Applied Machine Learning for Simulated Reprocessing Safeguards: Unsupervised Networks

Nathan Shoman[1], Benjamin Cipiti[1], Thomas Grimes[1], Ben Wilson[2], Randall Gladen[2]

[1]Sandia National Laboratories
[2]Pacific Northwest National Laboratory
E-mail: nshoman@sandia.gov

Abstract

A goal of the International Atomic Energy Agency (IAEA) is to deter the spread of nuclear weapons through detection of nuclear material and technology misuse. Detecting diversion of nuclear material from large bulk handling facilities, such as a reprocessing plant, is a goal that can prove to be both challenging and resource intensive as it often requires destructive analysis of numerous samples taken from various locations across the facility. The IAEA has sought out methods to develop an integrated system of instrumentation and data processing to reduce this burden. The goal of this work is to leverage machine learning (ML) methods to improve the effectiveness and efficiency of safeguards by utilizing higher uncertainty measurements, such as process monitoring and Non-Destructive Assay measurements, which are not extensively used in traditional safeguards methods. This work is part of a series of two documents that consider the use of ML to improve one aspect of safeguards, namely nuclear material accountancy. This part considers unsupervised networks that are used to detect anomalous behavior that could be indicative of material loss. The unsupervised approach is shown to exceed traditional methodologies but only after several practical barriers have been accounted for and resolved.

Keywords: safeguards; data science; machine learning; nuclear material accountancy; reprocessing

1. Introduction

The International Atomic Energy Agency (IAEA) was established as an organization within the United Nations to promote the peaceful use of nuclear power [1]. One function of the IAEA is the implementation of safeguards for member states. The goal of safeguards is the timely detection of diversion of significant quantities (SQs) of nuclear material for weapon purposes and deterrence of such diversion by the risk of detection. Nuclear material accountancy (NMA) is one method used by the IAEA to implement safeguards. NMA can be thought of as an audit of nuclear facilities that verifies reported quantities of material to ensure they have not been diverted. This is accomplished through several methods such as sampling and process monitoring. Safeguards can be further complemented by other systems such as containment and surveillance (C/S), particularly for large throughput facilities.

Existing NMA systems are well understood and have been implemented at numerous facilities. However, NMA often requires low uncertainty destructive assay (DA) measurements to reach timeliness goals. These measurements are often time consuming and expensive as they must be performed in an analytical laboratory. Other types of measurements, such as process monitoring (PM) and non-destructive assay (NDA), can be used for remote monitoring to lead to lower costs, but often have relatively high uncertainties. Machine learning (ML) has revolutionized many fields and offers promise in safeguards related tasks like anomaly detection. This work hypothesizes that ML could more effectively leverage underutilized measurements with higher uncertainties (e.g. NDA and PM) to improve costs associated with NMA.

2. Background

International safeguards are implemented to guard against diversion of significant quantities of nuclear material. This is defined by the IAEA as the approximate amount of nuclear material for which the possibility of manufacturing a nuclear explosive device cannot be excluded, which for plutonium is 8 kg [2]. One simple approach for the NMA component of international safeguards is item counting. Here, simple counting of discrete items is used to account for items that contain nuclear materials (e.g. fuel assemblies). When

combined with statistics and random sampling, item accounting is indeed the preferred method for facilities where material is most often found in discrete items. However, the focus of this work is large facilities where material is often in bulk form (e.g. powders or solutions) that require methods beyond simple item accounting [3]. The goal of this work is to develop machine learning approaches to improve material accountancy of these large facilities. It is then important to accurately describe traditional methods such that the proposed machine learning based framework can be fairly compared to the current state-of-the-art.

2.1 Traditional Nuclear Material Accounting

Material Unaccounted For (MUF) [4] is a core component of NMA. MUF is a quantitative balance between flows of material into and out of a facility. Usually, facilities will have multiple material balances that are divided up to reach certain timeliness goals or due to physical constraints within a facility (e.g., separate buildings). MUF is calculated at regular intervals defined by the material balance period (MBP). Subject matter expertise is used to determine both the number and size of material balances in addition to the material balance period. The MUF calculation at a given time t with measurement locations i and total number of locations for a given measurement n is given by Equation 1.

$$MUF_T = \left(\sum_{i=1}^{n_I} I_{i,t-1} + \sum_{i=1}^{n_{inp}} Tin_{i,t} - \sum_{i=1}^{n_{out}} Tout_{i,t} \right) - \sum_{i=1}^{n_I} I_{i,t} \quad (1)$$

The individual terms in the equation are as follows:

- $\sum_{i=1}^{n_{inp}} Tin_{i,t}$: Total input transfers
- Transfer terms are often streams of material which should then be time integrated. The total transfer term would then become $\sum_{T=t-1}^{t} \sum_{i=1}^{n_{inp}} Tin_{i,t}$
- $\sum_{i=1}^{n_{out}} Tout_{i,t}$: Total output transfers
- $\sum_{i=1}^{n_I} I_{i,t}$: Total of all inventories at time t
- $\sum_{i=1}^{n_I} I_{i,t-1}$: Total of all inventories at time *t-1*

The expectation is that $MUF_t = 0$ when no material has been removed as all material has been accounted for. However, measurements always have some associated error, which causes a non-zero MUF even during normal conditions.

2.2 Measurement Error

No measurement is perfect and therefore is accompanied by some degree of uncertainty. Safeguards measurements are often characterized by a multiplicative error model as described in Equation 2.

$$M_{i,t} = G_{i,t}\left(1 + S_i + R_{i,t}\right) \quad (2)$$

Where

$$S_i \sim N(0, \delta_S^2)$$

$$R_{i,t} \sim N(0, \delta_R^2)$$

The terms above are defined as follows:

- $M_{i,t}$: Measured quantity of interest at location i at time t
- $G_{i,t}$: True quantity of interest (unobservable) at location i at time t
- S_i : Short-term systematic (i.e. epistemic) error
- Arises from measurement conditions or settings that remain unchanged from some period of time
- Difficult to reduce
- Example: Error due to calibration curve
- $R_{i,t}$: Random error (i.e. aleatory)
- Varies unpredictable under repeated conditions
- Can be reduced through repeated measurements
- Example: Counting statistics
- δ : Relative standard deviation

The random and systematic errors are assumed to be independent Gaussian random variables with zero mean and variances δ_S^2 and δ_R^2. Measurement errors are approximately normally distributed according to Equation 3. The specific values of the variances depend on the measurement technology that is used. The IAEA has published the International Target Value (ITV) guidelines [5] which provides expected performance metrics and variances.

$$M_{i,t} \sim N\left(G_{i,t}, G_{i,t}^2(\delta_R^2 + \delta_S^2)\right) \quad (3)$$

Measurement error plays an important role in the performance of anomaly detection for material losses. Generally, a material loss can be thought of as a mean shift in the normally distributed material balance, as expressed in Equation 4. A key goal of NMA is to detect this shift.

$$N\left(\mu_t \rightarrow \mu_t\ , \sigma_{MB}^2\right) \quad (4)$$

The body of statistics literature contains a range of different tests that can be used for change detection such as the one shown in Equation 4. However, all approaches are generally subject to limitations arising from measurement error as expressed in Equation 5. The probability of detection of a mean shift in a known, normal distribution (i.e. true positive) approaches the probability of false alarm (i.e. false positive) as the variance increases.

$$\lim_{\sigma_{MB} \to \infty} P\left(Alarm \vee N\left(\mu_t\ , \sigma_{MB}^2\right)\right) = P\left(Alarm \vee N(\mu_t, \sigma_{MB}^2)\right) \quad (5)$$

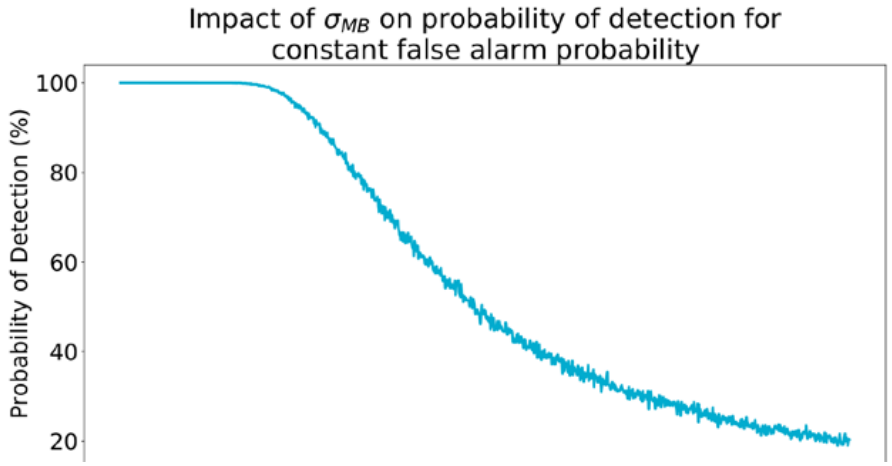

Figure 1: Probability of detection shown as a function of uncertainty for a constant false alarm probability.

