Contract No:

This document was prepared in conjunction with work accomplished under Contract No. DE-AC09-08SR22470 with the U.S. Department of Energy (DOE) Office of Environmental Management (EM).

Disclaimer:

We put science to work.™

Savannah River National Laboratory™

OPERATED BY SAVANNAH RIVER NUCLEAR SOLUTIONS

A U.S. DEPARTMENT OF ENERGY NATIONAL LABORATORY • SAVANNAH RIVER SITE • AIKEN, SC

Event Definition for the Automated Detection of Nuclear Proliferation Activities

T. L. Danielson

A. A. Kail

J. A. Pike

April 2020

SRNL-STI-2020-00155, Revision 0

SRNL.DOE.GOV

DISCLAIMER

Printed in the United States of America

Prepared for
U.S. Department of Energy

SRNL-STI-2020-00155
Revision 0

Keywords: *Artificial Intelligence*
Machine Learning
Proliferation

Retention: *Permanent*

Event Definition for the Automated Detection of Nuclear Proliferation Activities

T. L. Danielson
A. A. Kail
J. A. Pike

April 2020

Savannah River National Laboratory™
OPERATED BY SAVANNAH RIVER NUCLEAR SOLUTIONS

Prepared for the U.S. Department of Energy under contract number DE-AC09-08SR22470.

REVIEWS AND APPROVALS

AUTHORS:

T. L. Danielson, Environmental Modeling Date

A. A. Kail, Integrated Computing Date

J. A. Pike, Advanced Modelling, Simulation and Analytics Date

APPROVALS:

D. A. Crowley, Manager, Environmental Modeling Date

J. A. Horne, Acting Manager, Integrated Computing Date

P. L. Lee, Manager, Advanced Modelling, Simulation and Analytics Date

EXECUTIVE SUMMARY

In FY2020, Savannah River National Laboratory (SRNL) in collaboration with the Discovery Analytics Center (DAC) at Virginia Polytechnic Institute and State University (VT) began developing a demonstration prototype system that uses multiple machine learning and data analytic methods on large-scale open data sources to identify new, developing, or undeclared nuclear programs. One of the most challenging aspects of applying machine learning techniques to such a problem is the high likelihood of extremely sparse data from disparate sources. To overcome this challenge, the current work will use a strategic combination of supervised, semi-supervised, and unsupervised learning techniques to ingest and fuse data streams to make a forecast of nuclear activities in a targeted geospatial location. Identifying potential data sources and training supervised learning algorithms is dependent upon the development of a robust foundation of targeted event domains that fundamentally define the nuclear activities of interest. This report documents the definition of a hierarchical structure for both nuclear activity and event domains that will be used to guide the research team in development or use of existing semantic dictionaries that are instrumental to searching, parsing, and categorizing events for the forecasting system's use.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

AGR	Gas-cooled Reactor
BWR	Boiling Water Reactor
DAC	Discovery Analytics Center
DOE	Department of Energy
FBR	Fast Neutron Reactor
IAEA	International Atomic Energy Agency
LWGR	Light Water Graphite Reactor
NRC	Nuclear Regulatory Commission
PHWR	Pressurized Heavy Water Reactors
PNNL	Pacific Northwest National Laboratory
PWR	Pressurized Water Reactors
SNM	Special Nuclear Materials
SRNL	Savannah River National Laboratory
SRPPF	Savannah River Plutonium Processing Facility
SRS	Savannah River Site
US-NRC	United States Nuclear Regulatory Commission
VT	Virginia Polytechnic Institute and State University

1.0 Introduction

In FY2020 the Savannah River National Laboratory (SRNL) in collaboration with the Discovery Analytics Center (DAC) at Virginia Polytechnic Institute and State University (VT), and funded by the Department of Energy's (DOE) Defense Nuclear Nonproliferation Research and Development, began developing a demonstration prototype system that uses multiple machine learning and data analytic methods on large-scale open data sources to identify new, developing, and/or undeclared nuclear programs. This report documents the definition of events of interest and a hierarchical structure that will be used in the forecasting system. This report will also guide the research team in development or use of existing semantic dictionaries that are instrumental to searching, parsing, and categorizing events [1] [2] [3]. In addition, various geolocations around the United States are identified for sensitivity and null case testing.

A previous successful forecasting system [4] [5] that this project is adapting for use depended on rapidly evolving events in days and weeks where numerous events provide considerable training and learning data. Nuclear weapons program development is inherently a long evolving program over several years where many aspects overlap with nuclear power production without numerous examples for training data. In this case, each milestone along the path of a weapons program becomes an event of interest to the ultimate concluding event of a weapon prototype. A program milestone would occur only once in a geolocation such that after the milestone, the capability or activity can occur at will. For example, many locations have nuclear power reactors where spent fuel accumulates. To create the capability to reprocess the fuel, years of program development, design, and construction would occur before startup but once the facility is operating, fuel reprocessing can occur at-will. The event of interest would be reprocessing spent nuclear fuel, but several related events are expected to be defined or developed over time to improve the ability to differentiate or categorize discovery of nuclear related events that are associated with the reprocessing facility. Each step along the way of development and program decision making involves people and activities that make up smaller, identifiable events more readily found in open sources. The accumulation of the related event discoveries is hypothesized to indicate when the primary event of interest is under development, or exists, and remains unreported in open sources. These related events become the relevant events of interest that can be used in this forecasting system.

To test the hypothesis, this project is focused on one small aspect of a nuclear weapons program, where a significant opportunity presented itself with the announcement of the proposed Savannah River Plutonium Processing Facility (SRPPF). Here, core weapon component manufacturing is planned to start by 2030. DOE announced this project on May 10, 2018 via a release "Joint Statement from Ellen M. Lord and Lisa E. Gordon-Hagerty on Recapitalization of Plutonium Pit Production"

> *To achieve DoD's 80 pits per year requirement by 2030, NNSA's recommended alternative repurposes the Mixed Oxide Fuel Fabrication Facility at the Savannah River Site in South Carolina to produce plutonium pits while also maximizing pit production activities at Los Alamos National Laboratory in New Mexico. [6]*

The facility is planned to start up in 2030 which offers two possible courses of development, one based on open source data and predictability before the announcement and the other based on accumulation of data after the announcement that would detail the many related events that could be used for developing detailed related event definitions for a forecasting system. In this project we will focus first on predictability before the announcement and consider how to accumulate information for the latter as time and resources allow.

This document is anticipated to be updated as the project evolves, and necessary refinements are identified.

The following sections discuss the initial approach and define the structure of the numerous nuclear related events of interest, geospatial locations that will be used for evaluating system performance, and performance criteria.

2.0 Defining Nuclear Related Events of Interest

Numerous nuclear related events routinely occur in most areas of the United States because of uranium enrichment operations, reactor fuel manufacturing, power production, nuclear medicine, and research as well as DOE production and legacy facilities. Among other challenges, detecting the inception of a new facility (i.e., not yet announced) in a massive open data environment is inherently difficult due to a high signal-to-noise ratio that results from continuously evolving information about existing production facilities. In other words, the signal belongs to a sparse data set that is derived from potentially disparate data sources, where no single nuclear related event discovery is anticipated to be a direct predictor. Therefore, the system will necessarily accumulate events over the time period before any first public announcement and cumulatively estimate the likelihood (to be developed) of an incipient, developing, or existing, undeclared, facility or program. As such, the machine learning methods and models will rely on the use of ontological-type hierarchies to support taxonomies for categorizing all discovered events. Substantial work has been done to develop ontologies on various parts of the nuclear fuel cycle and component fabrication. [1] [2] [3] Wherever possible, existing ontologies will be adapted and used for this project.

