Enhancing Nuclear Safeguards with Time Series Sketching-Based Nuclear Material Loss Tracking

Hao Huang[1], Scott Evans[1], Philip Honnold[2], Nathan Shoman[2]

[1] *GE Vernova Advanced Research, Niskayuna, NY, USA*

[2] *Sandia National Laboratories, Albuquerque, NM, USA*

hao.huang1@gevernova.com, evans@gevernova.com, phonnol@sandia.gov, nshoman@sandia.gov

Abstract

Detecting nuclear material loss or leak events is a critical challenge for nuclear safeguards and material accountancy in the recycling of used nuclear fuel. The MAYER (Multi-Sensor Assimilation Yielding Enhanced Reliability) project, part of the ARPA-E CURIE (Converting UNF Radioisotopes Into Energy) program, aims to develop a framework for tracking material loss in nuclear facilities by integrating multi-sensor data with predictive modeling. In this paper, we introduce MAterial Loss Tracking via Time Series Sketching (MALTS), a deep learning-based method designed to detect material loss events across the nuclear fuel recycling system. MALTS enhances the accuracy and robustness of nuclear material loss tracking by employing time series sketching to capture essential patterns while filtering out sensor noise, resulting in more stable predictions despite sensor noise effects. This approach also improves the time efficiency of material loss tracking by reducing the dimensionality of high-frequency sensor data, thereby enhancing computational scalability and enabling real-time inference. To further provide insights into the leak, MALTS ranks anomalous channels by post-processing results with a pretrained vision-language model (VLM) that considers the system flow diagram, generating a sorted list of anomalous channels from upstream to downstream. The initial leak location is identified as the first upstream channel. Experimental results demonstrate MALTS's effectiveness and efficiency in accurately identifying unseen nuclear material loss events and pinpointing initial leak locations, making it suitable for deployment within the MAYER digital twin framework for nuclear material safeguards.

1. Introduction

Ensuring the security and accountability of nuclear materials is a strategic priority for the Advanced Research Projects Agency–Energy (ARPA-E) and the U.S. Department of Energy (DOE). Developing solutions that would enable facility reprocessing operators to keep the tracking and measurements of used fuel flowing in the pipes of the facilities themselves has the potential to yield estimated savings in the hundreds of millions of dollars annually in operating costs while enabling more used fuel to be recycled with enhanced safeguards. Currently, fuel reprocessing operators have to temporarily shut down plant operations to perform a full inventory of nuclear materials within the plant. Additionally, there are planned scheduled outages for these physical inventory checks.

The CURIE (Converting UNF Radioisotopes Into Energy) program, established by the DOE through ARPA-E, aims to develop advanced technologies for recycling used nuclear fuel (UNF) and improving nuclear safeguards *CURIE program by ARPA-E* (2022). A critical component of CURIE is the MAYER project (Multi-Sensor Assimilation Yielding Enhanced Reliability) *MAYER ARPA-E* (2023), which focuses on creating digital twins to enhance nuclear material accountancy and loss detection. MAYER integrates data from multiple sensors within nuclear facilities, using predictive models to identify material loss or leak events accurately Honnold et al. (2024).

State-of-the-art methods for nuclear material accountancy, such as Material Unaccounted For (MUF) and Sequential Imbalance Testing for MUF (SITMUF), rely heavily on statistical analysis of material balances, necessitating a sufficient observation period to ensure reliable assessments Shoman & Moosir (2023). However, these methods often struggle with issues such as sensor noise and incomplete data, which can impede their ability to detect nuclear material loss or leaks in real time. Machine learning (ML) approaches have shown promise in enhancing anomaly detection by learning complex patterns from historical data. Nevertheless, supervised ML models require labeled loss events, which are frequently unavailable or incomplete in historical records, thereby limiting their applicability Shoman & Honnold (2022). Unsupervised ML algorithms present an attractive alternative, as they do not require prior knowledge of potential safeguard anomalies. Despite this advantage, previous research has highlighted the limitations of popular unsupervised algorithms, particularly their susceptibility to measurement noise in safeguards data streams Shoman & Honnold (2022).

To address these challenges, we propose MAterial Loss Tracking via Time Series Sketching (MALTS), a deep learning-based prediction method designed to detect nuclear material loss events across all monitoring inventory channels in the targeted system. Utilizing time series sketching, MALTS enhances sensor noise robustness and time efficiency by reducing the dimensionality of high-frequency sensor data, enabling real-time inference. It captures essential patterns while filtering out measurement noise, improving the model's ability to generalize to unseen loss events. Additionally, the model provides insights into detected leak events by ranking anomalous channels through post-processing with a pretrained vision-language model (VLM) that considers the system flow diagram, generating a sorted list of anomalous channels from upstream to downstream. The initial leak location is identified as the first upstream channel.

This paper presents the design and experimental results of our proposed MALTS. We evaluate the model's performance using simulated sensor data from Sandia National Laboratories and demonstrate MALTS's effectiveness and efficiency in accurately and promptly detecting nuclear material loss events. The proposed approach aligns with MAYER's goal of creating a resilient and adaptable digital twin for enhanced nuclear safeguards. Our work enables fuel reprocessing operators to perform real-time tracking of a plant's nuclear materials inventory, helping to prevent unplanned shutdowns and reduce the duration of planned shutdowns, potentially saving operators hundreds of millions of dollars annually.

