DISPARATE DATA INTEGRATION FOR ADVANCED FACILITY MONITORING

KAREN A. MILLER Los Alamos National Laboratory Los Alamos, New Mexico, USA Email: kamiller@lanl.gov

KENDRA L. VAN BUREN, JANETTE R. FRIGO, MAX ZEYEN, JOSHUA P. SACKOS, PAUL M. MENDOZA Los Alamos National Laboratory Los Alamos, New Mexico, USA

AUSTIN P. WRIGHT University of California, Berkeley Berkeley, California, USA

Abstract

The International Atomic Energy Agency (IAEA) must ensure the peaceful use of all nuclear materials with a budget that has been compared in size to that of the police department in Vienna. This includes, for example, coverage of nearly 1,300 nuclear facilities spread around the globe and verification of over 200,000 significant quantities of nuclear material. The amount of information the IAEA collects is on an upward trajectory, and data overload is poised to be an everincreasing stress on the IAEA's ability to perform its safeguards mission. Los Alamos National Laboratory (LANL) has been investing over the past several years in experimental studies within a number of its unique facilities to characterize activity patterns and operational modes using automated methods for disparate data integration. Building on the success of these preliminary studies, there is currently an effort funded by the Laboratory-Directed Research and Development (LDRD) Program at Los Alamos to develop a testbed at one of the laboratory's radiological facilities for advancing this work, specifically aimed at safeguards-relevant data streams. The overarching goal of this work is to develop and experimentally validate methods to improve the efficiency and effectiveness of safeguards verification at nuclear facilities, which will allow the IAEA to better utilize the data it is already collecting. The paper describes work that has been completed to date as well as implications for future areas of research.

1. INTRODUCTION

The world is at an inflection point where our ability to collect data far outpaces our ability to make use of it, and the proliferation of sensors and so-called Internet of Things (IoT) is changing the way society uses data in profound ways. The International Atomic Energy Agency (IAEA) is fundamentally a data-driven organization, and the amount and complexity of data the agency collects to fulfill its safeguards mandate continues to grow [1,2]. The IAEA is taking proactive steps to modernize its information technology infrastructure [3] at the same time that unprecedented advances in data science and automation are transforming the technology landscape in commercial industry. The convergence of these enabling factors provides a catalyst for renewed focus on developing new tools and methodologies for improving the IAEA's productivity.

The objective of safeguards is to deter the spread of nuclear weapons through early detection of the misuse of nuclear material or technology. Increased use of continuous, unattended monitoring systems has opened the door for the IAEA to move beyond static signatures to characterizing dynamic activities in aggregate over time and space. The hypothesis of this work is that characterizing "normal" activity patterns in multisource data will provide a baseline for the IAEA to detect anomalies and give them enhanced capability to detect misuse or diversion of nuclear materials. The project pulls inspiration from modern smart home technologies and commercial IoT sensing and applies it to the challenge of monitoring and verification of nuclear facilities. Building off the idea that the whole is greater than the sum of its parts, the goal is to develop automated methods for combining disparate data streams representative of those collected by the IAEA to improve detection of undeclared activities at safeguarded facilities.

The potential for this type of integrated analysis to improve facility monitoring has long been recognized. In the early 1990s, the Safeguards Group at Los Alamos developed a neural-network-based method called Video Time and Radiation Analysis Program (VTRAP) for identifying patterns and detecting anomalies in combined containment and surveillance (C&S) and non-destructive assay (NDA) data [4]. While VTRAP was not operationalized at the time, technological advances in data storage and processing have dramatically improved the feasibility of implementing the underlying approach.

More recently, Los Alamos has been investing in research to advance next-generation monitoring systems using the laboratory as a testbed [5,6]. As part of the initiative, our one-year project was targeted at building the experimental components of a testbed for safeguards technology. It required a multidisciplinary team of experts in areas such as NDA systems, distributed sensing, computer science, and machine learning. The project began in November 2017 and included the following tasks:

- Design a sensor network for recognizing targeted activities;
- Design the associated data management architecture;
- Perform scoping measurements using available sensors; and
- Conduct preliminary data analysis.