Put simply, smaller mean shifts relative to the variance are more difficult to detect as they often get lost in the noise. This is shown more concretely in Figure 1 where the detection probability for an arbitrary, fixed material loss is quantified for a fixed false alarm probability and various levels of measurement uncertainty.

Finding strategies for reducing the material balance uncertainty has been a historical target for safeguards R&D given the impact on detection of material loss. One possible improvement would be to reduce the measured quantity size which would require a more frequent material balance period. This requires some optimization as too frequent material balance closures will result in higher false alarm probabilities [6]. Improving measurement uncertainty, which also reduces material balance uncertainty via smaller δ_S^2 and δ_R^2, is currently what drives the use of expensive DA measurements.

2.3 Sequential Material Balance Testing

Discussion so far has focused on a single material balance at a specific point in time. However, timely detection of potential material losses, a key goal of the IAEA, often requires multiple sequential material balances. For example, consider the case of a single yearly material balance where a diversion is initiated near the beginning of the year. Consequently, it would be months before the loss could be detected. Sequential material balances also have the added benefit of reducing the uncertainty of any single balance while noting that there are some restrictions on frequency of material balance closure.

Each individual material balance is comprised of potentially many normally distributed measurements which imply the material balance will also be normally distributed. As each single material balance has some mean and variance, a sequence of MBs can be expressed as a multivariate normal in Equation 6.

$$\overrightarrow{MB_t} \sim MVN(\overrightarrow{\mu_t}, \Sigma_t) \quad (6)$$

where

$$\overrightarrow{\mu_t} = (\mu_1, \mu_2, \ldots, \mu_t)^T$$

Recall that each individual material balance at time t is a function of the total current and previous inventories ($\sum_{i=1}^{n_I} I_{i,t}$ and $\sum_{i=1}^{n_I} I_{i,t-1}$). This results in a temporally correlated material balance sequence. However, the Standardized Independent Material Unaccounted For (SITMUF) transformation [7] can be used to decorrelate the sequence by accounting for the analytically determined covariance (using propagation of variance), Σ, and the conditional expectation with a few assumptions. Although not covered extensively here, traditional NMA relies on the residual between the observed MUF value and the conditional expectation of MUF. A sequential test, namely Page's trend test [8] [9], is used to detect trends in the material balance sequence residual. Under normal conditions, the SITMUF sequence should be approximately zero owing to a good conditional expectation. Material losses lead to poor expectations and larger residuals.

2.4 Machine Learning

Machine Learning (ML) refers to algorithms that perform a task without being explicitly programmed to do so. ML has seen a large surge in interest and is now embedded into many aspects of our daily lives. Although arguably less popular than domains such as computer vision, anomaly detection has benefited greatly from improvements in ML. Given the limitations described in previous sections, namely the dependence of traditional NMA on measurement uncertainty, it would be desirable to develop a ML framework that could sidestep the limitation. Specifically, a notable

improvement would be the use of lower cost, but higher uncertainty process monitoring (PM) and non-destructive assay (NDA) measurements to detect material loss. Such a framework would require framing material loss as an anomaly detection problem. This contrasts with traditional NMA which attempts to detect diversions through direct quantification of nuclear material (i.e. MUF).

There are many different anomaly detection algorithms that have been proposed as there is no universal solution for all problems. Consequently, this work represents only one potential, but informed solution for applied ML to improve nuclear material accountancy. Specifically, this work considers supervised regression with an unsupervised anomaly detection problem. The supervised regression problem requires the ground truth to learn an approximate function for some task. In this case, the regression task is to learn the behavior of parts of the PUREX reprocessing facility. Then, an unsupervised anomaly detection algorithm is used to detect unusual behavior. This class of anomaly detection algorithm does not require specific labelled examples of anomalies and instead relies on some proxy metric to describe normality. Unsupervised methods are particularly desirable for safeguards applications where it can be difficult or impossible to provide examples of all credible material loss pathways. This also facilitates a more direct comparison with the existing benchmark (Page's trend test on SITMUF) which also has no requirement with regards to examples of material loss.

In contrast, supervised approaches do require explicit, labelled examples of material loss, but do offer some potential advantages. For example, supervised approaches enable for direct optimization of material loss detection rather than specification of a proxy problem. Direct optimization through specific examples of material loss could also lead to better feature representation in supervised approaches leading to improved performance for known, high consequence loss pathways. Supervised approaches may prove useful, but were not considered in this work.

2.4.1 Related Work

Several previous works have attempted to develop improved strategies for guarding against material losses by developing novel approaches. One example is the Multi-Isotope Process Monitor (MIP) [10] wherein existing process monitoring measurements were combined with pattern recognition techniques in an attempt to develop more effective detection of material loss at large throughput facilities. MIP used principal component analysis (PCA) [11] to reduce the dimensionality of gamma a spectra to learn new representations that express most of the signal variance. Then, PCA statistics such as Q-residual could be used to detect anomalies. The approach used by MIP was limited by the linear reduction in dimensionality. Other works [12] [13] have sought to improve on commonly used trend tests used on SITMUF. For example, there have been prior attempts to use autoregressive moving average (ARMA) [14] models with SITMUF in an effort to detect material loss [12]. A companion work has also considered the application of supervised deep learning anomaly detection to detect material loss. That proposed work leverages a few examples of material loss in attempt to generally improve anomaly detection [15].

3. Problem Statement

Traditional statistics for nuclear material accountancy have a strong reliance on low measurement uncertainty to produce good probabilities of detection as shown previously in Figure 1. ML methods excel at finding subtle changes in signals that could indicate anomalous behavior. Ideally, a ML-based framework could utilize higher uncertainty (and potentially unattended) measurements to detect material loss at the same performance level as traditional approaches.

3.1 Process Modelling

Obtaining data from actual nuclear facilities is often impractical due to cost and limited availability. The Separation and Safeguards Performance Model (SSPM) [16] [17] PUREX flowsheet has been used to provide synthetic training, test, and validation data for the techniques described in this work. The model was developed for systems-level analysis of safeguards design for various facilities including UREX+, PUREX, gaseous enrichment, fuel fabrication, electrochemical reprocessing, and more. The model uses MATLAB Simulink to track elemental and isotopic material flows through various unit operations. Measurement blocks are used to simulate different types of measurements such as PM, NDA, and DA. Several common statistical tests used by the IAEA are also integrated into the model.

A PUREX SSPM flowsheet, based on a generic facility [18], is shown in Figure 2. The grey blocks represent the processing vessels throughout the plant and contain significant detail about inventories, timing of operations, filling/emptying sequences, etc. Signals connecting the blocks contain mass flow information for all nuclear material and bulk flows. The blue blocks represent measurement points which feed the traditional material balance calculation. The shaded regions (red, blue, and green) correspond to various prediction regions where neural networks are used to learn the area's behavior.

3.2 Baseline Machine Learning Approach

This work is motivated by the universal approximation theorem [19] which states that an arbitrary-width single layer neural network can approximate any well-behaved function. It is important to note that the theorem does not comment on the learnability of such well-behaved function.

