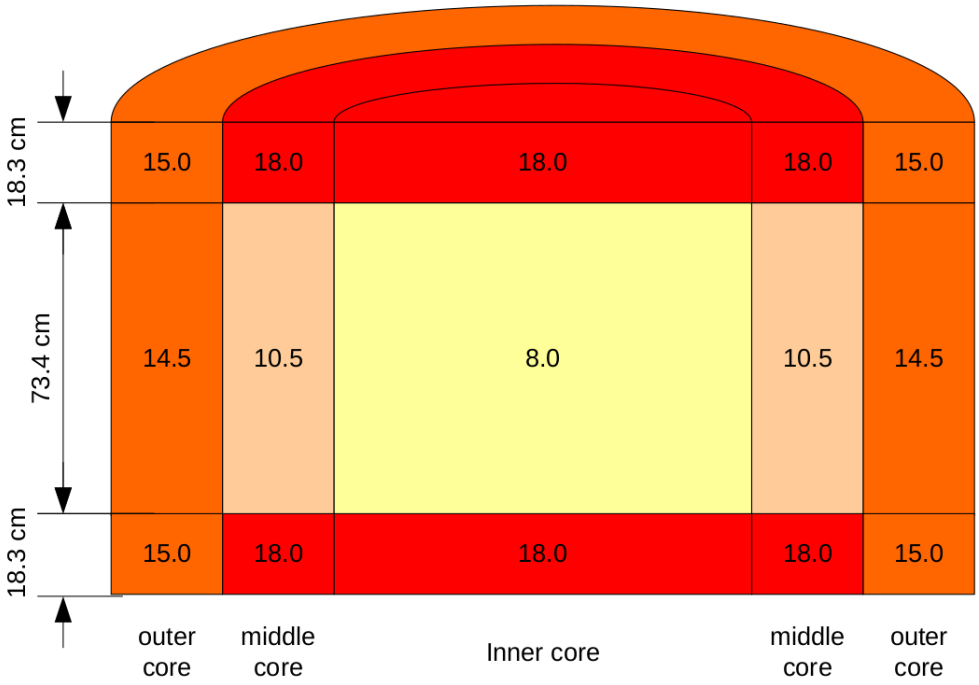

Figure 2.4: AFR-100 enrichment zoning scheme in weight % U-235 [59]

The LFR-FCR in unity operation was selected for this study due to the longer lasting fuel cycle. While the conversion ratio 0 reactor does contain a larger fraction of fissionable transuranic isotopes, the fuel cycle only lasts about 1.5 EFPYs compared to 5 EFPYs with the unity core. Due to the rare-event nature of antineutrino detection, antineutrino-based safeguards have the most potential with longer-lasting fuel cycles [9].

To sustainably burn fuel, the core is divided into three zones, as seen in Figure 2.5. The isotopic composition and assembly-specific locations of these zones were optimized to last the 5 EFPY while maintaining a flat radial power distribution for the duration [59]. The composition within these zones as well as other parameters are listed in Table 2.4.

Table 2.3: Properties of the AFR-100 reactor design [59]

Parameter	
Power	250 MW_{th}
Fuel Cycle Length	30 EFPY
Primary Coolant	Sodium
Number of Fuel Pins per Assembly	91
Number of Fuel Assemblies	150
Pin Outer Radius	1.501 cm
Cladding Thickness	0.052 cm
Fuel Pin Pitch	1.606 cm
Active Core Height	110 cm
Composition [Weight %]	
Zr	10
U	90
Enrichment	8 - 18 %U-235
Smear Density	10.94 g/cc
Specific Power Density	10.5 MW/t
Initial Heavy Metal Inventory	23.9 t
Average Discharge Burnup	101 MWd/kg

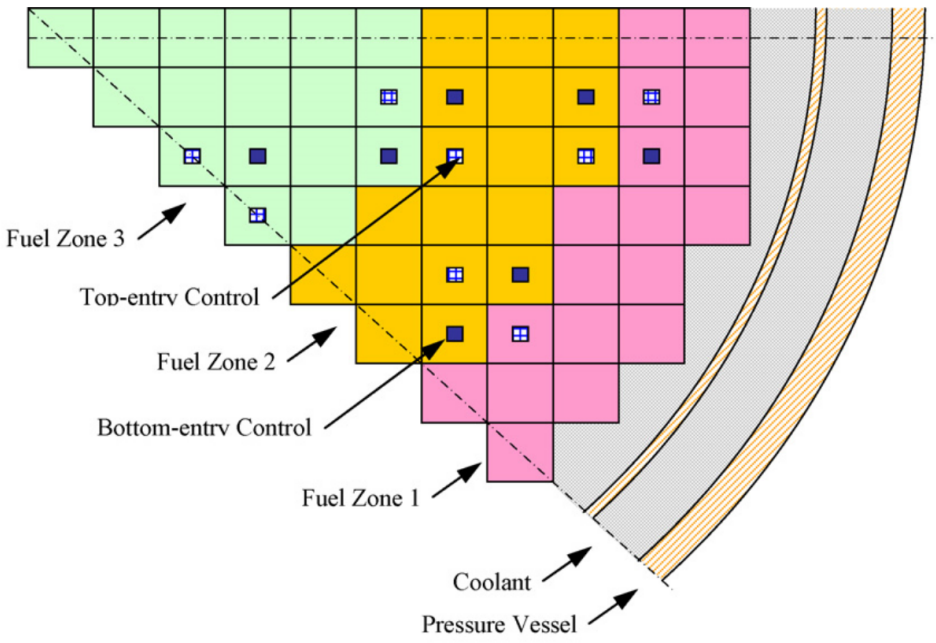

Figure 2.5: LFR-FCR1 radial core zoning [61]

Table 2.4: Properties of the LFR-FCR1 reactor design [61] [60]

Parameter	Zone 1	Zone 2	Zone 3	Core
Power	-	-	-	2400 MW_{th}
Fuel Cycle Length	-	-	-	5 EFPY
Primary Coolant	-	-	-	Lead
Number of Fuel Pins per Assembly	-	-	-	441
Number of Fuel Assemblies	128	124	197	349
Pin Outer Radius	-	-	-	0.752 cm
Cladding Thickness	-	-	-	0.063 cm
Fuel Pin Pitch	-	-	-	0.978 cm
Active Core Height	-	-	-	130 cm
Composition [Weight %]				
Zr	10.00	15.00	19.00	-
TRU	15.00	14.17	13.50	-
U	75.00	70.83	67.50	-
Enrichment	-	-	-	0.71 %U-235
Smear Density	11.53 g/cc	10.70 g/cc	10.12 g/cc	-
Specific Power Density	-	-	-	40.60 MW/t
Initial Heavy Metal Inventory	25.6 t	21.5 t	15.2 t	62.2 t
Average Discharge Burnup	60 MWd/kg	86 MWd/kg	93 MWd/kg	77 MWd/kg

CHAPTER 3
METHODOLOGY

3.1 Confidence in Antineutrino Yield

The crux of the RETINA system is the antineutrinos collected by the detector. The simulated antineutrinos collected, however, will contain some uncertainty based on the statistical nature of radiation decay as well as the propagated uncertainty throughout the system [62]. While some uncertainty is unavoidable, there are many forms of uncertainty that can be mitigated [63]. Our confidence in these antineutrino yields is defined by the uncertainty propagated up to the antineutrino yield dataset used in deviation detection. The following subsections outline our method for deriving antineutrino yields and their associated uncertainties for reference case reactors, or reactors that are operating as reported to the IAEA.

3.1.1 Reactor Modeling

SERPENT2 [64] was used for high-fidelity reactor modeling calculations. This code, developed at VTT Technical Research Centre of Finland, uses Monte Carlo methods for three-dimensional, continuous-energy particle transport. SERPENT2 was chosen as the reactor modeling tool based on its built-in neutronics and depletion calculation ability as well as its quick computation speed through parallelization [65].

While SERPENT2 has many useful capabilities, in this work, we are primarily interested in highly accurate reactor inventories and isotopic fission rates throughout the entirety of the fuel cycle. Reactor inventories are used to determine the assemblies that depict a proliferation threat based on significant plutonium quantities [8]. Isotopic fission rates are the most vital component of reactor modeling since they correlate to the antineutrino spectra being emitted from the core [66]. As outlined in section 2.2, different fissionable

isotopes will produce a different amount of antineutrinos at various energies [42]. Since the RETINA system in practice will need to detect diversion scenarios at the end of a reactor's lifetime just as well as with a fresh core, there is a special importance on fairly accurate inventories and isotopic fission rates at the end of the core fuel cycle [8].

Each reactor of interest, the FHR, AFR-100, and LFR-FCR1, were modeled in SERPENT2 for the entirety of their fuel cycles. Their geometries are shown in Figure 3.1 as modeled in SERPENT2. For each Monte Carlo run, there were 200000 source neutrons per cycle, 200 active cycles, and 20 inactive cycles. The burnup steps for each reactor design are listed in Table 3.1 and were normalized to maintain reactor power density. The burnup steps varied between the reactor designs depending on the core lifetimes and how drastically the inventory changed. The isotopic fission rates were calculated using the analog reaction rate mode in SERPENT2, in which sampled events were counted and recorded.

The isotopic fission reaction rate, FR_i, for the isotopes of interest (Table 3.2) were isolated for a desired burnup step and stored in a matrix along with their listed variances, $\sigma^2_{FR_i}$ using Jupyter Notebook. The listed variances only account for statistical variation through the Monte Carlo approach. All other modeling variances, such as reactor inventory through burnup, were assumed to be negligible compared to the statistical uncertainty [67].

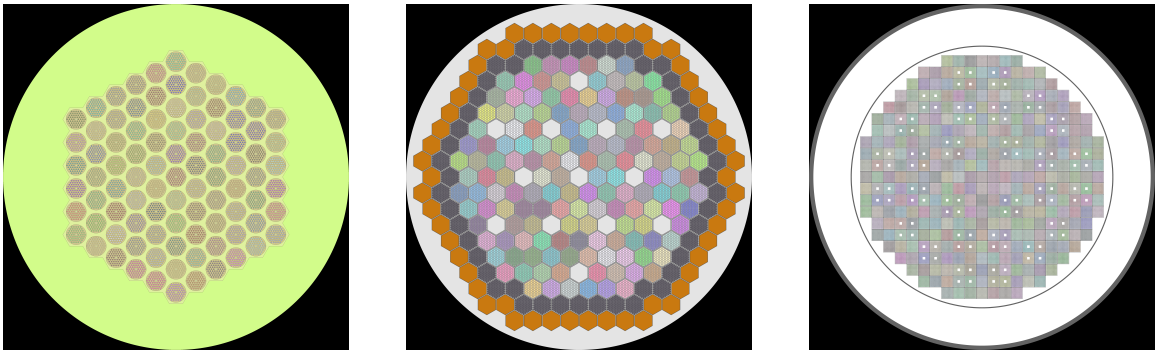

Figure 3.1: Reactor geometries for the FHR (left), AFR-100 (middle), and LFR-FCR1 (right)

Table 3.1: Burnup steps used for reactor modeling in cumulative effective full power days (EFPD)

Burnup steps [EFPD]		
FHR	AFR-100	LFR-FCR1
0.1	0.1	0.1
1	1	1
15	10	180
30	365	365
90	1825	1800
180	3650	-
270	4836	-
365	5475	-
395	5749	-
-	7300	-
-	7756	-
-	9125	-
-	10950	-

Table 3.2: Isotopes of interest for reactor modeling

Isotope
U-233
U-234
U-235
U-236
U-238
Np-237
Pu-238
Pu-239
Pu-240
Pu-241
Pu-242
Am-241
Am-242
Am-242m
’ Am-243
Cm-243
Cm-244

3.1.2 Antineutrino Translation

The Huber-Mueller model [42] was used to convert the isotopic fission rates into antineutrino spectra being emitted from the core. The Huber-Mueller spectra values, in unit yield per fission and MeV, along with relative uncertainties are labeled in libraries by isotope, incident neutron energy, and antineutrino energy. The values were first multiplied by the energy bin width, 0.5 MeV, to translate the values into integrated counts, in units yield per fission. The libraries were further simplified through assuming all fissions occurred from intermediate neutrons [68] (400 keV). Since not all isotope libraries have thermal-induced neutron values, such as U-238, or fast-induced neutron values, such as Pu-239, this greatly simplifies the data while still including the isotopic antineutrino spectra variations. These final values in energy bins, b, ranging from 2 MeV to 9 MeV in 0.5 MeV increments will be listed as antineutrino yields, $AY_{i,b}$, with their relative uncertainty, $\frac{\sigma_{AY_{i,b}}}{AY_{i,b}}$.

To retain isotopic uncertainty information propagated through the final uncertainty, the isotopic fission rate matrix was multiplied by the altered Huber-Mueller library to form a dictionary of isotopic antineutrinos generated per second within the source reactor core, $ASR_{i,b}$. Antineutrinos in this form, however, have no memory; they all have the same probability of interaction with the detector independent of the fissioning isotope. With this characteristic in mind, the isotopic antineutrino generation rate is truncated into the summation antineutrino source rate, ASR_b, as seen in Equation 3.1.

$$ASR_b = \sum_{i}^{isotopes} [ASR_{i,b}] = \sum_{i}^{isotopes} [AY_{i,b} * FR_i] \tag{3.1}$$

The uncertainties for all contributing factors to the antineutrino generation rate must be propagated forward. While the antineutrino yield libraries and isotopic fission rates are the only explicit values used to determine the antineutrino generation rate, we often see the antineutrino generation rate as a linear function of reactor thermal power, P_{th} [11]. This relationship between thermal power and antineutrino generation indicate any relative

change in the reactor power would lead to the near-identical change in the antineutrino generation rate. This trend is captured by including a relative reactor power uncertainty, $\frac{\sigma_{P_{th}}}{P_{th}}$, of 0.007. This value is based on the 0.7% capability of current flow meter technology for nuclear reactors [66]. The relative antineutrino source rate variance, $(\frac{\sigma_{ASR_b}}{ASR_b})^2$, is then calculated following Equation 3.2.

$$(\frac{\sigma_{ASR_b}}{ASR_b})^2 = \frac{1}{ASR_b^2} \sum_{i}^{isotopes} [(\frac{\sigma_{P_{th}}}{P_{th}})^2 + (\frac{\sigma_{AY_{i,b}}}{AY_{i,b}})^2 + (\frac{\sigma_{FR_i}^2}{FR_i^2})] \tag{3.2}$$

While the uncertainty must propagate forward to conclude the final system confidence, the influence of fission rate uncertainty, antineutrino yield uncertainty, and reactor thermal power uncertainty must be stored for a comprehensive uncertainty analysis. To account for their significance to the final uncertainty, fractional values were calculated to represent their contribution to the propagated antineutrino generation rate, as seen in Equation 3.3 and Equation 3.4 for the fission rate fraction, FRF_i, and power fraction, PF_i, respectively. With two of the three contributing sources of uncertainty stored, we can calculate the final fraction, the antineutrino yield fraction, by exclusion. As seen in the normalized contribution function, Equation 3.5, any contributing factor can be isolated if the other fractions and antineutrino generation rate are set.

$$FRF_i = \frac{(\frac{\sigma_{FR_i}^2}{FR_i^2})}{(\frac{\sigma_{P_{th}}}{P_{th}})^2 + (\frac{\sigma_{AY_{i,b}}}{AY_{i,b}})^2 + (\frac{\sigma_{FR_i}^2}{FR_i^2})} \tag{3.3}$$

$$PF_i = \frac{(\frac{\sigma_{P_{th}}}{P_{th}})^2}{(\frac{\sigma_{P_{th}}}{P_{th}})^2 + (\frac{\sigma_{AY_{i,b}}}{AY_{i,b}})^2 + (\frac{\sigma_{FR_i}^2}{FR_i^2})} \tag{3.4}$$

$$1 = \frac{1}{\sigma_{ASR_b}^2} \sum_{i}^{isotopes} [FRF_i * \sigma_{ASR_{i,b}}^2 + PF_i * \sigma_{ASR_{i,b}}^2 + (1 - FRF_i - PF_i) * \sigma_{ASR_{i,b}}^2] \tag{3.5}$$

3.1.3 Detector Effects

The antineutrino generation rate must then be processed to account for actual antineutrinos collected from the source. The spectra translation follows Equation 3.6 [8], in which SDR_b is the antineutrino source detection rate for energy bin b, L is the distance between the detector and the reactor, n_{prot} is the number of target protons within the detector, $IBD_{xsec,b}$ is the inverse-beta decay cross section averaged over each energy bin b, ϵ is the detector efficiency, and ASR_b is the antineutrino generation rate for energy bin b. The number of target protons were calculated following Equation 3.7, in which ρ_{prot} is the proton density of the detector, V_{det} is the detector volume, m_{det} is the detector mass, and G_{det} is the detector specific gravity. The detector values of interest are listed in Table 2.1. The distance between the detector and the reactor is dependent on the size of the reactor. The smaller reactor designs, the FHR and AFR-100, can likely accommodate for a closer antineutrino detector compared to the larger LFR-FCR1. For this reason, 17 m were used for the smaller reactor cases while 25 m were used for the larger reactor cases [8].

$$SDR_b = \frac{1}{4\pi L^2} n_{prot} IBD_{xsec,b} \epsilon * ASR_b \tag{3.6}$$

$$n_{prot} = \rho_{prot} * V_{det} = \rho_{prot} * \frac{m_{det}}{G_{det}} \tag{3.7}$$

Since the detector translation equation is relatively simple with values multiplying, the propagation of error equation [69] can be used to derive Equation 3.8 for the relative antineutrino source detection rate variance, $(\frac{\sigma_{SDR_b}}{SDR_b})^2$. In this equation, $\frac{\sigma_L}{L})^2$ is the relative distance variance, $(\frac{\sigma_{n_{prot}}}{n_{prot}})^2$ is the relative number of target protons variance, $(\frac{\sigma_{IBD_{xsec,b}}}{IBD_{xsec,b}})^2$ is the inverse-beta decay cross section variance, $(\frac{\sigma_\epsilon}{\epsilon})^2$ is the detection efficiency variance, $(\frac{\sigma_{IBD_{fit,b}}}{IBD_{xsec,b}})^2$ is the inverse-beta decay fit variance, and $(\frac{\sigma_{ASR_b}}{ASR_b})^2$ is the antineutrino source rate variance.

$$(\frac{\sigma_{SDR_b}}{SDR_b})^2 = 4(\frac{\sigma_L}{L})^2 + (\frac{\sigma_{n_{prot}}}{n_{prot}})^2 + (\frac{\sigma_{IBD_{xsec,b}}}{IBD_{xsec,b}})^2 + (\frac{\sigma_\epsilon}{\epsilon})^2 + (\frac{\sigma_{IBD_{fit,b}}}{IBD_{xsec,b}})^2 + (\frac{\sigma_{ASR_b}}{ASR_b})^2 \quad (3.8)$$

The distance between the detector and reactor will likely be known up to 1 cm (1σ uncertainty) based on well documented construction dimensions and placing the detector against the containment building wall to minimize standoff distance [51]. While the variation in scintillation mass is not listed in EJ-309 documentation [70], we can assume a 1% relative number of proton uncertainty is a conservative estimate based on frequent detector calibration periods [51]. Current inverse-beta decay cross sections are reported to be accurate to 1.4% (1σ) on average [71]. The relative detector intrinsic efficiency uncertainty is set to a conservative 1% as well to account for poorly-characterized background and system degradation [53]. Lastly, the inverse-beta decay cross sections had to be fit and integrated to match the energy bins set by the Huber-Mueller spectra libraries. A simple second-order polynomial was used to fit the cross section dataset and the integrated fitting error was returned via the scipy.integrate.quad function in Jupyter Notebook.

3.1.4 Background

The uncertainties up to this point are based on systematic variations that can be controlled and studied further for mitigation; for example, the Monte Carlo modeling can use more particles and run longer and the antineutrino libraries can be further studied to lower the isotope-specific antineutrino spectra to lower the overall uncertainty. That said, there are some aspects of radiation detection, such as background and counting statistics, that are unavoidable. While we can partially shield from background radiation and can be relatively certain in background events collected [50], any background events will add on to the uncertainty from counting statistics [62].

The background counts were based on results from the PROSPECT experiment at HFIR

[52]. The background values had to be modified to fit the desired energy bins and scaled to account for the different fiducial mass of the RETINA detector. A simple ratio between the RETINA detector, which is designed to be 5 tons, and the PROSPECT AD-I detector, which is 1.48 tons, was used to scale the background linearly. While the background is specific to HFIR, the RETINA system in application would ideally be calibrated to the location during a reactor-off period and/or would be verified through simulation [50].