A previous PNNL lead study developed a detailed schematic diagram of the nuclear fuel cycle and weapons development process which is known as the "Wacker chart" after the lead author (reference unavailable). A very condensed overview of the Wacker chart is illustrated in Figure 2-1, where the processes for making special nuclear materials are on the left (i.e., the nuclear fuel cycle) and those for making the weapons are on the right (i.e., the weapons development processes). Those processes and activities associated with the making of Special Nuclear Materials (SNM)[1] are considered to occur in the "early detection period." Figure 2-2 shows the summary nuclear fuel cycle and weapons development process schematic. The red circled item represents where the proposed SRPPF falls on the chart under item 18.1.3.

The Wacker chart represents one subset of all possible pathways to building a nuclear weapons capability and is useful here to illustrate the specific, narrow focus of this project and will serve as an information guide for defining an ontological structure of nuclear activity domains that illustrate the preliminary information relationships in the nuclear proliferation conceptual domain. The chart will also help guide interpretation and adaptation of ontologies that exist for specific parts of the chart. One should note that existing ontology work may not cover everything needed for this project nor are the computational approaches being limited to computational ontology. Any needed gaps in existing ontologies will be addressed by the research team when discovered.

The proposed SRPPF has several advantages that will facilitate long-term success in the forecasting system's development. First, the researchers have access to data streams that pre-date the first announcement of the mission and therefore, the algorithms' temporal detection capabilities (i.e., detection at the earliest possible stages of mission development) can be quantified and refined. Next, the proposed SRPPF has been widely publicized and therefore, if all data streams are allowed and processed, the detection of the mission should not be particularly challenging. However, having such publicly available announcements regarding the mission and its evolution will allow the researchers to incrementally restrict the algorithms' access and quantify the capabilities using progressively more sparse and disparate data streams, thereby providing a proof-of-concept in more data-restrictive environments. Finally, the proposed SRPPF is a relatively specific activity that acts as only one step in the broader activity of weapons production and therefore will allow fine tuning of the forecasting system to detect activities at such a level of specificity.

---

[1] SNM is defined by Title I of the Atomic Energy Act of 1954 as plutonium, uranium-233, or uranium enriched in the isotopes uranium-233 or uranium-235, but does not include source material. The definition includes any other material that the NRC determines to be special nuclear material although the NRC has not declared any other material as SNM.

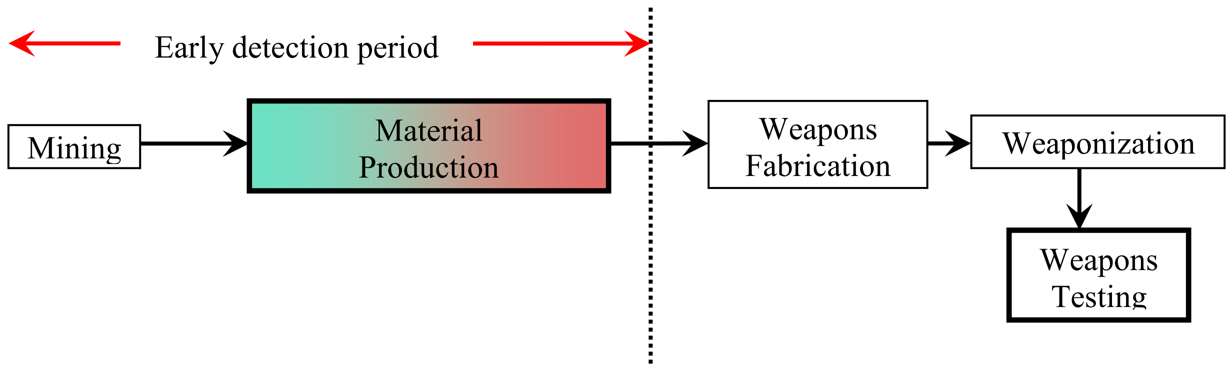

Figure 2-1. Condensed overview of the nuclear weapons production pathway

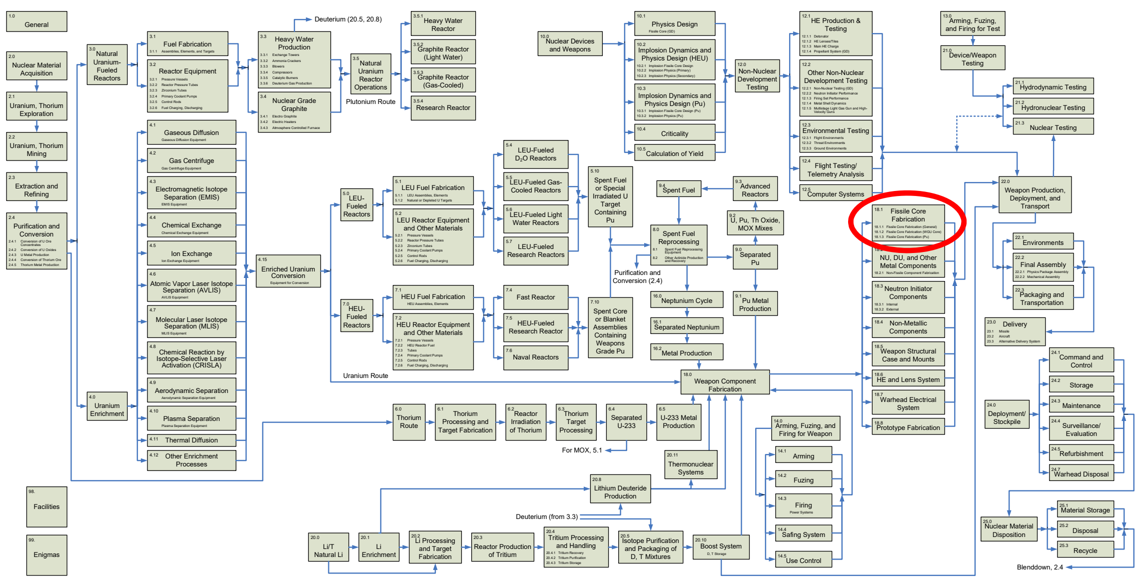

Figure 2-2. Summary Nuclear Fuel Cycle and Weapons Development Process

2.1 Activity Domains of Interest

Before describing in detail the activity domains of interest and the event domains of interest, a brief definition of the relationship between activities and events is merited. Activities are topical domains which describe a collection of functionally related nuclear activities whose features are sought in the open source. Events are domains containing features of activities. Therefore, the discovery of an event (as will be outlined in Section 2.2) in the open source provides an indicator that a nuclear activity of interest is occurring. Events and activities are unique, but an event can be a feature of multiple activities. Therefore, it is a unique combination of events that defines an activity. In other words, the target activity domain's unique events can be obtained by taking the difference between the set of events that define the target activity and the set of events that define all other activities.

