SCIENCE & GLOBAL SECURITY
2020, VOL. 28, NO. 1, 28–47
https://doi.org/10.1080/08929882.2020.1716461

Data Science in Support of Radiation Detection for Border Monitoring: An Exploratory Study

Christopher Hobbs[a], Peter McBurney[b], and Dominic Oliver[b]

[a]Centre for Science & Security Studies, Department of War Studies, King's College London, Strand, London, UK; [b]Department of Informatics, King's College London, Strand, London, UK

ABSTRACT
Radiation detection technology is widely deployed to identify undeclared nuclear or radiological materials in transit. However, in certain environments the effective use of radiation detection systems is complicated by the presence of significant quantities of naturally occurring radioactive materials that trigger nuisance alarms which divert attention from valid investigations. The frequency of nuisance alarms sometimes results in the raising of alarming thresholds, reducing the likelihood that systems will detect the low levels of radioactivity produced by key threat materials such as shielded highly enriched uranium. This paper explores the potential of using data science techniques, such as dynamic time warping and agglomerative hierarchical clustering, to provide new insights into the cause of alarms within the maritime shipping environment. These methods are used to analyze the spatial radiation profiles generated by shipments of naturally occurring radioactive materials as they are passed through radiation portal monitors. Applied to a real-life dataset of alarming occupancies, the application of these techniques is shown to preferentially group and identify similar commodities. With further testing and development, the data-driven approach to alarm assessment presented in this paper could be used to characterize shipments of naturally occurring radioactive materials at the primary scanning stage, significantly reducing time spent resolving nuisance alarms.

ARTICLE HISTORY
Received 29 May 2019
Accepted 16 November 2019

Introduction

International concerns over the clandestine movement of sensitive nuclear and radiological materials across borders have existed for more than 70 years.[1] These have driven the development and deployment of detection systems that identify and assist in the interception of illicit radioactive materials. In the post-9/11 period, there has been a dramatic increase and internationalization of efforts to detect and intercept radioactive materials, precipitated by a perceived increase in the potential for acts of nuclear

CONTACT Christopher Hobbs christopher.hobbs@kcl.ac.uk Centre for Science & Security Studies, Department of War Studies, King's College London, Strand, London, WC2R 2LS, UK

terrorism by non-state actors.[2] Over the past two decades billions of dollars have been spent on the deployment of radiation detectors to counter these perceived threats.[3] U.S.-led programs, such as the Container Security Initiative (CSI) and the Megaports Initiative, have supported the installation and operation of radiation detection systems at seaports, airports, road, and rail border crossings in more than 50 countries.[4]

However, despite significant investment in this area, technological challenges remain in the sensing and characterization of key threat materials, such as highly enriched uranium (HEU), as it is believed that adversaries are likely to shield and ship HEU and plutonium in small quantities thereby reducing their radioactive emissions.[5] Detection challenges are arguably most acute in the maritime supply chain where large scale shipments of naturally occurring radioactive materials (NORM) and commercial radioactive goods regularly trigger nuisance alarms.[6] Examples include ceramics, fertilizers, granite, radiopharmaceuticals and industrial sources.[7] The time and attention devoted to investigating these nuisance alarms diverts attention from other activities and slows material flows and trade. High levels of nuisance alarms can also lead to increased alarming thresholds, reducing the ability of systems to detect materials with low radiation levels such as shielded HEU.[8]

A 2019 U.S. Department of Energy (DOE) study estimated that 1.66% of all containers at seaports triggered an alarm because of NORM.[9] This represents a significant number of containers given the high volume of global seaborne trade and the U.S. objective of scanning 100% of imported containers.[10] For example, the throughput of the world's busiest port, Shanghai International in China, surpassed 40 million twenty-foot equivalents units (TEU) in 2017, equivalent to more than 100,000 containers a day.[11] At the DOE rate, that would amount to approximately 1800 nuisance alarms daily. At other large ports, NORM could produce hundreds of nuisance alarms in a single day.

This paper explores the potential of using data science techniques such as dynamic time warping (DTW) and agglomerative hierarchical clustering to reduce maritime nuisance alarms in the initial scanning stage. Tens of millions of Radiation Portal Monitor (RPM) occupancies are created and recorded every year, including hundreds of thousands of alarming records.[12] However, systematic analysis of this data has been limited to an inferential analysis on the stream of commerce and the pace of site operations.[13] The application of data science tools to this information, in support of the real-time assessment of alarms, has not been studied.