The gross background detection rate, BDR_b, can be added to the source detection rate outlined in Equation 3.6 to form a total antineutrino detection rate. For deviation analysis, more useful values are the cumulative antineutrino counts collected based on the number of detectors, n_{det}, and collection period, $t_{collect}$, as seen in Equation 3.9 with associated variance in Equation 3.10. Once again, an uncertainty fraction, SF_b, was used to store the significance of the source detection rate on the final antineutrino spectra as seen in Equation 3.11. The influence of the background detection rate can be deduced by subtracting the source detection rate fraction from 1. The RETINA system is assumed to have 3 antineutrino detectors and to have a collection period of 3 months.

$$\mu_b = n_{det} * t_{collect} * (SDR_b + BDR_b) \tag{3.9}$$

$$\sigma^2_{\mu_b} = (n_{det} * t_{collect})^2 * (\sigma^2_{SFR_b} + \sigma^2_{BDR_b}) \tag{3.10}$$

$$SF_b = \frac{(\frac{\sigma^2_{SFR_b}}{SFR^2_b})}{(\frac{\sigma^2_{SFR_b}}{SFR^2_b}) + (\frac{\sigma^2_{BDR_b}}{BDR^2_b})} \tag{3.11}$$

3.1.5 Counting Statistics

Lastly, for antineutrino deviation analysis, the counting statistic nature of radiation detection will play a role when the RETINA system is placed next to a reactor core. In practice,

the onsite RETINA system will collect some measured value of antineutrino counts, n_b, with a counting statistic uncertainty, σ_{n_b}. This value and uncertainty can be folded into the current RETINA system simulated structure by using Equation 3.12. For a confidence analysis, however, there is no measured value to derive the counting statistic. The counting statistic uncertainty can be assumed to some extend using the simulated spectra, as seen in Equation 3.13, but will break from Gaussian statistics due to the quantity-dependency of counting statistics.

$$Deviation_b = (\mu_b \pm \sigma_{\mu_b}) - (n_b \pm \sigma_{n_b}) \tag{3.12}$$

$$Deviation_b = \mu_b \pm \sqrt{(\sigma_{\mu_b}^2 + n_b)} - n_b \tag{3.13}$$

To elaborate, the propagated simulated uncertainty follows Gaussian statistics so there is an equal probability (68% [18]) that the number of antineutrinos collected will fall either above or below 1σ deviation from μ_b. If the counting statistics, considered $\sqrt{n_b}$ [18], are based off these values, however, the larger values would be considered more uncertain than the lower values. For example, if μ_b was set to be 100 and σ_{μ_b} was set to 10, there is a 68% that the antineutrino counts fall between 90 and 100 as well as 100 and 110. From collection, however, the standard deviation would be $\sqrt{90}$ for 90 counts and $\sqrt{110}$ for 110 counts.

This distribution can be accurately modeled through Monte Carlo methods. To do so, sample points are generated from a normal distribution with μ_b set as the mean and σ_{μ_b} as the standard deviation. From this dataset, the sample points can be randomized from a new normal distribution that matches that sampled point's current value, as seen in Equation 3.14 with the processed sample points, s_b. The final values of these points would represent the distribution expected through the propagated uncertainty and counting statistics.

$$n_b \sim N(\mu_b, \sigma_{\mu_b})$$
$$s_b \sim N(n_b, \sqrt{n_b}) \tag{3.14}$$

While the stochastic method can be used to form an approximation of the exact distribution, the significance of the propagated uncertainty and counting statistics to the final confidence becomes convoluted. Specific error bars would need to be formed on both sides of mean and evaluated independently for each component of the uncertainty analysis. Since these results would not be as interpretable as standard Gaussian statistics, the Monte Carlo results were fit to a Gaussian to generate a distribution that closely reflects the original distribution.

The Gaussian fit helped better shape the distribution, but the initial Monte Carlo sampling became too computationally intensive. Determining the influence of each aspect would require a significant amount of computational resources compared to a statistical approach. To counteract this computational demand, the Monte Carlo approach was bypassed for adding the variances of the propagated uncertainty and the mean counting statistics. This result would be slightly skewed due to the counting statistics not fitting the mean exactly but is useful to determine our confidence in antineutrino yields.

3.1.6 Figures of Merit

The confidence in antineutrino yield is evaluated by the following metrics:

1. Component Significance to Confidence: what RETINA system components are the most influential towards our antineutrino yield uncertainty and to what extent?

2. Reactor Significance to Confidence: how do different reactor design parameters, such as power level and initial reactor inventory, impact the antineutrino yield confidence

for different factors of uncertainty?

3. Temporal Significance to Confidence: how do does the reactor burnup influence the antineutrino yield confidence for different factors of uncertainty?

3.2 Confidence in Diversion Scenario Detection

The purpose of the RETINA system is to safeguard nuclear reactors. Before implementation, the system will need to prove it can effectively and consistently detect SNM being removed from a core within a reasonable amount of time. Due to the innate wide distribution of antineutrino yields, a probability threshold is set to classify diversion scenarios, or scenarios in which SNM has be removed from the core, from reference scenarios, or scenarios in which the core includes the reported inventory. The RETINA system's detection confidence is defined as the probability that SNM has been removed from the core based on the antineutrino yield. The confidence in the system is tested for a variety of diversion scenarios to determine RETINA system sensitivity limits. This work is outlined in Dunbrack *et al.* [72].

3.2.1 Diversion Scenarios

Diversion scenarios were initially selected by Stewart [51]. The AFR-100 reactor design was chosen to test RETINA system diversion confidence. This reactor design is of particular interest for sensitivity limits since it is considered a small reactor ($< 300MWe$ [6]), meaning there will be a considerably small antineutrino yield for diversion analysis compared to a large reactor, and the reactor breeds a significant quantity (SQ) of plutonium by the end of the reactor's lifecycle [59]. An SQ of plutonium, 8 kg, is defined as the amount of nuclear material in which manufacturing a nuclear weapon cannot be excluded [73]. While an SQ of lower enriched uranium (LEU), 75 kg U-235, could also be removed from the AFR-100 core in a diversion scenario, the smaller SQ of plutonium would allow an adversary to take a smaller quantity of SNM from fewer assemblies. A variety of diversion

scenarios were simulation to account for adversaries' different priorities and resources. Regardless of the scenario, a total of 1 SQ of plutonium was simulated being removed from the core.

The AFR-100 reactor design is outlined in Figure 3.3. As previously mentioned in subsection 2.4.2, the reactor initially burns U-235 as the fuel source in the outer portion of the core. As the reactor ages, plutonium is generated throughout the core based on U-238 absorbing a neutron. Due to the larger U-238 inventory in the inner core, an SQ of plutonium is first available in a single assembly towards the center of the core at 15.75 EFPY. If timing and assembly accessibility are important diversion factors to an adversary, a scenario in which 1 SQ of plutonium is removed from a center assembly (upper left of Figure 3.3), labeled scenario 1, is ideal.

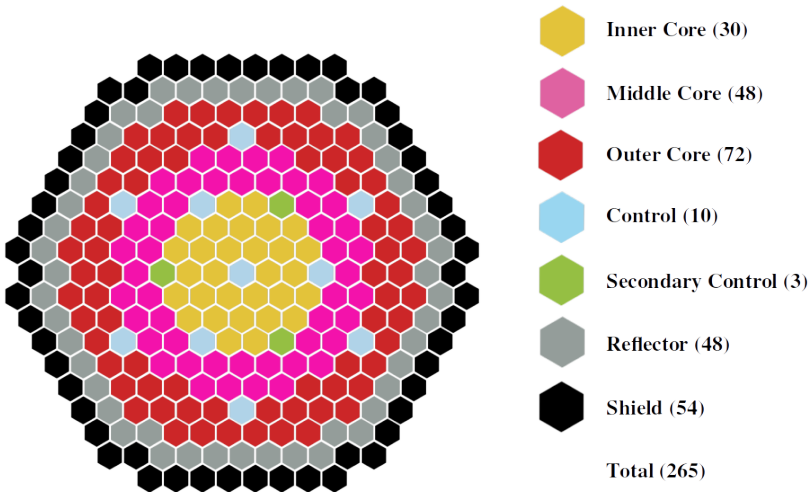

Figure 3.2: AFR-100 core layout with the numbers in parentheses corresponding to number of assemblies in a particular core region

Alternatively, there are diversion benefits if an adversary is able to pull from two assemblies and wait until a later burnup step. From breeding plutonium inwards, there is a higher fission rate around the center of the core at the end stages of the reactor lifecycle. If fuel is removed from this region, this could significantly alter the isotopic fission rate and is more likely to shift the antineutrino yield detected by the RETINA system. The outer regions, while still containing an SQ of plutonium, have a lower plutonium fission rate so are less

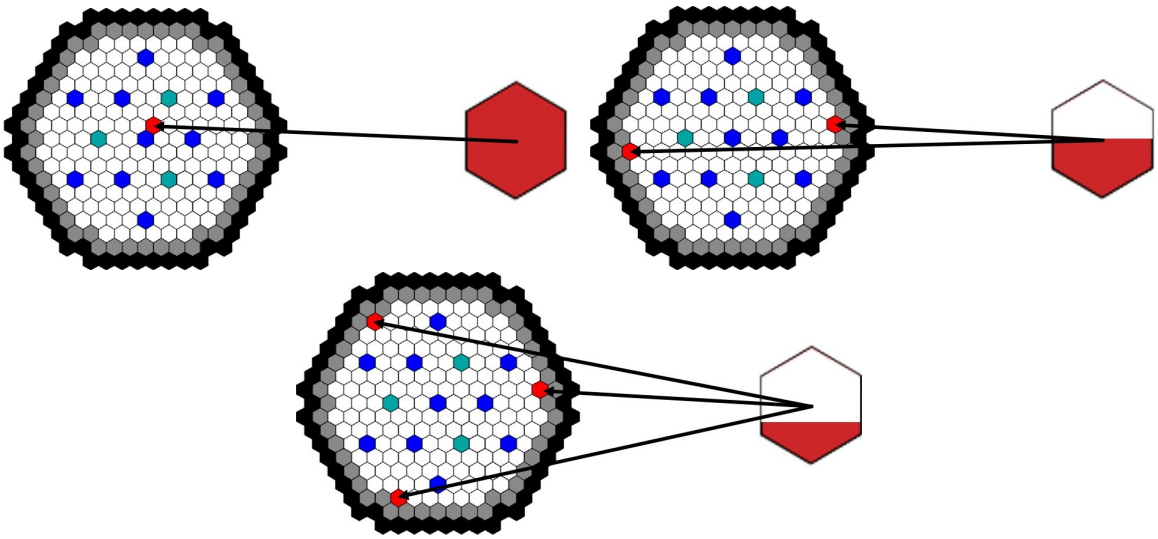

Figure 3.3: Diversion scenarios assembly location 1 (upper left), 2 (upper right), and 3 (lower center) in which 8 kg of plutonium total was removed from the core. Figures adapted from Stewart [51]

likely to alter the antineutrino yield if plutonium is removed. While an adversary would have to remove 1/2 of an SQ of plutonium from each assembly since no single assembly has enough plutonium for a single-assembly-diversion, the plutonium vector in these assemblies are more pure compared to the inner region. An ideal plutonium vector has a maximal amount of Pu-239 and Pu-241 and a minimal amount of Pu-238 and Pu-240. For plutonium processing, Pu-238 is an undesired hazard since it is not fissile and generates an energetic alpha particle in its fairly active decay. Pu-240 is also undesirable since it has a significant spontaneous fission rate, meaning any bomb containing a large percentage of Pu-240 would fissile out on detonation. The plutonium vectors for each scenario are listed in Table 3.3. Scenario 2 is set at 21.25 EFPY and illustrated in the upper right portion of Figure 3.3.

The final diversion scenario location variation, scenario 3, is the most ideal for an adversary. If an adversary is able to access three cores, in which 1/3 of an SQ of plutonium is removed, the scenario could happen early on in the fuel lifecycle, 13.25 EFPY, while still including a low outer plutonium fission rate and pure plutonium vector retention. This scenario is seen in the bottom center of Figure 3.3.

Table 3.3: Plutonium isotopic composition for various diversion scenario locations

Plutonium Removed	Scenario Vector 1	Scenario Vector 2	Scenario Vector 3
Pu-238	0.31	0.15	0.08
Pu-239	93.93	97.00	98.15
Pu-240	5.55	2.80	1.75
Pu-241	0.20	0.05	0.02
Pu-242	0.01	<0.01	<0.01

Along with diversion scenario location, an adversary can customize a scenario with fuel replacement. Fuel replacement refers to the nuclear material that an adversary can place in the diverted assembly to fill the void. While there are endless possibilities for fuel replacement, natural uranium (NU), labeled fuel replacement a, is seen as a reasonable fuel replacement since it is likely easily available to an adversary with assembly access and contains fissionable material. In a more ideal diversion scenario to the adversary, LEU would better reflect the fissile inventory of the assembly. LEU with an enrichment based on the average of that region was selected for fuel replacement b.

The last diversion scenario variable altered was power manipulation. While replacement fuel will mask the missing plutonium fuel to some extent, there will be a change in the antineutrino yield due to different amount and isotopic inventory of the fuel. This alteration will lead to either an increase or decrease in the overall antineutrinos collected by the RETINA system. This behaviour, however, in the overall counts can be mitigated through power manipulation. For example, increasing the reactor power correlates to more fuel burning, which correlates to generating more antineutrinos. If an adversary wanted to mask a diversion scenario that would typically lead to a decrease in overall antineutrino counts, they can increase the reactor power to generate more antineutrinos and match the RETINA system's expected antineutrino yield. For diversion scenario variations, nominal and manipulated power levels were simulated, following the labeling convention listed in Table 3.4.

Table 3.4: Diversion scenarios labels by assembly diversion locations (1, 2, and 3), fuel replacement (a for NU and b for LEU), and power manipulation (nominal and manipulated)

Assembly Location	Fuel Replacement	Power Manipulation
Scenario 1	a	Nominal
Scenario 1	a	Manipulated
Scenario 1	b	Nominal
Scenario 1	b	Manipulated
Scenario 2	a	Nominal
Scenario 2	a	Manipulated
Scenario 2	b	Nominal
Scenario 2	b	Manipulated
Scenario 3	a	Nominal
Scenario 3	a	Manipulated
Scenario 3	b	Nominal
Scenario 3	b	Manipulated

Nominal power level refers to the power level matching the expected power level perfectly while manipulated power level refers to the optimal power level to match the expected antineutrino yield. The optimal power manipulation value was set through Equation 3.15, in which a hyperparameter, x, is set to minimize the Chi-square value, χ^2, for an expected antineutrino yield, n_b, a diverted antineutrino yield, n_b', and the propagated uncertainty at nominal power, σ_{norm}. This equation reflects the Chi-square goodness of fit test, in which the deviated antineutrino yield is directly compared to the expected yield for variation, but includes a penalty term, $(\frac{x}{\sigma_{norm}})^2$. This penalty term is added to account for reactor power variation possibly setting off other safeguard measures. The strength of this term is dependent on how drastically the adversary alters the antineutrino spectra, x, and uncertainty of this spectra, σ_{norm}.

$$\chi^2 = (\sum_b \frac{(n_b - (1+x)n_b')^2}{n_b}) + (\frac{x}{\sigma_{norm}})^2 \tag{3.15}$$

3.2.2 System Variation

Diversion scenario detection confidence has been studied in Stewart [51]. The diversion scenarios were simulated using a joint-toolkit system of MC^2-3 (MCC) [74] and REBUS-3 [75]. MCC utilizes cross section libraries and well-defined reactor geometries to produce spatial and energy dependent effective microscopic cross sections for full-core neutronics. The deterministic code, REBUS, was then used for neutronics, which directly reports the isotopic fission rate, and for depletion calculation, which updates the macroscopic cross sections then altered by MCC for the next burnup steps. This process, outlined in Figure 3.4, was iterated for every 0.25 EFPY steps for the entirety of the AFR-100 fuel lifecycle (30 EFPY) [51].

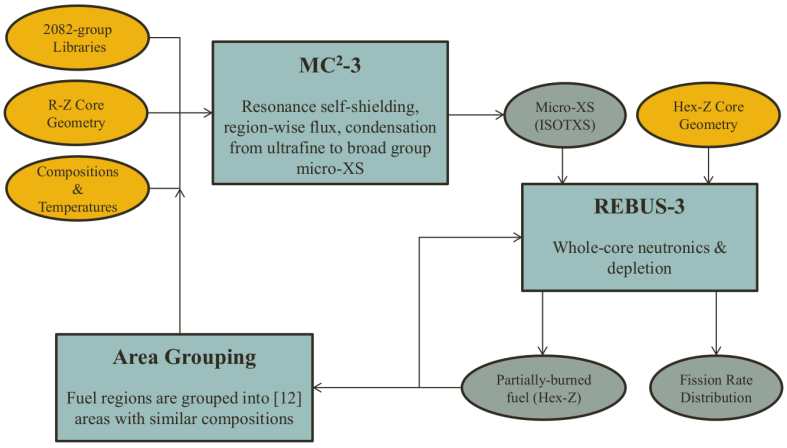

Figure 3.4: Reactor modeling calculation scheme with rectangles indicating codes, ovals indicating data sets, gold indicating user-defined data, and grey indicating calculated data. Figure from Stewart [51]

Uncertainty within this process was complex due to the iterative nature of the joint-toolkit system. Any deviation in microscopic cross sections can shift the isotopic fission rate which then would alter the cross sections in the next burnup step. Accounting for this complex iterative behaviour, a Monte Carlo stochastic method, sampling from thermal power variations, microscopic cross sections, and burnup histories, was developed for system convergence. After sampling from 1000 histories, the AFR-100 isotopic fission rates

were well converged. This process is described in more detail in Stewart [51].

While this process was sufficient in Stewart, SERPENT2 (SERPENT) was used for reactor modeling and isotopic fission rate calculations for the antineutrino yield confidence study. SERPENT, as described in subsection 3.1.1, is a Monte Carlo-based tool that includes both a neutronics and depletion tool. The SERPENT-resulting antineutrino yield associated with the RETINA system is in agreement with the MCC/REBUS-resulting antineutrino yield (Figure 3.5). Considering the MCC/REBUS system to be a more comprehensive system, based on smaller burnup steps and a more generalized uncertainty propagation system, all SERPENT-driven values fall well within 1σ for the burnup of interest (Figure 3.6). The MCC/REBUS derived antineutrino yields are assumed for the entirety of the diversion detection confidence portion of this work.

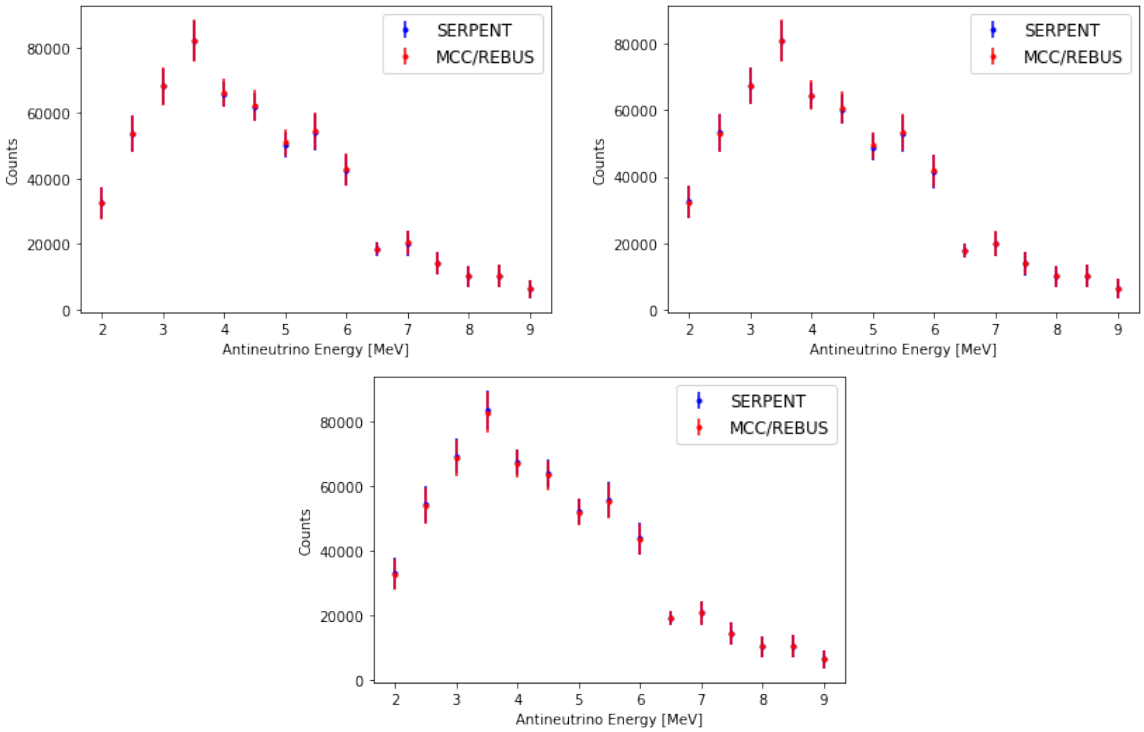

Figure 3.5: Antineutrino yield and associated uncertainty (1σ) for an AFR-100 reactor core at 15.75 EFPY (upper left), 21.25 EFPY (upper right), and 13.25 EFPY (bottom center) using SERPENT and MCC/REBUS toolkits for reactor modeling

The diversion detection analysis system also has been altered since the work outlined in