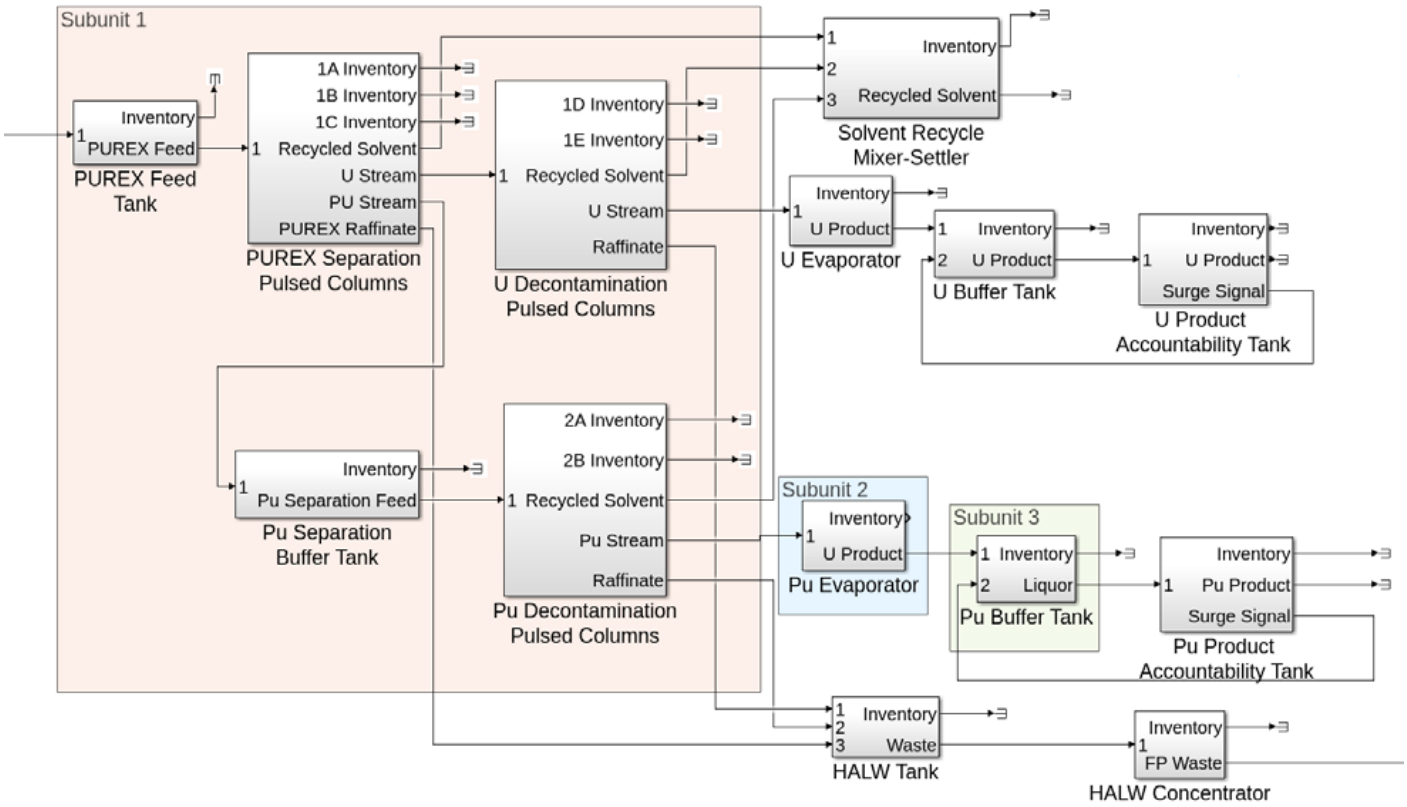

Figure 2: SSPM PUREX Model. Several labelled and shaded regions represent different areas of MBA2 that were learned by individual neural networks (i.e. subunits).

Nonetheless the hypothesis of this work is that a neural network should be able to learn the behavior of a large throughput nuclear facility, specifically a PUREX reprocessing facility. A material loss should appreciably change facility behavior such that the neural network will no longer provide accurate predictions. In turn, this will lead to discrepancies between observations and predictions that could be used to detect and possibly locate anomalous behavior (i.e. material loss). The hypothesis is summarized below in Figure 3.

The proposed unsupervised ML approach requires two steps. The first step is the prediction step where the neural network learns the behavior of a certain facility process (or area of processes) under normal conditions. Ideally, the neural network should be able to learn this behavior by way of the universal approximation theorem. Then, as the facility changes under anomalous conditions, the neural network predictions should no longer agree with observations. Figure 4 shows an example of this behavior wherein poor predictions are made during a window of anomalous behavior.

PUREX facilities have some operations that are time dependent and are not well suited to traditional feed forward neural networks, which have no temporal capacity. Consequently, the prediction step utilizes Long short-term memory (LSTM) [20] networks to complement traditional neural networks to capture the temporal properties of certain signals. For example, PUREX facilities have several mixing tanks that are dependent on material that has entered previously. Specific neural network architectures and data representation have a strong impact on accurate predictions. This work found that the LSTM networks trained well and produced good predictions when temporal behavior is captured by passing a window of history as input.

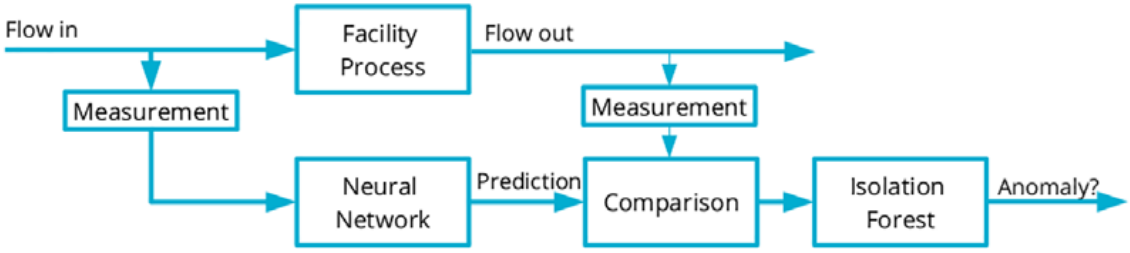

Figure 3: Proposed setup for applied ML for NMA

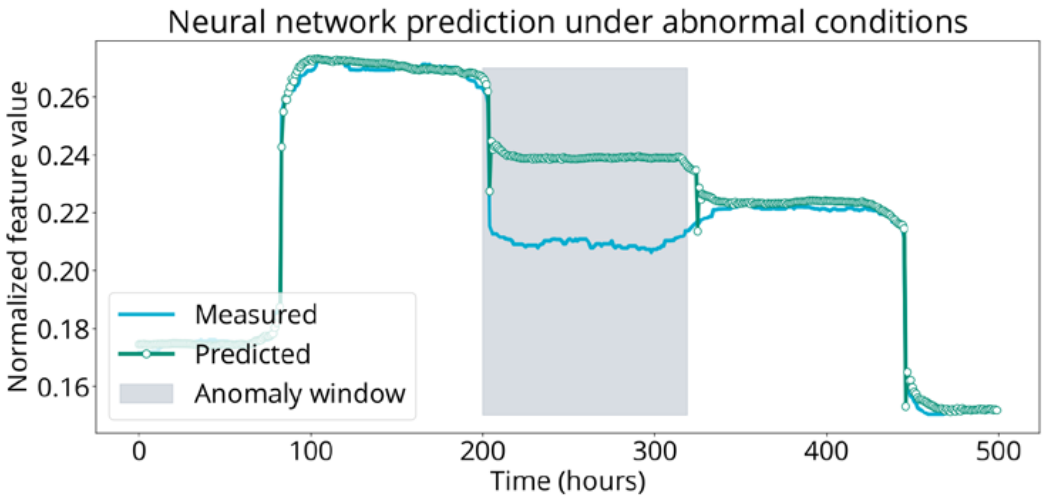

Figure 4: Neural network prediction during abnormal conditions.

