

TRACKING MATERIAL TRANSFERS AT A NUCLEAR FACILITY WITH PHYSICS-INFORMED MACHINE LEARNING AND DATA FUSION

Ken Dayman

Oak Ridge National Laboratory

Jason Hite

Oak Ridge National Laboratory

Riley Hunley

Oak Ridge National Laboratory

Nageswara S. V. Rao

Oak Ridge National Laboratory

Christopher Greulich

Oak Ridge National Laboratory

Michael Willis

Oak Ridge National Laboratory

James Ghawaly

Oak Ridge National Laboratory

Dan Archer

Oak Ridge National Laboratory

Jared Johnson

Oak Ridge National Laboratory

ABSTRACT

Operations at nuclear fuel cycle facilities are complex sequences of subtasks that include material production, destruction, or chemical processing; packing and unpacking; and transfer into, out of, or between areas or buildings within the facility. Identifying these tasks and tracking temporal patterns in their occurrences might indicate and help characterize long-term operations at the facility. We have developed methods to selectively identify transfers of special nuclear material using a network of low-frequency acoustic and NaI detectors deployed around a multiuse research reactor and radiochemical processing facility. The testbed for the methods is the collocated High Flux Isotope Reactor and Radiochemical Engineering Development Center at Oak Ridge National Laboratory, which have been instrumented with multiple measurement modalities under the Multi-Informatics for Nuclear Operation Scenarios project. By studying the physics of neutron-emitting nuclear material in a heavily shielded cask, we developed a physics-informed feature construction procedure that enables specific identification of shielded material using spectral data when typical analysis procedures, such as peak analysis, are not applicable. Data streams from six measurement stations are fused using a sequence model based on long, short-term memory recurrent neural networks. In the initial results, trained on synthetic data and transferred to real data collected at the testbed, these recurrent models successfully identified known sequences of material transfer detections throughout the multisensor network that are indicative of material transfers while also differentiating these patterns from spurious false alarms typical of complex, multiuse facilities.

INTRODUCTION

Detecting, characterizing, and tracking operations and activities at nuclear fuel cycle facilities are complex processes that typically include collecting, storing, and analyzing material assays, radiation measurements, material transfer and processing logs, security records (e.g., controlled-area access logs), and security sensor data (e.g., video footage). Often these data are analyzed by security teams in a resource-intensive, largely manual review process. The growing ubiquity of small footprint sensors and resource-efficient data storage, transfer, and analysis capabilities can potentially revolutionize how material protection and security is ensured in the future by enabling automated

analysis of activities at nuclear fuel cycle facilities using multisensor systems and artificial intelligence (AI) analysis methods.

The US Department of Energy sponsored Multi-Informatics for Nuclear Operations Scenarios (MINOS) project aims to develop multimodal data analysis capabilities for detecting and characterizing operations at a real-world, multiuse nuclear facility. Under the MINOS project, a testbed has been established at Oak Ridge National Laboratory (ORNL) consisting of a high-power research reactor—the High-Flux Isotope Reactor (HFIR)—and a collocated radiochemical processing facility—the Radiochemical Engineering Development Center (REDC). The testbed has been instrumented in and around HFIR and REDCX with multiple sensors across several measurement modalities including radiation, effluent, electromagnetic, seismic, and infrasound. These sensors collect and store persistent data at a range of sampling rates and resolutions, and the resulting multimodal dataset is used to identify measurable signatures associated with physical phenomena to develop and demonstrate multimodal data-analytic solutions to various nuclear operation detection and characterization problems.

Previous work developed and demonstrated approaches to identify vehicle-borne transfers of special nuclear material around the testbed. These approaches used a combination of low-frequency acoustic data to identify vehicles and used count rate–based radiation alarms to identify radioactive sources in proximity to the sensors. Although this initial work demonstrated the potential of this data-fusion approach, numerous false positives were observed owing to lack of selectivity of the count rate analysis. This paper highlights progress made to improve the techniques used to detect and track transfers of nuclear material between HFIR and REDC and focuses on the processing and analysis of the radiation measurement modality to increase selectivity of detections for shielded special nuclear material and develop a method to track highly variable temporal patterns of these detections in the testbed’s multisensor network.

THE TESTBED

This section provides a brief overview of the testbed and associated instrumentation from which collected data is used in the development and demonstration of the autonomous, AI-based analysis and data-fusion methods.