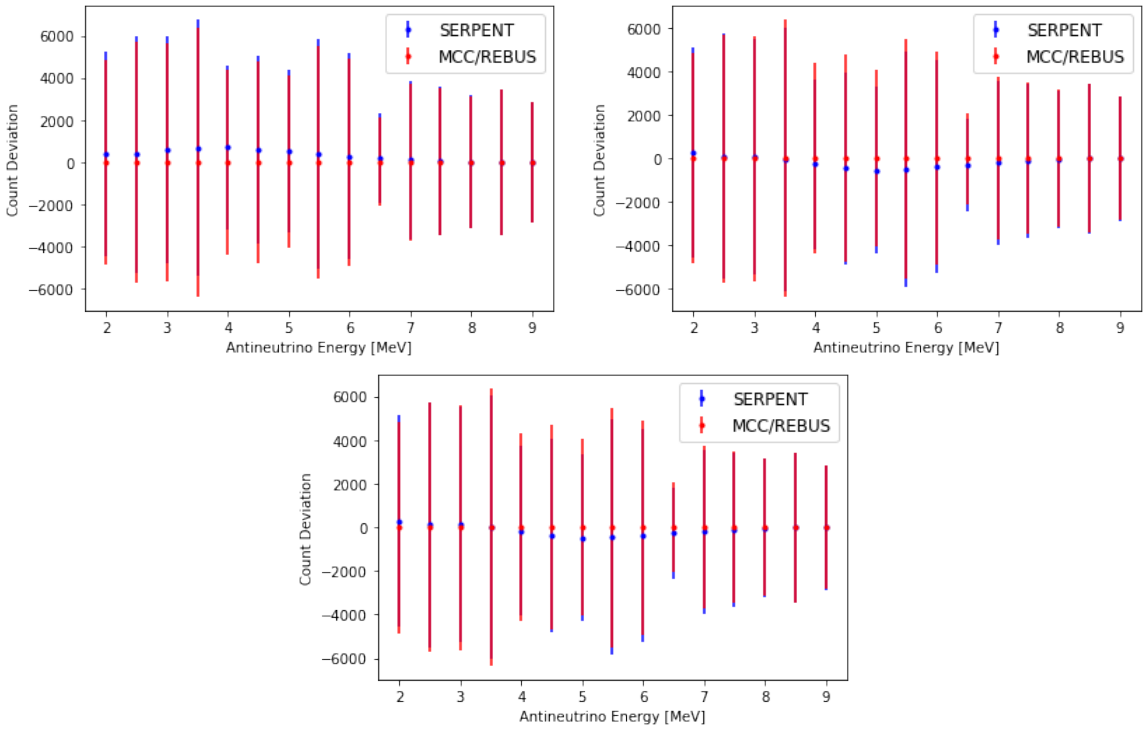

Figure 3.6: SERPENT-based antineutrino yield deviation from MCC/REBUS-based antineutrino yield for an AFR-100 reactor core at 15.75 EFPY (upper left), 21.25 EFPY (upper right), and 13.25 EFPY (bottom center)

Stewart [51]. For diversion detection variance, the previous system included the variance of a Chi-square goodness of fit test for the deviation distribution (Equation 3.16) [76]. In this equation, T_0 is the χ^2 value defined in Equation 3.15 for a given diversion scenario and reference case. A cumulative distribution function, Φ, is then taken in terms of a critical false positive rate value based on the Chi-square distribution, T^{α}_{crit}, the expected Chi-square value, and the Chi-square standard deviation, $2\sqrt{T_0}$ (Equation 3.17). The value of T^{α}_{crit} depends on the allowable false positive rate, or the rate in which the RETINA system reports a diversion scenario when one has not taken place, set by the user. The resulting value, β, represents the false negative rate, or the rate in which the RETINA system does not report a diversion scenario when one has taken place. Our detection confidence in the RETINA system is defined by the probability of the RETINA system to not report a false negative, described in Equation 3.18.

$$T \sim N(T_0, 2\sqrt{T_0}) \tag{3.16}$$

$$\beta = \Phi(\frac{T^{\alpha}_{crit} - T_0}{2\sqrt{T_0}}) \tag{3.17}$$

$$Detection\ Confidence = 1 - \beta \tag{3.18}$$

This deviation analysis method does not account for complex bin-to-bin relationships. Initially, from the starting deviation Chi-square values (Equation 3.15), the system only accounts for the deviation from each antineutrino yield bin to the antineutrino yield bin of the reference case. There are cases, such as with reactor power manipulation, when this general deviation is at a minimum. Inadvertently, however, this minimization causes this oscillation effect in which the upper some energy bins overestimate yields while other bins underestimate yields (Figure 3.7). A more complex method is needed to process the antineutrino yields for bin-to-bin comparisons.

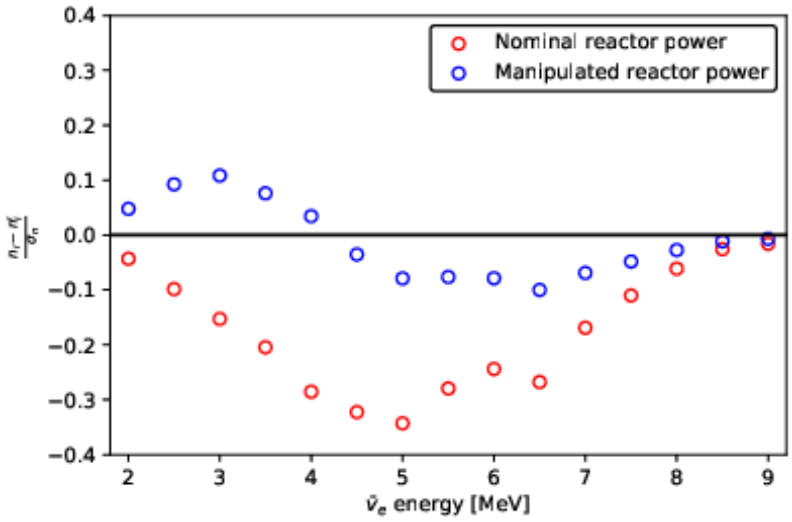

Figure 3.7: Oscillation effect from goodness-of-fit function minimization for diversion 1a

3.2.3 Support Vector Machine Modeling

Machine learning (ML) models have been used in many fields to process complex datasets [77]. A support vector machine model (SVM) is used in this work to classify the antineutrino yield as a diversion scenario or a reference scenario. Specifically, a soft-marginalized SVM is applied to allow for some misclassification to generate a more robust system. The model iterates over different kernel weights, **w**, to minimize Equation 3.19. For the antineutrino dataset, since the yields go from 2 MeV to 9 MeV in 0.5 MeV increments, the weights are stored in a 15 feature vector. The left hand portion of the equation, $C * \Sigma_{i=1}^{N} max(0, 1 - L(y, f(\mathbf{x}_i)))^2$, is a penalty term for any misclassified values. The extent of this penalty depends on the hyperparameter, C, which scales the loss function relative to the right hand term. The loss function includes the actual data classifier, y, and the kernel classifier, $f(\mathbf{x}_i)$ [78]. There are many forms of kernel classifiers, such as linear, polynomial degree 2, polynomial degree 3, radial basis function, and sigmoid [79]. The kernel type is selected to fit the shape of the dataset. Since the shape of the dataset is unknown, all 5 types of kernels are fit to diversion scenario 1a for selection. A squared hinge loss function was chosen for the system loss function.

$$\min_{w}[C * \Sigma_{i=1}^{N} max(0, 1 - L(y, f(\mathbf{x}_i)))^2 + \frac{1}{2}||\mathbf{w}||^2] \tag{3.19}$$

The right hand portion, $\frac{1}{2}||\mathbf{w}||^2$, describes the spread of the soft-margin. Ideally, the diversion and reference antineutrino yields will be fairly different. If the clusters are far apart, however, there are many ways to draw the SVM classification boundary while still classifying the data perfectly. In Figure 3.8, the classification boundaries could be placed right next to either the red or green cluster while still providing perfect classification. The second weighted term, $||\mathbf{w}||$, minimizes the distance between the boundary and the clusters. The entirety of the equation is differentiable so with each model iteration the weights become more inline with the optimal classification boundary.

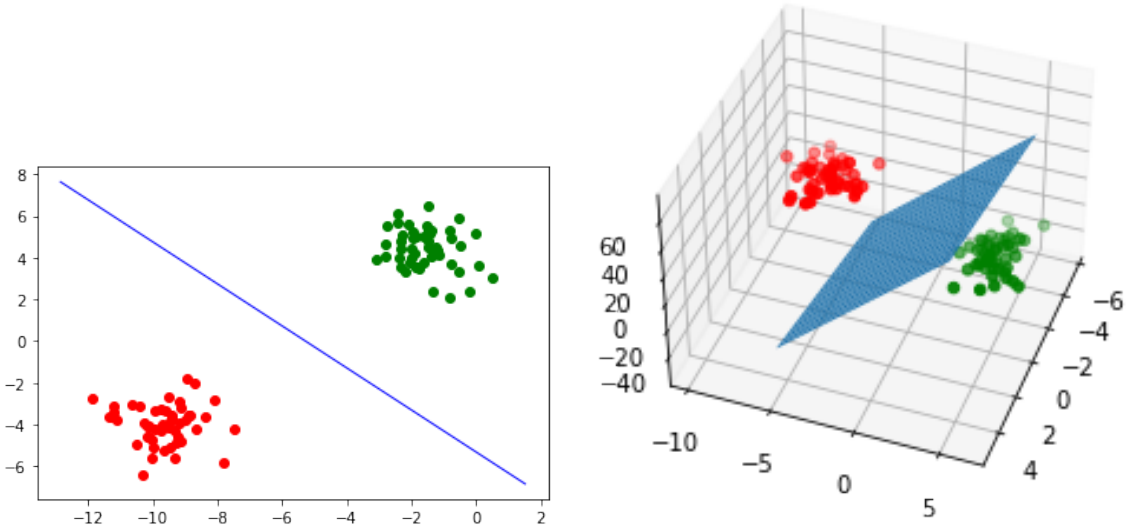

Figure 3.8: Visual of a support vector machine applied in 2 (left) and 3 (right) dimensions

Once the SVM weights are set, test data is processed through the SVM model for classification values. While a classification value of -1 represents a diverted scenario spectra and 1 represents a reference scenario spectra, most values fall somewhere in between the boundaries. The classification values for both the diversion scenario and reference scenario are placed in a histogram and fit to Gaussian distributions to mitigate the stochastic fluctuation of sampling. A $T^{\alpha}_{cv,crit}$ classification boundary is set to match a user-defined false

positive rate limit. Similar to the statistical Chi-square approach, a cumulative distribution function (Equation 3.20) is used to isolate the false negative rate of the diversion-based Gaussian distribution. Unlike the Chi-square approach, however, the reference case has an associated false negative rate (Equation 3.21) that must be subtracted out. The confidence is defined as the probability that the RETINA system detected a diversion scenario above the reference standard. Without this definition distinction in this formula (Equation 3.22), the system diversion detection confidence would always be greater than 0 regardless of any applied diversion scenarios.

$$\beta_{cv} = \Phi(\frac{T^{\alpha}_{cv,crit} - \mu_{cv}}{\sigma_{cv}}) \tag{3.20}$$

$$\beta'_{cv} = \Phi(\frac{T^{\alpha}_{cv,crit} - \mu'_{cv}}{\sigma'_{cv}}) \tag{3.21}$$

$$Detection\ Confidence = Confidence_{cv} - Confidence'_{cv} = \beta'_{cv} - \beta_{cv} \tag{3.22}$$

3.2.4 Data Selection

To use ML, a model must be trained for a sampled dataset. This dataset can easily retain the antineutrino yield distribution information through selective sampling. These samples, however, would be widely deviating from the mean compared to the actual deviation between the reference and scenario cases. As illustrated in Figure 3.9, the propagated uncertainty (1σ) spans far further than the actual deviation between the distributions. The SVM would likely not converge to a meaningful classification boundary if there is considerable variation in the dataset [80]. To aid in model convergence, samples are generated from a more idealistic distribution from source and background counting statistics. The difference in antineutrino spectra is then more significant compared to the overarching uncertainty in

the system, as seen in Figure 3.10, While there are other sources of uncertainty, counting statistics represent an unavoidable hurtle that the RETINA system must surpass for diversion detection confidence.

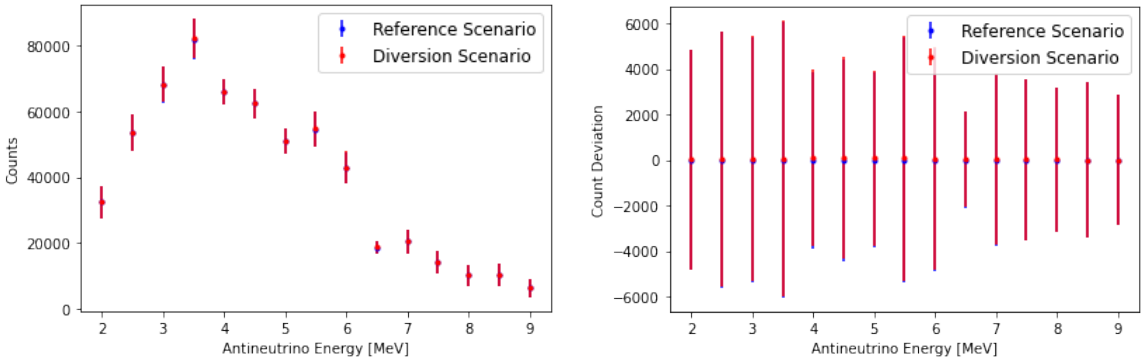

Figure 3.9: Antineutrino yield (left) and reference-based deviation (right) for diversion scenario 1a nominal assuming all sources of uncertainty

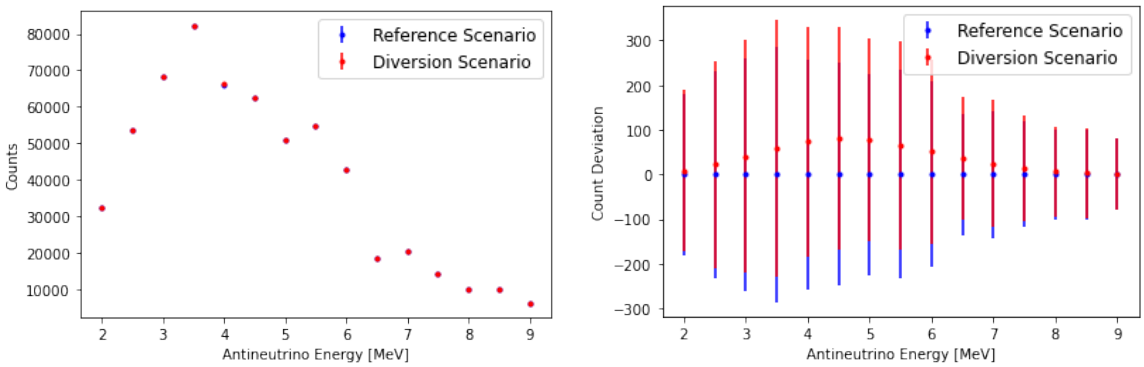

Figure 3.10: Antineutrino yield (left) and reference-based deviation (right) for diversion scenario 1a nominal assuming only counting statistics as the source of uncertainty

The generated dataset was then further processed to scale the importance of each feature. With each energy bin being a feature to the model, the spectra needs to be compressed to indicate significant deviations. For example, deviating by 1000 counts might not be significant in the 3.5 MeV bin, where around 80,000 counts are expected, but this deviation might be significant in the 9, where around 5,000 counts are expected. To scale the features according to their expected deviation and uncertainty, the standard score, $\frac{s_i - n'_i}{\sigma_s}$, in which s_i is the sample antineutrino yield, n'_i is the expected antineutrino yield, and σ_s is the sample

uncertainty, was taken for every count. Features that match the expected value would be near 0 while the highly-deviating values would be near 1.

There is also customization that can be made in ML model training. The SVM models are most well-prepared for datasets that reflect the selected training data. For example, if an SVM model was trained from diversion scenario 1a nominal data, it would likely classify diversion scenario 1a nominal testing data with more confidence than any other scenario. The RETINA system, however, should be prepared for all diversion scenarios. To make a more robust system, datasets from multiple scenarios can be used for SVM model training. While the complex training data scheme would better prepare the RETINA system for a variety of scenarios, the system would be more generalized to fit all scenarios.

Model training-testing selection was divided into four groupings: individualized models, power-ignorant models, unseen models, and power-ignorant unseen models. The individualized models represent the most idealistic RETINA system. The system utilizes 12 SVM models that are individually trained and tested for a specific diversion scenario. In practice, however, there would be no way of knowing which model is accurate without an explicit diversion scenario. Power-ignorant models include training from a grouped dataset of both nominal and manipulated power scenarios while testing from the individual 12 scenarios. As seen in Figure 3.7, manipulated power scenarios cause an antineutrino yield oscillation effect. The ML models might rely on this oscillation pattern too much for nominal power scenarios and too little for manipulated power scenarios. By grouping the power datasets together, the models would be equally-trained for both oscillation spectra and general deviation spectra.

The final two model variations, unseen models and power-ignorant unseen models, involve training from the majority of the scenarios to generate a robust system towards all possible diversion scenarios. Considering there are infinitely many ways to divert SNM from a reactor core, the RETINA system will need to safeguard not only diversions expected by the system, but scenarios never simulated for preparation, or 'unseen' scenarios.

The unseen model grouping trains the SVM models from 11 of the 12 scenarios and then tests the model against the withheld scenario. The final model grouping, the power-ignorant unseen models, train from all scenarios excluding the scenario of interest and its alternating power scenario. The model is then tested against only the scenario of interest. By withholding the power-varying scenario equivalent, the model is equally trained from both nominal and manipulated diversion scenarios. Compared to the unseen models, in which there is an imbalance with 11 training datasets, this model variation proportionally considers the oscillation pattern of the manipulated power spectra and the general deviation of the nominal power spectra.

The unseen scenarios require a 16th model feature, a timestamp, since the assembly location varying scenarios occur at 13.25, 15.75, and 21.25 EFPY. The SVM models will need to discriminate between the reference scenario spectra and the diverted scenario spectra for all three reactor ages. To transfer this information through the SVM models, the reactor age was added in as an extra feature and normalized by all timestamp options to fall between 0 and 1. There, 13.25 EFPY had a feature value of 0 and 21.25 EFPY had a feature value of 1. Linearly interpolating, 15.75 had a feature value of 0.3125. The corresponding feature values were added to the spectra vectors for every reference scenario and diversion scenario for their respective timestamp.

3.2.5 Figures of Merit

The IAEA has well-defined limits for reactor safeguard implementation. For the RETINA system to be used for low-probability scenarios, the system must have a minimum detection confidence of 0.2 (with a 0.05 false positive rate limit) within 3 months of diversion [73]. The confidence in diversion detection is evaluated by the following metrics:

1. Scenario Significant to Confidence: can the RETINA system detect any simulated diversion scenario above current IAEA detection limits?

2. Model Significance to Confidence: can the RETINA system become robust to effectively detect unseen diversion scenarios?

CHAPTER 4
RESULTS AND DISCUSSION

4.1 Confidence in Antineutrino Yield

The RETINA system was applied to a variety of reactor designs and inventories for component analysis. The final uncertainties correspond to our confidence in the RETINA system in certain situations. In the following subsection, the component, reactor, and temporal significance are outlined and discussed in terms of antineutrino yield confidence.

4.1.1 Component Significance

Each uncertainty component was propagated forward and stored to assess leading causes of uncertainty that limit the antineutrino yield confidence. The AFR-100 reactor design with a fresh core is used as a surrogate for all other reactor designs in regards to component significance. Explicit processed data for the AFR-100 fresh core antineutrino yield confidence analysis can be found in Appendix A.

The initial component of the RETINA system simulation is reactor simulation and antineutrino translation. While isotopic fission rate is an important aspect to generate the antineutrino spectra, not all isotopes are equally important to our final antineutrino yield confidence. As discussed in subsection 2.1.1, on average, a fissioning U-235 atom produces about 1.9 detectable antineutrinos while a fissioning Pu-239 atom produces about 1.5 detectable antineutrinos. With antineutrinos as the final information carriers, the uncertainty from one fissioning U-235 atom is considered about 25% more important than than the uncertainty from one fissioning Pu-239 atom. This weighted factor is accounted for by analyzing the isotopic contribution to antineutrino flux, as seen in Figure 4.1.

The antineutrino yield was then converted by the antineutrino yield uncertainty dictio-

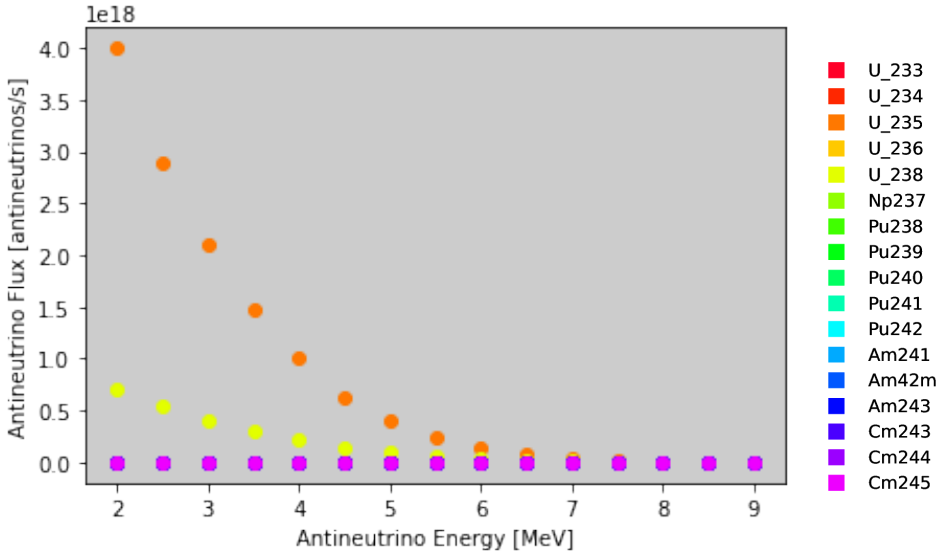

Figure 4.1: U-235 and U-238 as the significant contributing isotopes to the antineutrino flux for a fresh AFR-100 reactor core

nary, Figure 4.2, to produce the antineutrino yield uncertainty. The relative antineutrino yield uncertainty, normalized by total contribution, is seen in Figure 4.3. While U-235 fission was the most significant contributor to the antineutrino flux, the large U-238 fractional uncertainty at low (< 3 MeV) and high (> 6.5 MeV) antineutrino energies led to the total U-238 uncertainty contribution being much higher than the U-235 uncertainty contribution at low and high antineutrino energies. Since these two isotopes produced the most amount of antineutrinos for this scenario, they were the most significant isotopes that contributed to our final uncertainty.