The remaining sections describe progress to date as well as implications for future areas of research. As safeguards technology has evolved to meet new challenges over the past fifty years, one constant that has not changed is the central and indispensible role of highly-skilled IAEA analysts in reaching safeguards conclusions. Within that context, the aim of this research is not to replace the analysts with automation but to maximize their value by freeing more time for verification tasks that require subject-matter expertise, reasoning, and critical thinking.

2. THE FACILITY

The first step of the project was to evaluate candidate facilities for the testbed. The team chose LANL's Safeguards Training Facility (TA-66) as the target facility based on a variety of attributes. Importantly, TA-66 is a Category 3 radiological facility. It includes office spaces, conference rooms, and a laboratory used for NDA training courses that feature hands-on nuclear material measurements [7]. Materials used regularly in the facility include, for example, cans of uranium, plutonium, and mixed uranium-plutonium oxide (MOX) as well as fresh fuel assemblies.

Typical activities at TA-66 include general office work, seminars, preparation for NDA courses, and hosting NDA courses. TA-66 was chosen, in part, because the operational modes are relatively simple, allowing for a phased approach to developing methods. The goal is to minimize ambiguity in initial studies focused on demonstrating the fundamental concepts. If this is successful, the concepts and approaches can then be applied to facilities with more complex activities.

3. DATA STREAMS

The IAEA analyses data from State declarations, on-site inspections, fielded equipment, and other sources to draw its safeguards conclusions [8]. This project integrates data streams that are representative of those that the IAEA might collect for a typical nuclear facility. Specifically, the data streams include nuclear material inventory records, an NDA system, and a network of custom-designed sensors that serve as a proxy for surveillance data.

3.1. Nuclear Material Inventory Records

LANL's Source Tracker software is used to monitor locations and quantities of nuclear materials [9]. It is akin to physical inventory listings provided through State declarations. The database logs transactions associated with nuclear material movements between material balance areas and buildings. For each transaction, it records information such as date and time of transaction, transaction type (e.g., check in, check out, transfer, etc.), location, source identification number, source activity, and material type (uranium, plutonium, MOX, etc.).

Exploratory analysis comparing historical Source Tracker data to past training course schedules provides confidence that inventory records will serve as a rich data stream for characterizing activities at TA-66. For example, the total number of transactions per day is typically larger on course and course setup days compared

to other times. The daily temporal patterns of when transactions occur may also provide some differentiation between course and course setup days, with course days having a more structured and repeatable schedule than course setup days. Finally, the material type may provide finer granularity on course topics.

3.2. Non-destructive Assay System

NDA systems are used routinely by the IAEA for both quantitative and qualitative measurements of nuclear materials. We have installed an NDA system comprised of standard IAEA hardware and software in the TA-66 laboratory space to monitor neutron counts and material movements. The setup includes a ³He-based slab detector called the Portable Handheld Neutron Counter (PHNC) [10] with an Advanced Multiplicity Shift Register (AMSR) [11] for data acquisition. These components are paired with a standalone laptop running Multi-Instrument Collect (MIC) [12] for data collection in unattended mode. Data is pulled off the laptop manually at regular intervals. The system is currently setup to collect singles and doubles neutron counts every two seconds. This produces a continuous time series of neutron measurements as shown in Figure 1.



FIG. 1. Sample data showing singles neutron counts as a function of time for a day with no activity (top), a course setup day (middle), and a course day (bottom). The series were smoothed with a Savitzky-Golay filter to reduce noise.