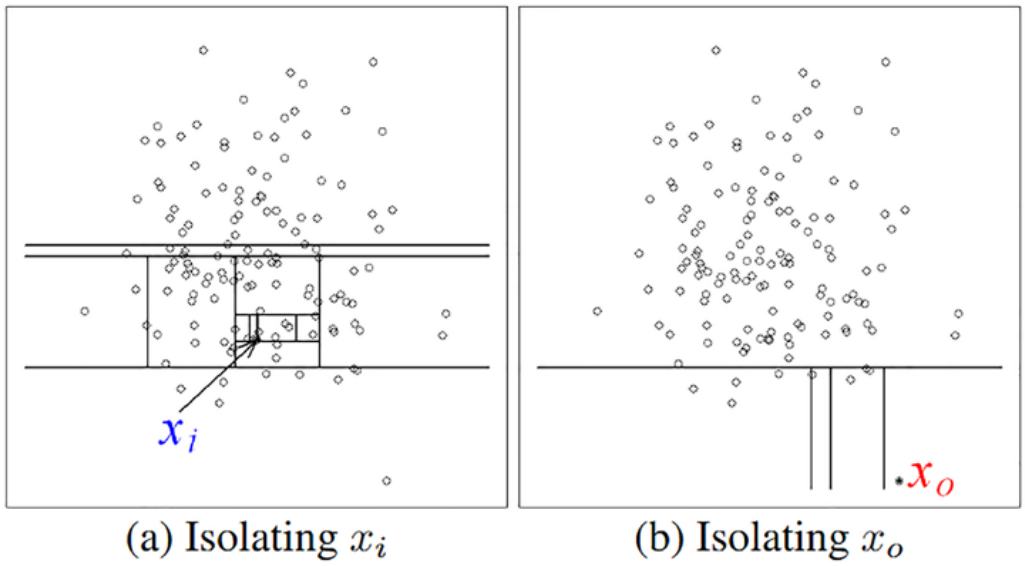

Figure 5: Isolation forest uses recursive splitting to measure the abnormality of a point. This figure shows a normal point, xi, which takes many splits to isolate it from the larger population. In contrast, the abnormal point x_0 requires fewer splits. [21].

The difference between the prediction and observed value, which in this work will be referred to as reconstruction error, is arbitrary due to imperfect predictions even under normal conditions. For example, the neural network used in the prediction step can never calculate predictions with full accuracy which always results in some non-zero prediction error. A second step is required to translate these arbitrary reconstruction errors into alarms and probabilities of detection. Identification of anomalous behavior is complex as PUREX facilities have large multidimensional datasets that arise from measurements at multiple locations each with several features. Instead of using a simple static threshold to detect anomalous behavior (e.g. alarm if a reconstruction error is greater than some scalar value), isolation forest is used here in combination with a classification window to define an alarm condition.

Isolation forest [21] is an unsupervised (requiring no examples of abnormal behavior) anomaly detection algorithm. The key intuition behind isolation forest is that anomalies should be few and different from normal data. The algorithm proceeds by selecting an observation, randomly selecting a feature, then randomly selecting a split value between the minimum and maximum. This process occurs recursively until the observation has been isolated from the larger dataset. Put simply, isolation forest will generate a list of logical criteria that make a particular observation appear unique. The criteria (i.e. splittings) can be represented as a tree structure. Gathering multiple sets of criteria results in a

forest, hence the name isolation forest. The path length of an observation averaged over several random trees is used as a proxy for normality. Points with path lengths below a threshold (as abnormal points should take less logical criteria to isolate) are considered anomalous. A visual intuition for isolation forest is shown in Figure 5.

Isolation forest has several hyperparameters that can be optimized through a grid search. These include the number of trees, maximum number of samples to train each estimator, and maximum number of features to train each estimator, which for this work, are set to 100, 15000, and 5 respectively. An additional hyperparameter, namely, the rate of contamination in the training dataset (i.e. percent of data estimated to be anomalous), cannot easily be discovered through a grid search.

This work generally assumes a 2% contamination rate even though the entire dataset is normal. Effectively, this forces classification of 2% of the training dataset as anomalous. The normal points that are classified as anomalous represent observations with the highest applied errors (i.e. errors drawn from the distribution tails). As classifications alone are insufficient for detecting anomalous behavior (as some normal points are classified as off normal), an alarm criterion on the classification is required. Using prior knowledge that material loss should be rare, it can be assumed that isolation forest will only infrequently produce false positives (i.e. points that are classified as abnormal but are normal). An example of isolation forest output for different anomalies is shown in Figure 6.

Note that there are some classifications being made as normal in the protracted anomaly shown in Figure 6. This is a function of a particular set of measurement realizations. As anomalies become more protracted and closer to the uncertainty bounds, the off-normal classifications become increasingly sparse. Eventually, the algorithm will no longer classify anomalies as off normal as the anomaly magnitude decreases below the uncertainty bounds. This creates the need for a classification window. Off-normal classifications that are dense should represent an anomaly; therefore, a certain number of off-normal classifications in a particular window should cause an alarm. For example, if 10 out of the last 15 classifications are off-normal then the alarm condition has been reached.

It should be noted that there is some dependency between the classification window and contamination rate. Although this work used a 2% contamination rate, there are a range of possible values (1-6%) that still resulted in good detection levels. However, it is important to adjust both parameters (contamination rate and classification window) in parallel. Often, a higher contamination rate still resulted in the same detection probabilities, but higher false alarm rates. Consequently, the classification window requires adjustment in conjunction with the contamination rate.

The proposed unsupervised machine learning approach can be summarized as follows:

- Stage 1: Neural networks are used to predict behavior of several locations within PUREX facility
- Stage 2a: Isolation forest uses reconstruction errors (i.e. prediction - observation) from all subunits as input to classify behavior as normal or off-normal
- Stage 2b: A threshold is applied over recent outputs from stage 2 (i.e. isolation forest). If there are many off-normal classifications recently then an alarm condition is reached

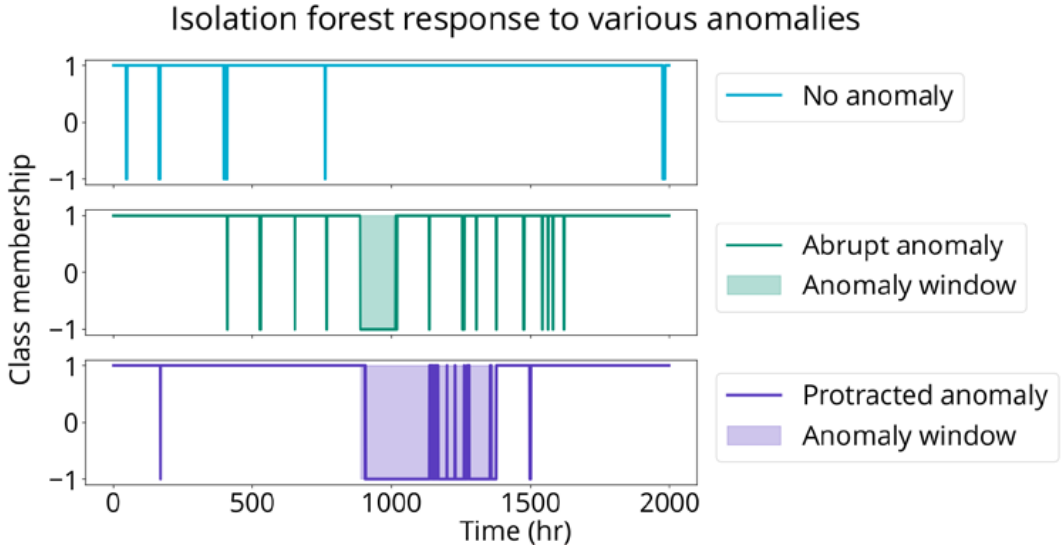

Figure 6: Isolation forest response to different anomalies. Class 1 is “Normal”, and class -1 is “Off-normal”.

4. Experimental Setup

The first step of this anomaly detection framework is to generate several datasets from the SSPM PUREX model to train and test the two-stage machine learning pipeline. The SSPM runs simulated randomized input fuel entering the facility to reflect real-world operation resulting in additional material flow and inventory variation. In practice, actual facilities will have a distribution of possible inputs and outputs rather than a single fixed input and output which results in a more difficult anomaly detection problem.

Each dataset generated by the SSPM model contains about 100 different runs to obtain good performance statistics for traditional approaches to benchmark against and to ensure sufficient data is available for training the machine learning algorithms. The runs are 6480 hours long (270 days), which is about one operating year for a PUREX facility. Ideally, this machine learning approach will operate directly on signals of interest (e.g. gamma spectra), however, the computational overhead of calculating tens of thousands of gamma spectra is large. Instead, this work considered mass with applied errors noting that this is not a direct evaluation of algorithmic performance. However, the use of mass to assess performance is a reasonable proxy as mass and gamma spectroscopy are related by a constant. The dataset evaluated in this work contained features representing ^{134}Cs, ^{137}Cs, ^{154}Eu, ^{241}Am, and ^{241}Pu in most cases. These features represent quantities that could realistically be observed at a PUREX facility.

Individual datasets for each location have a shape of [100 x 6480 x 5] where 100 is the number of runs (operational facility years), 6480 is the time in hours (assumed to be 270 operational days per year), and 5 is the number of features contained. The machine learning pipeline required several datasets to perform training and performance evaluations which is detailed below.