Figure 2-3 illustrates a hierarchical structure of nuclear activity domains that are the foundation of the event definitions where each activity domain represents a set of event domain features. The level of specificity of a given domain changes with its rank, illustrated here by the vertical positioning (i.e., least specific to most specific from top to bottom). The top-level domain, "Nuclear Activities", is made up by two sub-domains: "Weapons" and "Other". Because the long-term vision of the current work includes a forecasting system that is applicable to any activity sub-domain of "Weapons", "Fissile Core Fabrication" is allocated to a more general sub-domain "Weapons Component Fabrication". The singular goal of the current work, however, is to discover only the events which are unique to "Fissile Core Fabrication" (i.e., a weapons component) in the open source in order to make an inferential forecast of the SRPPF at the earliest possible point in time. Note that the various sub-domains in the hierarchy that have dotted outlines are being used only to illustrate that additional sub-domains may exist at a given rank, but anything other than what is labeled falls outside the scope of the current work.

As can be extracted from Figure 2-2, there is a substantial amount of potential overlap in the start-to-finish high-level workflow used to arrive at "Fissile Core Fabrication" compared to the workflow used to arrive at other nuclear activities that do not result in any weapons-related endpoint. To distinguish "Fissile Core Fabrication" from the occurrence of other nuclear activities (i.e., a classification problem), it will be necessary to train machine learning algorithms with data from all domains where the specificity of the learned features is correlated with the rank of the sub-domain. Following this concept, the sub-domain "Other" is populated with sub-domains that are anticipated to overlap with the target sub-domain of "Fissile Core Fabrication". Reiterating, a sub-domain's specificity increases as sub-domain rank increases and therefore, a higher rank in any branch of the hierarchy implies a higher degree of anticipated overlap between domains. For example, here, the sub-domain "Commercial Power" is presumed to have more overlap with "Fissile Core Fabrication" than the sub-domain "Medical, Research, Industrial" and thus, it has a higher rank which matches that of "Fissile Core Fabrication".

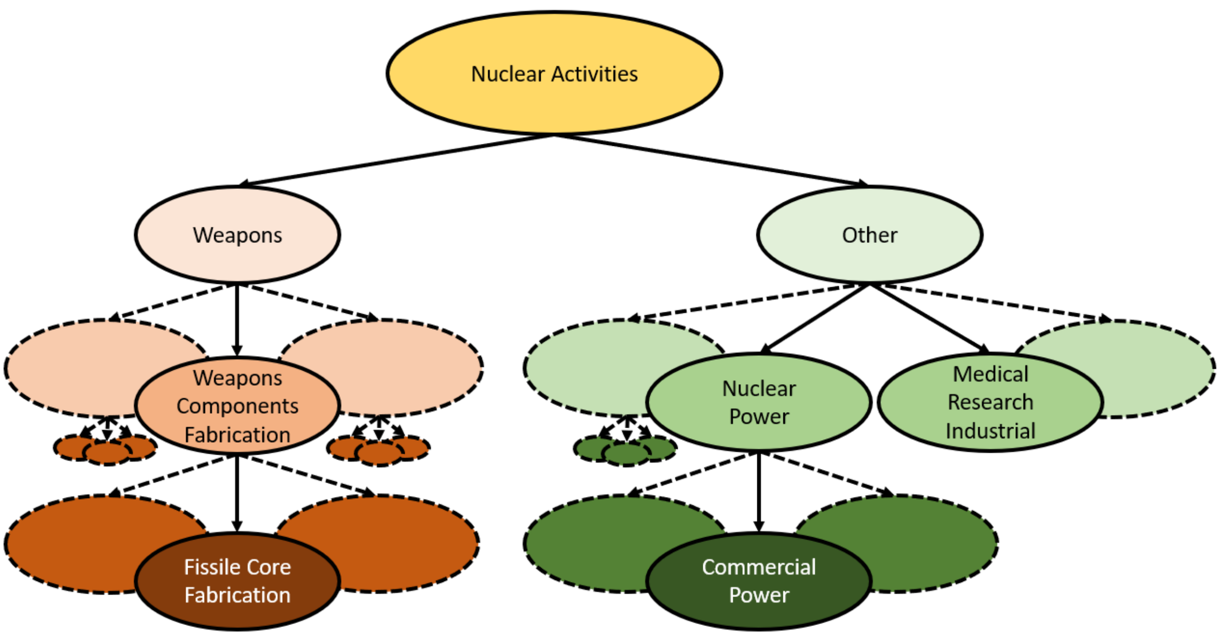

Figure 2-3. Hierarchical structure for defining the activity domains of interest.

The intended use of this hierarchical structure is two-fold. First, it allows for the specification of well-defined topical domains such that an algorithmic approach can be used for targeted data ingestion. Second, it defines the classification problem that will eventually lead to a forecast with a calculated degree of confidence. Through the collection of features from the open source, such a hierarchical structure would, at a minimum, allow a true positive binary classification for the occurrence of "Fissile Core Fabrication" if all the sub-domain's unique features could be discovered. Such success is unlikely, and therefore machine learning models will be tuned to appropriately ingest and fuse features from the open source for making such conclusions.

In the following sub-sections, further description of the highest rank activity sub-domains will be provided and a selection of potentially detectable features (i.e., events) of the activities will be highlighted. The examples of detectable features that are presented are not to be taken as an all-inclusive set, but rather a cursory and superficial subset to assist in illustrating the kinds of events/indicators the forecasting system will be equipped to find.

2.1.1 Commercial Power[2] and Weapons

The activity "Commercial Power" precisely refers to the design, construction, or operation of a commercial nuclear reactor with a mission entirely centered around providing electrical power to a country's energy grid. Many similarities exist among the design features, material processing flowsheets, and nuclear fuel cycles of common commercial nuclear reactors and those used for producing nuclear weapons grade materials as part of a nuclear proliferation effort. Under certain design criteria, use of a nuclear reactor can be multi-purpose (i.e., both production of electricity and weapons grade materials are accomplished) and/or can potentially be converted to production reactors. Therefore, a high-level overview of nuclear reactors is provided to highlight key details that will assist in distinguishing between the activities "Fissile core fabrication" and "Commercial Power". In this overview, emphasis is placed on identifying important non-nuclear materials components, nuclear materials, re-processing, timeline of fuel cycle(s), and reaction byproducts (e.g., fission products, activation products, etc.).

While the specific design features of commercial nuclear reactors continuously evolve to address needs pertaining to safety and efficiency, the most common characterizations include: 1) Pressurized Water Reactors (PWRs), 2) Boiling Water Reactors (BWRs), 3) Pressurized Heavy Water Reactors (PHWRs), 4) Gas-cooled Reactor (AGR), 5) Light Water Graphite Reactors (LWGRs), and 6) Fast Neutron Reactors (FBRs). In all cases, the basis for electrical power generation is the same: Energy released from the continuous fission of atoms transfers heat to water (or a gas) to produce steam that drives a turbine which in turn produces electricity. The primary components of commercial nuclear reactors include the fuel, moderator, control rods, coolant, pressure vessel, a steam generator, and containment.