An introduction to DTW and its utility in analyzing time-series data is provided and applied to a dataset of RPM spatial radiation profiles for alarming occupancies triggered by containerized shipments of NORM.[14]

DTW is performed on spatial radiation profiles generated by alarming containers as they pass through an RPM. This accounts for variations in container speed so that alarming containers can be meaningfully compared and their degree of similarity calculated. The relative dissimilarity of these warped profiles is then assessed using agglomerative hierarchical clustering, creating groups of similar objects. These are illustrated as dendrograms and show a preferential clustering of common NORM commodities. The key results of this study are then summarized, with suggestions for future work.

Radiation detection at maritime facilities – current and proposed systems

At maritime facilities containerized cargo may be inspected for the presence of nuclear and radiological materials when passing through a point of entry or exit, or for transshipments during loading and unloading.[15] Typically, multi-stage detection protocols are employed, with alarming containers undergoing a series of increasingly intrusive inspections until the source of the alarm is determined. A key component of radiation detection systems are RPMs. These are passive, non-intrusive detectors which are widely employed in large-scale primary scanning at maritime and other facilities. In a primary inspection, a container passes slowly through an RPM, which detects gamma emissions and neutron radiation over the course of its occupancy.

The presence of neutrons automatically triggers an RPM alarm and follow-up investigation, although such events are relatively rare in most environments. In contrast there are many potential sources of gamma radiation, which necessitates a gamma count threshold.[16] If the total gamma radiation received is above the count threshold an alarm is triggered. The spatial gamma profile of a container is also captured by the RPM and visually examined by officials for anomalies. For example, a sharp spike in the gamma profile might indicate a strong radioactive point source.[17] RPM gamma and neutron information is assessed in conjunction with the shipment manifest, which lists the declared content of the container and other information.

The spectral resolution of existing RPMs is limited and, following an alarm, a secondary inspection may be performed to characterize the radiation source. In a secondary inspection a container is typically moved to a secure area before a manual external inspection is conducted with passive handheld radioisotope identification devices (RIIDs).[18] Gamma and X-ray images of the container may also be taken during a secondary inspection to identify dense materials that might be shielding radioactive emissions. If a secondary inspection is inconclusive then a tertiary inspection may be launched, whereby the container is opened and unpacked. Disruptions to cargo flows increase significantly when secondary and tertiary inspections are triggered. While a primary inspection can be conducted in less than a minute, a secondary inspection might take tens of minutes, and a tertiary inspection could take several hours.[19]

Challenges in assessing alarms and variations in practice

It is intrinsically difficult to identify and characterize low radioactive threat materials at locations where commercial radioactive goods are transited and where significant quantities of NORM are present. Although the presence of radiation can be detected during the short time it takes for container to pass through an RPM, radioisotope identification, an important step in resolving nuisance and other alarms, currently requires at least a secondary inspection. Secondary inspections can also be inconclusive as RIIDs may fail to detect the relatively low activity of NORM containing commodities.[20] These challenges are compounded by fluctuating local factors relating to the ambient environmental conditions and detector setup, which can affect the radiation received, further complicating alarm assessment. Local factors include the weather, the separation of the RPM detectors, their volume and "cross-talk" from nearby containers.[21]

Practical deployment of detection systems is shaped by the priorities of national detection programs, the environment, available resources, and commercial considerations. These priorities are frequently in conflict and a difficult compromise between them will often have to be reached. As an example, RPM alarming thresholds and the triggering of secondary and tertiary inspections will vary based on how the operator balances the need to maintain throughput rates with their security priorities.[22] At some facilities security considerations dominate, while at others the management of nuisance alarms largely dictate how detection systems are configured and operated.[23] This can result in considerable variation in operations and approach from facility to facility. For example, a recent study by one of the authors revealed that RPM nuisance alarm rates ranged from more than 10% to less than 1% across facilities. Even greater variation was seen in secondary inspection rates following initial RPM alarms, which were triggered at some facilities in over 90% of cases, and at others in less than 1% of cases.[24]

Efforts to improve the operation of detection systems for border monitoring

There has been considerable research and development aimed at improving the speed, detection and characterization of radioactive materials at maritime (and other) facilities. A significant proportion of this work has been focused on the initial scanning stage, as the ability to accurately identify the cause of alarms at this stage would reduce the need for secondary and tertiary inspections. For example, tens of millions of dollars have been invested on spectroscopic portal monitors (SPM) which offer the promise of both high-sensitivity and spectral resolution.[25] In theory, SPMs could provide both simultaneous radiation detection and nuclide identification during initial scanning.[26] However, SPMs have struggled to meet operational requirements in certain environments, in particular at maritime facilities, and their deployment to date has been limited.[27]