After antineutrino translation, detector effects and background were included to predict antineutrinos collected by the RETINA system. From the values listed in subsection 3.1.3, there was no energy dependency on the uncertainty contribution from detector components. The relative detector effect uncertainty contribution is shown in Figure 4.4. The inverse-beta decay was the most significant component of system uncertainty. The number of target protons and intrinsic efficiency uncertainty were near-identical in their significance to the propagated uncertainty. Figure 4.5 depicts the detected background expected after a 3 month collection period using the RETINA system at HFIR.

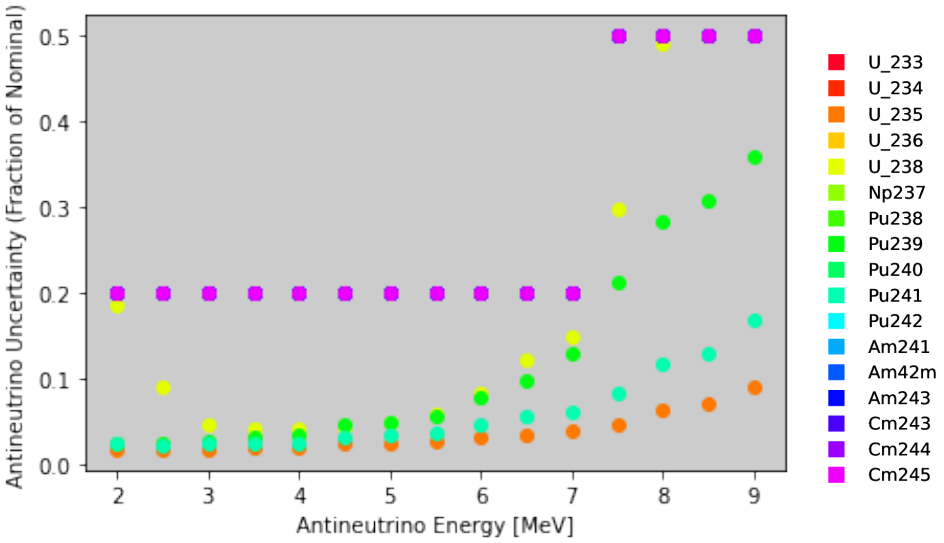

Figure 4.2: Processed Huber-Mueller antineutrino yield uncertainty dictionary

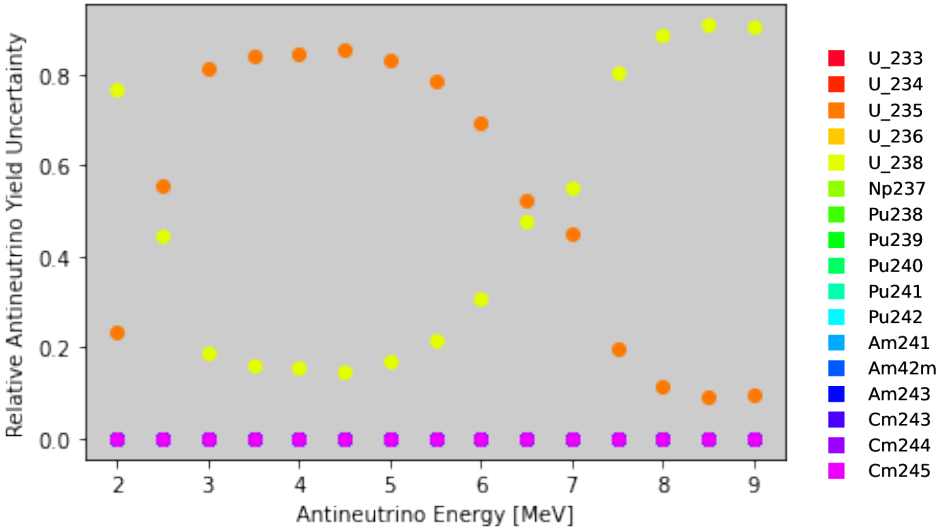

Figure 4.3: U-235 and U-238 as the significant contributing isotopes to the relative antineutrino emission rate uncertainty for a fresh AFR-100 reactor core

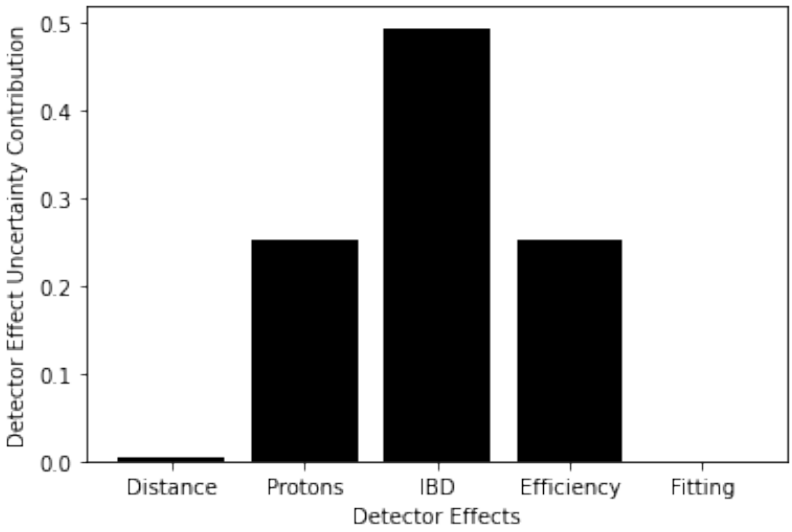

Figure 4.4: Inverse-beta decay, target proton quantity, and intrinsic efficiency as the significant contributing factor to the relative detector effect uncertainty

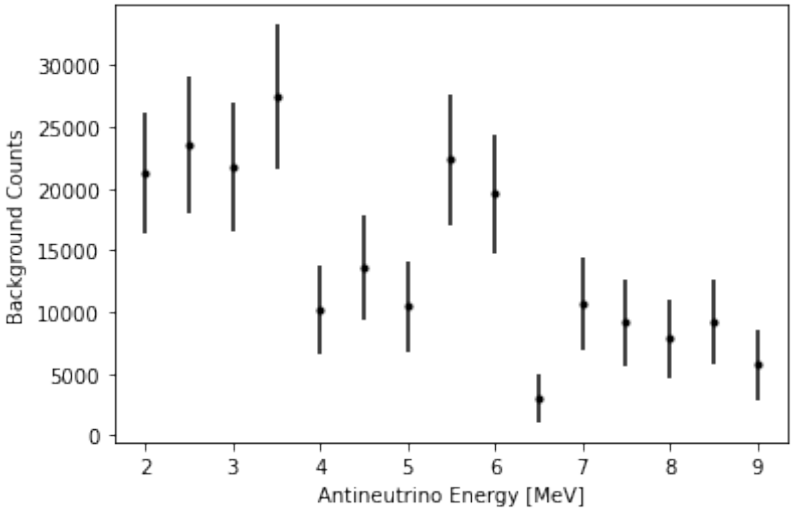

Figure 4.5: Background counts and associated uncertainty (1σ) for a RETINA system at HFIR

Counting statistics were initially applied via a Monte Carlo sampling process. The samples were then fit to a Gaussian for better interpretability. Alternatively, the variances between the propagated uncertainty and mean counting statistics were added for simplicity. The results for the 5 MeV antineutrino energy bin and a skewed example are shown in Figure 4.6. As seen in the 5 MeV histogram, the skewed nature of applying counting statics to a varying n_b in the $\sim 47,500$ region has little impact compared to the propagated uncertainty. As the region gets towards comparable $\sqrt{n_b}$ and σ_{μ_b}, the Gaussian and added variance fits tend to deviate from the actual distribution. In the example, with a $\mu_b = 4$ and a $\sigma_{\mu_b} = 2$, the Gaussian and added variance fit were virtually identical with values of $s = 4.0007$ and $\sigma_s = 2.8284$ for the Gaussian fit and $s = 4$ and $\sigma_s = 2.8345$ for the added variance fit. The fits matched to even more precise values at in the realistic regions of the antineutrino yield datasets.

The final antineutrino yield and associated uncertainty, accounting for both the propagated uncertainty and counting statistics, is shown in Figure 4.7. Since all uncertainties were stored in some manner, the contribution of every variance previously discussed was evaluated for their significance for all 15 energy bins. The uncertainties stored include the reactor thermal power uncertainty, the fission rate uncertainty from simulation, the antineutrino yield uncertainty from the Huber-Mueller spectra libraries, the uncertainty due to detector effects, the background uncertainty, the uncertainty from counting statistics for background events, and the uncertainty from counting statistics from source events. The leading 5 uncertainties can be seen in Figure 4.8. The other 2 uncertainties, power uncertainty and fission rate uncertainty, were much less significant compared to the other 5. Their values can be seen in Appendix A.

The background was the most significant contributing factor to the system uncertainty for all antineutrino energies. This was an expected trend, to some extent, considering the large error bars on the background counts and the large portion of events that come from background. A majority of the source antineutrinos are collected in around the 3.5/4 MeV

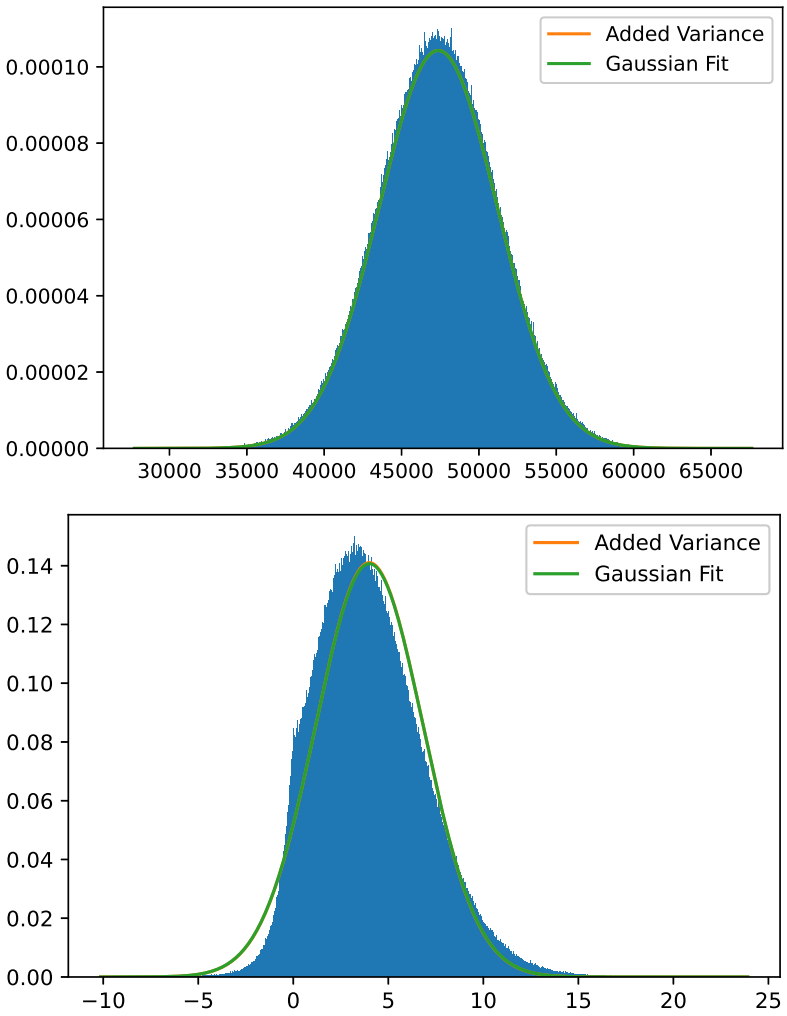

Figure 4.6: Histogram of 50,000 sampled 5 MeV antineutrino events for a fresh AFR-100 reactor core (top) and a skewed example (bottom).

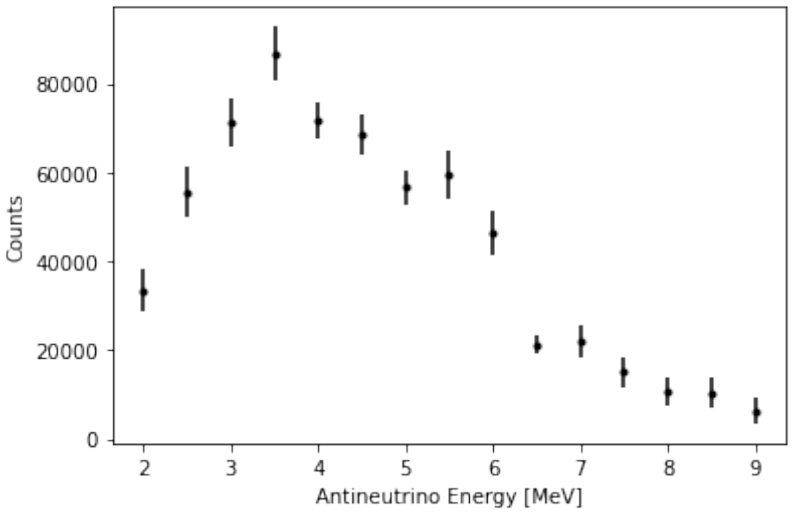

Figure 4.7: Antineutrino yield and associated uncertainty (1σ) for a fresh AFR-100 reactor core

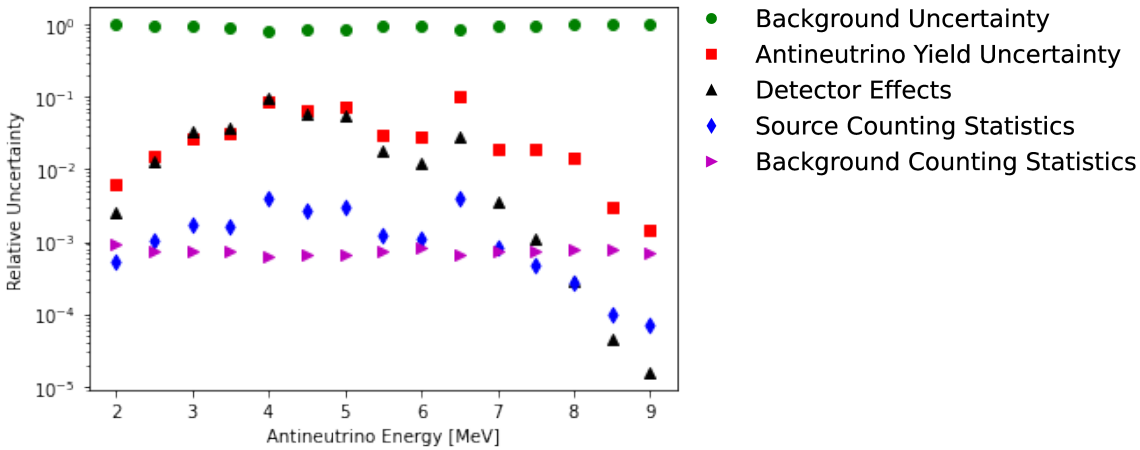

Figure 4.8: Five leading causes of uncertainty for the RETINA system 17 m away from a fresh AFR-100 reactor core

energy bin. For this reason, the background uncertainty is less significant around those bins compared to source-component uncertainties, such as those uncertainties from antineutrino yield libraries and detector effects. The signal-to-background, however, on average is fairly low ($\sim 1.5 : 1$), indicating that the background uncertainty will play a significant role in our yield confidence across the whole measured spectra.

4.1.2 Reactor Significance

The 3 reactor designs outlined in section 2.4 were simulated and processed through the RETINA system to determine reactor-specific variations in the antineutrino yield confidence. The total antineutrino yields and their associated uncertainties are shown in Figure 4.9. As expected, the overall antineutrino yields increase with reactor power. To get a better sense of reactor design alters the antineutrino spectra regardless of operating power, the spectra were normalized by power for comparison, as seen in Figure 4.10. The differences at this point in the spectra are from variations in the reactor inventories, background significance, and reactor-detector standoff distances.

The confidence in antineutrino yield depends not only on the uncertainty by antineutrino quantity, but the relative impact of that antineutrino quantity uncertainty. For example, the RETINA system can be fairly confident in a collection near the LFR-FCR1 core even if it counts $5,000$ fewer antineutrinos than expected. But if the same $5,000$ count deviation was to occur at the FHR core, our confidence in that yield would significantly diminish. Fractional uncertainties, as seen in Figure 4.11, can be used to evaluate our confidence in the antineutrino yield and compare across the different reactor designs. Overall, the FHR had the largest fractional uncertainties ranging from about 9% to 48% while the LFR-FCR1 had the smallest fractional uncertainties ranging from about 3% to 35%

As previously stated, after power normalization, reactor inventory variation could cause spectra uncertainty variations. From initial inventories, most of the relative antineutrino yield uncertainties come from U-235 in the FHR and AFR-100 spectra and from Pu-239

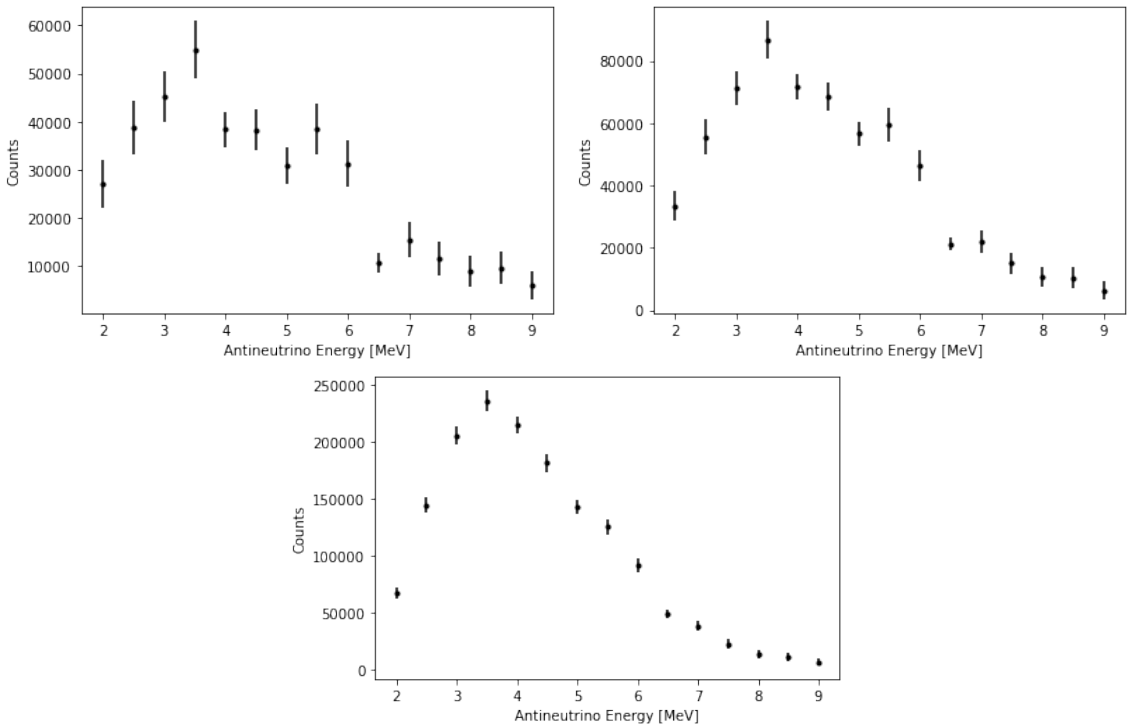

Figure 4.9: Antineutrino yields for the RETINA system 17 m away from a fresh FHR core (upper left), AFR-100 core (upper right), and 25 m away from a fresh LFR-FCR1 core (lower center)

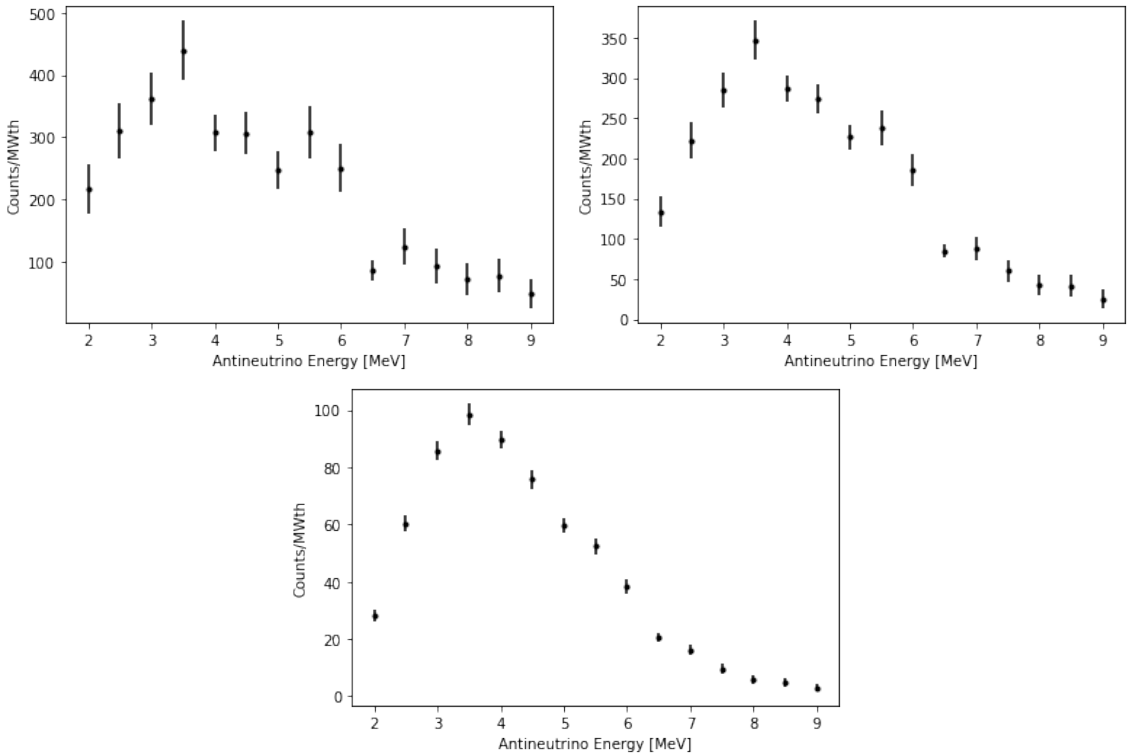

Figure 4.10: Antineutrino yields normalized by reactor power for the RETINA system 17 m away from a fresh FHR core (upper left), AFR-100 core (upper right), and 25 m away from a fresh LFR-FCR1 core (lower center)