The fuel, moderators, and coolants for the various types of reactors are shown in Table 2-1. Thorium is also a viable fuel, though it is not currently in use in any known commercial or large-scale facility. Each type of fuel requires varying degrees of differences in the nuclear fuel cycle. A schematic of the open and closed nuclear fuel cycle for various fuels (i.e., natural U or UO_2/PuO_2) is shown in Figure 2-4. Note that if natural UO_2 is used, the enrichment step is by-passed and the converted natural Uranium goes directly to pellet fabrication. This example highlights how the detection of specific events might allow one to infer the specific nature of the activity that is occurring. That is, if no enrichment is detected, it could be an indicator that 1) no weapons are being created in that specific location and 2) that the reactor is possibly a PHWR or an AGR. This example is only to illustrate the point – in practice, several indicators would be needed for inferential forecasting of activities. Going one step further, if the production or acquisition of heavy water (i.e., D_2O) is detected, this would provide further evidence of a PHWR. If, on the other hand, a signal can be obtained from the acquisition of CO_2 or graphite (minus enrichment), this would provide evidence in the direction of an AGR. Another potential indicator is the extent to which enrichment is being performed, if at all. Most nuclear reactors dedicated to commercial power require enrichment to 3.5-5% U-235, whereas weapons will be enriched to greater than 90%. Indicators pertaining to the enrichment may

[2] Summary information related to commercial nuclear power has been extracted from resources gathered from the World Nuclear Association information library [14].

potentially be embedded in the overall makeup of the enrichment facility (e.g., its overall size or material cost) about which the forecasting system would seek to identify surrogate data.

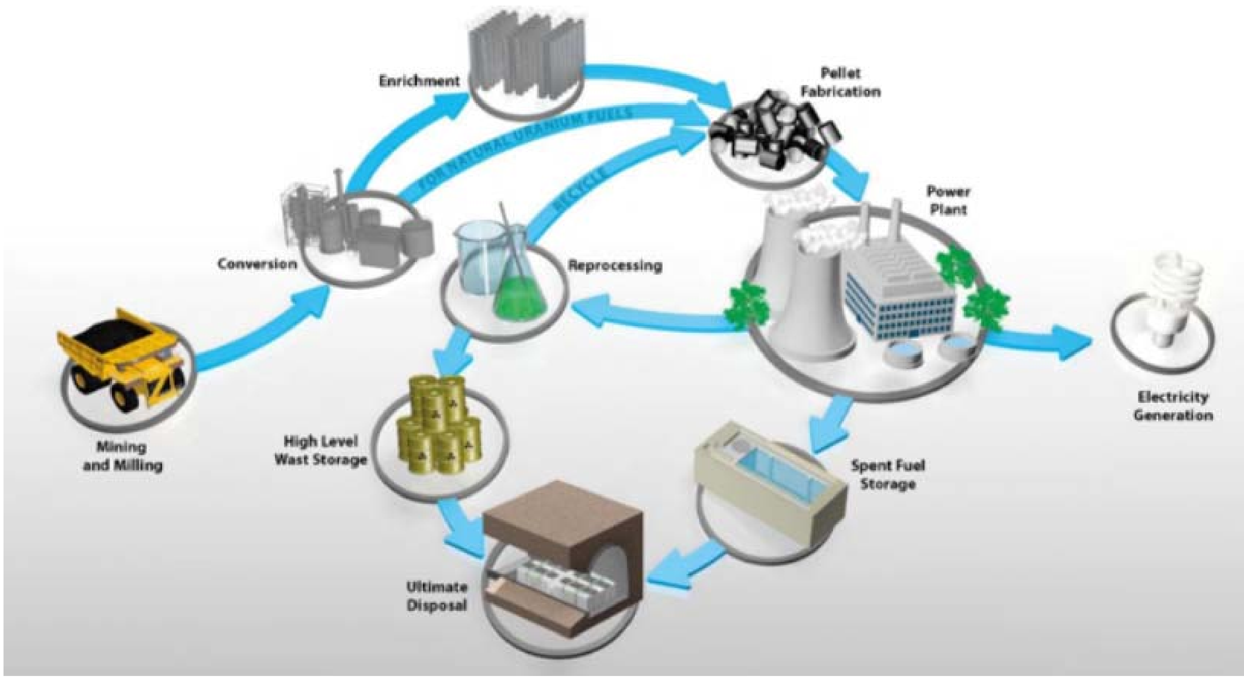

Figure 2-4. Nuclear fuel cycle for commercial nuclear reactors.

Table 2-1. Reactor types, fuels, moderators, and coolants.

Reactor Type	Fuel	Moderator	Coolant
PWR	Enriched UO_2	H_2O	H_2O
BWR	Enriched UO_2	H_2O	H_2O
PHWR	Natural UO_2	D_2O	D_2O
AGR	Natural U (metal) or Enriched UO_2	CO_2	Graphite
LWGR	Enriched UO_2	H_2O	Graphite
FBR	PuO_2 and UO_2	Liquid Na	none

During pellet fabrication, fuel pellets are fabricated and arranged in long tubular fuel rods, which most typically use a zirconium alloy, zircaloy, due to its high neutron transparency coupled with its capability of withstanding the potential high temperatures and pressures of the reactor environment. Because of its nuclear applications, zirconium is subject to controls on trading – especially "nuclear grade" zirconium. Generally, the quantity of zircaloy that is needed for a fuel assembly is on the order of 100 kg. "Non-nuclear grade" zirconium generally has elevated concentrations of hafnium present which decreases the overall performance of the alloy but is not subject to the same trade controls. Given the controls placed on these non-nuclear materials, and the quantities needed, some indicator of this commerce could potentially be found in the open source. If "non-nuclear grade" zirconium is purchased and hafnium impurities need to be extracted, some fingerprint of this materials processing scheme may be identified. For reactors that run at lower temperatures (i.e., reactors *not* used for commercial power generation), alternative alloys such as aluminum can be used. If there is a known nuclear reactor, but no indicator of zirconium acquisition, this may indicate that the reactor is not intended for commercial power.

The reactor vessel and steam generators require specialized corrosion resistant materials capable of withstanding high neutron flux, high temperatures, high pressures, and the presence of nuclear transmutation products (e.g., helium). With each evolving generation of reactor, the materials science continues to evolve providing increased resistance to radiation damage. Existing reactors generally use ferritic steels clad with austenitic stainless steel for the reactor vessel and Inconel alloys for the steam generator components. In new generations of reactors, oxide dispersion strengthened steels have gained interest for use in both. The overall containment (i.e., structural surrounding of the reactor) is comprised of highly reinforced concrete that can be up to 1-meter thick and may have detectable indicators in commerce and transit as well. Indicators of the acquisition of these large reactor components with specialized materials may be detected at several stages including the time of purchase and/or transit from the manufacturer to the reactor location. If a steam generation system is detected, this is indicative of the use of a PWR, or alternatively, that a BWR is not present because BWRs do not use steam generators.

Newly established reactors with new fuel require a neutron source in order to initiate the reaction. Often, this is made of beryllium in combination with an alpha emitter such as polonium-210, radium-226, or Plutonium-238, Americium-241, or with a source that undergoes spontaneous fission such as Californium-252. The detection of the acquisition of these specialized materials could be an indicator for the age of a reactor. Generally, the refueling of the reactor is performed at 12, 18, or 24-month intervals. For some designs, the plant must go to complete shut-down, while others can be refueled at reduced loading. If refueling events can be detected, this can provide insights to the reactor's age and the type of reactor that is operating. Spent fuel is sent for re-processing after a cool-down period. Depending on the type of fuel, the re-processing schemes can vary and result in different waste products and chemical and material acquisition events that, if detected, can provide insights to the type of reactor and its purpose. As can be seen in Figure 2-2, the activities that occur after the fuel is removed from the reactor provide some of the most significant evidence toward the intentions of the reactor (i.e., proliferation or not-proliferation). A once through fuel cycle with no reprocessing is a strong indicator for a power reactor. Closed fuel cycles create more possibilities without prior knowledge and therefore potential indicators of waste streams, and reprocessing methodologies would be necessary. It should be noted that the forecasting system is intended to identify fissile core fabrication *before* a core is produced and therefore, if detection occurs based on indicators from ongoing activities at the back end of the fuel cycle, there is a high likelihood that the forecast is too late. On the other hand, if the planning or construction of facilities that are used on the back end of the fuel cycle can be detected, the forecasting system would prove successful.