3.3. Surveillance System

Containment and surveillance (C&S) is highly complementary to nuclear material accountancy tools such as inventory records and NDA systems. It plays an important role in maintaining continuity of knowledge on nuclear materials between on-site inspections. Recognizing that recent advances in commercial sensing technologies may have implications for safeguards, as part of this work, we are exploring the possibility of using very low power, inexpensive sensors as an alternative to technologies like the IAEA's Next Generation Surveillance System (NGSS) [13].

Leveraging capabilities developed as part of LANL's stormwater runoff monitoring project [14], the team designed and built a set of custom sensor nodes for monitoring the general activity level in common areas at TA-66 such as entryways and the building's large conference room. Each node consists of a RoboDyn microphone, Wingoneer light sensor, SenseHat with an inertial measurement unit (IMU) and environmental sensor for temperature and humidity, and a Raspberry Pi computer. Photographs of the individual components

and an assembled unit are shown in Figure 2. The total cost of components for each sensor node is less than \$100, and we have installed seven of them at TA-66.



FIG. 2. Photographs of the individual components of the surveillance sensors and an assembled sensor node.

The sensor nodes can be deployed as standalone systems where data is collected manually from each node or as a wireless sensor network (WSN). Using wireless transmission, each node can send data directly to a wireless router that, in turn, sends data to the data collection computer where it is synchronized with a network time protocol (NTP) timestamp. Standard wireless security protocols such as Wired Equivalent Privacy (WEP) and Wi-Fi Protected Access (WPA or WPA2) allow for data encryption over the WSN.

4. DATA MANAGEMENT

The large volume and variety of data collected for this work demands a robust data management approach. Our goal was to design a data architecture that is easy to maintain and facilitates automation. The first step was to setup a centralized and secure collaborative workspace for the research team through a Git repository on one of LANL's internal networks. Git is a version control system for tracking changes to documents or source code and coordinating work among multiple contributors to a project. The Git repository contains a number of Python scripts and Jupyter notebooks generated for the project to automatically clean the various types of raw data files (i.e., parsing data files and removal of corrupt entries or fields that are irrelevant for analysis).

The cleaned data is uploaded to a NoSQL MongoDB database [15] using a representational state transfer application programming interface (REST API). MongoDB does not come with a RESTful API, so that functionality is provided by the Python Eve Framework [16]. The database is accessible on LANL's Darwin supercomputer for data processing and analysis, and it resides in a team share folder that requires user authentication for access control.

5. PRELIMINARY RESULTS

The dataset we used for preliminary analysis consists of fifteen weeks of concurrent neutron detector and Source Tracker data. Five courses were held at TA-66 during this period. While most of the courses are similar, there are some notable differences in duration and content between courses. The extreme example includes one course in our test dataset on statistics that did not include any nuclear material measurements. These variations provide realistic conditions for quantifying the sensitivity of approaches to perturbations that are difficult to replicate in simulations or synthetic data.

For the purpose of this analysis, we consider three classes of activities for a given day: (1) course days, (2) course setup days, and (3) no activity days (i.e., weekends and all other workdays). In order to build basic classification models, we reduced the dimensionality of the raw data. To do this, we first extracted features from

KAREN A. MILLER et al.

the data. This entails deriving a set of characteristics from a time-series distribution to represent the whole sequence. In the case of the neutron data shown in Figure 1, we extracted basic descriptive features such as the mean, standard deviation, skewness, minimum, and maximum. We also considered features that describe the spiked and highly-concentrated nature of the signals such as entropy, the high-frequency Fourier power, the fourth moment, and the number of data points above a relative threshold. Finally, we added features we found to be descriptive of the unique characteristics of the classes, like the dip in signal during the typical lunch break. Each of these features compresses a full day distribution into a single value that can be used for classification. Because the Source Tracker data is much more sparse than the neutron data, the features represent values contained in the logs aggregated over each day (e.g., total number of transactions involving plutonium sources).