1. First stage (neural network prediction) is trained with a 0.75/0.25 split for training and validation
2. First stage generated training dataset of normal residuals
3. Second stage (isolation forest) is trained
4. Final datasets reflecting different scenarios are used to evaluate performance

The final step in the pipeline is evaluation of a normal dataset to determine a false alarm probability and several anomalous datasets to quantify detection performance. Specifically, four different anomalous scenarios of increasing difficulty are considered. Although not fully described here note that scenario 1 is the easiest to detect while scenario 4 is the most difficult. All datasets used in this evaluation had errors applied according to the multiplicative error model described in Section 2.2 to represent real world conditions more accurately.

5. Identified Performance Factors

Real world application introduces several challenges that impact model performance. It is important to identify and resolve these issues given the high consequence environment of safeguards. The impact of several specific factors and traditional machine learning requirements are explored in the following.

5.1 Facility Discretization

Early work centered on training a single neural network that could learn the facility behavior at all measured locations. However, several unique facility operations prevented effective training of this large neural network. This resulted in the adoption of a discretized approach that segmented the facility into smaller regional neural networks. The following sections describe the primary facility regions (referred to as "subunits") that resulted from the facility segmentation. This discretization of the facility had no observable performance impacts on the supervised regression task. That is, information contained in each subunit was sufficient for the task of predicting the feature value at the next time step.

5.1.1 Subunit 1: Pulsed Separation Columns

The first segment of the facility encompasses a region from the head end, where dissolved nuclear fuel enters the process, to the output of the decontamination column, where a purified plutonium solution leaves. Figure 2 shows this area highlighted in red.

This area consists of several continuous processing operations that are straightforward for a neural network to learn. Specifically, a bi-layer LSTM with a running history of input material is used to predict the output of the decontamination columns. The length of the input history was selected to be 200-hours, which was based on empirical performance as measured by mean-squared error (MSE) on the next time step prediction. The running history approach is expressed in Equation 7 where $f(x_t^n)$ is approximated by the bi-layer LSTM. Equation 7 denotes features as n and time as t. Use of LSTM layers is key to capture temporal dependencies between the inputs and outputs of this facility region.

$$
\begin{aligned}
&{}_1199 \leq t < \infty \\
&\overrightarrow{x_t^n} = [x_{t-199}^n, \ldots, x_t^n] \\
&\hat{x}_{t+1}^n = f\left(\overrightarrow{x_t^n}\right)
\end{aligned}
\tag{7}
$$

5.1.2 Subunit 2: Pu Evaporator

The second segment of the facility encompasses a single unit operation, namely the evaporator, or "Pu Evaporator" as shown in the blue region of Figure 2. This operation

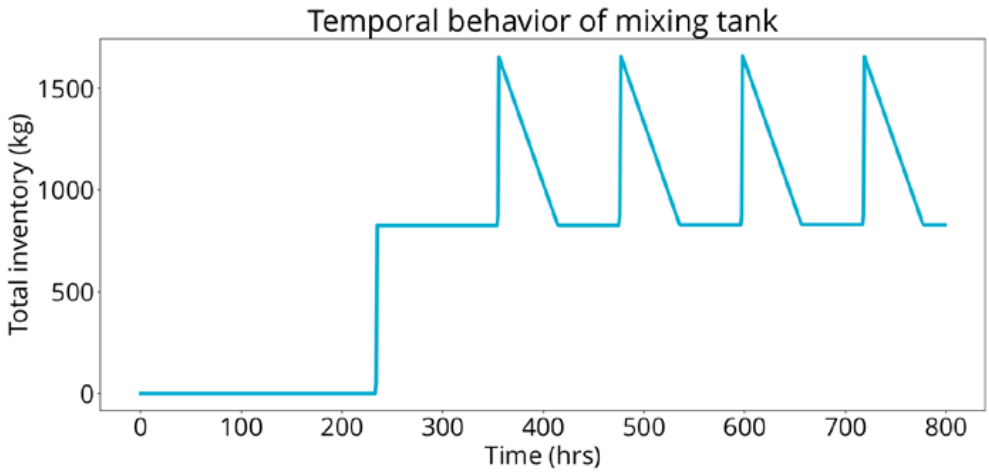

Figure 7: Mixing tank inventory

requires unique consideration as the signal is converted from continuous to discrete. During normal operation the evaporator accumulates solution until a setpoint is reached. Then, the evaporator reduces water content of the accumulated solution and outputs a discrete batch of material that is processed in a following operation.

The previous approach of a running history (described in Section 5.1.1) as input for a LSTM network is not appropriate for this area of the facility. For example, consider if the running history approach is used while the evaporator is accumulating solution and the setpoint had not been reached. The neural network would attempt to predict the previous output batch while the inventory reflects a different product. This area of the facility uses fixed window of time that precisely map the accumulated solution to the corresponding product. Here, a single layer feed-forward neural network is used to predict the output product given the total accumulated material

5.1.3 Subunit 3: Pu Buffer Tank

The final segment consists of another single unit operation, a mixing tank, listed as “Pu Buffer Tank” in the green region of Figure 2. This operation serves as a surge tank for the “Pu Accountability Tank”. During normal operation, this tank fills indefinitely, unless a surge signal is sent, in which case it empties. This behavior is regular and is shown in Figure 7. For this operation, the running history similar to what is described in Section 5.1.1 is used. A bi-layer LSTM network is again used; however, several hand-engineered features are required for the LSTM to make accurate predictions.

Note in Figure 7 that the tank output is a combination of two quantities: the previous batch of material to arrive in the tank and the residual tank inventory at the previous time step. Thus, a running average of the mixing tank inventory must be estimated. This running average can be roughly approximated as $(x_{t-1} + x_t)/2$. Additionally, the actual tank level measurement must also be included. During normal operations, there are slight variations in the input and output batch size, which will impact the running average calculation. Rather than adjusting the running average feature by hand, the LSTM can learn to adjust it during training provided that the tank level measurement, which is a function of the input and output sizes, is provided.

5.2 Training data availability

Machine learning algorithms often require large training datasets to demonstrate adequate performance at test time. One important factor driving required training data is model size. As the number of trainable weights and biases increases so does the training dataset size requirements. Safeguards data is often difficult to obtain, and real-world constraints could lead to little available training data. Therefore, it is important to consider how the dataset size impacts performance using concrete metrics. This section ignores measurement error (described in Section 5.4) to isolate the impact of available training data. As such, it is important to note these results would not reflect real-world results as there are other performance factors in addition to training data availability.

A parametric study is conducted to consider the impact of training dataset size on machine learning performance (probability of detection). It is important to note that both stages of the proposed approach (prediction and classification) require training data. Further, as the prediction stage is used to train the classification stage, there is a compounded effect of reduced training data. The classification stage will not only suffer from less training data, but poorer quality data as the prediction stage also degrades.

The baseline assumes 100 operational years of training data. This is chosen to be sufficiently large to ensure that training performance is driven by the machine learning algorithm hyperparameters which enabled fine tuning. It should be noted that this is not simply 100 iterations of the same operational year (i.e. same pattern with different errors applied), but unique simulations. There are a wide

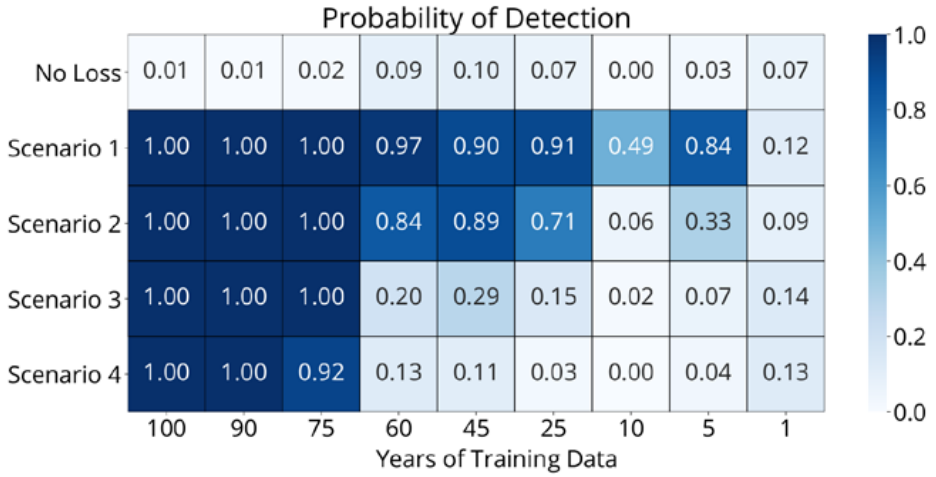

Figure 8: Probability of detection for several material loss scenarios with varied training dataset sizes. A probability of 1.00 indicates a 100% probability of detection whereas a probability of 0.00 indicates no probability of detection.

range of potential facility patterns due to the many combinations of input fuel that could be selected that are adequately captured in large datasets.