Commercial nuclear power has the potential to provide a substantial portion of the total energy grid to the population in the region. By examining the typical energy demands and the trends through time (i.e., increasing or decreasing energy demand), it may be possible, along with additional input data, to identify the presence of a reactor, the type of reactor, and its intended purpose.

2.1.2 Medical, Research, and Industrial Applications

The activities that occur in "Medical, Research, Industrial" can be broadly categorized as activities that rely on nuclear materials reprocessed from irradiated targets or spent reactor fuel. In most cases, the overlap with either "Commercial Power" or "Fissile core fabrication" is presumed to be negligible and/or superseded by "Commercial Power". Consequently, less specific training on sub-domain features will be required for successful and accurate classification and/or prevention of false positive "Fissile core fabrication" forecasts. Thus, only a brief description of activities in "Medical, Research, Industrial" is provided, but a more detailed description can be found at the US Nuclear Regulatory Commission (US-NRC) [7] and International Atomic Energy Agency (IAEA) websites [8].

2.1.2.1 Medical

Nuclear medicine makes use of radionuclides in the diagnosis and/or treatment of disease. More specifically, it encompasses two categories, diagnostic medical imaging and interventional medicine or radiotherapy. In diagnostic medical imaging, radiopharmaceuticals, or pharmaceutical drugs containing radioactive isotopes, are taken internally and an external detector captures images from the radiation emitted. In interventional nuclear medicine, radionuclides are used as the source of treatment for diseases such as hyperthyroidism, thyroid cancer, and blood disorders. The treatment is provided intravenously or orally either directly into involved organs, or nearby. In many cases, both categories use radionuclides that can be produced as by-products of uranium or plutonium fission, or with research reactors and accelerators.

Commonly used isotopes in diagnostic medical imaging applications include:

- flourine-18
- gallium-67
- krypton-81m
- rubidium-82
- nitrogen-13
- technetium-99m
- indium-111
- iodine-123
- xenon-133
- thallium-201

Commonly used isotopes in interventional medical applications include:

- yttrium-90
- iodine-131
- lutetium-177
- samarium-153
- strontium-99
- caesium-137
- cobalt-60
- iridium-192
- iodine-125
- palladium-103
- ruthenium-106

2.1.2.2 Research

The activity "Research" refers to the development or application of peaceful nuclear technologies that are not intended to be, or are not yet, commercially viable. A broad and comprehensive list of nuclear research activities might fall under the following descriptions:

- Accelerator and beam technologies
- Nuclear reactor physics, chemistry, instrumentation, and control
- Radiation shielding, detection, and measurement

- Nuclear materials and fuels, radiochemistry, and reprocessing
- Radioactive waste management
- Fusion energy research
- Environmental science
- Radiometric dating
- Isotope production
- Material irradiation
- Teaching/Training
- Nuclear data measurements

2.1.2.3 Industrial

Industrial applications of nuclear material are wide and varied and, similar to medical, are not directly related to the production of the material, but instead in the downstream use of material produced during the nuclear fuel cycle. Many industry uses are now considered unsafe and no longer in use, such as the use of uranium or thorium as a pigment in ceramics and glassware, or radium in watch faces and airplane gauge dials. Many of these applications are subject to strict requirements by the US-NRC and IAEA. Some common industrial uses of nuclear material include the following applications:

- Industrial radiography
- Irradiation
- Well logging
- Gauging devices
- Power source
- Calibration of instruments
- Various uses of depleted uranium

Commonly used radionuclides in these applications include:

- Irridium-192
- Cobalt-60
- Cesium-137
- Americium-241
- Beryllium-9
- Radium-226
- Uranium-235
- Uranium-238
- Tritium
- Plutonium-238

2.2 <u>Event Domains of Interest</u>

To reiterate from Section 2.1, events are features of activities that are sought in the open source to allow inferential forecasting of the planned, ongoing, or past occurrence of "Nuclear Activities". More specifically, the activities "Fissile core fabrication", "Commercial Power", and "Medical, Research, and Industrial" are defined by a unique combination of features that belong to four event domains: "Acquisition Events", "Political/Diplomatic Events", "Economic Events", and "Population/Personnel Events". Each event domain is defined by a taxonomical structure (illustrated in Figure 2-5; an equivalent illustration could be shown for each of the other two activity domains with no alteration to the event domains as shown) which provides a roadmap for information extraction from open data sources and subsequent event classification. The need to define a complete taxonomy for each nuclear activity (i.e., from the top-level event domains to bottom rank very specific events such as the acquisition of a specific quantity of a specific material) will be circumvented through the leveraging of machine learning techniques. Here, each event domain will have one sub-rank of events defined and therefore remains quite general. Such a level of generalization has the advantage of preventing "overfitting" of any specific activity. Notably, the optimal number of sub-rank events will be a high impact research question that the team will seek to answer in the prototype system's testing.

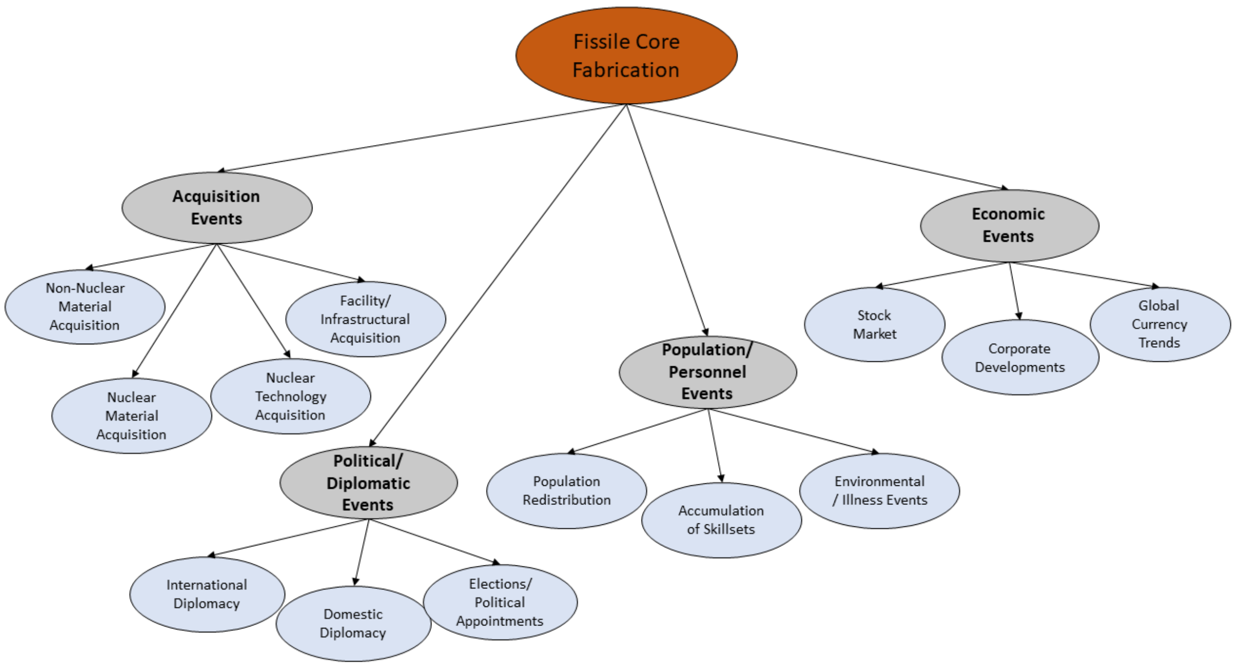

Figure 2-5. Illustration of the relationship of events to activities.