Dimensionality reduction techniques were then used to visualize relationships in the feature space to illustrate how useful the features are in separating out the three classes. Canonical correlation analysis (CCA) is a technique for dimensionality reduction that resulted in better discrimination between classes than the more common principal component analysis (PCA) technique. Instead of maximizing internal variance within a single dataset, CCA maximizes the correlation between two different datasets, which maximizes the separability by correlating the features with a one hot encoding of the labels. The results are shown in Figure 3. The plot shows good distinction between the days with no activity versus course and setup days. The course days that appear in the cluster of no activity days represent the statistics course where no nuclear material was used, which aligns with expectations for these two datasets. The separability between setup and course days is less distinct but shows promise for preliminary results.



CCA Cross Decomposition with Labels

FIG. 3. Canonical correlation analysis (CCA) of the neutron detector and Source Tracker datasets.

Once the feature space was built up, the next step was to perform proof-of-concept testing with a variety of classification models to explore their utility to the facility monitoring problem. In machine learning, classification techniques are used to learn a function that maps input parameters (e.g., neutron and Source Tracker data) onto known outputs (e.g., course, setup, or no activity days) so that as new data is collected, we can accurately predict the state of operations. We explored a range of models including support vector machine (SVM), decision tree classifier, random forest, and nearest neighbors in the CCA subspace. Confusion matrices were calculated to show the disaggregated accuracies found by cross validation of the models for each

IAEA-CN-267

classification scenario. These are shown in Figure 4. Each confusion matrix shows the prediction probability for each prediction versus true label pair. The results further demonstrate how much easier it is to distinguish between activity versus no activity days than course versus setup days using only the neutron and Source Tracker data.



FIG. 4. Confusion matrices for the four classification models tested using the feature space for the neutron detector and Source Tracker datasets.

The feature space we calculated is essentially independent of time. We used a tool called Matrix Profile [17] to explore some of the temporal information contained in the data. For the time series analysis, we only considered data from the neutron detector. Matrix Profile uses a technique called all pairs similarity to identify recurrent temporal patterns between days. These patterns are used to build representative signatures called the Matrix Profile Index for each category in the classifier. Figure 5 shows the Matrix Profile Index for our no activity day, course day, and setup day categories. These signatures show distinct peaks for the separate classes and show promise of more time-dependent models in future work.



FIG. 5. Matrix Profile Index aggregate category distributions for no activity days (top), course days (middle), and setup days (bottom).

6. CONCLUSIONS

Characterizing activity patterns through the integration of disparate data streams has the potential to provide the IAEA with enhanced capability to detect misuse and diversion within safeguarded facilities. This approach would represent a paradigm shift from a focus on the known to discovery of the unknown. The primary outcome of this work was development of the experimental infrastructure necessary to support research in disparate data integration for advanced facility monitoring. Encouraging results from exploratory data analysis using supervised machine learning methods and time series analysis showcased the feasibility of using classification models with safeguards-relevant data streams. Preliminary results provide confidence that accurate models for activity characterization under realistic operational conditions can be developed. Building on successes and lessons learned in the first year of the project, follow-on studies will focus on continuing to improve the data collection and management architecture in parallel with improvements to predictive models. The aim is to create a feedback loop between data scientists and subject-matter experts in relevant safeguards technologies to design concepts and approaches that are optimized for the application.

In the short term, integration of data from the surveillance sensors installed at TA-66 into the data analysis is expected to improve model accuracy significantly. Based on the preliminary results, further exploration of how to leverage temporal signatures also shows great promise. Unsupervised machine learning techniques and anomaly detection are also natural offshoots of this work to enhance detection of undeclared activities.

The longer-term vision for this work is to provide the IAEA with decision support tools to increase productivity and to help deter would-be proliferators through the threat of early detection. The aim of our integrated approach is to allow the IAEA to do more with the data they are already collecting by leveraging advances in data science as a force multiplier for complex problems.