The parametric study considered the joint performance of reduced training data on the machine learning pipeline. That is, both stages are trained on reduced training datasets. For example, when 10 years of training data is used, the first prediction stage is trained on 10 years' worth of data. Then, 10 years' worth of predictions are generated and used to train the second stage. This essentially doubles the amount of training data that would be required in practice. This study also incorporated early stopping during training of the prediction stage to ensure that any performance losses are due to the inability of a smaller dataset to represent the test distribution rather than less training time. Results of this parametric study are shown in Figure 8.

Unsurprisingly, lower quantities of training data have a larger impact on the more difficult to detect scenarios. These scenarios tend to be relatively large changes compared to the uncertainty arising from measurement error. The more difficult scenarios have much lower performance while seeing sharp drop offs at certain quantities of training data. This is largely due to inflexible alarm threshold. Recall that the alarm criteria specified in this work is defined by a certain number of off-normal classifications within a window of time. Small changes in performance resulting in fewer off-normal classifications could result in large changes in alarm probabilities. For example, consider a threshold criterion of 30 off-normal classifications, sampled at a rate of once per hour, in a 50-hour window. A small degradation in performance that results in an average of 27 off-normal classifications when also sampled at the same rate in a 50-hour window translates to many fewer alarm triggers.

Poorer performance of the prediction stage also results in degraded detection performance for some of the subunit areas. The average prediction error for normal behavior generally increased with shrinking training dataset sizes (shown in Figure 9). A key assumption in this work is that a neural network can adequately learn facility operations. Increasing prediction errors from an inability to learn facility behavior results in a more difficult classification task.

It is interesting to note some subunits, which correspond to specific unit operations, are more susceptible to reduced training data than others. This phenomenon is not well understood and a target for future work. However, one possibility is that more complex operations require more training data, which is supported by machine learning literature. Often data requirements scale with both algorithm size and task complexity. Subunit 1 is a relatively simple area of the facility (pulsed separation columns) that degrades little with decreased training data whereas subunit 2 is more complex (evaporator) has a significant decrease in performance (i.e. higher errors).

5.3 Facility transients

Routine operations at bulk nuclear facilities can sometimes lead to transients that are not malicious in nature. These changes in behavior could make it difficult to detect anomalous behavior that occurs at the same time. However, a successful detection algorithm should be able to recover after the transient has ended and regain performance. Recovery of the proposed machine learning pipeline is evaluated by generating datasets representative of two different facility transients. The performance is reported as the reconstruction error (i.e. prediction-observation), which has strong correlations to probability of detection.

Facility transients can be grouped into several categories despite the numerous different potential scenarios. This work considers two different types of the transient. The first includes scenarios that change facility behavior but do not result in a new baseline. An example in this category might be small changes to product composition as a result of a

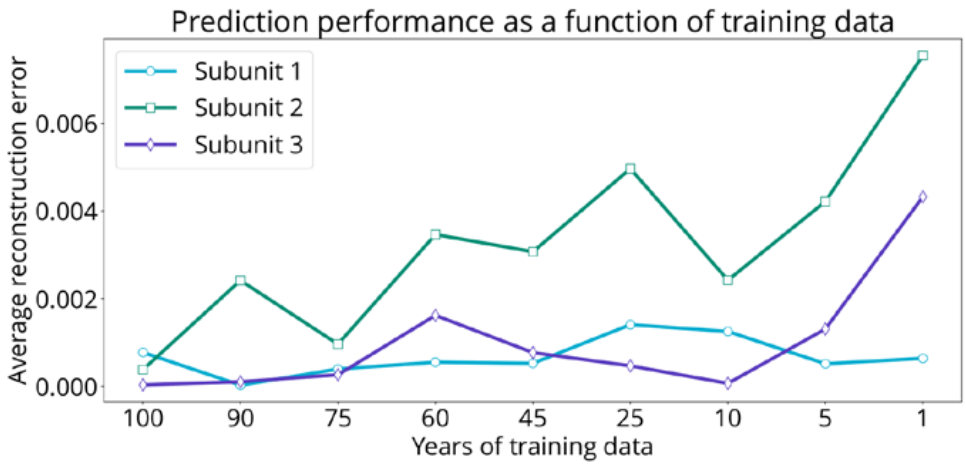

Figure 9: Impact of training data on prediction performance.

vessel leak. The second category includes scenarios that do significantly alter facility which results in a new baseline. Examples here would include transients that cause surge vessels to have a new equilibrium or changes to operational timing.

The first transient considered is in the first category where there is a temporary change in the behavior of the particular area. Figure 10 shows changes in the neural network prediction during and after this transient. The prediction is desirable during the transient in that the prediction error is high compared to normal behavior. However, after the transient has passed, the prediction no longer returns to a baseline performance level. This appears to indicate that the training data no longer represents conditions in the facility and that there has been some prolonged change in facility behavior. This is problematic as anomaly detection performance in the period of time after the transient will decrease.

The second transient represents a scenario where facility behavior is explicitly changed moving forward. For example, this could be damage to a unit operation or represent sensor failure. Figure 11 shows the prediction error during and after a simulated transient. Again, desirable behavior is shown during the transient where there is a sharp increase in prediction error. However, the prediction performance does not return to pre-transient levels. This is somewhat expected as the training data, which is represented to the algorithm as normal, no longer accurately represents the facility. Action would be required to adjust the prediction stage such that it reflects the updated normal conditions.

Both types of facility transients result in poor prolonged prediction performance, which would negatively impact probability of detection. While problematic, there exist some strategies within the machine learning literature to retrain algorithms in an online environment. Future work should target strategies to mitigate the impact of facility transients which are likely to occur in real-world scenarios.

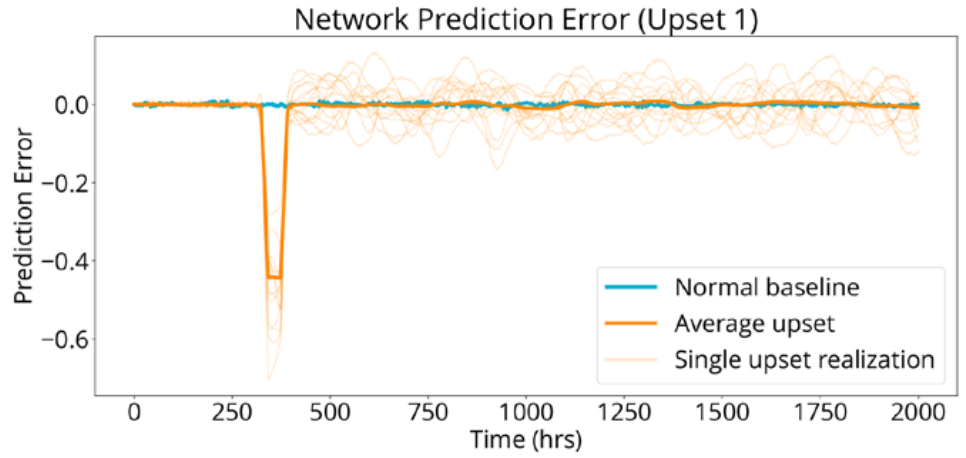

Figure 10: Prediction error of a neural network for a single isotope during a facility transient that changes the baseline of a facility section.