2. RELATED WORKS

Anomaly detection in time-series data is a critical component of nuclear material accountability, where timely and accurate detection of material loss is essential for safety and regulatory compliance. Over the years, several deep learning architectures, including Recurrent Neural Networks (RNNs), Temporal Convolutional Networks (TCNs), and Transformers, have been employed to address these challenges. Each method offers unique strengths and limitations that influence their suitability for specific tasks within the nuclear domain.

Recurrent Neural Networks (RNNs) have traditionally been favored for sequential data processing due to their ability to maintain a hidden state that captures information from previous time steps. This feature allows them to handle long-term dependencies within time-series data Hochreiter & Schmidhuber (1997). However, traditional RNNs suffer from vanishing and exploding gradient problems, which limit their effectiveness in capturing long-range dependencies crucial for detecting subtle anomalies in nuclear material flows. Variants such as extended Long Short-Term Memory (xLSTM) Beck et al. (2024) and Gated Recurrent Units (GRUs) Ravanelli et al. (2018) were introduced to mitigate these issues, but they remain computationally intensive and slow during inference, posing challenges for real-time monitoring in nuclear facilities. Additionally, their high computational requirements make them difficult to deploy on resource-constrained devices often used in nuclear environments.

The Shallow RNN (ShaRNN), proposed by Dennis et al. (2019), offers a more efficient approach by using shallow RNN layers while maintaining accuracy for time-series classification tasks Dennis et al. (2019). Although ShaRNN reduces computational resource usage, it still inherits the limitations of RNNs in handling long-range dependencies and sequence length, which are critical for real-time anomaly detection in nuclear material accountancy Cho et al. (2014); Graves (2013).

Temporal Convolutional Networks (TCNs) have emerged as a promising alternative to RNNs, offering the ability to model long-range dependencies without the vanishing gradient issues Bai et al. (2018b). TCNs utilize dilated convolutions to expand the receptive field, efficiently capturing temporal patterns necessary for monitoring nuclear material flows. Their parallelization capability during training enhances computational efficiency compared to RNNs. The ModernTCN model, introduced by Luo and Wang (2024), represents a state-of-the-art approach in time-series analysis, enhancing the basic TCN architecture with a purely convolutional structure Luo & Wang (2024). This allows ModernTCN to process long-range temporal dependencies effectively while maintaining stable gradients during training, making it suitable for the complex data streams in nuclear facilities.

Despite their advantages, TCNs can be computationally expensive, especially with high-dimensional or high-frequency data typical in nuclear monitoring systems, as large kernel sizes are needed to capture long-range dependencies Lea et al. (2017). Although ModernTCN and other TCN variants improve performance over earlier models, they still face scalability challenges with massive time-series datasets, a common scenario in nuclear material monitoring Bai et al. (2018a); Yu et al. (2024).

Transformers have revolutionized machine learning by capturing long-range dependencies through self-attention mechanisms, eliminating the need for recurrent connections Vaswani et al. (2017). Initially designed for natural language processing, Transformers have been successfully applied to time-series anomaly detection Xu et al. (2021). Their ability to process sequences in parallel and model both local and global dependencies makes them well-suited for detecting anomalies in large-scale datasets, such as those encountered in nuclear material tracking. Recent works, such as the Transformer-based model for multivariate time-series anomaly detection, highlight their potential for capturing complex temporal patterns Tu et al. (2024).

However, despite their success, Transformer-based models are computationally expensive and require large amounts of data for training, posing challenges in resource-limited nuclear environments. They also typically require longer training and inference times, which can be a limitation in scenarios needing fast, real-time anomaly detection. Furthermore, Transformers are prone to overfitting when dealing with noisy datasets, making them less reliable in situations where sensor data is characterized by high levels of noise Tuli et al. (2022); Zhao et al. (2024).

The rapid evolution of deep learning architectures, including RNNs, TCNs, and Transformers, has greatly advanced the field of anomaly detection in time-series data. However, these models face limitations, particularly regarding computational speed, scalability, and their reliance on high-frequency datasets, which pose challenges in the context of nuclear material loss tracking. Our proposed method, MALTS, presents a promising alternative by offering a more efficient, robust, and interpretable solution for anomaly detection. This makes it particularly well-suited for real-time applications in nuclear material accountability. By leveraging time series sketching, MALTS effectively tackles the unique challenges of noise and data dimensionality inherent in nuclear environments, thereby enhancing the reliability and efficiency of material loss tracking.

3. OUR METHODOLOGY

We aim to monitor nuclear material loss from multi-sensor time series data using a streaming approach based on time series sketching techniques Huang et al. (2024). Our method enhances processing speed and enables real-time monitoring while minimizing the need for extensive historical data storage. Unlike batch-processing methods that process entire datasets at once, time series sketching operates incrementally on small time series batches, dynamically updating the sketch to represent observed data in real time. This approach is not only fundamentally different but also more efficient than RNN- and CNN-based techniques, which often incur higher computational overhead or suffer from long-term memory loss. These distinctions will be explored further in subsequent sections.