Facilities

HFIR is an 85 MW_{th} high-performance research reactor used primarily to produce a constant neutron flux for neutron scattering applications and isotope production. REDC contains a variety of radiochemical infrastructure, including gloveboxes, hot cells, and industrial chemical processing equipment. Both HFIR and REDC are multiuse facilities with a variety of research and operational missions and activities occurring simultaneously. Two major efforts currently ongoing at HFIR/REDC are the production of ²³⁸Pu and Cf. In both cases, target material is purified and processed into targets at REDC and transferred to HFIR, where the targets are loaded during a shutdown, irradiated for one or more cycles (each cycle lasting approximately 28 days), cooled in the reactor pool at HFIR, transferred externally back to REDC, where the irradiated targets are disassembled, chemically dissolved and separated, and the product material (i.e., Pu or Cf) is purified.

Sensors

Radiation data are collected by a small network of external sensors placed in and around HFIR and REDC, and along roads used for transferring material between the facilities and for general access and work. Each measurement station includes NaI detectors, data processing and collection hardware (e.g., a multichannel analyzer), and computers to locally store and wirelessly transmit the data to a centralized server in a robust, air-conditioned enclosure. Figure 1 shows the position of these measurement stations around HFIR and REDC; the sensors are labeled as MUSE nodes for historical reasons. Nominally, data are reported and stored at one-second time intervals; however, in practice, data are time-integrated prior to our analysis. More information on the radiation detection network deployed at the testbed can be found in [1].



Figure 1. Map of radiation sensors distributed around HFIR and REDC to analyze transfers of material between facilities.

DEVELOPING A PHYSICS-INFORMED FEATURE

Multiple types of material transfers are common around and between HFIR and REDC. These include, but are not limited to, shipments of fresh fuel for HFIR, shipments of spent fuel offsite, transfers of fresh ^{237}Np or Cm targets from REDC to HFIR for irradiation, transfers of irradiated $^{237}\text{Np}/^{238}\text{Pu}$ or Cm/Cf targets from HFIR to REDC for processing, and transfers of ^{225}Ac from REDC to the main ORNL campus. The primary interest is identifying, characterizing, and tracking transfers of irradiated $^{237}\text{Np}/^{238}\text{Pu}$ or Cm/Cf targets and differentiating those from other material transfers and other movements common around the site (e.g., personnel, working, or construction vehicle movements). Previous research has focused on developing analyses based on low-frequency acoustic data to detect and characterize vehicles moving around the testbed [2], and these detections have been fused with radiation alarm data derived from analysis of gross count rates (e.g., by using common radiation detection algorithms such as the sequential probability ratio test [3] or k-sigma [4]). In initial analyses, this lightweight approach to radiation data analysis does alarm on the transfers of interest; however, numerous false positives occurred that we attribute to innocuous events such as Rn washout associated with rain. To this end, developing a more selective radiation analysis, based on spectral information, was pursued.

During transfers of irradiated Np and Cm targets, the material is contained in a large, heavily shielded cask colloquially called the “Q-Ball.” The Q-Ball includes concrete and steel structural elements and a water backfill to moderate and attenuate neutrons emitted by the source. The abundance of shielding virtually guarantees no photons will be detectable at their characteristic energies by external detectors owing to attenuation and down scattering. Traditional spectral analysis that relies on peak analysis is not applicable to this problem because photopeaks are not typically observed, as shown in Figure 2. Although peaks are not prominent when the loaded cask is proximate to the detector, an apparent shift is observed in the gross energy distribution toward higher energies, and we hypothesize this shift is associated with secondary radiation generated by neutron interactions within the cask itself. Such interactions tend to yield higher-energy gamma rays than typically found in common terrestrial sources.




Figure 2. Typical gamma-ray spectra collected when a loaded cask is present (blue) and absent (orange). The spectra have been normalized to the unit area to highlight shape differences.