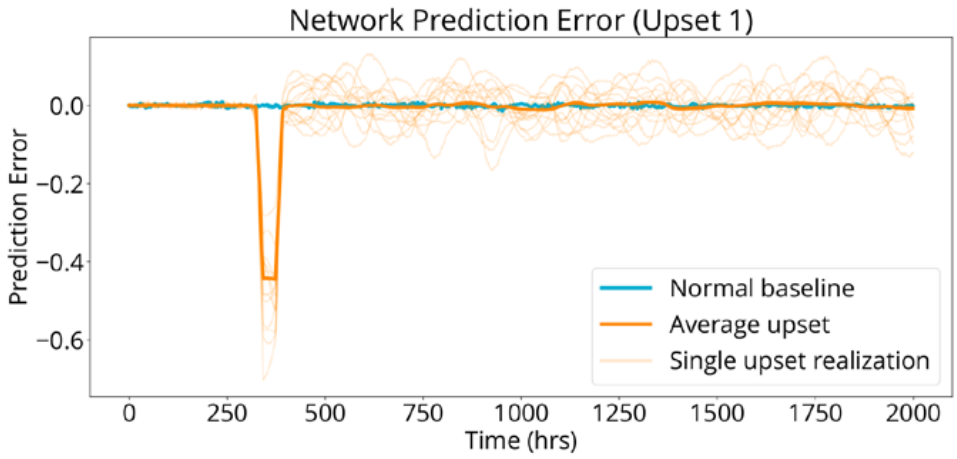

Figure 11: Prediction error of a neural network for a single isotope during a facility transient that changes facility behavior.

5.4 Measurement Error

Measurement error is a reality for the deployment of any real-world NMA system. Traditional NMA systems must implement specific strategies to detect losses when measurements are contaminated with error. For example, the SITMUF transformation can mitigate some impacts of measurement error by converting a MUF sequence to an uncorrelated sequence. Many common anomaly detection algorithms in the machine learning literature are prone to failure when used with error contaminated data. There are also few documented strategies on mitigating measurement error as most literature focuses on bias in supervised learning settings. For example, fairness which is an important area of research, seeks to remove human bias from collected datasets. However, this is fundamentally different from the multiplicative error model encountered in safeguards.

Fundamentally, detection of material loss (or any anomaly at all) is a mean shift detection problem. That is, given a normal distribution of features, can a shift in population mean be detected? This intuition forms the basis of several common anomaly detection algorithms. The proposed machine learning pipeline here also relies on a similar premise. Consider the training objective for the prediction stage; predict the facility behavior given some input. This is achieved through a mean squared error objective which attempts to minimize the difference between training examples and the prediction. It can be shown that the relationship in Equation 8 is true.

$$argmin_\theta \frac{1}{N}\sum_{i=1}^{N}(y_i - x_i\theta)^2 \equiv argmin_\theta E_{x\sim \hat{p_{data}}}[logp_{model}(x)] \quad (8)$$

The mean squared objective is essentially the negative log-likelihood (i.e. cross-entropy) between the empirical distribution and a Gaussian model (i.e. the learned distribution, assumed to be normal). Effectively, the prediction stage tries to learn a function that produces an output distribution as close as possible to the training distribution. Then, under anomalous conditions, the learned distribution is no longer representative indicating that a mean shift has occurred.

Earlier, this work showed that increases in material balance uncertainty reduces the probability of detection for a material loss. A similar phenomenon is at play for the mean shift detection problem. Increases in a distribution's variance reduces the probability that a mean shift can be detected. This can be shown using a variety of approaches including a simple application of Bayes' theorem to a more complex analysis of variance (i.e. ANOVA) procedure.

The previous section showed that there is a strong dependence on sufficient large training datasets to achieve satisfactory anomaly detection performance. It is reasonable to assume, given the limited amount of safeguards data, that multiple measurement campaigns might be required to create a dataset of sufficient size. Each measurement campaign will have its' own unique set of calibrations (i.e. systematic error), that when aggregated together, will result in a larger variance than any individual dataset as shown in Figure 12.

The aggregation of multiple measurement campaigns results in the machine learning pipeline essentially learning variation due to measurement error in addition to facility behavior. This leads to lower anomaly detection performance than traditional statistical methods used for safeguards. This phenomenon is particularly nuanced and discussed at length in a companion work [22].

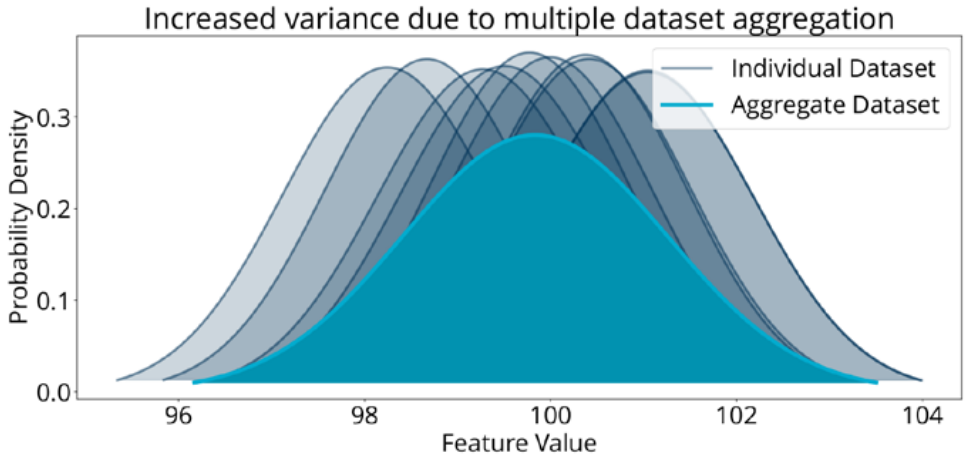

Figure 12: Probability density functions for multiple normal datasets.

5.5 Threshold selection

Recall that the second stage of this proposed machine learning approach requires some threshold for an alarm criterion. Isolation forest generates class labels when given the reconstruction error from the prediction stage. However, due to variation caused by measurement error, the classifications will never be perfect. Intuitively, material losses should generate more off-normal classifications (i.e. true positives) in each period as off-normal classifications made for normal observations (i.e. false positives) will be randomly distributed. One potential alarm condition would be requiring a specific number of off-normal classifications in a particular window of time. A common metric to tune safeguards thresholds (often defined by regulations to be 5%) is the false alarm probability (FAP, i.e. false positive rate).

Threshold optimization is underdefined in this case as there is one constraint (5% FAP) and two unknowns (window size and total classifications). This leads to multiple possible solutions for threshold criteria. In practice this has some impact on detection of abrupt material loss as shown in Figure 13. A parametric study is conducted that considered multiple combinations of windows and total classifications that resulted in a 5% FAP. The difference for most losses is insignificant but caused a 23% difference in detection probability for scenario 1, which is the most abrupt scenario. This reflects threshold where the total number of classifications is larger than the duration of the abrupt loss.

6. Ideal results

The previous section identified several factors that require attention to ensure adequate performance of the machine learning pipeline. Ideal (optimistic) performance can be quantified by accounting for these factors. The overall performance of the machine learning algorithm is compared to the traditional Page's trend test on SITMUF under near identical conditions in Figure 14. The ideal conditions used for Figure 14 made several assumptions:

- Sufficient training data available
- Optimal threshold selection
- No facility transients

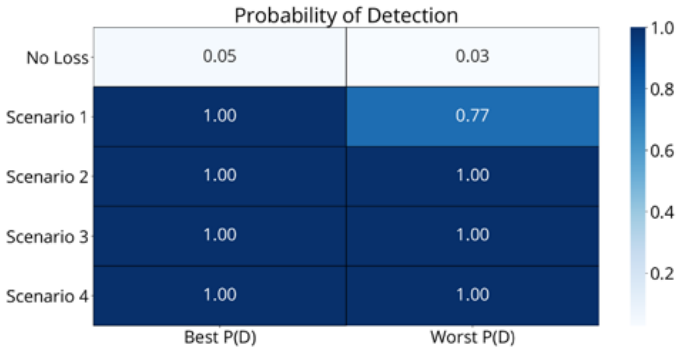

Figure 13: Probability of detection for several loss scenarios with varied thresholds.

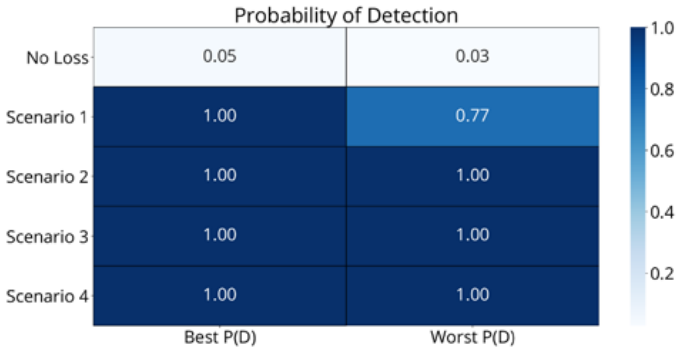

Figure 14: Detection probabilities for various loss scenarios.