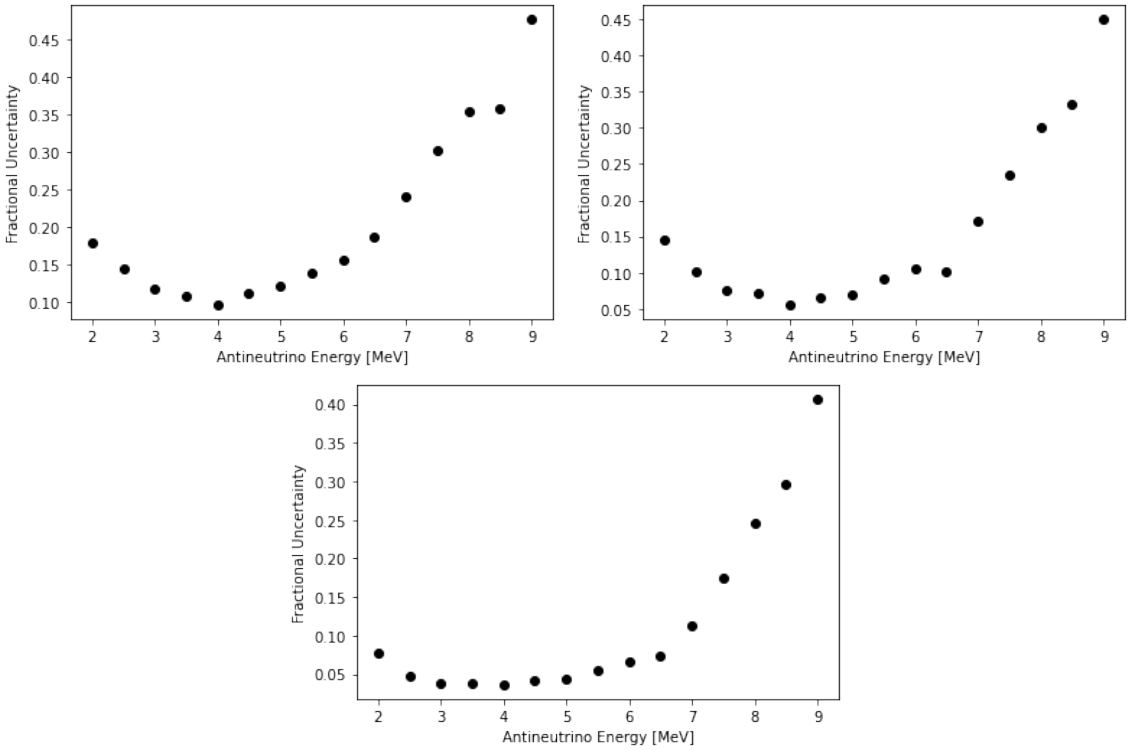

Figure 4.11: Fractional uncertainty for the RETINA system 17 m away from a fresh FHR core (upper left), AFR-100 core (upper right), and 25 m away from a fresh LFR-FCR1 core (lower center)

in the LFR-FCR1 spectra. The LFR-FCR1's initial inventory includes a large quantity of TRU and NU, which support the idea that U-235 would not play a large role in the spectra uncertainty. The LFR-FCR1 design, similar to the AFR-100 design, utilizes a faster neutron spectrum for U-238 fission. This design characteristic leads to a larger U-238 fission contribution, noticed at the lower and higher energy bins for the AFR-100 and LFR-FCR1 designs.

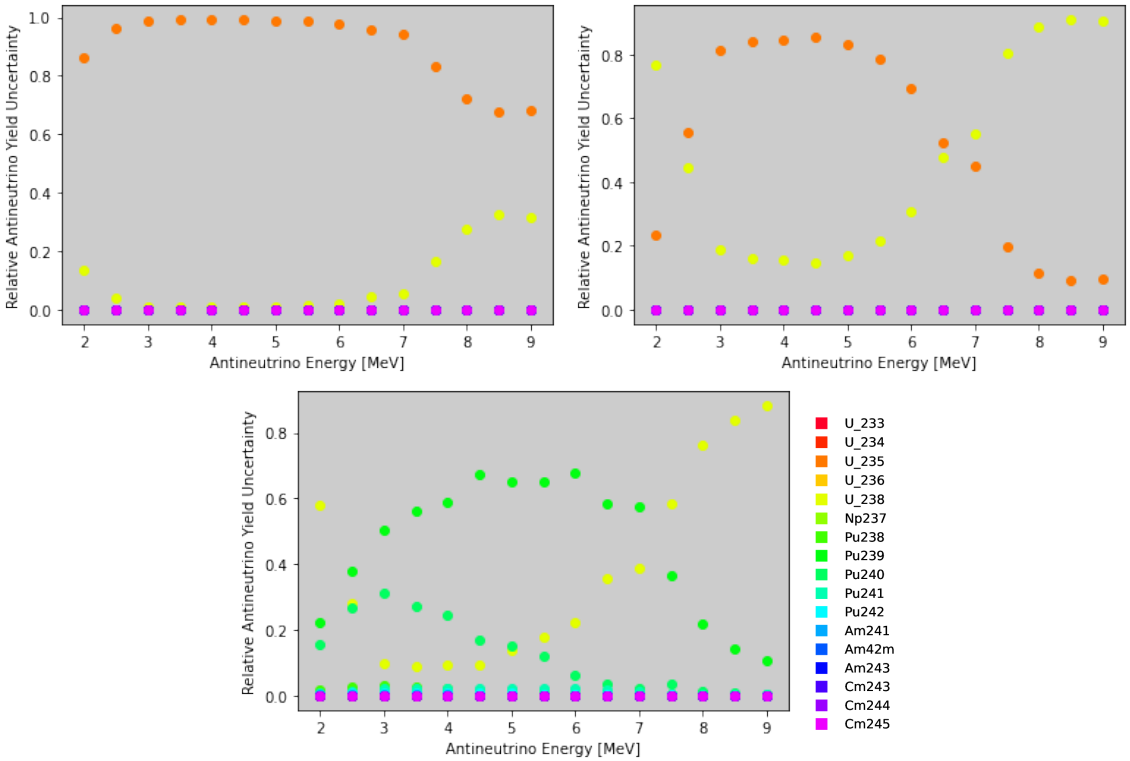

Figure 4.12: Relative antineutrino emission rate uncertainties for a fresh FHR core (upper left), AFR-100 core (upper right), and a fresh LFR-FCR1 core (lower center)

The 5 leading causes of uncertainties for the 3 reactor designs are shown in Figure 4.13. There are slight differences between the FHR and the AFR-100 uncertainty significant factors. One notable difference is the varying importance of counting statistics depending on antineutrino energy. The FHR design, with an approximate signal-to-background ratio of 0.95, a majority of the uncertainty from counting statistics is due to the background events. For the AFR-100 design, with an approximate signal-to-background ratio of 1.5, a majority of the uncertainty from counting statistics is due to the source events. The LFR-

FCR1 design, while located further away from the RETINA system (25 m compared to 17 m), has a much larger signal-to-background ratio of about 6.5. With a much larger signal source, the overwhelming background uncertainty is relatively lowered. This is especially significant around the 4 MeV antineutrino energy bin based on the larger antineutrino event rate.

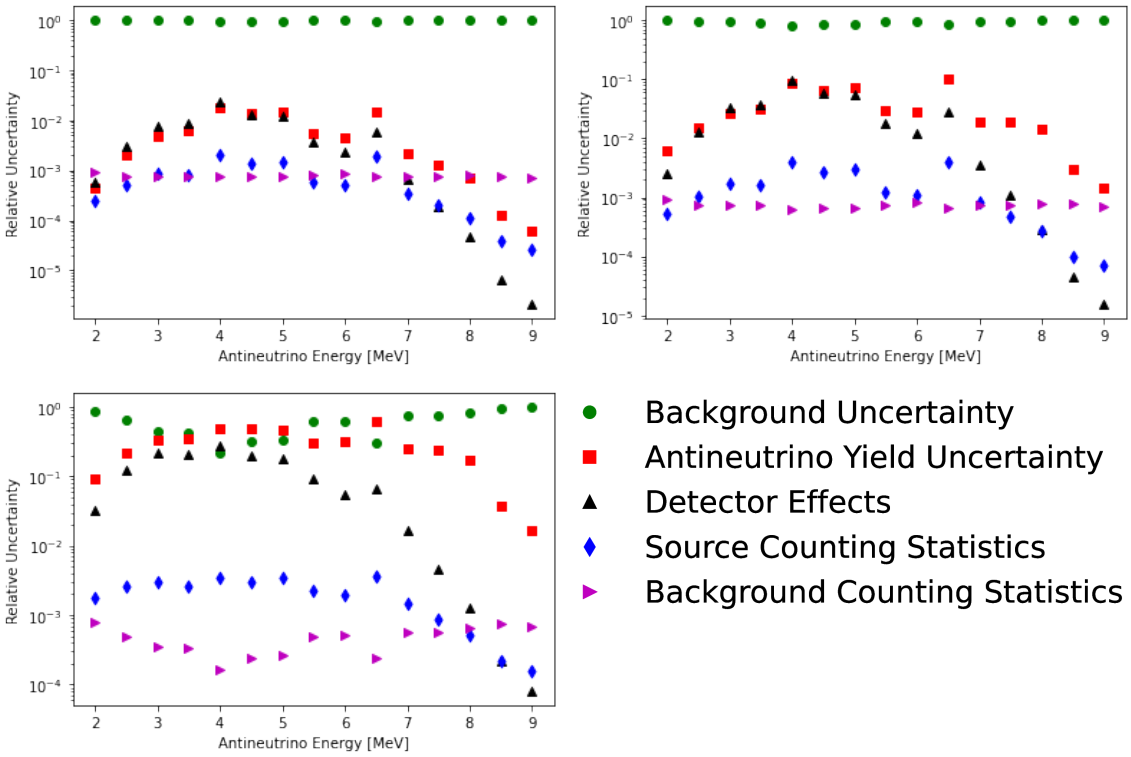

Figure 4.13: Five leading causes of uncertainty for the RETINA system 17 m away from a fresh FHR core (upper left), AFR-100 core (upper right), and 25 m away from a fresh LFR-FCR1 core (lower left)

4.1.3 Temporal Significance

The final component for the antineutrino yield confidence analysis is temporal effects. While many aspects of the RETINA system uncertainty remain motyl consistent for the entirety of the reactor's lifecycle, such as detector effects and background, the reactor fissile inventory will change with time. Depending on the reactor design, there can be significant changes in the reactor inventory. Since the AFR-100 has the longest fuel cycle (30 EFPY)

and is designed to breed plutonium as fuel, it was selected as the reactor design for this section. As seen in Figure 4.14, the age of the AFR-100 reactor has little influence on the overall antineutrino yield.

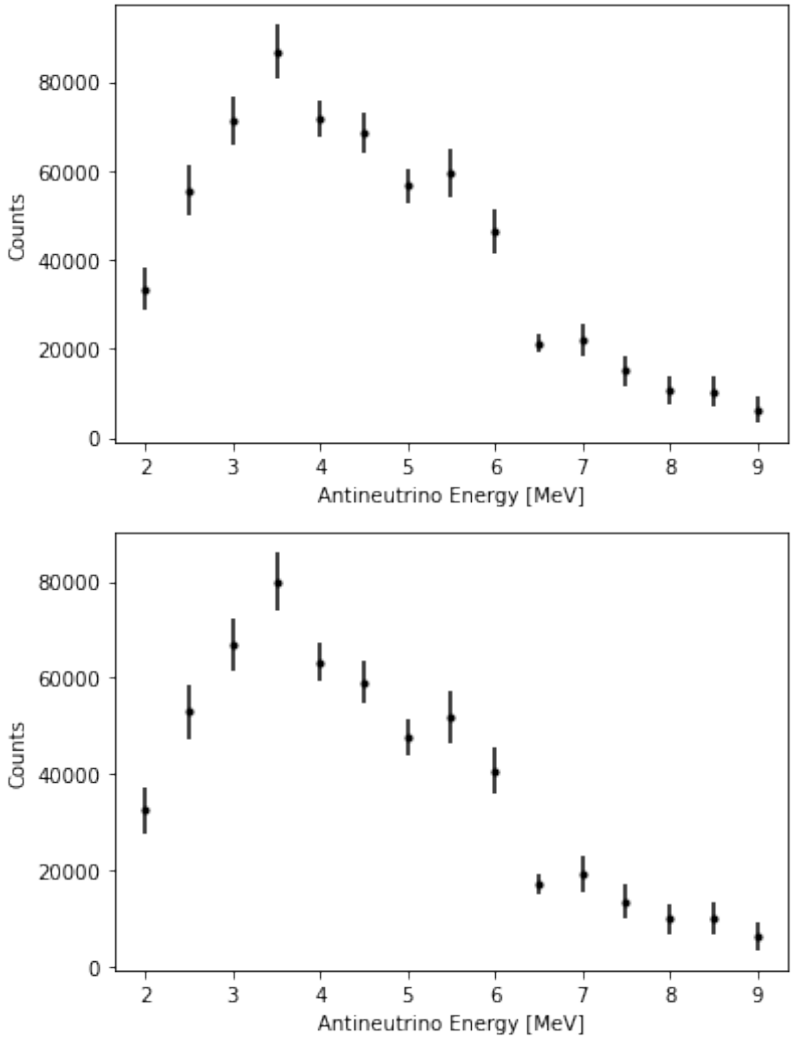

Figure 4.14: Antineutrino yields for the RETINA system 17 m away from an AFR-100 core at 0 EFPY (top) and 30 EFPY (bottom)

The antineutrino flux, however, changes drastically as the reactor ages (Figure 4.15). As expected, after 30 EFPYs of breeding plutonium, the antineutrinos being emitted from the reactor are largely due to fissioning plutonium. The U-235 still contributes a large portion of the antineutrino flux uncertainty. With 3 isotopes contributing relatively similar amounts of antineutrinos to the flux rather than the initial 2, the yield uncertainties become more

evenly distributed among the fissioning isotopes (Figure 4.16). Plutonium-241, a relatively insignificant fissioning isotope, can be seen in Figure 4.16 due to the larger antineutrino uncertainty library associated with that isotope.

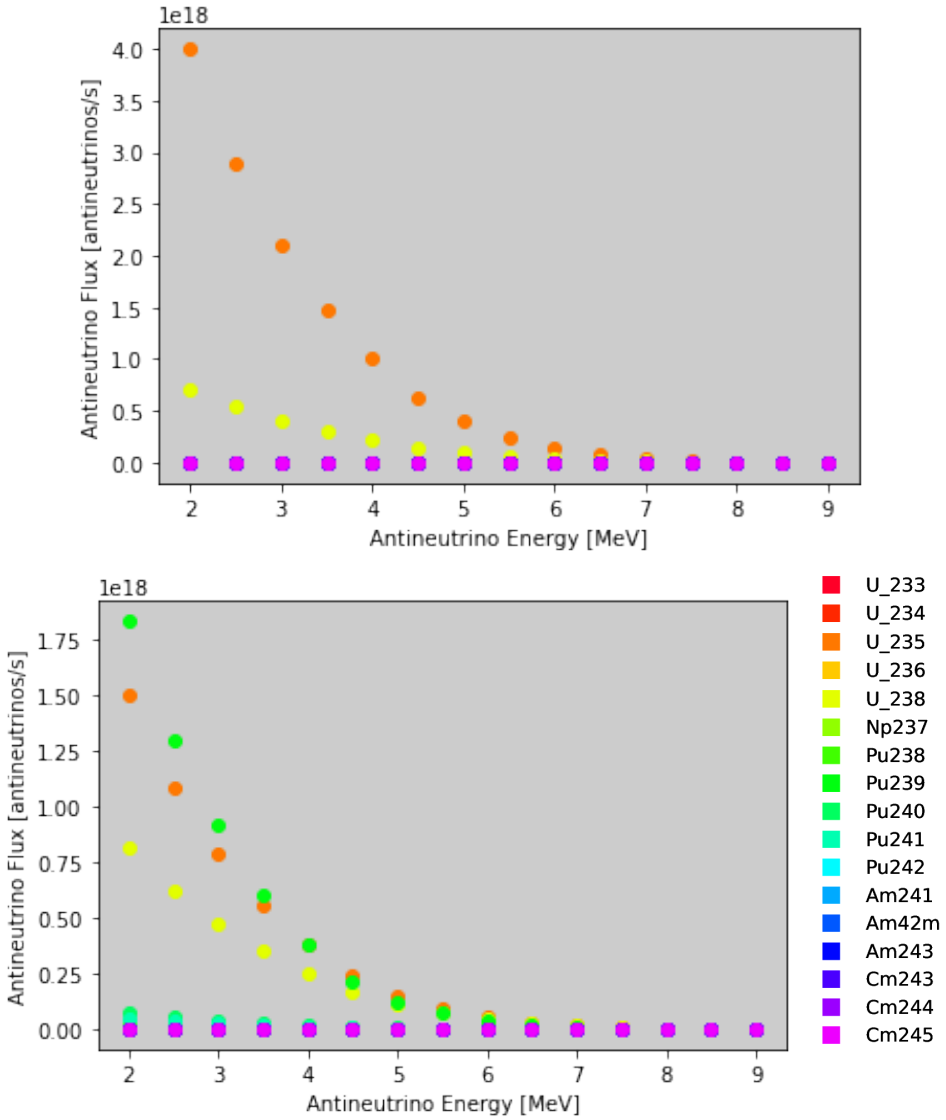

Figure 4.15: U-235, U-238 and eventually Pu-239 as the significant contributing isotopes to the antineutrino flux for an AFR-100 core at 0 EFPY (top) and 30 EFPY (bottom)

The overall uncertainty rankings for the confidence analysis, Figure 4.17, do not change orders depending on the age of the reactor. That said, there are slight deviations in the significance of antineutrino yield uncertainty compared to the other components. For example, at the 6.5 MeV bin for the fresh core, the detector effects surpass the importance of the an-

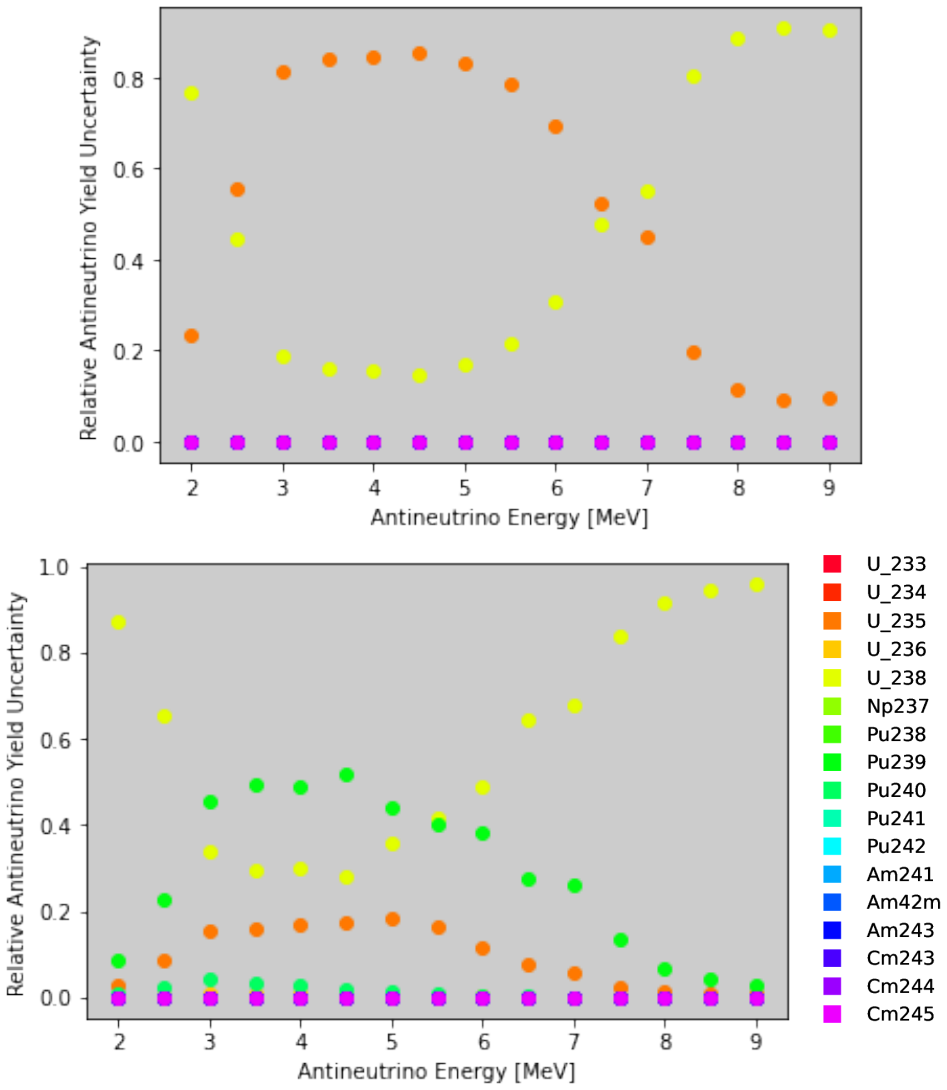

Figure 4.16: U-235, U-238 and eventually Pu-239 as the significant contributing isotopes to the relative antineutrino emission rate uncertainties for an AFR-100 core at 0 EFPY (top) and 30 EFPY (bottom)

tineutrino yield uncertainties in the nearby bins. In the 30 EFPY case, the antineutrino yield uncertainties around that region all remain more prominent than the detector effect uncertainties. Regardless of these slight variations, the background uncertainty remains the overwhelming contribution to lower antineutrino yield confidence in the RETINA system.