ACKNOWLEDGEMENTS

Funding for this work was provided by the Laboratory-Directed Research & Development (LDRD) Program at Los Alamos National Laboratory.

REFERENCES

- [1] IAEA, IAEA Safeguards Key Facts (2018), https://www.iaea.org/sites/default/files/16/06/sg_keyfacts.pdf
- [2] CLAUDE, F., COYNE, J., CALLE D., GAGNE, D., SIMMONS, A., WILSON, B., IANCU, I., "Data fusion at scale: strengthening safeguards conclusions through applied analytics," Proc. INMM Annual Meeting, Indian Wells, California, USA (2017).
- [3] MOSAIC The Modernization of Safeguards Information Technology: Completing the Picture, report, IAEA, Vienna, Austria, 2016.
- [4] MENLOVE, H.O., HOWELL, J.A., RODRIGUEZ, C.A., ECCLESTON, G.W., BEDDINGFIELD, D., SMITH, J.E., BAUMGART, C.W., "Integration of video and radiation analysis data," Proc. INMM Annual Meeting, Naples, Florida, USA (1994).
- [5] LOKHOV, A.Y., LEMONS, N., MCANDREW, T.C., HAGBERG, A., BACKHAUS, S., "Detection of cyberphysical faults and intrusions from physical correlations," ACM SIGKDD Workshop on ODD, San Francisco, California, USA (2016).
- [6] CASLETON, E., OSTHUS, D., VAN BUREN, K., Imputation for multisource data with comparison and assessment techniques, Appl. Stochastic Models Bus. Ind. 34 1 (2018) 44-60.
- [7] NPR Morning Edition, How do you find plutonium? Go to nuclear inspector school (2015), https://www.npr.org/2015/10/19/449031762/how-do-you-find-plutonium-go-to-nuclear-inspector-school
- [8] IAEA, Safeguards Implementation in 2017 (2018), https://www.iaea.org/sites/default/files/18/04/sgimplementation-2017.pdf
- [9] CLAASSEN, J.P., Source Tracker User's Manual, LA-UR-12-00576, report, Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 2012.
- [10] MENLOVE, H.O., Manual for the Portable Handheld Neutron Counter (PHNC) for Neutron Survey and the Measurement of Plutonium Samples, LA-14257-M, report, Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 2015.
- [11] ABHOLD, M.E., SWEET, M.R., BOURETT, S.C., HANSEN, W.J., AUDIA, J., KRICK, M.S., SALAZAR, S.D., STEWART, J.E., "Advanced multiplicity shift register with networking capability," Proc. INMM Annual Meeting, Phoenix, Arizona, USA (1999).
- [12] PELOWITZ, D., "An international standardized utility to collect data, MIC," Proc. INMM Annual Meeting, Orlando, Florida, USA (2004).
- [13] HADFI, G., JOHN, M., LIGOURI, C., MOESLINGER, M., MURRAY, J., ROCCHI, S., "Next Generation Surveillance System (NGSS): field implementation & associated developments," Proc. IAEA Safeguards Symposium, Vienna, Austria (2014).
- [14] FRIGO, J., AYERS, H., KULATHUMANI, V., HINZEY, S., SEVANTO, S., PRIOCOU M., YANG, X., MCCABE, K., SAARI, A., SENTZ, K., "Novel WSN hardware for long range low power monitoring," Proc. IEEE DCoSS, Ottawa, Ontario, Canada (2017).
- [15] MongoDB (2018), https://www.mongodb.com/
- [16] Python Eve Framework, Eve. The Simple Way to REST (2018), http://python-eve.org/
- [17] YEH, C.M., ZHU, Y., ULANOVA, L., BEGUM, N., DING, Y., DAU, H.A., SILVA, D.F., MUEEN, A., KEOGH, E., "Matrix Profile I: all pairs similarity joins for time series: a unifying view that includes motifs, discords, and shapelets," Proc. IEEE ICDM, Barcelona, Spain (2016).