To capture this spectral energy shift, we first integrate the measured data for one minute (recall data are reported with a one-second time resolution) and capture the overall spectral shape using an algorithm called Baseline Estimation and Denoising using Sparsity (BEADS) [5]. This algorithm treats the measured energy spectrum, y , as a superposition of a low-frequency baseline, f , sparse signal with sparse derivatives, x , and noise, w , as shown in Equation 1.

$$y = f + x + w \quad (1)$$

The sparse signal (i.e., peaks in the spectrum) are extracted by solving the convex optimization problem shown in Equation 2.

$$\hat{x} = \underset{x \in \mathbb{R}^N}{\operatorname{argmin}} \left\{ \frac{1}{2} \|H(y - x)\|_2^2 + \sum_{i=0}^M \lambda_i \sum_{n=0}^{N_i-1} \phi([D_i x]_n) \right\} \quad (2)$$

Here, H is a high-pass filter, ϕ is an asymmetric penalty function used to encourage sparsity, and D_i are finite-difference operators used to approximate derivatives of a discrete signal. As shown by the BEADS developers, Equation 2 can be solved efficiently to give a unique and globally optimal solution [5]. We have ported the original authors’ implementation to Python and modified the internal algorithm parameters to accommodate the lower energy resolution given by NaI detectors and to explicitly return the baseline estimate.¹ An example of this spectral decomposition is shown in Figure 3.




Figure 3. Example decomposition of a NaI spectrum using modified BEADS implementation.

Given the smooth baseline estimate (blue in Figure 3), we next apply a cumulative summation transform to remove effects of gross count rate—which can be affected by source-detector distance, time of day, or season—and highlight the baseline’s change in shape. Two such transformed baseline estimates are shown in the left pane of Figure 4, and the difference between the two is shaded. To summarize the energy distribution’s shift to higher energies when the shielded, neutron-emitting source is near the NaI detector, we compute a head-to-tail ratio—shown in Equation 3—to yield a univariate feature, in which values higher than the typical baseline indicate shifts toward higher energies, values lower than the baseline indicate shifts toward lower energies, and large peaks above baseline are associated with proximate neutron-emitting sources.

$$\beta = -\log \frac{\sum_{i=1}^H x_i}{\sum_{i=H+1}^N x_i} \quad (3)$$

In practice, the typical value of this derived metric slowly changes over time, and the signal can display high-frequency variability. Therefore, the robust identification of peaks in the metric, suggesting the presence of loaded casks used to transfer irradiated special nuclear material around the HFIR/REDC site, cannot be done with a simple linear threshold or moving average, but rather a moving median is used. The output of this complete workflow is shown in Figure 5.

¹ BEADS was originally developed to only extract peaks in gas and liquid chromatography.

- (1) To summarize the analysis workflow, for each detector station in the network (six total): data are integrated over time for one minute to create spectra;
- (2) the baseline for each spectrum is extracted with the modified BEADS algorithm;
- (3) the baseline is transformed to a cumulative distribution function;
- (4) a head-to-tail ratio is computed for each transformed function, and the time series of head-to-tail ratios create a univariate time series for each detector station; and
- (5) each pseudo-continuous time series is converted to a binary time series by defining alarm states using a moving median analysis to track the nominal metric value and alarming on significant excursions above or below the nominal value.




Figure 4. (Left) Baseline extracted from spectra collected with (orange) and without (blue) the proximate loaded Q-Ball cask and then transformed with a cumulative sum. The difference is highlighted in green. (Right) Head-to-tail ratio computed and plotted as a function of time. The large peak is associated with a loaded cask passing the detector.

TEMPORAL SEQUENCE ANALYSIS

During development, our proposed analysis was observed to be sensitive to transfers of irradiated Np and Cm targets in the heavily shielded Q-ball cask but would occasionally alarm during other events, especially rain and associated Rn washout. Additionally, because HFIR and REDC are complex, multiuse facilities, movements of radiative material that could trigger the analysis discussed previously do not always indicate a transfer of interest. To more uniquely identify transfers between the testbed facilities, we developed a model to automatically analyze the temporal pattern in alarms (red in Figure 5) to differentiate a transfer following one of several known routes from spurious events that could lead to false alarms when data from a single detection position are considered in isolation.

Transfers at the testbed have numerous sources of variation, so even data associated with transfers of the same material following the same route will be different. These sources include (but are not limited to) transfer start time, vehicle speed, source strength, human variability (e.g., delays at stop signs), stochastic radiation emission/detection, and transport between the source and a detector. Moreover, even though the spectral shape analysis discussed in the previous section performed well in initial analyses and is reasonably selective for shielded neutron sources, this indicator is imperfect and is subject to false positives and false negatives. To capture such variable and complex signals, a multilayer neural network model was developed.




Figure 5. Physics-informed univariate metric derived from radiation data collected by a network of six NaI sensors distributed along roads at the HFIR/REDC site. The noisy and gently varying baseline metric values are estimated with a moving median, and large excursions (i.e., prominent peaks) above and below baseline values are highlighted in red.