4.2 Confidence in Diversion Scenario Detection

The Chi-square goodness of fit test was applied to each diversion scenario listed in Table 3.4. After applying a distribution and processing through a cumulative distribution function, the confidence values were calculated and placed in Table 4.1. Due to the precision limits of Python, any confidence lower than 10^{-16} is considered 0. These results are included in Dunbrack *et al.* [72].

Table 4.1: Detection confidence through Chi-square approach

Scenario	Confidence
1a Nominal	2.21 x 10^{-2}
1a Manipulated	2.52 x 10^{-13}
1b Nominal	8.56 x 10^{-4}
1b Manipulated	6.34 x 10^{-12}
2a Nominal	5.63 x 10^{-12}
2a Manipulated	0
2b Nominal	0
2b Manipulated	0
3a Nominal	4.69 x 10^{-12}
3a Manipulated	0
3b Nominal	0
3b Manipulated	0

After the Chi-square approach was considered, sample points were generated from a random distribution. As previously mentioned in section subsection 3.2.4, the unprocessed spectra would not be ideal features due to the low variance between the diverted and reference scenarios, as seen in Figure Figure 4.18. The standard score of each sampled yield was used to increase the variance of these features (Figure Figure 4.19.

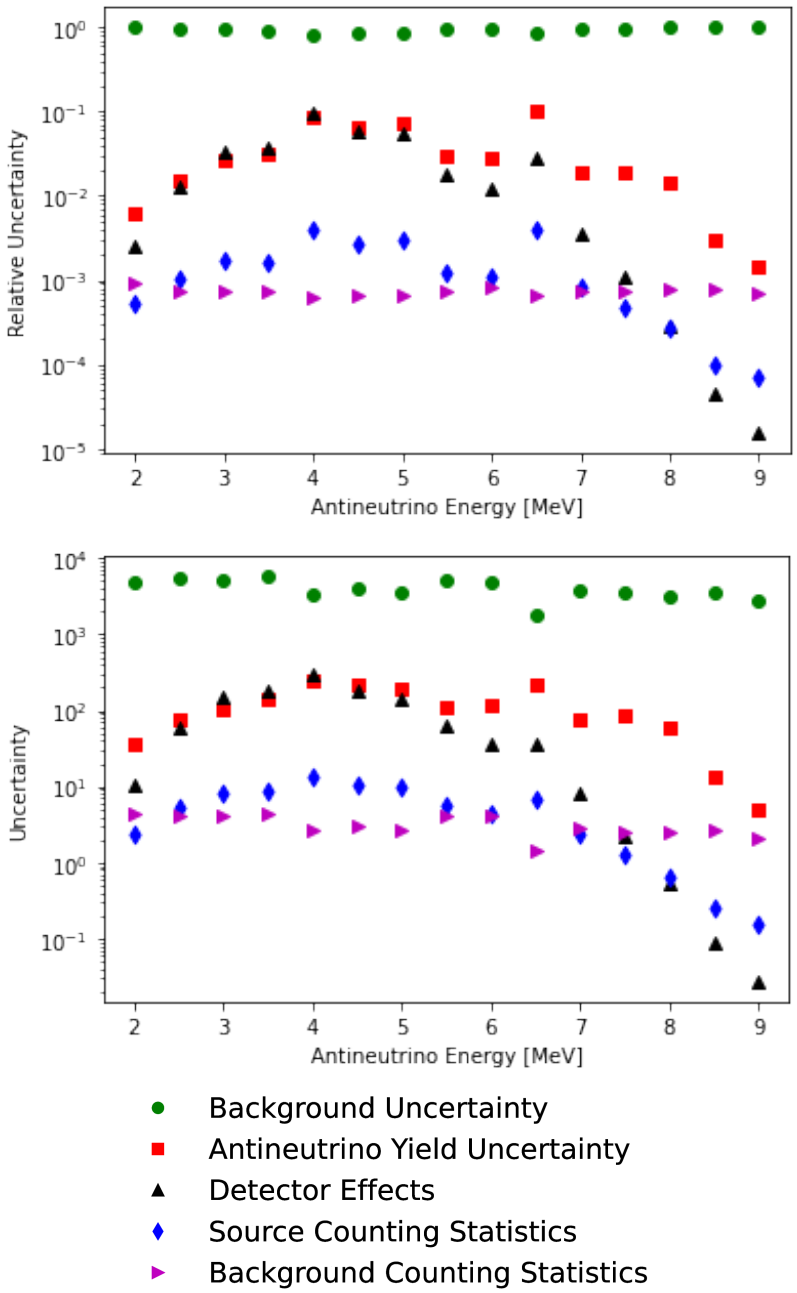

Figure 4.17: Five leading causes of uncertainty for the RETINA system 17 m away from an AFR-100 core at 0 EFPY (top) and 30 EFPY (bottom)

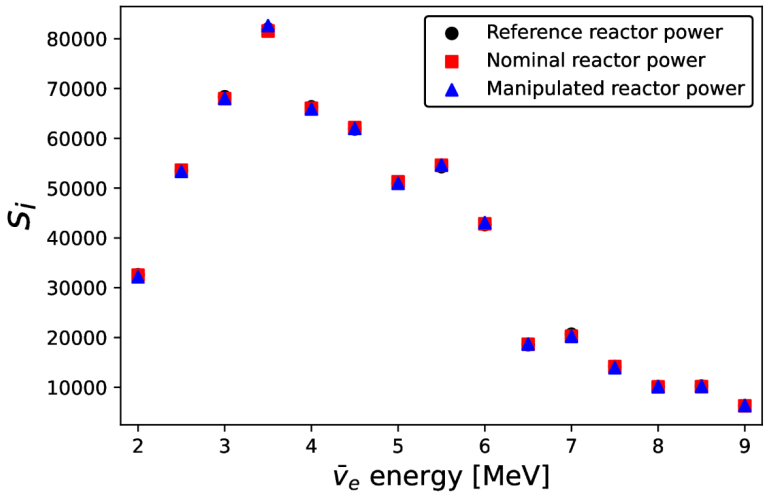

Figure 4.18: Low variance sampled points from antineutrino yield distribution

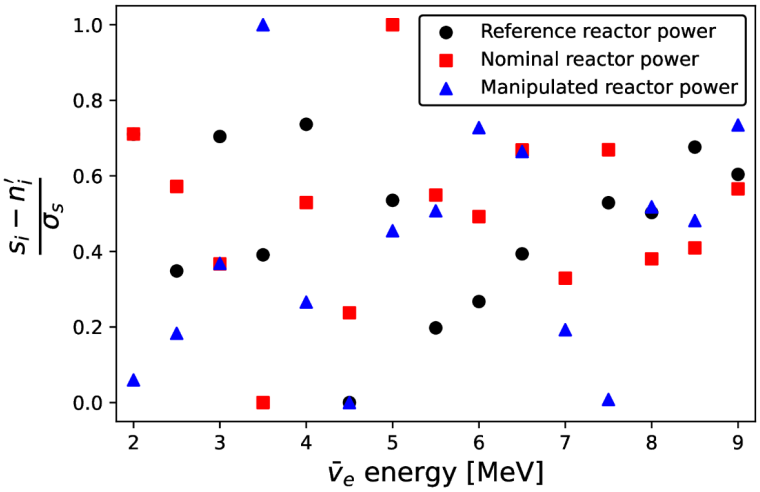

Figure 4.19: High variance sampled points from antineutrino standard score distribution

The SVM model was selected to process the standard score features. A variety of regularization hyperparameter values, ranging from 10^{-15} to 10^{0} and scaling by factors of 10, and various kernel functions, $f(\mathbf{x}_i)$, including linear, polynomial degree 2, polynomial degree 3, radial basis function, and sigmoid, were applied to an SVM to classify data from diversion scenario 1a nominal. The small regularization hyperparameter values typically required over 1000 iterations for SVM convergence. Small regularization hyperparameters were chosen to create a high bias system, or a system that is generalized to fit widely diverse datasets [81]. Although in some cases a high bias model can underfit data [82], a low C value was necessary to fit the wide distribution of the datasets. As the hyperparameter decreased, fewer iterations were required to converge the model with little effect on system performance. A regularization hyperparameter value of 10^{-12} was set for the SVM models since any smaller value had little impact on required iterations and performance. The linear kernel type was chosen since it led to relatively high detection confidence while performing significantly faster than the other kernel types. Equation 4.1 includes the updated SVM model design.

$$\min_{w}[10^{-12} * \Sigma_{i=1}^{N} max(0, 1 - y_i \mathbf{w}^T \mathbf{x}_i)^2 + \frac{1}{2}||\mathbf{w}||^2] \tag{4.1}$$

After the hyperparameters were selected, sample points for each scenario were processed through the ML model for a sensitivity analysis. A total of 10^6 samples were used in each model evaluation with an 80/20 train-test split. The resulting test classification values were binned in a histogram and fit to a Gaussian distribution, as seen in Figure 4.20. A decision boundary was then set at the classification value that would match a 0.05 false positive rate. The difference in area between the distributions above this decision boundary, outlined by Equation 3.20 - Equation 3.22, is the detection confidence. This process was repeated for 100 iterations to get a standard deviation for the models depending on sample fluctuations. Due to the stochastic nature of this process, it is possible to have negative detection confidence. If the distributions are near exact, the reference spectra could have one

extra sample over the classification boundary that skews the fitted Gaussian distribution. Since this is a nonsensical value, all negative confidence values are considered to be 0 from precision limits.

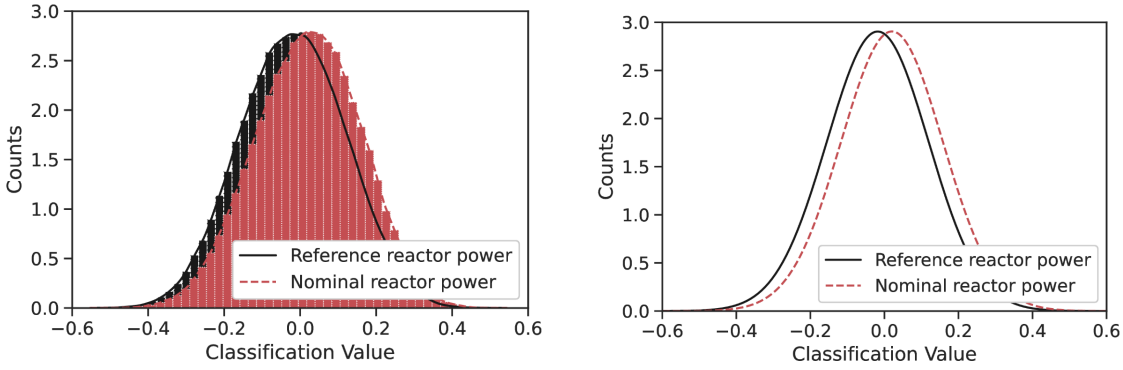

Figure 4.20: Histogram (left) and Gaussian fit (right) of SVM classification values from individualized model diversion 1a sampling

4.2.1 Scenario Significance

The resulting detection confidence values are found in Table 4.2. The scenario resulting in the highest detection confidence is scenario 1a nominal while the scenario resulting in the lowest detection confidence is scenario 3b manipulated. The most influential component of scenario significant is the assembly diversion location. The lowest confidence value for scenario 1, 2.08 x 10^{-2}, is higher than the highest confidence value for all other scenarios, 1.26 x 10^{-2}. The following most influential scenario factor is power manipulation. As seen in the confidence values for scenario 2, power manipulation can lower the confidence value by almost a factor of 10. The final scenario variation, fuel replacement, had the smallest impact on the overall confidence of the RETINA system. On average, LEU compared to NU fuel replacement decreased the confidence value by about 25%.

4.2.2 Model Significance

Overall, the SVM models performed significantly better than the previous Chi-square approach. The confidence values are listed in Table 4.3 and Table 4.4. While the highest SVM

Table 4.2: Detection confidence through SVM individualized model approach

Scenario	Individualized Models
1a Nominal	4.28 x $10^{-2} \pm$ 6.94 x 10^{-18}
1a Manipulated	2.08 x $10^{-2} \pm$ 1.04 x 10^{-17}
1b Nominal	3.63 x $10^{-2} \pm$ 6.94 x 10^{-18}
1b Manipulated	2.28 x $10^{-2} \pm 0$
2a Nominal	1.26 x $10^{-2} \pm$ 5.20 x 10^{-18}
2a Manipulated	3.82 x $10^{-3} \pm 0$
2b Nominal	1.01 x $10^{-2} \pm 0$
2b Manipulated	2.70 x $10^{-3} \pm 0$
3a Nominal	8.93 x $10^{-3} \pm 0$
3a Manipulated	3.46 x $10^{-3} \pm 0$
3b Nominal	6.65 x $10^{-3} \pm 0$
3b Manipulated	1.47 x $10^{-3} \pm 0$

confidence value of 4.28 x 10^{-2} is lower than the 0.2 confidence value needed for IAEA implementation, it is significantly higher than the 2.21 x 10^{-2} confidence value calculated from the Chi-square approach. Along with this better performance, the SVM models, in general, were able to detect a majority of the scenarios to some extent. In the previous Chi-square approach, half of scenarios were considered to have a 0 confidence value due to precision limits. Out of all 4 model variations and 12 diversion scenarios, only 4 of the scenarios were considered to have a 0 confidence value.

The individualized models performed significantly better than the other SVM models. Not only did these models improve the confidence values of the RETINA system, but they were the most consistent models. After 100 iterations, a majority of the individualized models converged to the exact same confidence value. The other models typically fluctuated with a standard deviation of about 10^{-3}.

The other three models varied in performance depending on the scenario. In general, the power-ignorant models performed second best compared to the individualized models. Similar to the individualized models, the power-ignorant models are trained for specific cases while remaining partially robust to the spectra oscillations of manipulated reactor

Table 4.3: Detection confidence through SVM model approach

Scenario	Individualized Models	Power-Ignorant Models
1a Nominal	$4.28 \times 10^{-2} \pm 6.94 \times 10^{-18}$	$4.21 \times 10^{-2} \pm 3.51 \times 10^{-4}$
1a Manipulated	$2.08 \times 10^{-2} \pm 1.04 \times 10^{-17}$	$1.38 \times 10^{-2} \pm 1.08 \times 10^{-3}$
1b Nominal	$3.63 \times 10^{-2} \pm 6.94 \times 10^{-18}$	$3.52 \times 10^{-2} \pm 6.37 \times 10^{-4}$
1b Manipulated	$2.28 \times 10^{-2} \pm 0$	$1.82 \times 10^{-2} \pm 1.08 \times 10^{-3}$
2a Nominal	$1.26 \times 10^{-2} \pm 5.20 \times 10^{-18}$	$1.24 \times 10^{-2} \pm 1.26 \times 10^{-3}$
2a Manipulated	$3.82 \times 10^{-3} \pm 0$	$2.78 \times 10^{-3} \pm 1.37 \times 10^{-3}$
2b Nominal	$1.01 \times 10^{-2} \pm 0$	$9.49 \times 10^{-3} \pm 1.38 \times 10^{-3}$
2b Manipulated	$2.70 \times 10^{-3} \pm 0$	$9.63 \times 10^{-4} \pm 1.42 \times 10^{-3}$
3a Nominal	$8.93 \times 10^{-3} \pm 0$	$8.19 \times 10^{-3} \pm 1.48 \times 10^{-3}$
3a Manipulated	$3.46 \times 10^{-3} \pm 0$	$1.97 \times 10^{-3} \pm 1.59 \times 10^{-3}$
3b Nominal	$6.65 \times 10^{-3} \pm 0$	$5.62 \times 10^{-3} \pm 1.31 \times 10^{-3}$
3b Manipulated	$1.47 \times 10^{-3} \pm 0$	$5.83 \times 10^{-4} \pm 1.36 \times 10^{-3}$

power. The third best performing models were the power-ignorant unseen models.

Table 4.4: Detection confidence through SVM model approach for unseen scenarios

Scenario	Unseen Models	Power-Ignorant Unseen Models
1a Nominal	3.84 x $10^{-2} \pm$ 6.10 x 10^{-4}	4.05 x $10^{-2} \pm$ 4.38 x 10^{-4}
1a Manipulated	1.34 x $10^{-2} \pm$ 1.26 x 10^{-3}	1.72 x $10^{-2} \pm$ 1.02 x 10^{-3}
1b Nominal	3.55 x $10^{-2} \pm$ 6.41 x 10^{-4}	3.62 x $10^{-2} \pm$ 5.81 x 10^{-4}
1b Manipulated	1.40 x $10^{-2} \pm$ 1.15 x 10^{-3}	1.46 x $10^{-2} \pm$ 1.11 x 10^{-3}
2a Nominal	1.18 x $10^{-2} \pm$ 1.32 x 10^{-3}	1.18 x $10^{-2} \pm$ 1.27 x 10^{-3}
2a Manipulated	3.27 x $10^{-3} \pm$ 1.31 x 10^{-3}	3.41 x $10^{-3} \pm$ 1.63 x 10^{-3}
2b Nominal	$0 \pm$ 1.77 x 10^{-3}	$0 \pm$ 1.45 x 10^{-3}
2b Manipulated	$0 \pm$ 1.38 x 10^{-3}	$0 \pm$ 1.38 x 10^{-3}
3a Nominal	7.91 x $10^{-3} \pm$ 1.33 x 10^{-3}	8.05 x $10^{-3} \pm$ 1.28 x 10^{-3}
3a Manipulated	2.54 x $10^{-3} \pm$ 1.33 x 10^{-3}	2.43 x $10^{-3} \pm$ 1.44 x 10^{-3}
3b Nominal	$0 \pm$ 1.75 x 10^{-3}	$0 \pm$ 1.81 x 10^{-3}
3b Manipulated	$0 \pm$ 1.51 x 10^{-3}	$0 \pm$ 1.64 x 10^{-3}

CHAPTER 5
CONCLUSIONS

5.1 Confidence in Antineutrino Yield

The uncertainty in the final antineutrino yields are considerably high, ranging from 3% to 48% depending on the energy bin and reactor of interest. Since background uncertainty was typically the most significant component of uncertainty, reactor designs with a low signal-to-background ratio had the largest fractional uncertainty for all energy bins. Regardless of the reactor design, the lowest fractional uncertainties remained around the 4 MeV bin due to the larger source term. Especially for the larger reactor, the LFR-FCR1, in which background uncertainty, antineutrino yield uncertainty, and detector effects have comparable significance in antineutrino yield confidence around the 4 MeV bin, the fractional uncertainty becomes more consistent at a low value.

The background uncertainty was the most impactful component to the low confidence in antineutrino yield. Considering background is an unavoidable aspect of the RETINA system, there are two ways to mitigate this hindrance: by increasing the signal-to-noise ratio or by gaining confidence in the detected background. As outlined in Equation 3.10, the propagated uncertainty scales with both the source uncertainty and the background uncertainty. If the background uncertainty cannot be improved, having a more significant source term uncertainty would make the background uncertainty relatively insignificant. Alternatively, the background uncertainty can be studied directly for uncertainty mitigation.

The second leading uncertainty, lowering the antineutrino yield confidence, is the antineutrino yield uncertainty. This value corresponds to the sum of the antineutrino yield libraries' uncertainties weighted by the isotopic fission rates. While the libraries for the less common reactor inventory actinides, such as Americium and Curium, include large

theoretical uncertainties ranging from 20% to 50%, not many of these actinides produce a notable amount of antineutrinos compared to U-235, U-238, and Pu-239. While U-235 or Pu-239 typically produce the most amount of antineutrinos in these simulated reactor cores, U-238 contributes the most to the antineutrino yield uncertainty at low and high energy bins due to the uncertainty library. The U-238 antineutrino yield uncertainty library, discussed in subsection 2.2.1, required high error bars for low and high energy bins based on later-derived values for U-238 fast fission. For the overall antineutrino yield confidence to increase, future research should focus on improving the antineutrino yield libraries for common fissile isotopes, such as U-235 and Pu-239, as well as for U-238 for lower and higher antineutrino energy bins.