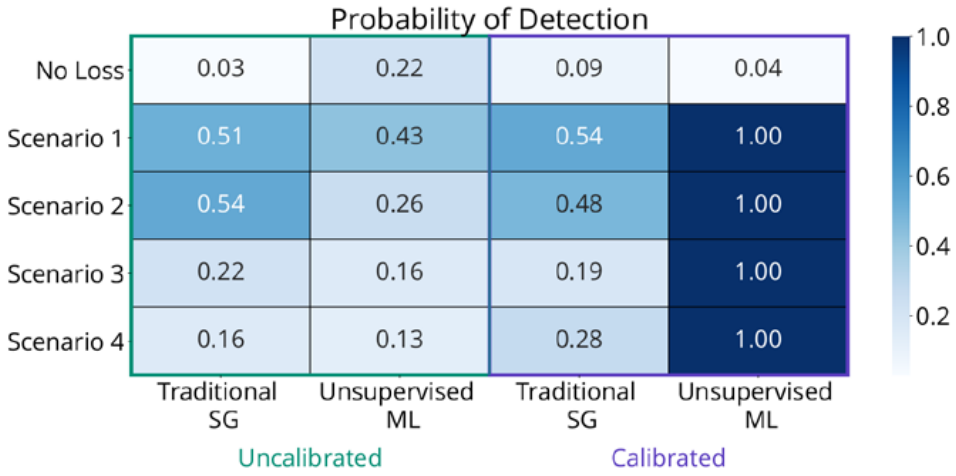

Figure 15: Detection probabilities for various loss scenarios.

Measurement error has many potential sources, and as such, could be difficult to resolve in real-world deployment scenarios. Consequently, it is not surprising that large portions of R&D for safeguards target reductions in measurement error. As there are no obvious data driven solutions to reduce the measurement error, Figure 15 considers the performance of both traditional statistical methods for safeguards and the unsupervised machine learning pipeline under "Uncalibrated" and "Calibrated" measurement conditions. Uncalibrated conditions are similar to current practices at facilities where sensors are placed and measured independently. The calibrated condition considers an experimental procedure wherein sensors are calibrated against each other (i.e. cross-calibrated). Here, the systematic errors for each sensor are the same non-zero value. For example, instead of having one sensor +1% biased and another being -2% biased, all sensors are biased at the same level.

The simulated calibration procedure has a large impact on the performance of the machine learning approach. Without it, performance is worse than traditional safeguards and very poor for most scenarios. It is interesting to note that the traditional safeguards approaches do not significantly benefit from this calibration procedure. This likely arises from implementation details for each approach. The machine learning algorithm is comparing signals from different locations in the facility, which is sensitive to mismatched biases specifically (i.e. large differences between sensors). However, the traditional safeguards approach is focused on quantifying MUF, which is sensitive to error in general.

7. Conclusions

This work proposed an unsupervised machine learning pipeline consisting of two steps to improve safeguards of bulk facilities. Several practical performance factors are crucial for real world performance. Several important findings are summarized below:

- Data representation is important to achieve adequate training performance
- Data availability has a significant impact on performance - Some facility operations are easier to learn than others and thus less susceptible to smaller training datasets
- Online training will likely be required after facility transients
- Measurement error has a significant negative impact on anomaly detection performance of the unsupervised machine learning approach
- The proposed alarm criteria are inflexible and can cause result in poor performance when not properly optimized. - A better criterion should be developed in future work.

Additionally, the generalization of the proposed pipeline was not studied in depth here. However, it is hypothesized that this approach will exhibit poor generalization even for facilities of the same type (I.e. other PUREX reprocessing facilities). The behavior learned through training will likely vary from facility to facility due to differences in equipment and facility layout. Applicability of common mitigation strategies for small datasets, such as transfer learning [23], to this problem remain unknown.

This work shows that unsupervised machine learning has the potential to out-perform traditional safeguards, but several requirements must be satisfied. There are several challenging limitations that are raised which make it unlikely that ML will wholly replace traditional safeguards in the near future. Data driven systems will likely complement existing safeguards systems until future work can resolve important barriers identified here.

8. Acknowledgements

This work was funded through the National Nuclear Security Administration's Office of International Nuclear Safeguards. The authors would also like to acknowledge the contributions of Michael R. Smith and Richard Fields (Sandia National Laboratories) who also contributed to this work. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2021-15805 J.

9. References

[1] United Nations, IAEA Statute, 1956.

[2] International Atomic Energy Agency, IAEA Safeguards Glossary, 2001 Edition, 2002.

[3] International Atomic Energy Agency, Establishing a System for Control of Nuclear Material for Nuclear Security Purposes at a Facility during Use, Storage and Movement, 2019.

[4] A. Goldman, R. Picard and J. Shipley, "Statistical Methods for Nuclear Material Safeguards: An Overview," Technometrics, vol. 24, pp. 267-275, 1982.

[5] K. Zhao, M. Penkin, C. Norman and S. Balsley, "International Target Values 2010 for Measurement Uncertainities in Safeguarding Nuclear Materials," ESARDA Bulletin, 2012.

[6] R. R. Picard, "Sequential analysis of material balances," Journal of Nuclear Materials Management, vol. 15, 1987.

[7] B. Jones, "Near real time material accountancy using SITMUF and a joint page's test: comparison with MUF and CUMUF tests," ESARDA Bulletin, vol. 15, pp. 20-26, 1988.

[8] E. S. Page, "Continuous Inspection Schemes," Biometrika, 1954.

[9] T. Burr and M. S. Hamada, "Revisiting Statistical Aspects of Nuclear Material Accounting," Science and Technology of Nuclear Installations, 2013.

[10] C. R. Orton, C. G. Fraga, R. N. Christensen and J. M. Schwantes, "Proof of concept experiments of the multi-isotope process monitor: An online, nondestructive, near real-time monitor for spent nuclear fuel reprocessing facilities," Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, vol. 672, pp. 38-45, 2012.

[11] I. T. Jolliffe, "Principal Component Analysis and Factor Analysis," in Principal Component Analysis, 1986, pp. 115-128.

[12] T. Burr and M. S. Hamada, "Smoothing and Time Series Modeling of Nuclear Material Accounting Data for Protracted Diversion Detection," Nuclear Science and Engineering, pp. 307-320, 2017.

[13] S. Prasad, D. Booth, M. Y. Hu and S. Deligonul, "The Detection of Nuclear Material Losses," Decision Sciences, vol. 26, no. 2, pp. 265-281, 1995.

[14] E. J. Hannan, Multiple Time Series, Wiley, 1970.

[15] R. Gladen, T. Grimes, B. Wilson, D. Jack, N. Shoman and B. B. Cipiti, "Neural Assessment of Non-Destructive Assay for Material Accountancy," ESARDA Bulletin, vol. 63, 2021 (Preprint).

[16] B. Cipiti and N. Shoman, "Bulk Handling Facility Modeling and SImulation for Safeguards Analysis," Science and Technology of Nuclear Installations, 2018.

[17] B. B. Cipiti, "Process Monitoring Considerations for Reprocessing," in Institute of Nuclear Materials Management, 2015.

[18] R. J. Jones, W. R. Kane and M.-S. Lu, Detailed Description of an SSAC at the Facility Level for Irradiated Fuel Reprocessing Facilities, 1986.

[19] G. Cybenko, "Approximation by Superpositions of a Sigmoidal Function," Mathematics of Control, Signals, and Systems, pp. 303-314, 1989.

[20] S. Hochreiter and J. Schmidhuber, "Long Short-term Memory," Neural Computation, 1997.

[21] F. T. Liu, K. M. Ting and Z.-H. Zhou, "Isolation Forest," in IEEE International Conference on Data Mining, 2008.

[22] N. Shoman and T. Burr, "Impact of Measurement Error on Deep Neural Networks for Nuclear Material," Nuclear Engineering and Design, 2021 (Preprint).

[23] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang and C. Liu, "A Survey on Deep Transfer Learning," in International Conference on Artificial Neural Networks, 2018.