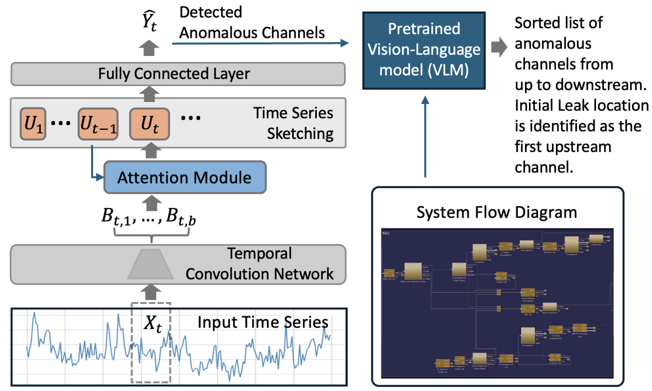

Figure 1. MALTS starts by extracting time series embeddings (B) with Temporal Convolutional Networks. These embeddings incrementally update the time series sketch (U) via an attention module. A fully connected layer predicts future measurements, where high residuals signal potential material loss and identify relevant inventory channels. A pretrained vision-language model (VLM), which understands the system flow diagram, further analyzes these channels to pinpoint the initial leak location.

3.1. Our proposed MALTS

MALTS utilizes time series sketching Huang et al. (2024), a technique that incrementally condenses high-dimensional time series data $X \in \mathbb{R}^{m\times n}$ into a more compact latent form $U \in \mathbb{R}^{m\times k}$, where $k \ll n$. The compressed representation U encapsulates all observed data to date and is used to predict incoming data. High prediction residuals act as indicators of potential material loss and help identify the associated inventory channels. Unlike traditional batch-processing methods, this approach continuously updates U as new data arrives, enabling efficient real-time monitoring without the need to store or process the entire historical dataset.

Step 1: Temporal Convolutional Network (*TCN*). In the field of time series analysis, capturing complex temporal dependencies is crucial for understanding the underlying patterns and dynamics of the data. To achieve this, we employ a Temporal Convolutional Network (TCN) Lea et al. (2017), a robust neural network architecture specifically designed to handle sequential data with temporal characteristics.

The TCN excels at extracting meaningful patterns from time series data by processing it through a series of 1D convolutional layers. Unlike traditional recurrent neural networks (RNNs), which process data sequentially, TCNs utilize convolutional operations to capture temporal dependencies across multiple time steps simultaneously. This capability allows TCNs to efficiently model temporal dependencies, making them particularly well-suited for our subsequent time series sketching step.

Our approach applies the TCN to sliding windows of time series data, as illustrated in the bottom left of Figure 1. Each sliding window represents a batch of consecutive time steps, providing a localized view of the data for the TCN to analyze. By moving the window across the entire time series, the TCN extracts features that encapsulate the temporal dynamics present within each segment. The output of the TCN is a set of extracted features, denoted as B_i, where i represents the index of the batch or window. These features are rich in temporal information and serve as a compact representation of the patterns identified by the TCN.

These extracted features B_i are then used as inputs to the subsequent time series sketching process (Step 2). The sketching process further refines the features, distilling them into a form that is both computationally efficient and robust to noise. While the TCN captures short-term fluctuations, the subsequent time series sketching step captures long-term trends, providing a comprehensive view of the data's temporal structure.

By utilizing a Temporal Convolutional Network, we effectively harness the power of convolutional operations to capture and represent complex temporal dependencies in time series data. This step is crucial for enabling real-time analysis and decision-making, as it reduces the dimensionality of the data while preserving its essential characteristics. This foundation sets the stage for the subsequent stages of our framework, ensuring that the extracted features are both informative and actionable for downstream tasks.

Step 2: Time Series Sketching. In our previous work Huang et al. (2024), we introduced a streaming time series sketching algorithm, detailed in Algorithm 1. This algorithm sequentially updates the time series sketch U using new time series embeddings B_i through the following process:

(1) Positional Encoding: To preserve the crucial temporal order information in time series analysis, we integrate a continuous positional encoding mechanism (Line 2) that maintains computational efficiency. We utilize Continuous Augmented Positional Embeddings (CAPE) Likhomanenko et al. (2021), which are recognized for their computational efficiency and robustness when handling variable-length inputs Zhai et al. (2023).

(2) Initialization: For the first batch (Lines 3–7), the sketch update operates on $(I + \frac{1}{b}\widetilde{B}_1\widetilde{B}_1^T)U_1$ instead of directly using $\frac{1}{b}\widetilde{B}_1\widetilde{B}_1^T U_1$ to facilitate stable convergence.

(3) Incremental Update: For subsequent batches (Lines 8–16), the sketch is updated using $W_0 U_{i-1}U_{i-1}^T U_i + \sum_{j=1}^{b} W_j \widetilde{B}_{i,j}\widetilde{B}_{i,j}^T U_i$, where W_0 and W_j represent weights for the prior sketch covariance and each new sample covariance, respectively. Here i denotes the batch index and j the sample index within each batch. The first term in the formula ensures that the historical information captured in the previous sketch U_{i-1} is preserved, while the second term allows the sketch to incorporate new data from the current batch. The weights W provide a mechanism to control the influence of the prior sketch and the new data. This weighting is crucial for balancing the retention of historical information with the integration of new observations. It allows the model to dynamically adjust the emphasis on past versus current data, which can be particularly important in environments where the relevance of historical data may change over time.

(4) Attention-Based Weighting: The weights in the incremental update formula are derived via an attention mechanism. Keys are projected from the prior sketch and sample covariances (Line 11), while the query is derived from the new batch covariance (Line 12) to dynamically adjust weight assignments.