The primary design features of the network are a one-dimensional convolution layer, followed by two long short-term memory (LSTM) layers, the output of which is aggregated by a single fully connected dense layer to provide the final output—whether a 20-minute time sequence of alarms derived from the 6-detector network corresponds to a transfer of Np/Pu or Cm/Cf between HFIR and REDC.

There are two immediate benefits to this model construction. First, applying convolution operations to the input signal removes the effects of variable start time because convolution is a shift-invariant operation; that is, the output of the convolutional layer does not change if the input signal is linearly shifted (i.e., the start time of the alarm states associated with a transfer is changed). As Bronstein showed, any shift-invariant function is circulant, and convolution is the only linear shift-invariant function [6]. Having effectively time aligned the beginning of a pattern of alarm states, an LSTM is used to determine if the noisy pattern is sufficiently consistent with a known transfer route. As stated, even the same transfer can give rise to different patterns in alarm states; however, the LSTM is invariant to time warping and well-behaved signal deformations [7], and the model should be insensitive to changes in the pattern (e.g., stretching caused by changes in vehicle speed).

To train the model, a large amount of labeled training data (i.e., example alarm patterns associated with and without Np/Pu or Cm/Cf transfers) was required. Despite persistent collections around the testbed, an insufficient number of transfers was observed. Therefore, synthetic alarm patterns were generated and used for training the model, which was then tested on the available real data collected at the testbed. To generate the required data, the procedure enumerated below was used.

- (1) Construct a graph representation of the radiation detection stations around the testbed with each station represented as a node. An edge between nodes i and j is constructed if a road is between the detector stations. The weight of the edge is proportional to the driving distance.

- (2) Choose a starting node, and with equal probability, move to any connected node not previously visited with probability p , and stop with probability $1 - p$.
- (3) Choose a constant vehicle speed from a uniform distribution ranging from 0.3 to 1.5 m/s, a start time within a 20-minute window, and a “visible distance” (i.e., from how far away can the detector detect the source, which is related to source strength) from a half-normal distribution with location 0 and scale 10 .²
- (4) Calculate time intervals in which the source is in view of each detector station.
- (5) Apply assumed true positive and false positive statistics to each 1-minute window and each detector station within the synthetic measurement.
- (6) Apply a “transfer of interest” label if the route chosen in step 2 is a known transfer route or its reverse and apply a “not transfer” label otherwise.
- (7) Repeat steps 2–6 until enough data is generated.

These training data were used to train the hybrid convolution-LSTM model described, which was then applied to all the data—collected during May 2020 at the testbed—in 20-minute increments (i.e., each 20-minute period was input into the model to assess if a transfer had occurred), and the results of this initial analysis are shown in Figure 6.

Each row of the figure shows the physics-informed univariate feature from one of the detection nodes shown in Figure 1. The highlighted regions show when the sequence model estimated that a transfer of Np/Pu or Cm/Cf targets between HFIR and REDC occurred, including on May 27, when a transfer was known to have occurred. At present, the causes of the alarms in the first half of the month are unknown; however, upon inspection of the raw radiation spectra, it is clear that *something* of interest occurred, and further work to attribute these alarms is ongoing. Considerable spikes might be observed outside of the highlighted regions identified by the sequence model, so it appears the sequence is successfully discriminating against many of the spurious movements of radioactive material typical of the HFIR and REDC site. Given the small number of alarms (four in a month), the automated analysis already represents a powerful capability to triage large volumes of sensor data and reduce the workload of human analysts.

SUMMARY

Previous work showed the viability of identifying vehicle-based transfers of heavily shielded nuclear material around a multiuse facility by fusing low-frequency acoustic data and gross radiation counts collected by a network of collocated audio and radiation sensors distributed around the site [2]. False alarm rates were observed using this approach, so a more selective radiation analysis workflow was developed using a combination of physics-informed feature construction and variable temporal pattern analysis.

² That is the random variable $Y = |X|$, where $X \sim \mathcal{N}(0, 10)$.




Figure 6. Suspected transfers of Np/Pu and Cm/Cf targets identified by the convolution-LSTM sequence model (red) overlaid on plots of the physics-informed spectral shape metric. Data were collected throughout May 2020 at the HFIR/REDC testbed. The alarm on May 27 (far right) is known to be a target transfer; work is ongoing to attribute the other alarm events.