Other efforts have sought to improve the ability of existing systems. For example, the International Atomic Energy Agency's (IAEA) mobile application Tool for Radiation Alarm and Commodity Evaluation (TRACE) provides a readily accessible database of NORM with information on the specific radionuclides found in each commodity.[28] This can help officials unfamiliar with the radiative properties of NORM to interpret alarms from RIIDs during secondary inspections by facilitating manual spectral matching to commodity information on the shipment manifest. Similarly, the approach presented here does not require new detectors or associated hardware.

A new approach – characterizing nuisance alarms using spatial gamma profiles

The approach proposed here could potentially improve detection efficiency and free up time to focus on difficult to characterize alarms, including potentially those caused by threat materials. It could also enable lower alarm thresholds by significantly increasing the capacity of maritime facilities to respond to primary RPM alarms. This would in turn increase the likelihood of detecting key threat materials such as shielded HEU.

The speed at which containers are driven through RPMs is variable, resulting in spatial gamma profiles of varying lengths and shapes, which are challenging to interpret. Widely varying speeds of containers moving through RPMs are common, as illustrated by analysis presented here, where occupancy periods (time spent inside the RPM) ranged from 5.2 to 33.8 seconds.[29] The impact of this is illustrated in Figure 1, which for a single container of NORM passed through a RPM multiple times, shows how variations in speed can significantly distort the gamma radiation profile.[30] This provides further evidence that RPM spatial gamma profile information is currently useful only in identifying clear anomalies, such as a strong unshielded radioactive source.

The data science technique of DTW is a useful method of accounting for the speed distortion of spatial gamma profiles so that they can be mathematically compared and a degree of dissimilarly between them calculated. It is hypothesized that the degree of dissimilarity will be small between the warped spatial gamma profiles from shipments of similar NORM commodities relative to shipments of other NORM commodities. For example, shipments of ceramics would have similar warped spatial gamma profiles, relative to say shipments of fertilizer. If true, this approach could be used to differentiate between NORM shipments, grouping similar alarming occupancies and linking them to specific commodity types. This is explored using hierarchal agglomerative clustering, a common technique for grouping objects based on their similarity. If validated, once a reference database of warped spatial gamma profiles for common NORM commodities is developed, this approach could be used to categorize nuisance alarms at the primary scanning stage.

Times series and the analysis of real-world data through dynamic time warping

RPM spatial profiles can be considered as time series, a sequence of gamma counts indexed in time as a container passes through a detector. Time series are most simply compared by calculating the Euclidean distance, the shortest path between points on two series that occur at the same time. This can be a useful metric for comparing the similarity between series if they are in phase, i.e. with similar events occurring at the same moment. However, as previously discussed, variations in container speed through an RPM mean that corresponding peaks and troughs will occur at different times, negating the use of this metric. For example, if the Euclidean distance were used to compare the RPM spatial gamma profiles in Figure 1, it would erroneously conclude that they were produced by different, as opposed to the same shipment.

Comparing out of phase time series data is a common exercise. It is used successfully in automatic speech recognition where there can be considerable variation in speed, pace, tempo, rhythm, and pronunciation. Similar challenges can also be found in signature matching, music and signal processing, and even in the rehabilitation of stroke victims.[31]

Dynamic time warping

To compare time series data with temporal variations, DTW calculates the best non-linear alignment between two series, matching similar features even if they are out phase.[32] The difference between Euclidean matching and non-linear DTW matching is illustrated visually in Figure 2 for two example timeseries of differing lengths. Here the black lines show how corresponding points are matched and subsequently compared using each