(5) Normalization for Stability: The normalization steps (Lines 6 and 15) are implemented to prevent the sketch magnitudes from increasing, thereby ensuring robust predictions and stable convergence.

Algorithm 1 Time series sketching

Input: A sequence of time series embedding batches $\{B_1, B_2, ...\}$, batch size b, starting vector $U_0 \in \mathbb{R}^d \sim \mathcal{N}(0, I)$, positional encoding function *pos*, key projection function *key*, query projection function *qry*.
Output: Sketch U_i at any time
Initialization : $U_1 \leftarrow U_0/\|U_0\|_2$
1: **for** each B_i **do**
2: Equip the input batch with their positional information by $\widetilde{B}_i \leftarrow B_i + pos(B_i)$
3: **if** $i = 1$ **then**
4: **while** U_1 not converge **do**
5: $U_1 \leftarrow U_1 + \frac{1}{b}\widetilde{B}_1\widetilde{B}_1^T U_1$
6: $U_1 \leftarrow U_1/\|U_1\|_2$
7: **end while**
8: **else**
9: $U_i \leftarrow U_{i-1}$
10: **while** U_i not converge **do**
11: calculate keys:
$K_0 \leftarrow key(U_{i-1}U_{i-1}^T U_i)$
$K_1 \leftarrow key(\widetilde{B}_{i,1}\widetilde{B}_{i,1}^T U_i)$
...
$K_b \leftarrow key(\widetilde{B}_{i,b}\widetilde{B}_{i,b}^T U_i)$
12: calculate query: $Q \leftarrow qry(\widetilde{B}_i\widetilde{B}_i^T U_i)$
13: calculate weights W from query Q and keys K
14: $U_i \leftarrow W_0 U_{i-1}U_{i-1}^T U_i + \sum_{j=1}^{b} W_j \widetilde{B}_{i,j}\widetilde{B}_{i,j}^T U_i$
15: $U_i \leftarrow U_i/\|U_i\|_2$
16: **end while**
17: **end if**
18: **end for**

Time series sketching reduces the dimensionality of high-frequency data, which inherently filters out noise by focusing on the most significant features and patterns. The attention mechanism in MALTS is tailored to work with the reduced data representation, allowing it to focus on the most relevant features without being overwhelmed by noise. This contrasts with the full attention mechanism in Transformers, which can sometimes amplify noise if not properly regularized.

Step 3: Real-Time Prediction and Leak Detection.

In the dynamic environment of real-time monitoring, accurately predicting system behavior and identifying anomalies as they occur is crucial for maintaining operational integrity. This step involves leveraging our predictive modeling to continuously assess the state of the system and detect potential leaks.

During each time step of the time series data transfer, the most recent data representation, known as the sketch U_i, is utilized. This sketch is a condensed form of the data that captures essential patterns and trends while filtering out noise. It serves as a compact yet informative input for the predictive model.

The sketch U_i is fed into a fully connected layer, a type of neural network layer that connects every input to every output, allowing for complex transformations and predictions. This layer processes the sketch to generate a prediction of the incoming measurement for the next time step. The prediction is not limited to a single channel but rather spans the entire space of monitoring channels, ensuring comprehensive coverage of all areas of interest within the system.

The core of this step lies in the comparison between the predicted and actual measurements. By calculating the residual, or the difference between these two values, the system can assess the accuracy of the prediction. A high residual indicates a significant deviation from expected behavior, which is flagged as a potential material leak. This deviation suggests that the system is experiencing an anomaly that could signify a leak or other malfunction.

The ability to predict across all monitoring channels is vital for effective leak detection. It ensures that no part of the system is overlooked and that any irregularities are promptly identified. This comprehensive approach allows for the observation and analysis of all channels, providing a holistic view of the security and accountability of nuclear materials.

By implementing real-time prediction and leak monitoring, the system can quickly respond to anomalies, minimizing the risk of undetected leaks and enabling timely interventions. This proactive monitoring strategy enhances the reliability and safety of the system, ensuring that potential issues are addressed before they escalate into more significant problems.

Step 4: Localizing Initial Leak Location. Once a potential leak event has been identified, the next crucial step is to pinpoint the initial location of the leak within the system. This process begins by taking the unsorted list of anomalous channels, which have been flagged during the leak detection phase (Step 3), and inputting them into a pretrained vision-language model (VLM).

The VLM is specifically designed to interpret and analyze complex system flow diagram, which represent the interconnected pathways and components of the system being monitored. By leveraging its understanding of both visual and textual data, the VLM can effectively map the relationships and dependencies between different channels within the system.

Upon receiving the list of anomalous channels, the VLM processes this information in the context of the system flow diagram. It evaluates the position and role of each channel within the overall system architecture, taking into account the direction of flow and the sequence of operations. This allows the VLM to generate a sorted list of anomalous channels, organized from upstream to downstream.

The sorting process is critical because it helps to trace the path of the anomaly back to its origin. In fluid or process systems, upstream channels are those that occur earlier in the flow sequence, closer to the source of the material or signal. By identifying the first channel in the sorted list, the system can accurately localize the initial leak location. This channel is considered the most likely point of origin for the leak, as it is the first point in the flow where the anomaly was detected.

This method of localizing the initial leak location not only enhances the accuracy of the detection process but also provides valuable insights for maintenance and repair teams. By knowing the exact starting point of a leak, teams can more efficiently address the issue, minimizing downtime and preventing further damage to the system.