Because the material of interest is heavily shielded during transfers, traditional photopeak analysis is not viable, so an approach to holistically analyze the spectrum was adopted. First, the baseline of the spectrum is extracted. To capture the shift in energy distribution when a shielded source is present, the extracted baseline is converted to a cumulative distribution function, and then a head-to-tail ratio is computed to yield a univariate metric capturing the relative position of the spectrum's energy distribution. Time series of this metric exhibit significant noise, occasional erratic changes, and slowly varying shifts in the nominal value. Therefore, a moving median approach is used to capture the moving baseline and identify significant departures from baseline. The net output of this workflow is a time series of binary variables for each of the radiation sensors in the network (see Figure 1), in which nonzero values correspond to significant shifts in the energy distribution in the spectra.

The temporal patterns of these binary values are analyzed to identify transfers that follow known routes and differentiate these from other spurious behavior expected at the site. Because transfers have variable start times, driving speeds, and source strengths, a model was designed to be shift-invariant and robust to time warping, and a neural network combining a convolutional layer with LSTM layers was adopted. To train this complex model, vehicle movements around the site and expected associated detection patterns were simulated by generating random walks around a graph abstraction of the site. Finally, after training the model on these simulated data, the model was applied to real data collected throughout May 2020 at the HFIR/REDC testbed. The known transfer of interest was successfully identified with, at most, three false positives during the month. Attribution of these

false alarms is ongoing, but initial inspection of the raw data suggests these alarms are not entirely nuisance alarms.

Further work will include more nuanced simulations of vehicle movements around the site and inclusion of detection patterns of additional measurement modalities, especially low-frequency acoustic and seismic measurements, to enable non-trivial multimodal data fusion using our existing sequence model architecture.

VALUE ADDED FOR NONPROLIFERATION AND SAFEGUARDS MISSIONS

Autonomously detecting and characterizing events at nuclear fuel cycle facilities—and combining these observations in sequence models to capture the context of the events—has the potential to greatly enhance material accountancy systems and benefit both domestic and international safeguards missions. For example, a system built using multiple sensors and modern data analytics could help provide positive assurance on all material within a facility by tracking all relevant activities and could even serve as a virtual second person for two-man rule movements. If realized, such a capability could greatly reduce the staffing and cost of complying with material security requirements.

Hard sensor data typical of nuclear safeguards and security applications has traditionally not been the focus of the wider data analytics community, but we believe the work presented here and within the larger MINOS effort demonstrates the viability of collecting and autonomously processing large volumes of hard sensor data that are immediately indicative of real but complex physical processes and phenomena associated with nuclear material and its processing.

ACKNOWLEDGMENTS

This work was funded by the Office of Defense Nuclear Nonproliferation Research and Development (NA-22), within the US Department of Energy’s National Nuclear Security Administration.

REFERENCES

- [1] A. D. Nicholson, D. E. Archer, I. Garishvili, I. R. Stewart, and M. J. Willis, “Characterization of gamma-ray background outside of the High Flux Isotope Reactor,” *J. Radioanal. Nucl. Chem.*, vol. 318, no. 1, pp. 361–367, 2018.
- [2] J. Hite *et al.*, “Automated Vehicle Detection in a Nuclear Facility Using Low-Frequency Acoustic Sensors,” in *Proceedings of the 23rd International Conference on Information Fusion*, 2020, no. July, pp. 1–6.
- [3] K. D. Jarman, L. E. Smith, and D. K. Carlson, “Sequential probability ratio test for long-term radiation monitoring,” *IEEE Trans. Nucl. Sci.*, vol. 51, no. 4 I, pp. 1662–1666, 2004.
- [4] P. E. Fehlau, J. C. Pratt, J. T. Markin, and T. Scurry Jr., “Smarter radiation monitor for safeguards and security,” in *24th Annual Meeting of the Institute of Nuclear Materials Management*, 1983.
- [5] X. Ning, I. W. Selesnick, and L. Duval, “Chromatogram baseline estimation and denoising using sparsity (BEADS),” *Chemom. Intell. Lab. Syst.*, vol. 139, pp. 156–167, 2014.
- [6] M. M. Bronstein, “Geometric Deep Learning: the Erlangen Programme of ML,” in *The Ninth International Conference on Learning Representations*, 2021.
- [7] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, “Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges,” 2021.