2.2.1 Acquisition Events

Acquisition events encompass all events where an asset, object, or skill, has been obtained within the location of interest. Acquisition can result through trade/commerce, engineering, synthesis, construction, or theft. The detection of the possession of an asset, object, or skill implies acquisition. Given this definition, essentially all activities in Figure 2-5 can be classified as acquisition events. There are four secondary rank events to an acquisition event:

1) Non-nuclear material acquisition
2) Nuclear material acquisition
3) Nuclear technology acquisition
4) Facility/Infrastructural acquisition

2.2.1.1 Non-Nuclear Material Acquisition

A non-nuclear material acquisition encompasses any acquisition of materials or resources that are non-nuclear in nature, but that could be of potential use in a proliferation effort. Non-nuclear materials acquisition can occur through trade/commerce, theft, science/engineering, construction, or synthesis. Examples of non-nuclear materials acquisitions are:

1) Computers/programable logic controllers
2) Piping/Electrical wiring
3) Concrete/Reinforced Concrete
4) Chemicals

2.2.1.2 Nuclear Material Acquisition

A nuclear material acquisition encompasses any acquisition of materials that are nuclear in nature. That is, any material which undergoes radioactive decay. Nuclear materials acquisition can be obtained through reprocessing, trade/commerce, theft, or synthesis in a nuclear reactor. Examples of such materials include:

1) Uranium/Uranium containing compounds
2) Tritium
3) Plutonium
4) Tc-99m

2.2.1.3 Nuclear Technology Acquisition

Nuclear technology acquisition encompasses any acquisition of technology that could potentially be related to known nuclear technologies. Nuclear technology can be acquired through trade/commerce, theft, or science/engineering. Examples of nuclear technology include:

1) Nuclear reactor design
2) Nuclear weapons design
3) Nuclear materials reprocessing facility design and process flow-sheeting
4) Chemical manufacturing/synthesis facility design and process flow-sheeting

2.2.1.4 Facility/Infrastructural Acquisition

Facility/Infrastructural acquisition encompasses any acquisition of new facilities or infrastructure. Facility/Infrastructural acquisitions can be obtained through engineering, construction, and purchase. Examples of facility/infrastructural acquisition include:

1) Nuclear materials re-processing/Chemical synthesis/Nuclear reactor facilities construction
2) Plant startup
3) Facility re-purposing
4) Construction of roadways

2.2.2 Political/Diplomatic Events

Political/Diplomatic events encompass all events concerning domestic governmental actions and/or international governmental actions/interactions. Political/Diplomatic events can result from elections, legislation, coups d'etat, public outreach, collaboration, or dissent. There are three secondary rank events to political/diplomatic events:

1) International Diplomacy
2) Domestic Diplomacy
3) Elections/Political Appointments

2.2.2.1 International Diplomacy

International diplomacy encompasses any interaction, positive or negative, among governmental actors of two or more nations. International diplomacy can occur through trade deals, immigration standards, economic sanctions, expression of alliance, expressions of dissent/hostility, or past/present military activities. Examples of international diplomacy include:

1) Trade embargos/tariffs
2) United Nations interventions
3) Joint military exercises
4) Presidential meetings

2.2.2.2 Domestic Diplomacy

Domestic diplomacy encompasses any internal interaction, positive or negative, between legislators or between legislators and citizens. Domestic diplomacy can occur through passing of laws and budgets, political community outreach, or perceived governmental oppression. Examples of domestic diplomacy include:

1) Increased military spending
2) High-level diplomatic visit to a region
3) Protest of governmental actions
4) Civilian internment

2.2.2.3 Elections/Political Appointments

Elections/Political appointments encompasses any actions within a nation that result in, or have the potential to result in, changes to governmental leaders. Elections/Political appointments can result from elections, political appointees, coups d'etat (or attempted coups d'etat), removal from office (e.g., impeachment), resignation, retirement, or death. Examples of elections/political appointments are:

1) Civil war resulting in governmental overthrow
2) Family inheritance of a high-ranking political position
3) Election of incumbent president
4) Election of non-incumbent president

2.2.3 Population/Personnel Events

Population/Personnel events encompass all events concerning the distribution and activities of a nation or nations' population. Population/Personnel events can result from increased interest in a specific field of study, regional population growth/contraction, or job growth. There are three secondary rank events to population/personnel events:

1) Population redistribution
2) Accumulation of skillsets
3) Environmental/illness events

2.2.3.1 Population Redistribution

Population redistribution encompasses any changes associated with the population of a nation, or specific region within the nation. Examples of population redistribution are:

1) Growing population in the northwest
2) Shrinking population in a city
3) Decreasing population of a nation
4) Increasing population of a state

2.2.3.2 Accumulation of Skillsets

Accumulation of skillsets encompasses any changes associated with a growing interest in a particular field of study within a nation. Accumulation of skillsets can occur when a growing trend of a particular interest is found in a regionally specific population. Examples of accumulation of skillsets are:

1) Increasing numbers of bachelor's degrees in nuclear engineering
2) Increasing numbers of doctorate degrees in science, technology, engineering, and mathematic fields
3) A growing network of experienced scientists in a specific city
4) Increase in plant operator-type job positions

2.2.3.3 Environmental/Illness Events

Environmental/Illness events encompasses all environmental or medical issues that are directly related to the production, storage, or disposal of materials or chemicals. Environmental/Illness events can occur from a discrete sudden event such as an accident, or from the long-term operation of a facility. Examples of illness events are:

1) Radiation poisoning
2) Elevated levels of mercury in water supplies
3) Atmospheric detection of radioactive isotopes from fallout

2.2.4 Economic Events

Economic events encompass all events concerning global commerce. Economic events can result from corporate expansion and/or existing presence to specific regions, fluctuations in commodity prices resulting from increased demand or decreased supply, or fluctuations in currency value due to changes in a nation's gross domestic product. There are three secondary rank events to economic events:

1) Stock market
2) Corporate developments
3) Global currency trends

2.2.4.1 Stock Market

Stock market encompasses any fluctuations in the stock market that may be traced back to a particular market sector, corporation, commodity, or global event. Examples of stock market are:

1) A decrease in the S&P 500 due to rumors of impending economic recession
2) An increase in oil prices due to increased demand
3) An increase in a concrete manufacturer's stock price

2.2.4.2 Corporate Developments

Corporate developments encompass all events concerning major corporate expansions. Corporate developments can occur when a corporation expands to a new geospatial region or market sector, merges with another corporation, files for bankruptcy, or goes out of business. Examples of corporate developments are:

1) Two major companies merge
2) A major corporation purchases a smaller corporation and expands its geospatial presence
3) A major nuclear energy corporation acquires a large defense contract[3]

2.2.4.3 Global Currency Trends

Global currency trends encompass all events concerning fluctuations to currency values. Global currency trends include national currency prices, cryptocurrency prices, and the intentional or consequential manipulation of currency prices. Examples of global currency trends include:

1) Currency manipulation involving printing of bills
2) Mining of Bitcoin concentrated in a specific region
3) Purchasing debt of a foreign government

[3] Note that this is similar to an acquisition event, but all EOIs assume that the proliferation activities are state-sponsored. With this example, one might also suspect there could be a trigger of activity in "Diplomatic Events within a Nation" regarding increased defense spending. Though these two events are similar in nature, they offer opposing vantage points.