Algorithm 2 MALTS

Input: A sequence of sensor time series batches $\{X_1, X_2, ...\}$.
Output: Detected leak and its initial location, if applicable.

```
for each X_i do
    Extract time series features using TCN (Step 1)
    Update time series sketch using Algorithm 1 (Step 2)
    Project the updated sketch to predict the incoming sensor reading Ŷ_i and identify anomalous channels based on high prediction residuals (Step 3)
    Utilize a pretrained VLM to sort anomalous channels from upstream to downstream, accurately pinpointing the initial leak location (Step 4)
end for
```

The complete MALTS framework is outlined in Algorithm 2 and depicted in Figure 1. By integrating time series sketching with an attention-based refinement process, MALTS balances computational efficiency with resilience to sensor noise, ensuring reliable real-time monitoring of nuclear material loss. Additionally, its capability to output sorted anomalous channels and identify initial leak locations enhances detection insights and informs subsequent leak solutions, making MALTS well-suited for real-time anomaly detection in environments with noisy or incomplete sensor data.

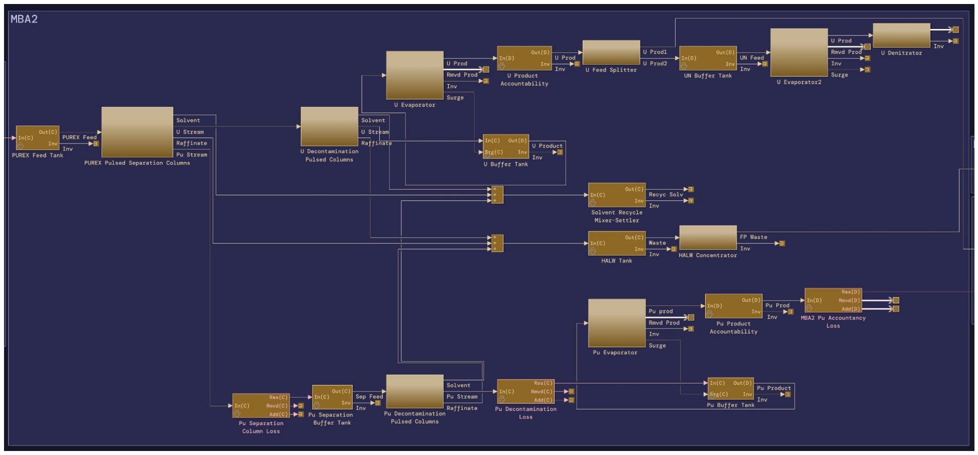

Figure 2. The simulated nuclear material flow diagram in a recycling system.

4. EXPERIMENTS

4.1. Problem and Dataset Description

Problem Introduction. Effective monitoring of nuclear material loss within reprocessing facilities is crucial for ensuring both nuclear safety and regulatory compliance. Traditional methods of Material Control and Accounting (MC&A) McMath et al. (2024); Taylor & Terentiev (1998) often face challenges in promptly detecting material loss due to the complex nature of nuclear processes and the inherent limitations of measurement systems. To address these challenges, the Multi-Sensor Assimilation Yielding Enhanced Reliability (MAYER) project *MAYER ARPA-E* (2023), under the ARPA-E CURIE program *CURIE program by ARPA-E* (2022), aims to develop advanced methodologies for real-time material loss detection in nuclear facilities. A key component of this initiative is the development of the Material Loss Tracking via Sketching (MALTS) model proposed in this paper, designed to enhance the detection and localization of material loss events by leveraging the observed multivariate time series data from multiple sensor channels.

Dataset Introduction.

The dataset utilized to evaluate the MALTS model was provided by Sandia National Laboratories and generated using the Material Performance Indicator Toolkit (MAPIT) Shoman & Moosir (2023). The dataset comprises multivariate time series data from 63 channels, each representing radiometric measurements at different locations within a simulated nuclear reprocessing facility (the flow diagram is shown in Figure 2). These channels collectively capture the dynamic behavior of nuclear materials as they move through various stages of the reprocessing workflow. The dataset contains both 'no-leak' and 'leak' data, representing normal operation and three simulated leak scenarios with different initial leak locations and propagation patterns, respectively. Each time series is sampled roughly every 30 minutes for 6480 hours, resulting in about 200,000 timesteps. The training set consists of 100 normal time series, while the testing set comprises 30 normal and 30 leak time series for each leak type. The data includes simulated measurement uncertainties and potential anomalies, providing a comprehensive testbed for evaluating the MALTS model's capability to detect and localize material loss events. By utilizing this MAPIT-generated dataset, the MALTS model's performance can be assessed in a controlled environment that closely mirrors the complexities and challenges of real-world nuclear fuel recycling systems. This evaluation is instrumental in demonstrating MALTS's potential for enhancing the reliability and efficiency of MC&A processes in nuclear facilities.

4.2. Experiment Setup

Baselines and Hyperparameter Setting. We compared our MALTS model with the following baseline methods: *MTCN* Luo & Wang (2024): ModernTCN is a convolution-based time series model with an enhanced Temporal Convolutional Network (TCN) architecture, designed for general analysis with a better balance of performance and efficiency. *LSTM* Filonov et al. (2017): The Long Short-Term Memory (LSTM) model is trained on normal data to regress future values, treating prediction error as an anomaly degree. *TranAD* Tuli et al. (2022): TranAD is a transformer-based anomaly detection model designed for multivariate time series data, leveraging self-attention mechanisms to capture long-range dependencies and detect anomalies with high precision. For their hyperparameter settings, we adhered to the default structures specified in their original papers, ensuring a fair comparison across models.