Detector effects also played a significant role in the uncertainty propagated through the RETINA system. The leading causes of uncertainty within these effects include inverse-beta decay cross sections, number of protons within the scintillation medium, and the overall intrinsic efficiency of the system. The IBD cross section uncertainty is the most significant of these components, nearly twice the relative contribution compared to the target proton uncertainty and intrinsic efficiency uncertainty. While the uncertainty values for the detector effects are accurate based on literature review, a more complex detector modeling aspect should be applied to generate a more comprehensive understanding of how detector effects fluctuate with varying antineutrino energies and detector designs.

If the RETINA system is to safeguard small reactor cores, this background uncertainty must be significantly lowered. For the FHR, an 125 MW_{th} reactor design, the background is more significant than the other leading causes of uncertainty by a factor of 100 for high yield energy bins and by a factor of $1,000$ for low yield energy bin. Even for the AFR-100, a 250 MW_{th} reactor design, the background uncertainty is more significant by factors ranging from 10 to 100. With larger reactors, such as the LFR-FCR1, the source term uncertainty is large enough to mitigate the significance of background uncertainty compared to the antineutrino yield uncertainty and detector effects uncertainty.

As for the temporal significance in antineutrino yield confidence for antineutrino-based safeguards, the largest difference forms from the antineutrino yield uncertainty. The fissile inventory shifts overtime as fuel is produced and depleted. Over the 30 EFPY lifespan of the AFR-100 fuel, Pu-239 becomes the leading antineutrino-producing isotope, replacing U-235. This shift in antineutrino generation propagates through the antineutrino yield uncertainty portion of the RETINA system. There is also a slight shift in the role between background uncertainty and the other forms of uncertainty due to the change in antineutrino production. Since, on average, Pu-239 generates fewer antineutrinos than U-235 per fission, switching to a Pu-239 fuel source would alter the signal-to-background ratio. With U-235 and Pu-239 having similar antineutrino yield uncertainty libraries, this temporal effect has little influence on the overall antineutrino yield confidence of the RETINA system.

5.2 Confidence in Diversion Scenario Detection

The following conclusions are found in Dunbrack *et al.* [72]. While none of the SVM models were able to classify diversion scenarios above the IAEA-set confidence level, all of the models performed better than the Chi-square goodness of fit approach for non-zero confidence values. Also, scenarios that were previously considered 0 now have confidence values. From this change in antineutrino yield processing, all 12 scenarios can be ranked in order of risk. For example, diversion scenario 3b nominal, with a confidence value of 6.65 x 10^{-3} for individualized SVM models, is easier to detect compared to diversion scenario 3b manipulated, with a confidence value of 1.47 x 10^{-3} for individualized SVM models. The previously used Chi-square approach would consider these scenarios equivalent in risk (with a 0 detection confidence value).

The most significant scenario parameter that influenced the diversion detection confidence was the number of diverted assemblies and their locations. The RETINA system could best isolate diversion scenarios in which 1 SQ of plutonium was diverted from 1 assembly. The next most meaningful diversion variable is reactor power manipulation. If

the reactor operator is an adversary and attempts to mask the shift in antineutrino yield through reactor power manipulation, they can significantly reduce the confidence in diversion scenario detection. The final variable in scenario selection was fuel replacement. The RETINA system is less confidence in scenarios with LEU fuel replacement compared to NU. In summary, the RETINA system is most prepared for scenarios in which only one assembly is diverted, the reactor is operating at nominal thermal power, and the diverted Pu-239 is replaced with NU. The RETINA system is least prepared for scenarios in which multiple assemblies are diverted, the reactor is operating at a manipulated thermal power, and the diverted Pu-239 is replaced with LEU.

The individualized models performed significantly better than than the power-ignorant models, the unseen models, and the power-ignorant unseen models. This is expected since each model is custom to the scenario of interest and does not need to be robust to a complex system of spectra. The power-ignorant models require a more generalized model that can process the information from the deviation of nominal scenarios and from the oscillation of manipulated power scenarios. The unseen models require similar complexity to become robust to any deviation from the reference case. The power-ignorant unseen models need both oscillation generalization as well as robustness for unexpected spectra.

The power-ignorant unseen models were not only robust enough to reasonably detect diversion scenarios, but they were more confidence in the detected diversion scenarios in certain cases compared to the power-ignorant. Power-ignorant unseen models are better performing than the power-ignorant models since there are more diverse diversion scenarios within the training dataset, making a more robust system.

The power-ignorant unseen models also led to higher confidence values than the unseen models. While the unseen models train from a larger variety of scenarios compared to the power-ignorant unseen models, the power-ignorant models have a more balanced dataset in regards to power manipulation. For example, for the unseen model tested against scenario 1a nominal, the model should utilize generalize deviation rather than an oscillation effect

since there is no oscillation in the nominal case. The model, however, is trained from 6 datasets with the oscillating spectra and only 5 datasets with no oscillation spectra. This imbalance in oscillation-based datasets caused the models to focus primarily on the less common pattern.

Both the unseen models and power-ignorant unseen models included confidence values of 0. These harder-to-detect scenarios, as supported by the individualized models, fall below the precision limit of this method. Since a total of 10^6 samples were generated for each model, the unseen models and power-ignorant unseen models, including 11 and 10 training datasets, respectively, have fewer samples from each scenario compared to the other two models. With fewer samples from each distribution, the classification boundaries are less likely to reach the true optimal position.

5.3 Future Work

There are many components of the RETINA system that can be improved. Aside from fundamental library improvements, such as with antineutrino yield libraries and inverse-beta decay cross section libraries, the are more RETINA system-focused improvements that can be made directly to the safeguards system. This future work includes applying the system to different reactor designs, utilizing more realistic multi-physics models, and implementing new ML techniques to improve our final confidence in the system.

The RETINA system could be used to safeguard an evolving nuclear industry with novel fuel types and complex facilities. Bulk fuel, for example, would require new safeguards for implementation [83]. An FHR with TRISO fuel, a previously mentioned fuel type in subsection 2.4.1, can be simulated in SERPENT2 with a variety of simplified diversion scenarios for confidence analysis. A simplified diversion scheme can be applied since assembly location would no longer be a user-controlled parameter. The RETINA system could also be applied to a facility with many small reactors. While the current RETINA system is unable to effectively safeguard small reactors following the IAEA confidence limits,

the system could have value in safeguarding a facility that contains many small reactors.

The simulation capability of the RETINA system can also be improved to account for more realistic and complex behaviors within the system. For example, there are uncertainties, such as those associated with reactor burnup and thermal power fluctuation, that would directly influence the isotopic fission rate. Rather than including these uncertainties on deterministically through propagation of error, a Monte Carlo method can be used to sample for power level and burnup history distributions for a more all-encompassing isotopic fission rate. This Monte Carlo approach has already been applied to some degree for cross sections in SERPENT through Total Monte Carlo (TMC). The detection translation component can also be more comprehensive by applying high-fidelity reactor modeling through GEANT4 simulation. Although there is a loss in interpretability through a Monte Carlo approach, the RETINA system is more realistic with complex relationships.

A modern ML technique, temporal difference learning, can be applied to the RETINA system to improve the final system diversion detection confidence. Since there are an infinite number of reactor diversion scenarios possible for simulation, a reinforcement learning (RL) model agent can be used to determine the optimal scenarios to remove SNM from a reactor core without being detected by the RETINA system. The RETINA system can then be trained to effectively detect these scenarios. With the RETINA system altered for these specific scenarios, the RL agent can then generate a new set of reactor diversion scenario simulations to avoid detection. This back-and-forth computational process, referred to as temporal difference learning, would provide a single, fundamental confidence value as well as identify possible antineutrino-based safeguards concerns for researchers and policy makers.

Appendices

APPENDIX A

SUPPLEMENTAL MATERIAL FOR ANTINEUTRINO YIELD CONFIDENCE

	U233	U234	U235	U236	U238	Np237	Pu238	Pu239	Pu240	Pu241	Pu242	Am241	Am242	Am242m	Am243	Cm243	Cm244	Cm245
Fissions	4.516190e+11	3.101479e+14	4.497990e+18	2.954330e+16	1.162460e+18	7.602470e+15	3.404290e+15	1.891190e+18	2.603550e+16	1.559900e+16	9.290770e+13	1.921310e+14	4.479630e+11	7.183630e+13	7.187650e+12	1.348230e+12	4.040800e+12	8.973380e+11
U233	2.039597e+23	5.836213e+24	5.890995e+26	5.416980e+25	3.517426e+26	2.585365e+25	1.705023e+25	3.928848e+26	5.150059e+25	3.804195e+25	3.029431e+24	4.136326e+24	2.023086e+23	2.572052e+24	7.803250e+23	3.497747e+23	5.959583e+23	2.858384e+23
U234	5.836213e+24	1.670005e+26	1.685681e+28	1.550044e+27	1.006495e+28	7.397903e+26	4.878845e+26	1.124222e+28	1.473666e+27	1.088553e+27	8.668578e+25	1.183591e+26	5.788967e+24	7.359809e+25	2.232864e+25	1.000864e+25	1.705307e+25	8.179135e+24
U235	5.890995e+26	1.685681e+28	1.701504e+30	1.564593e+29	1.015943e+30	7.467344e+28	4.924641e+28	1.134774e+30	1.487498e+29	1.098771e+29	8.749946e+27	1.194701e+28	5.843306e+26	7.428893e+27	2.253823e+27	1.010259e+27	1.721314e+27	8.255909e+26
U236	5.416980e+25	1.550044e+27	1.564593e+29	1.438699e+28	9.341957e+28	6.866489e+27	4.528383e+27	1.043465e+29	1.367808e+28	1.010359e+28	8.045887e+26	1.098570e+27	5.373128e+25	6.831131e+26	2.072470e+26	9.289691e+25	1.582810e+26	7.591602e+25
U238	3.517426e+26	1.006495e+28	1.015943e+30	9.341957e+28	6.066045e+29	4.458641e+28	2.940430e+28	6.775570e+29	8.881634e+28	6.560597e+28	5.224464e+27	7.133381e+27	3.488952e+26	4.435682e+27	1.345724e+27	6.032107e+26	1.027771e+27	4.929481e+26
Np237	2.585365e+25	7.397903e+26	7.467344e+28	6.866489e+27	4.458641e+28	3.277173e+27	2.161264e+27	4.980153e+28	6.528144e+27	4.822144e+27	3.840065e+26	5.243150e+26	2.564436e+25	3.260298e+26	9.891292e+25	4.433696e+25	7.554284e+25	3.623248e+25
Pu238	1.705023e+25	4.878845e+26	4.924641e+28	4.528383e+27	2.940430e+28	2.161264e+27	1.425333e+27	3.284363e+28	4.305248e+27	3.180157e+27	2.532486e+26	3.457807e+26	1.691221e+25	2.150135e+26	6.523211e+25	2.923979e+25	4.981977e+25	2.389497e+25
Pu239	3.928848e+26	1.124222e+28	1.134774e+30	1.043465e+29	6.775570e+29	4.980153e+28	3.284363e+28	7.568085e+29	9.920488e+28	7.327967e+28	5.835551e+27	7.967748e+27	3.897042e+26	4.954509e+27	1.503129e+27	6.737662e+26	1.147986e+27	5.506066e+26
Pu240	5.150059e+25	1.473666e+27	1.487498e+29	1.367808e+28	8.881634e+28	6.528144e+27	4.305248e+27	9.920488e+28	1.300409e+28	9.605734e+27	7.649428e+26	1.044438e+27	5.108368e+25	6.494529e+26	1.970350e+26	8.831943e+25	1.504817e+26	7.217528e+25
Pu241	3.804195e+25	1.088553e+27	1.098771e+29	1.010359e+28	6.560597e+28	4.822144e+27	3.180157e+27	7.327967e+28	9.605734e+27	7.095468e+27	5.650403e+26	7.714950e+26	3.773398e+25	4.797314e+26	1.455438e+26	6.523892e+25	1.111563e+26	5.331371e+25
Pu242	3.029431e+24	8.668578e+25	8.749946e+27	8.045887e+26	5.224464e+27	3.840065e+26	2.532486e+26	5.835551e+27	7.649428e+26	5.650403e+26	4.499640e+25	6.143721e+25	3.004907e+24	3.820292e+25	1.159023e+25	5.195234e+24	8.851819e+24	4.245583e+24
Am241	4.136326e+24	1.183591e+26	1.194701e+28	1.098570e+27	7.133381e+27	5.243150e+26	3.457807e+26	7.967748e+27	1.044438e+27	7.714950e+26	6.143721e+25	8.388517e+25	4.102841e+24	5.216151e+25	1.582508e+25	7.093471e+24	1.208610e+25	5.796836e+24
Am242	2.023086e+23	5.788967e+24	5.843306e+26	5.373128e+25	3.488952e+26	2.564436e+25	1.691221e+25	3.897042e+26	5.108368e+25	3.773398e+25	3.004907e+24	4.102841e+24	2.006708e+23	2.551231e+24	7.740080e+23	3.469432e+23	5.911338e+23	2.835245e+23
Am242m	2.572052e+24	7.359809e+25	7.428893e+27	6.831131e+26	4.435682e+27	3.260298e+26	2.150135e+26	4.954509e+27	6.494529e+26	4.797314e+26	3.820292e+25	5.216151e+25	2.551231e+24	3.243510e+25	9.840358e+24	4.410866e+24	7.515385e+24	3.604591e+24
Am243	7.803250e+23	2.232864e+25	2.253823e+27	2.072470e+26	1.345724e+27	9.891292e+25	6.523211e+25	1.503129e+27	1.970350e+26	1.455438e+26	1.159023e+25	1.582508e+25	7.740080e+23	9.840358e+24	2.985428e+24	1.338195e+24	2.280064e+24	1.093583e+24
Cm243	3.497747e+23	1.000864e+25	1.010259e+27	9.289691e+25	6.032107e+26	4.433696e+25	2.923979e+25	6.737662e+26	8.831943e+25	6.523892e+25	5.195234e+24	7.093471e+24	3.469432e+23	4.410866e+24	1.338195e+24	5.998359e+23	1.022021e+24	4.901902e+23
Cm244	5.959583e+23	1.705307e+25	1.721314e+27	1.582810e+26	1.027771e+27	7.554284e+25	4.981977e+25	1.147986e+27	1.504817e+26	1.111563e+26	8.851819e+24	1.208610e+25	5.911338e+23	7.515385e+24	2.280064e+24	1.022021e+24	1.741355e+24	8.352030e+23
Cm245	2.858384e+23	8.179135e+24	8.255909e+26	7.591602e+25	4.929481e+26	3.623248e+25	2.389497e+25	5.506066e+26	7.217528e+25	5.331371e+25	4.245583e+24	5.796836e+24	2.835245e+23	3.604591e+24	1.093583e+24	4.901902e+23	8.352030e+23	4.005870e+23

Figure A.1: Fission-covariance matrix for a simulated fresh core AFR-100 reactor

	Energy Bin Center [MeV]	[Yield/fission]
0	2.0	0.599203
1	2.5	0.432151
2	3.0	0.315090
3	3.5	0.220164
4	4.0	0.150242
5	4.5	0.094564
6	5.0	0.059291
7	5.5	0.036656
8	6.0	0.021156
9	6.5	0.011557
10	7.0	0.005975
11	7.5	0.002538
12	8.0	0.000980
13	8.5	0.000337
14	9.0	0.000135

	Energy Bin Center [MeV]	[Yield/fission]
0	2.0	0.513922
1	2.5	0.364303
2	3.0	0.257950
3	3.5	0.169379
4	4.0	0.107282
5	4.5	0.060332
6	5.0	0.034412
7	5.5	0.020283
8	6.0	0.010988
9	6.5	0.005523
10	7.0	0.002644
11	7.5	0.001010
12	8.0	0.000364
13	8.5	0.000113
14	9.0	0.000041

Figure A.2: Antineutrino spectra for U-235 (left) and Pu-239 (right)

	Uncertainty (Fraction of Nominal)
0	0.018
1	0.019
2	0.019
3	0.020
4	0.021
5	0.024
6	0.026
7	0.028
8	0.032
9	0.035
10	0.039
11	0.048
12	0.064
13	0.072
14	0.092

	Uncertainty (Fraction of Nominal)
0	0.026
1	0.026
2	0.028
3	0.032
4	0.035
5	0.046
6	0.049
7	0.056
8	0.079
9	0.099
10	0.131
11	0.214
12	0.284
13	0.308
14	0.359

Figure A.3: Relative antineutrino spectra uncertainty for U-235 (left) and Pu-239 (right)

Energy Bin Center [MeV]	Yield
2.0	4.513061e+18
2.5	3.276681e+18
3.0	2.395856e+18
3.5	1.670079e+18
4.0	1.135523e+18
4.5	7.090002e+17
5.0	4.425189e+17
5.5	2.746526e+17
6.0	1.587793e+17
6.5	8.700542e+16
7.0	4.529715e+16
7.5	1.977941e+16
8.0	7.867868e+15
8.5	2.912113e+15
9.0	1.272523e+15

Energy Bin Center [MeV]	Relative Variance
2.0	0.001238
2.5	0.000471
3.0	0.000264
3.5	0.000278
4.0	0.000300
4.5	0.000386
5.0	0.000465
5.5	0.000585
6.0	0.000964
6.5	0.001730
7.0	0.002622
7.5	0.010157
8.0	0.030376
8.5	0.041288
9.0	0.054028

Figure A.4: Total antineutrino generation rate (left) and relative antineutrino generation rate variance (right) for a simulated fresh core AFR-100 reactor

Energy Bin Center [MeV]	Variance Fraction
2.0	0.000044
2.5	0.000117
3.0	0.000211
3.5	0.000204
4.0	0.000194
4.5	0.000155
5.0	0.000133
5.5	0.000108
6.0	0.000067
6.5	0.000039
7.0	0.000027
7.5	0.000007
8.0	0.000003
8.5	0.000002
9.0	0.000002

Energy Bin Center [MeV]	Variance Fraction
2.0	0.017184
2.5	0.044776
3.0	0.079556
3.5	0.075833
4.0	0.070709
4.5	0.055587
5.0	0.046659
5.5	0.037089
6.0	0.022563
6.5	0.012630
7.0	0.008346
7.5	0.002138
8.0	0.000709
8.5	0.000515
9.0	0.000396

Figure A.5: Variance fraction for the fission rate uncertainty (left) and power uncertainty (right) for a simulated fresh core AFR-100 reactor

	0
0	3.307007e-44
1	1.203711e-43
2	2.541692e-43
3	4.303011e-43
4	6.487669e-43
5	9.095667e-43
6	1.212700e-42
7	1.558168e-42
8	1.945969e-42
9	2.376105e-42
10	2.848574e-42
11	3.363377e-42
12	3.920514e-42
13	4.519985e-42
14	5.161790e-42

Figure A.6: Inverse-beta decay cross sections fit to match desired energy bins

	Background Event Rate [Hz/bin]	Background Event Rate 1-Sig Error [Hz/bin]
0	0.000899	0.000204
1	0.000994	0.000234
2	0.000918	0.000222
3	0.001158	0.000249
4	0.000429	0.000152
5	0.000576	0.000178
6	0.000442	0.000154
7	0.000944	0.000225
8	0.000829	0.000203
9	0.000129	0.000083
10	0.000451	0.000157
11	0.000387	0.000147
12	0.000331	0.000133
13	0.000389	0.000145
14	0.000244	0.000120

Figure A.7: Expected background event rate and error fit to match desired energy bins

	Detection Rate
2.0	0.000491
2.5	0.001299
3.0	0.002005
3.5	0.002366
4.0	0.002425
4.5	0.002123
5.0	0.001767
5.5	0.001409
6.0	0.001017
6.5	0.000681
7.0	0.000425
7.5	0.000219
8.0	0.000102
8.5	0.000043
9.0	0.000022

Figure A.8: Antineutrino detection rate for a simulated fresh core AFR-100 reactor

	Mean Counts	Uncertainty (1-Sigma)	Fractional Uncertainty (1-Sigma)
2.0	32895.851918	4848.934671	0.147403
2.5	54220.305249	5620.891899	0.103668
3.0	69135.620431	5390.974359	0.077977
3.5	83358.465788	6068.937562	0.072805
4.0	67512.617197	3905.267984	0.057845
4.5	63832.585929	4453.099950	0.069762
5.0	52249.983607	3854.948196	0.073779
5.5	55651.139059	5419.711259	0.097387
6.0	43665.704565	4892.487230	0.112044
6.5	19143.752041	2114.013549	0.110428
7.0	20721.121996	3753.625787	0.181150
7.5	14322.927910	3523.951053	0.246036
8.0	10223.033756	3169.244547	0.310010
8.5	10214.062157	3443.782321	0.337161
9.0	6288.629806	2851.829265	0.453490

Figure A.9: Antineutrino yields for a 3 month collection period with associated uncertainties for a simulated fresh core AFR-100 reactor