3.0 Geospatial Locations of Interest

The prototype forecasting system will be applied to the region surrounding the target (i.e., the Central Savannah River Regional Area, South Carolina and Georgia, the entire Southeastern US, etc.). In addition, other regions around the country are identified in this section that have known similar facilities, no similar facilities but other nuclear related facilities, and regions known to be devoid of nuclear related facilities for training and sensitivity.

The following sections evaluate and propose regions of interest.

3.1 Regions of known nuclear activity

Two primary regions could be used as close surrogates for the proposed SRPPF: the regions surrounding Los Alamos, New Mexico, and the region surrounding Denver, Colorado. The former encompasses Los Alamos National Laboratory that currently fabricates fissile cores and the latter encompasses a closed fissile core fabrication facility, the Rocky Flats Plant. The Rocky Flats Plant has been closed long enough such that obtaining historically equivalent data useful to this project may be impractical. If that is found to be the case, Los Alamos extending to the entire state of New Mexico will be used on its own.

The extent of nuclear activity in the US is illustrated in Figure 3-1 which includes the location of all DOE facilities including legacy sites [9] and all sites listed by the NRC [10] and IAEA [11] [12] [13]. Sites shown may be active, decommissioned, decommissioning, or under construction. Contours radiate from each location in 60-mile increments and black dots represent the locations of cities with a population greater than 38,000 but limited to the largest five cities per state. Notably, except for Hawaii, Montana, and North Dakota, every state has a DOE facility (including legacy sites) or some other type of past, ongoing, or planned nuclear activity suggesting that almost any metropolitan area would be an acceptable target for testing in a complex data environment. Figure 3-2 shows an equivalent map that includes only DOE and Legacy Management sites. Comparison of Figure 3-1 and Figure 3-2 reveals a smaller selection of locations that will have nuclear sites that are not DOE or Legacy sites. These locations include Alabama, Arkansas, Delaware, Indiana, Kansas, Maine, Maryland, Minnesota, New Hampshire, North Carolina, Rhode Island, Vermont. Many of these states may be useful test cases, as facilities close to the state borders may have a significant influence and create a unique data environment. Any of these states may be used to test the forecasting system with nuclear facilities that are void of any current or past weapons related sites.

3.2 Regions known to be void of nuclear activity

Figure 3-1 reveals three states with no nuclear facilities: Hawaii, Montana, and North Dakota. A few regions around the cities of Billings, Montana; Fargo and Bismarck, North Dakota may be distant enough from nuclear sites such that there could be a low enough frequency of discovery of nuclear events to serve as a geospatially null test. Hawaii is the preferred test location which is also physically isolated.

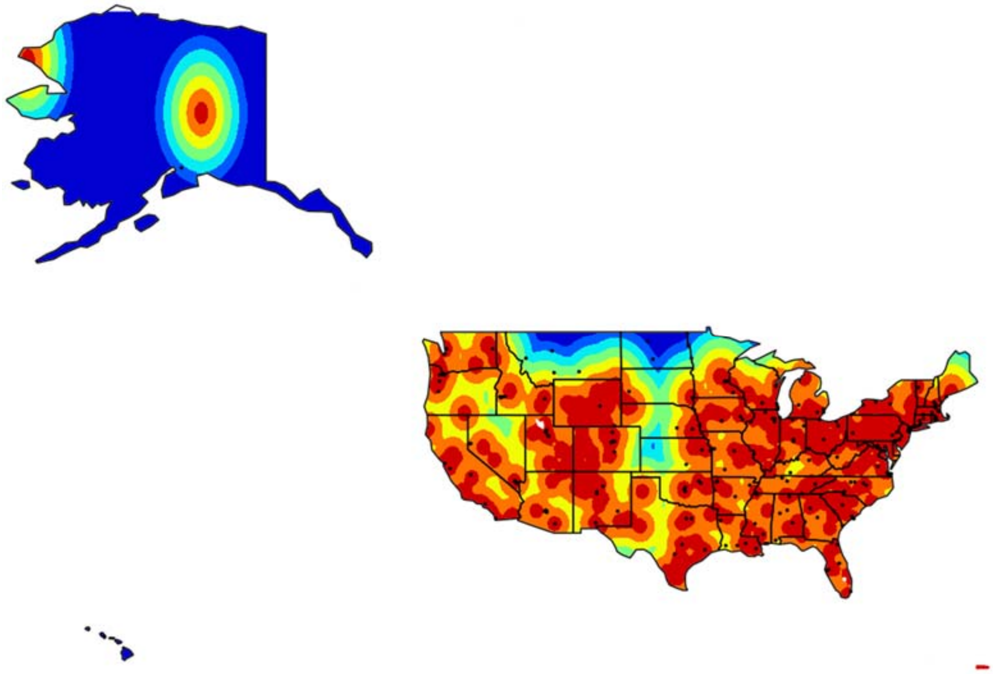

Figure 3-1. Heat Map showing 60 Mile Increments from All Nuclear Facilities

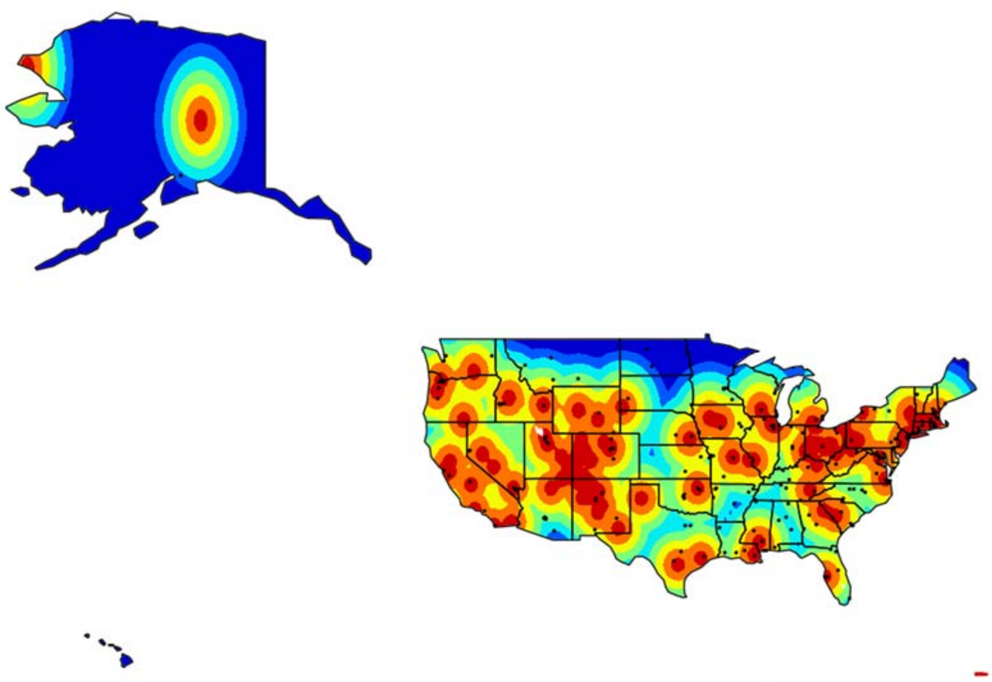

Figure 3-2. Heat Map showing 60 Mile Increments from DOE Nuclear Facilities Including Legacy Facilities