Evaluation Metrics.

We employed the following metrics to evaluate the performance of the models:

Table 1. The average AUPR (and standard deviation) in leak detection comparisons.

Method	Loss scenario 1			Loss scenario 2			Loss scenario 3		
	noise = 1e-5	noise = 1e-3	noise = 1e-2	noise = 1e-5	noise = 1e-3	noise = 1e-2	noise = 1e-5	noise = 1e-3	noise = 1e-2
MTCN	1.00(0.00)	0.95(0.01)	0.88(0.07)	1.00(0.00)	0.92(0.02)	0.87(0.09)	1.00(0.00)	0.90(0.03)	0.85(0.09)
LSTM	1.00(0.00)	0.93(0.02)	0.87(0.08)	1.00(0.00)	0.90(0.04)	0.85(0.10)	1.00(0.00)	0.87(0.07)	0.84(0.09)
TranAD	1.00(0.00)	0.96(0.01)	0.92(0.05)	1.00(0.00)	0.95(0.01)	0.90(0.08)	1.00(0.00)	0.92(0.03)	0.87(0.07)
MALTS	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	0.99(0.01)

Table 2. Ranks (and standard deviation) of truth initial leak location by each method.

Method	Loss scenario 1			Loss scenario 2			Loss scenario 3		
	noise = 1e-5	noise = 1e-3	noise = 1e-2	noise = 1e-5	noise = 1e-3	noise = 1e-2	noise = 1e-5	noise = 1e-3	noise = 1e-2
MTCN	1.00(0.00)	1.08(0.28)	1.11(0.32)	1.00(0.00)	1.12(0.32)	1.15(0.37)	1.00(0.00)	1.30(0.46)	1.50(0.59)
LSTM	1.00(0.00)	1.10(0.30)	1.15(0.35)	1.00(0.00)	1.13(0.34)	1.22(0.42)	1.00(0.00)	1.50(0.47)	1.70(0.90)
TranAD	1.00(0.00)	1.05(0.21)	1.08(0.28)	1.00(0.00)	1.09(0.30)	1.14(0.36)	1.00(0.00)	1.30(0.45)	1.40(0.49)
MALTS	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.00(0.00)	1.03(0.17)

1. *Leak Detection*: We use the Area Under the Precision-Recall Curve (AUPR) as a metric, which is not sensitive to class distribution and ranges from 0 to 1, with 1 indicating perfect detection. This metric is particularly useful in scenarios with imbalanced data, such as rare leak events.
2. *Initial Leak Localization*: We report the ranking of the true initial leak channel within the sorted channel array for each algorithm. Ideally, a rank of 1 indicates that the method successfully identifies the true initial leak channel as the highest-ranked.

We conducted 60 trials for each setting and documented the average performance and standard deviation to ensure statistical reliability of the results. These metrics provide a comprehensive assessment of the models' ability to not only detect leaks but also accurately localize the initial leak source, which is critical for timely intervention and mitigation in nuclear facilities.

4.3. Analysis Comparison

Leak Detection. We first report the performance of each algorithm in detecting material leaks. All baselines (MTCN, LSTM, TranAD) and our proposed *MALTS* train a regression model using leak-free data and identify time series with high prediction residuals during inference as anomalies or leaks. To evaluate the anomaly detection results, we use the Area Under the Precision-Recall Curve (AUPR), which ranges from 0 to 1, with 1 indicating perfect detection. AUPR is chosen because it is not sensitive to class distribution, making it particularly suitable for scenarios with imbalanced data, such as rare leak events.

Table 1 presents the leak detection performance for the three types of leaks using different algorithms. We observe that when the noise level is small (i.e., 1e-5), all algorithms exhibit comparably high accuracy. However, as the noise level increases, our proposed MALTS demonstrates significantly better accuracy and stable performance with a low standard deviation. On average, MALTS achieves up to 10% better AUPR when the noise level is 1e-3, and 18% better AUPR when the noise level is 1e-2. This superior performance is attributed to MALTS's ability to effectively filter out noise and capture essential patterns in the data by using the advanced time series sketching technique, enhancing its robustness in challenging conditions. This robustness is crucial for real-world applications where sensor noise can obscure critical signals.

Initial Leak Localization. Besides evaluating leak detection, another critical aspect is to find the initial material leak location to take subsequent action to resolve the issue. Table 2 presents the quality comparison of initial leak localization. Each method generates a sorted list of anomalous channels, and we recorded the rank of the ground truth initial leak channel for each method. Table 2 reveals that our proposed MALTS exhibits the best average performance, outperforming all other methods across the three leak cases. Overall, MALTS shows a $20\%+$ average improvement, consistently ranking the true initial leak location as the highest score in almost all cases and settings.