	Power Uncertainty	Fission Rate Uncertainty	AN Yield Uncertainty	Suite Uncertainty	Background Uncertainty	Source Counting Uncertainty	Background Counting Uncertainty
2.0	1.222392e-04	3.130209e-07	0.006991	0.002283	0.989205	0.000494	0.000905
2.5	6.300106e-04	1.645362e-06	0.013438	0.011864	0.972350	0.000972	0.000744
3.0	1.626698e-03	4.315967e-06	0.018816	0.030743	0.946431	0.001632	0.000747
3.5	1.792663e-03	4.830283e-06	0.021842	0.033784	0.940313	0.001519	0.000744
4.0	4.575065e-03	1.253429e-05	0.060115	0.085740	0.845130	0.003761	0.000665
4.5	2.731261e-03	7.619768e-06	0.046396	0.050531	0.897116	0.002532	0.000687
5.0	2.548919e-03	7.246958e-06	0.052073	0.046693	0.895162	0.002812	0.000704
5.5	8.198942e-04	2.387469e-06	0.021284	0.015023	0.960976	0.001134	0.000760
6.0	5.262432e-04	1.567561e-06	0.022796	0.009610	0.965242	0.001005	0.000819
6.5	1.267358e-03	3.892440e-06	0.099070	0.023042	0.872333	0.003602	0.000682
7.0	1.568020e-04	4.997129e-07	0.018630	0.002847	0.976895	0.000713	0.000758
7.5	4.692224e-05	1.636182e-07	0.021901	0.000859	0.976040	0.000417	0.000736
8.0	1.236556e-05	4.683401e-08	0.017435	0.000228	0.981307	0.000239	0.000779
8.5	1.883795e-06	8.431230e-09	0.003655	0.000035	0.995446	0.000086	0.000775
9.0	6.890171e-07	3.639708e-09	0.001737	0.000013	0.997476	0.000063	0.000710

Figure A.10: Contributing uncertainties for a simulated fresh core AFR-100 reactor

APPENDIX B

SUPPLEMENTAL MATERIAL FOR DIVERSION SCENARIO DETECTION CONFIDENCE

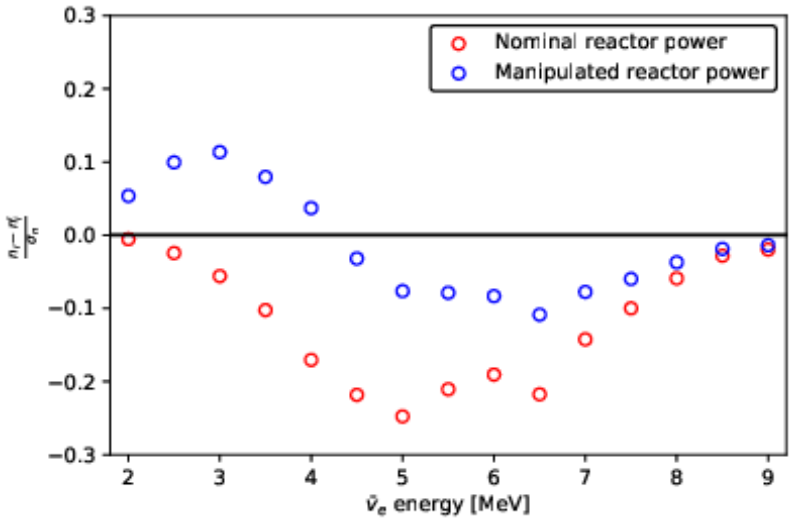

Figure B.1: Oscillation effect from goodness-of-fit function minimization for diversion 1b

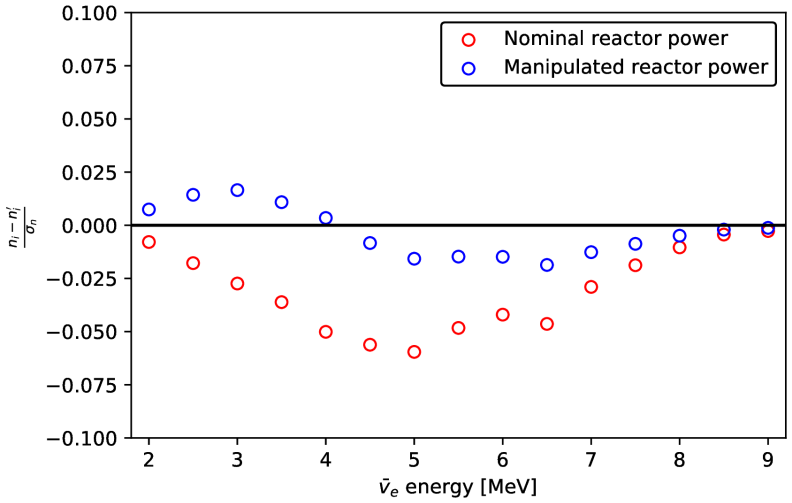

Figure B.2: Oscillation effect from goodness-of-fit function minimization for diversion 2a

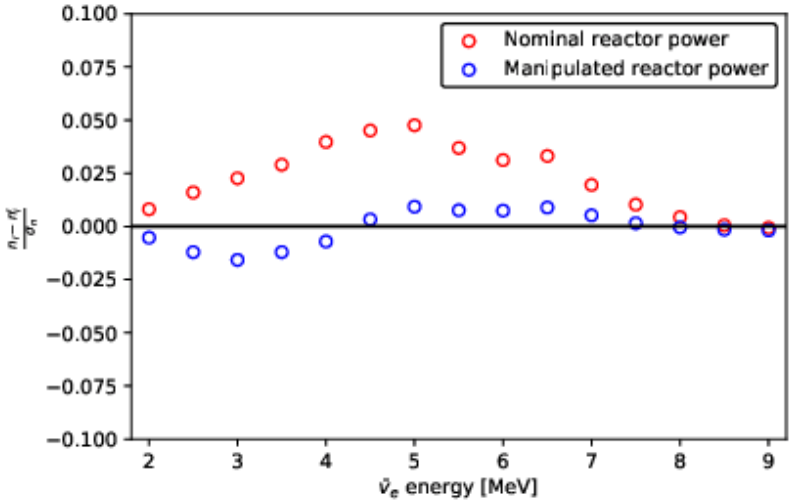

Figure B.3: Oscillation effect from goodness-of-fit function minimization for diversion 2b

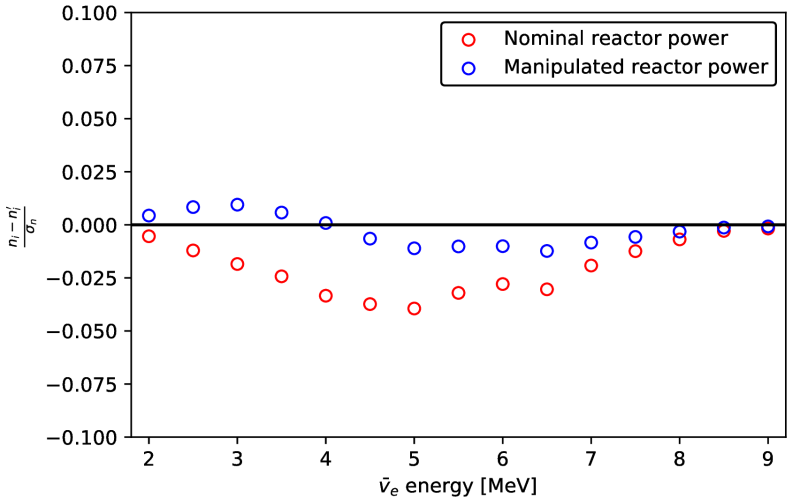

Figure B.4: Oscillation effect from goodness-of-fit function minimization for diversion 3a

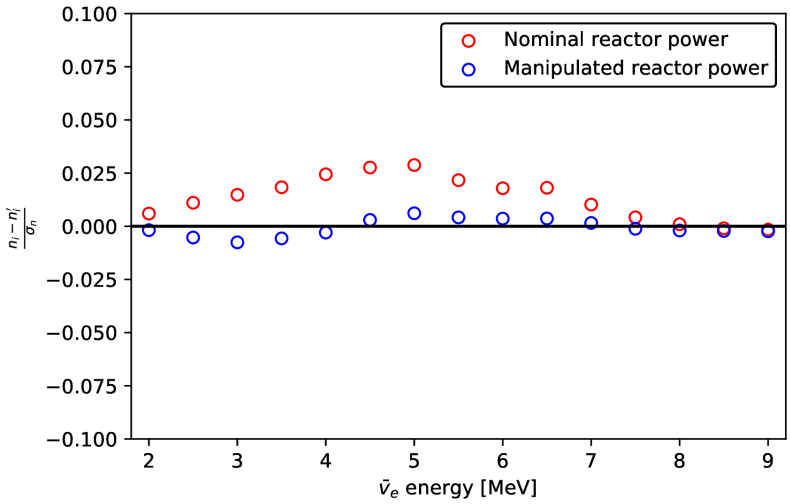

Figure B.5: Oscillation effect from goodness-of-fit function minimization for diversion 3b

REFERENCES

[1] *Energy - united nations sustainable development*, Nov. 2021.

[2] *World energy needs and nuclear power*, Nov. 2021.

[3] *Fact sheet: Advanced reactors, nuclear security and non-proliferation concerns*, Oct. 2021.

[4] Burek *et al.*, *Reinforcing the global nuclear order for peace and prosperity: The role of the iaea to 2020 and beyond*, May 2008.

[5] *The iaea and the non-proliferation treaty*, Jun. 2016.

[6] IAEA, "Options to enhance proliferation resistance of innovative small and medium sized reactors," International Atomic Energy Agency, Vienna, Tech. Rep. No. NP-T-1.11, 2014.

[7] E. Galdoz *et al.*, *Remote monitoring in safeguards: Security of information and enhanced cooperation*, Jul. 2011.

[8] C. Stewart, A. Abou-Jaoude, and A. Erickson, "Employing antineutrino detectors to safeguard future nuclear reactors from diversions," *Nature Communications*, vol. 10, no. 1, 2019.

[9] A. Bernstein, Y. Wang, G. Gratta, and T. West, "Nuclear reactor safeguards and monitoring with antineutrino detectors," *Journal of Applied Physics*, vol. 91, no. 7, pp. 4672–4676, 2002.

[10] Y. Kim, "Detection of antineutrinos for reactor monitoring," *Nuclear Engineering and Technology*, vol. 48, no. 2, pp. 285–292, 2016.

[11] N. S. Bowden, "Reactor monitoring and safeguards using antineutrino detectors," *Journal of Physics: Conference Series*, vol. 136, no. 2, p. 022 008, 2008.

[12] O. Akindele *et al.*, *Nu tools – final report*, Dec. 2021.

[13] V. Li, "Far-field monitoring of reactor antineutrinos for nonproliferation," *INMM 60th Annual Meeting*, Jul. 2019.

[14] E. Christensen, P. Huber, and P. Jaffke, "Antineutrino reactor safeguards: A case study of the dprk 1994 nuclear crisis," *Science amp; Global Security*, vol. 23, no. 1, pp. 20–47, 2015.

[15] M. Askins *et al.*, *The physics and nuclear nonproliferation goals of watchman: A water cherenkov monitor for antineutrinos*, 2015.

[16] A. Bernstein *et al.*, "Nuclear security applications of antineutrino detectors: Current capabilities and future prospects," *Science amp; Global Security*, vol. 18, no. 3, pp. 127–192, 2010.

[17] V. Rubchenya and J. Äystö, "Production of neutron rich isotopes in fission. a study for rnb facilities," *Nuclear Physics A*, vol. 701, no. 1-4, pp. 127–132, 2002.

[18] G. F. Knoll, *Radiation Detection and Measurement, Fourth Edition*. John Wiley & Sons, 2010.

[19] P. Vogel, L. Wen, and C. Zhang, "Neutrino oscillation studies with reactors," *Nature Communications*, vol. 6, no. 1, 2015.

[20] B. R. Sehgal, Ed., *Chapter 5 - Fission Product Release and Transport*. Boston: Academic Press, 2012, pp. 425–517, ISBN: 978-0-12-388446-6.

[21] Y. Wang, "Neutrino detectors: Present and future," *Physics Procedia*, vol. 37, pp. 22–33, 2012.

[22] "Neutrinos could shed light on why the universe has so much more matter than antimatter," *Nature*, vol. 580, no. 7803, pp. 305–305, 2020.

[23] F. Vannucci, "Interactions of neutrinos with matter," *Progress in Particle and Nuclear Physics*, vol. 95, pp. 1–47, Jul. 2017.

[24] S. Bilenky, "Neutrino oscillations: From a historical perspective to the present status," *Nuclear Physics B*, vol. 908, pp. 2–13, 2016, Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.

[25] G. Giacomelli, "The standard model of particle physics. neutrino oscillations," *Radiation Measurements*, vol. 44, no. 9-10, pp. 826–833, Oct. 2009.

[26] M. Gonzalez-Garcia, M. Maltoni, and T. Schwetz, "Global analyses of neutrino oscillation experiments," *Nuclear Physics B*, vol. 908, pp. 199–217, 2016.

[27] A. Formozov, *On the role of radiative losses in energy scale of large liquid scintillator and water cerenkov detectors*, 2018.

[28] O. Tomalak and R. Hill, "Theory of elastic neutrino-electron scattering," *Fermilab*, Feb. 2020.

[29] M. D. Bowen and P. Huber, "Inverse beta decay and coherent elastic neutrino nucleus scattering – a comparison," 2019.

[30] M. Kandemir and A. Cakir, "Comparison of plastic antineutrino detector designs in the context of near field reactor monitoring," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 927, pp. 353–361, 2019.

[31] P. Huber and T. Schwetz, "Precision spectroscopy with reactor antineutrinos," *Physical Review D*, vol. 70, no. 5, 2004.

[32] V. Fischer and E. Tiras, "Water-based liquid scintillator detector as a new technology testbed for neutrino studies in turkey," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 969, p. 163 931, 2020.

[33] G. Bak *et al.*, "Measurement of reactor antineutrino oscillation amplitude and frequency at reno," *Physical Review Letters*, vol. 121, no. 20, 2018.

[34] A. Conant, "Antineutrino spectrum characterization of the high flux isotope reactor using neutronic simulations," Ph.D. thesis, Georgia Institute of Technology, 2019.

[35] M. Estienne *et al.*, "Updated summation model: An improved agreement with the daya bay antineutrino fluxes," *Physical Review Letters*, vol. 123, no. 2, 2019.

[36] T. A. Mueller *et al.*, "Improved predictions of reactor antineutrino spectra," *Physical Review C*, vol. 83, no. 5, 2011.

[37] Dadwyer, *Dadwyer/oklo: A toolkit for modeling nuclides and nuclear reactions.*

[38] K. Schreckenbach, G. Colvin, W. Gelletly, and F. Von Feilitzsch, "Determination of the antineutrino spectrum from 235u thermal neutron fission products up to 9.5 mev," *Physics Letters B*, vol. 160, no. 4, pp. 325–330, 1985.

[39] F. von Feilitzsch, A. Hahn, and K. Schreckenbach, "Experimental beta-spectra from 239pu and 235u thermal neutron fission products and their correlated antineutrino spectra," *Physics Letters B*, vol. 118, no. 1, pp. 162–166, 1982.

[40] A. Hahn, K. Schreckenbach, W. Gelletly, F. von Feilitzsch, G. Colvin, and B. Krusche, "Antineutrino spectra from 241pu and 239pu thermal neutron fission products," *Physics Letters B*, vol. 218, no. 3, pp. 365–368, 1989.

[41] N. Haag *et al.*, "Experimental determination of the antineutrino spectrum of the fission products of 238u," *Physical Review Letters*, vol. 112, no. 12, 2014.

[42] P. Huber, "Determination of antineutrino spectra from nuclear reactors," *Physical Review C*, vol. 84, no. 2, 2011.

[43] I. Machikhiliyan, "The danss neutrino spectrometer: The results of reactor antineutrino studies," *Phys. Part. Nuclei*, vol. 53, pp. 546–551, 2022.

[44] V. Klimov, V. Kopeikin, L. Mikaelyan, K. Ozerov, and V. Sinev, "Neutrino method remote measurement of reactor power and power output," *Atomic Energy*, pp. 123–127, 1994.

[45] N. Bowden *et al.*, "Experimental results from an antineutrino detector for cooperative monitoring of nuclear reactors," *Nuclear Instruments and Methods in Physcis Research A*, vol. 572, pp. 985–998, 2007.

[46] M. G. T. Lasserre, "Double chooz, a search for the neutrino mixing angle theta-13," 2006.

[47] Y. Kuroda *et al.*, "A mobile antineutrino detector with plastic scintillators," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 690, pp. 41–47, 2012.

[48] J. Ashenfelter *et al.*, "The prospect reactor antineutrino experiment," *Nuclear Instruments and Methods*, Aug. 2019.

[49] M. Andriamirado *et al.*, "Improved short-baseline neutrino oscillation search and energy spectrum measurement with the prospect experiment at hfir," *Phys. Rev. D*, vol. 103, p. 032 001, 3 Feb. 2021.

[50] J. Ashenfelter *et al.*, "Background radiation measurements at high power research reactors," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 806, Jan. 2016.

[51] C. Stewart, "Antineutrino-based safeguards for ultra-high burnup fast reactors," Ph.D. thesis, Georgia Institute of Technology, 2016.

[52] J. Ashenfelter *et al.*, "The prospect physics program," *Journal of Physics G: Nuclear and Particle Physics*, vol. 43, no. 11, Oct. 2016.

[53] C. Stewart and A. Erickson, "Antineutrino analysis for continuous monitoring of nuclear reactors: Sensitivity study," *Journal of Applied Physics*, vol. 118, no. 16, p. 164 902, 2015.

[54] C. Forsberg, P. Peterson, and P. Pickard, *The advanced high-temperature reactor (ahtr)*, 2001.

[55] S. R. Greene *et al.*, *Pre-conceptual design of a fluoride-salt-cooled small modular advanced high temperature reactor (smahtr)*, Feb. 2011.

[56] P. C. Durst *et al.*, "Nuclear safeguards considerations for the pebble bed modular reactor (pbmr)," *Idaho National Laboratory*, 2009.

[57] H. Mohamed, D. Kotlyar, and Y. Shaposhnik, "Coupled neutronic-thermal-hydraulic analysis of a small fhr core with pin-type fuel assemblies," *Physor*, 2014.

[58] H. Mohamed and G. Parks, *Modeling a pin-type fuel assembly for a small fluoride salt cooled high temperature reactor (fhr)*, 2014.

[59] C. Grandy *et al.*, "Advanced fast reactor - 100 (afr-100) report for the technical review panel," 2014.

[60] A. Nikiforova, P. Hejzlar, and N. E. Todreas, "Lead-cooled flexible conversion ratio fast reactor," *Nuclear Engineering and Design*, vol. 239, no. 12, pp. 2596–2611, 2009.

[61] N. Todreas and P. Hejzlar, "Flexible conversion ratio fast reactor systems evaluation," 2008.

[62] V. Bulaevskaya and A. Bernstein, "Detection of anomalous reactor activity using antineutrino count evolution over the course of a reactor cycle," *Journal of Applied Physics*, vol. 109, no. 11, p. 114 909, 2011.

[63] M. Andriamirado *et al.*, "Prospect-ii physics opportunities," Feb. 2022.

[64] J. Leppanen, *Serpent - a monte carlo reactor physics burnup calculation code*.

[65] *Past, present and future challenges of developing the serpent monte carlo code*, Jul. 2019.

[66] Z. Djurcic, J. A. Detwiler, A. Piepke, V. R. Foster, L. Miller, and G. Gratta, "Uncertainties in the anti-neutrino production at nuclear reactors," *Journal of Physics G: Nuclear and Particle Physics*, vol. 36, no. 4, p. 045 002, 2009.

[67] D. Rochman and C. Sciolla, *Nuclear data uncertainty propagation for a typical pwr fuel assembly with burnup*, Apr. 2015.

[68] *Neutron energy: Classification of neutrons*, Oct. 2021.

[69] Libretexts, *Propagation of error*, Aug. 2020.

[70] *Neutron/gamma psd ej-301, ej-309*.

[71] V. Vyrodov *et al.*, "Precise measurement of the cross section for the reaction +p $\rightarrow$ e++n at the bourges reactor," *ZhETF Pisma Redaktsiiu*, p. 161, Feb. 1995.

[72] M. Dunbrack, C. Stewart, and A. Erickson, "High bias machine learning for antineutrino-based safeguards for small reactors," *Annals of Nuclear Energy*, vol. 169, May 2022.

[73] IAEA, "Iaea safeguards glossary: 2001 edition," International Atomic Energy Agency, Vienna, Tech. Rep. 3, 2002.

[74] C. Lee, Y. Jung, and W. Yang, 2018.

[75] *Software: Rebus-3 (fuel cycle / depletion codes)*, 2016.

[76] M. Blennow, P. Coloma, P. Huber, and T. Schwetz, "Quantifying the sensitivity of oscillation experiments to the neutrino mass ordering," *Journal of High Energy Physics*, vol. 2014, no. 3, Feb. 2014.

[77] X. Qi, S. Silvestrov, and T. Nazir, "Data classification with support vector machine and generalized support vector machine," *AIP Conference Proceedings*, 2017.

[78] F. Rossi and N. Villa, "Support vector machine for functional data classification," *Neurocomputing*, vol. 69, no. 7-9, pp. 730–742, 2006.

[79] M. Singla and K. K. Shukla, "Robust statistics-based support vector machine and its variants: A survey," *Neural Computing and Applications*, vol. 32, no. 15, pp. 11 173–11 194, 2019.

[80] L. Bottou and C.-J. Lin, "Support vector machine solvers," *Large-Scale Kernel Machines*, 2007.

[81] P. Mehta *et al.*, "A high-bias, low-variance introduction to machine learning for physicists," *Physics Reports*, vol. 810, pp. 1–124, 2019.

[82] T. Hastie, J. Friedman, and R. Tisbshirani, *The elements of Statistical Learning: Data Mining, Inference, and prediction*. Springer, 2017.

[83] O. A. Akindele, A. Bernstein, and E. B. Norman, "Antineutrino monitoring of thorium reactors," *Journal of Applied Physics*, vol. 120, no. 12, p. 124 902, 2016.