4.0 Initial Performance Metrics

Weapons program development is expected to proceed in confidence which is analogous to the planning for a public demonstration until it becomes widely known or announced. EMBERS performance was measured as the difference in the time of prediction and that of the very first report from any open source (i.e., the lead time) – this project will use a similar metric. The metric will identify the number of days between the time a forecast can be generated, with a minimum of at least 50% confidence, and the first public announcement of the proposed SRPPF, May 10, 2018.

The occurrence of many of the events and activities will be treated as binary classification problems (i.e., has the event or activity occurred or not). Therefore, to quantitatively assess the performance of many of the models, precision and recall will be used. Precision will quantify the ability of the models to identify *only* the relevant data points and is expressed as:

$$Precision = \frac{True\ Positives}{True\ Positives + False\ Positives}$$

Recall will quantify the ability of the models to find *all* relevant cases within a data set and is expressed as:

$$Recall = \frac{True\ Positives}{True\ Positives + False\ Negatives}$$

For the different models within the overall forecasting system (e.g., models used to formulate events, fuse data streams, forecast the occurrence of a nuclear activity, etc.), threshold values for positive versus negative will likely vary and will be tuned to the event type. For example, for some event types, it may be necessary, or desirable, to increase recall (i.e., reduce the threshold for a positive case) at the expense of precision to avoid missing relevant data points. In other cases, it may be more important increase the threshold for a positive case, thereby decreasing false positives and increasing precision at the expense of recall such that only relevant data is gathered. An initial minimum goal of 0.60 for both precision and recall for the prototype system. One may note that false negatives are less desirable than false positives in a finished system. A positive result requires other methods to verify the result, thus, numerous false positives is burdensome; whereas, a false negative will miss a desired result. As such, high recall is more important than high precision. Ultimately, this becomes a balance to minimize verification work in a completely functional system. In addition, the user will set these parameters for a particular application. For this prototype system, an initial modest goal is set for evaluation purposes. If first year results prove to meet these minimums, the subsequent development will be evaluated to improve the performance to optimize results and reduce false negatives to as low as practical.

All forecasts will include an audit trail (i.e., a detailed listing of the events, the originating data sources, etc.) of how the system arrived at its forecast.

If the results reveal singular data sources produce forecasts well in advance of the public disclosure with very high confidence, the data source will be screened from the models. High dependency on only one data source reduces the ability to translate the models and system to other locations.

5.0 References

[1] G. S. Davidson and e. el., "Data Mining for Ontology Development," Sandia National Laboratory, SAND2010-3708, June 2010.

[2] P. A. Pope and A. M. Spearing, "SNF Reprocessing Ontology: Design, Creation, Evaluation, & Use," Los Alamos National Laboratory, LA-CP-15-20237, April 28, 2015.

[3] R. Ward and A. Sorokine, "Nuclear Power Plant Ontology in OWL Format," Oak Ridge National Laboratory, ORN-LLTR-2012467.

[4] S. Muthiah , P. Butler , R. P. Khandpur , P. Saraf , N. Self , A. Rozovskaya , L. Zhao , J. Cadena , C. T. Lu , A. Vullikanti , A. Marathe , K. Summers , G. Katz , A. Doyle , J. Arredondo , D. K. Gupta , D. Mares and N. Ramakrishnan, "EMBERS at 4 years: Experiences operating an Open Source Indicators Forecasting System," arXiv:1604.00033, 2016.

[5] N. Ramakrishnan , P. Butler , S. Muthiah , N. Self , R. Khandpur , P. Saraf , W. Wang , J. Cadena , A. Vullikanti , G. Korkmaz , C. Kuhlman , A. Marathe , L. Zhao , T. Hua , F. Chen , C. T. Lu , B. Huang , A. Srinivasan , K. Trinh , L. Getoor , G. Katz , A. Doyle , C. Ackermann , I. Zavorin , J. Ford , K. Summers , Y. Fayed , J. Arredondo , D. Gupta and D. Mares, arXiv: 1402.7035, 2014.

[6] "Joint Statement from Ellen M. Lord and Lisa E. Gordon-Hagerty on Recapitalization of Plutonium Pit Production," 10 May 2018. [Online]. Available: https://www.energy.gov/nnsa/articles/joint-statement-ellen-m-lord-and-lisa-e-gordon-hagerty-recapitalization-plutonium-pit .

[7] United States Nuclear Regulatory Commission, [Online]. Available: https://www.nrc.gov/.

[8] International Atomic Energy Agency, 2020. [Online]. Available: https://www.iaea.org/.

[9] "LM Sites," Department of Energy Office of Legacy Management, 2020. [Online]. Available: https://www.energy.gov/lm/sites/lm-sites.

[10] "United States Nuclear Regulatory Commission," 12 9 2019. [Online]. Available: https://www.nrc.gov/reading-rm/doc-collections/datasets/.

[11] "Research Reactor Database," International Atomic Energy Agency, 2019. [Online]. Available: https://www.iaea.org/resources/databases/research-reactor-database-rrdb.

[12] "Nuclear Fuel Cycle Information System," International Atomic Energy Agency, 2019. [Online]. Available: https://infcis.iaea.org/NFCIS/About.cshtml.

[13] "Power Reactor Information System," International Atomic Energy Agency, 2019. [Online]. Available: https://www.iaea.org/resources/databases/power-reactor-information-system-pris.

[14] World Nuclear Association, 2020. [Online]. Available: https://world-nuclear.org/information-library.aspx.

Distribution:

J. S. Bollinger, james02.bollinger@srnl.doe.gov
P. Butler, pabutler@vt.edu
G. R. Cefus, gregory.cefus@srnl.doe.gov
D. A. Crowley, dave.crowley@srnl.doe.gov
T. L. Danielson, Thomas.Danielson@srnl.doe.gov
C. C. Herman, connie.herman@srnl.doe.gov
J. A. Horne, jeff.horne@srs.gov
R. D. Jeffcoat, ron.jeffcoat@srnl.doe.gov
A. A. Kail, Andrew.Kail@srnl.doe.gov
P. L. Lee, patricia.lee@srnl.doe.gov
B. Mayer, bmayer@cs.vt.edu
N. Muralidhar, nik90@vt.edu
S. M. Oswald, sandi.oswald@srnl.doe.gov
J. A. Pike, jeff.pike@srnl.doe.gov
D. L. Wilson, david.wilson@srnl.doe.gov
Records Administration (EDWS)