Figure 3 illustrates an output example of our MALTS from both the data-driven part and the VLM part. The data-driven neural network provides an unsorted list of detected anomalous channels, which is then fed into the VLM that takes into account the prior-known system flow diagram, and outputs a sorted list from upstream to downstream. The initial leak location is identified as the first upstream channel of the sorted list. This two-step process not only enhances the accuracy

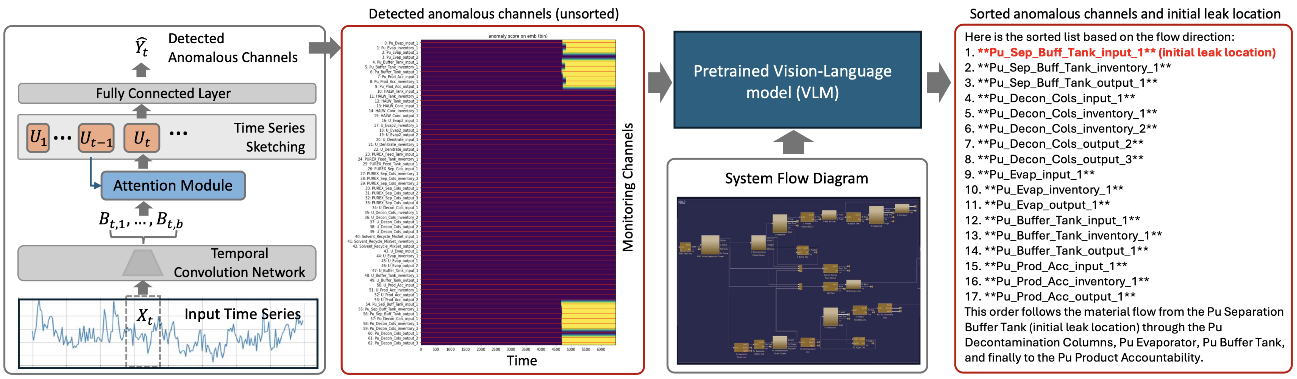

Figure 3. Output example of our MALTS.

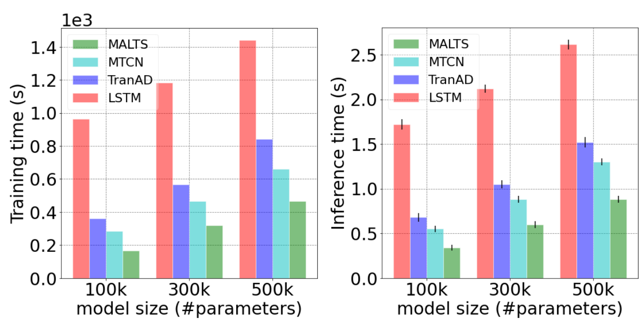

((a)) Training time comparison. ((b)) Inference time comparison.

Figure 4. Training and inference runtime comparison with different model sizes.

of leak detection and localization but also provides actionable insights for operators to quickly address the source of the leak.

4.4. Running Time Comparison

To demonstrate the running efficiency of our MALTS, Figures 4(a) and 4(b) present comparisons of training and testing times on a logarithmic scale. For a fair and comprehensive comparison, we evaluate the running time of all methods across varying model sizes based on parameter count. Specifically, we implement three versions of each algorithm, approximately with 100k (small), 300k (medium), and 500k (large) parameters. Our MALTS outperforms LSTM by a factor of three in terms of speed and is approximately twice as fast as TranAD and 50% faster than MTCN. This significant reduction in computational time is achieved without compromising accuracy, making MALTS a practical choice for real-time applications in nuclear material loss tracking. The efficiency gains are particularly crucial in operational settings where a rapid response to detected leaks is essential for maintaining safety and compliance, thereby broadening the applicability of MALTS in various nuclear safeguarding scenarios.

5. DISCUSSION

The core architecture of MALTS—particularly its use of time series sketching for noise-robust, efficient streaming anomaly detection—is inherently adaptable to other Prognostics and Health Management (PHM) domains such as aerospace, energy systems, and advanced manufacturing. The integration of a vision-language model (VLM) for physics-informed localization could similarly be extended to any system with known topology (e.g., pipeline networks, power grids). Future work will explore MALTS's performance on public PHM benchmarks, including the NASA Turbofan Degradation dataset and the IMS Bearing dataset, to assess cross-domain applicability.

We are working toward releasing a de-identified version of the dataset and open-sourcing a modular implementation of MALTS post-publication. To support adaptability across facilities, we are developing a transfer learning framework that allows fine-tuning MALTS with minimal data from new configurations, reducing retraining burden. We also intend to study repeatability under varying operational conditions, including maintenance cycles and human operator interactions, as part of ongoing field testing within the MAYER digital twin environment.

6. CONCLUSION

In this paper, we introduced the MAterial Loss Tracking via Time Series Sketching (MALTS) model, a novel approach designed to enhance the detection and localization of nuclear material loss events in reprocessing facilities. Our work is part of the broader MAYER project under the ARPA-E CURIE program, which aims to improve nuclear safeguards through advanced methodologies.

The MALTS model leverages time series sketching to effectively filter out sensor noise and capture essential patterns in multivariate sensor data, thereby improving the robustness and accuracy of anomaly detection. Our experimental results, conducted using a comprehensive dataset from Sandia

National Laboratories, demonstrate that MALTS outperforms popular machine learning models in both leak detection and initial leak localization, particularly under challenging conditions with high sensor noise levels.

Furthermore, MALTS exhibits significant computational efficiency, achieving faster training and inference times compared to baseline models. This efficiency, combined with its high accuracy, makes MALTS a practical solution for real-time monitoring and safeguarding of nuclear materials, potentially reducing operational costs and enhancing safety.

The integration of MALTS into the MAYER digital twin framework represents a significant step forward in the development of resilient and adaptable systems for nuclear material accountancy. By enabling real-time tracking and localization of material loss events, our approach supports facility operators in maintaining continuous operations and minimizing the impact of unplanned shutdowns.

Future work will focus on further refining the model's capabilities, including its adaptability to different facility configurations and its integration with other predictive maintenance tools. Additionally, we aim to explore the application of MALTS in other nuclear domains where real-time anomaly detection is critical.

In conclusion, the MALTS model offers a promising advancement in nuclear safeguards, providing a robust, efficient, and scalable solution for the complex challenges of material loss tracking in nuclear reprocessing facilities.

7. ACKNOWLEDGMENT

The information, data, or work presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0001688. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

REFERENCES

Bai, S., Kolter, J. Z., & Koltun, V. (2018a). Convolutional sequence modeling revisited.

Bai, S., Kolter, J. Z., & Koltun, V. (2018b). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.

Beck, M., Pöppel, K., Spanring, M., Auer, A., Prudnikova, O., Kopp, M., … Hochreiter, S. (2024). xlstm: Extended long short-term memory. *arXiv preprint arXiv:2405.04517*.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Curie program by arpa-e. (2022). `https://arpa-e.energy.gov/news-and-events/news-and-insights/us-department-energy-awards-38-million-projects-leading-used-nuclear-fuel-recycling-initiative`. (Accessed: 2022-10-21)

Dennis, D., Acar, D. A. E., Mandikal, V., Sadasivan, V. S., Saligrama, V., Simhadri, H. V., & Jain, P. (2019). Shallow rnn: accurate time-series classification on resource constrained devices. *Advances in neural information processing systems*, *32*.

Filonov, P., Kitashov, F., & Lavrentyev, A. (2017). Rnn-based early cyber-attack detection for the tennessee eastman process. *arXiv preprint arXiv:1709.02232*.

Graves, A. (2013). Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, *9*(8), 1735–1780.

Honnold, P., Shoman, N., & Cipiti, B. (2024). Increased complexity in simulating measurement systems.

Huang, H., Shah, T., Evans, S., & Yoo, S. (2024). Energy efficient streaming time series classification with attentive power iteration. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 38, pp. 12574–12582).

Lea, C., Flynn, M. D., Vidal, R., Reiter, A., & Hager, G. D. (2017). Temporal convolutional networks for action segmentation and detection. In *proceedings of the ieee conference on computer vision and pattern recognition* (pp. 156–165).

Likhomanenko, T., Xu, Q., Synnaeve, G., Collobert, R., & Rogozhnikov, A. (2021). Cape: Encoding relative positions with continuous augmented positional embeddings. *Advances in Neural Information Processing Systems*, *34*, 16079–16092.

Luo, D., & Wang, X. (2024). Moderntcn: A modern pure convolution structure for general time series analysis. In *The twelfth international conference on learning representations* (pp. 1–43).

Mayer arpa-e. (2023). `https://arpa-e.energy.gov/programs-and-initiatives/search-all-projects/monochromatic-assays-yielding-enhanced-reliability-mayer`. (Accessed: 03-2023)

McMath, G., Lousteau, A., & Smith, S. (2024). Nuclear material accounting and control measurements. In *Non-destructive assay of nuclear materials for safeguards and security* (pp. 709–717). Springer.

Ravanelli, M., Brakel, P., Omologo, M., & Bengio, Y. (2018). Light gated recurrent units for speech recognition. *IEEE*

Transactions on Emerging Topics in Computational Intelligence, *2*(2), 92–102.

Shoman, N., & Honnold, P. (2022). *Limitations for data-driven safeguards at enrichment facilities.* (Tech. Rep.). Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

Shoman, N., & Moosir, P. M. (2023). *Open-source software for material accountancy analysis* (Tech. Rep.). Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

Taylor, S., & Terentiev, V. (1998). *Us national nuclear material control and accounting system* (Tech. Rep.). Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States).

Tu, F.-F., Liu, D.-J., Yan, Z.-W., Jin, X.-B., & Geng, G.-G. (2024). Stft-tcan: A tcn-attention based multivariate time series anomaly detection architecture with time-frequency analysis for cyber-industrial systems. *Computers & Security*, *144*, 103961.

Tuli, S., Casale, G., & Jennings, N. R. (2022). Tranad: Deep transformer networks for anomaly detection in multivariate time series data. *arXiv preprint arXiv:2201.07284*.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.

Xu, J., Wu, H., Wang, J., & Long, M. (2021). Anomaly transformer: Time series anomaly detection with association discrepancy. *arXiv preprint arXiv:2110.02642*.

Yu, L.-r., Lu, Q.-h., & Xue, Y. (2024). Dtaad: Dual tcn-attention networks for anomaly detection in multivariate time series data. *Knowledge-Based Systems*, *295*, 111849.

Zhai, S., Likhomanenko, T., Littwin, E., Busbridge, D., Ramapuram, J., Zhang, Y., ... Susskind, J. M. (2023). Stabilizing transformer training by preventing attention entropy collapse. In *International conference on machine learning* (pp. 40770–40803).

Zhao, M., Peng, H., Li, L., & Ren, Y. (2024). Multivariate time series anomaly detection based on spatial-temporal network and transformer in industrial internet of things. *Computers, Materials & Continua*, *80*(2).