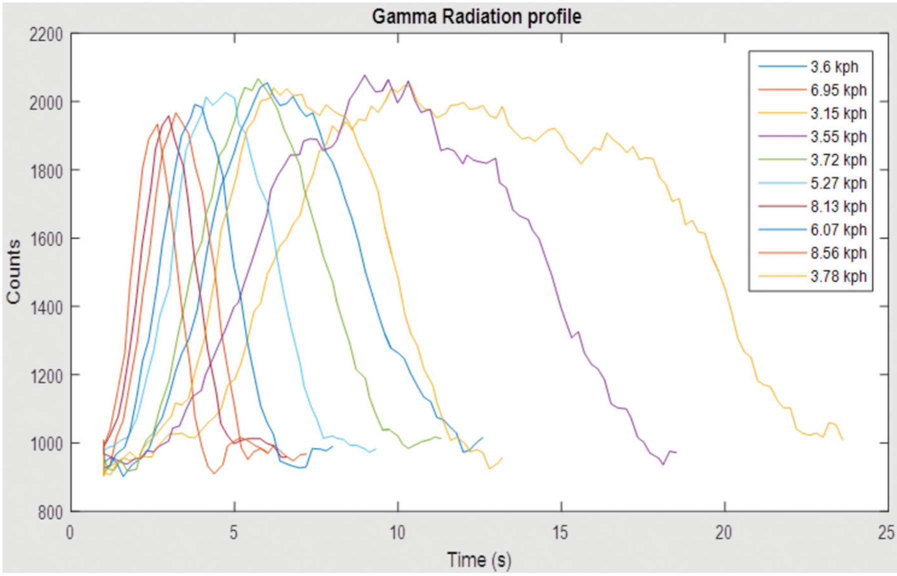

Figure 1. Gamma radiation profiles for an alarming container passed through the same RPM at different speeds (Reproduced from IAEA).

technique. In this example, DTW is better at matching the corresponding peaks and troughs.

DTW functions by modifying or warping time series data along an optimal path which considers all local compressions, shifting, and minimizing the cumulative distance between aligned points. Consider two time series $T_A = (t_{A1} \ldots t_{Ai} \ldots t_{AN})$ and $T_B = (t_{B1} \ldots t_{Bj} \ldots t_{BM})$, of different lengths N and M, with points occurring at different times. These can be used to construct a matrix, $C \in \mathbb{R}^{N\ x\ M}$, containing the distance between every point in series X with respect to every point in series Y.

$$C \in \mathbb{R}^{NxM} : c_{i,j} = \|t_{Ai} - t_{Bi}\|, i \in [1:N], j \in [1:M]$$

There are several ways to calculate the distance between points across different series. This study used the Manhattan or city-block metric, due to its ability to account for differences between similar series.[33] For example, small differences between RPM spatial gamma profiles could potentially be significant for low activity shielded HEU purposefully hidden inside a shipment of NORM. Rather than taking the shortest Euclidean distance between two points, the Manhattan metric measures the distance along axes at right angles. For example, in a plane with point 1 at (x_1, y_1) and point 2 at (x_2, y_2), the Manhattan distance is $|x_1 - x_2| + |y_1 - y_2|$.

The correspondence between the points in different series is then established through a warping path $\Phi = (\phi_{t,}\ _{\psi t})$ where $t = 1, \ldots, T$, under the following constraints, namely:

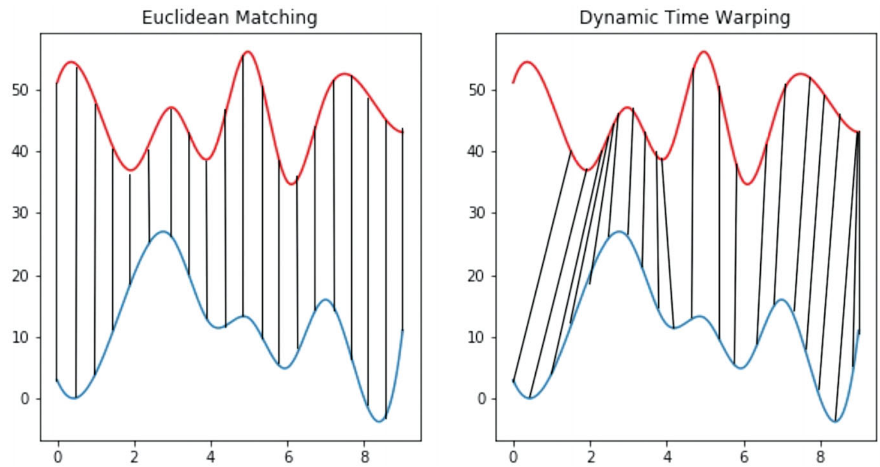

Figure 2. Comparison of Euclidean matching (left) and non-linear dynamic time warping matching (right) for two example time series.

1. *Start-point constraint.* The warping curve is anchored at the origin: $\phi_1 = \psi_1 = 1$;
2. *End-point constraint.* A global alignment is required, and the mapping covers both time series completely: $\phi_T = N$, $\psi_T = M$;
3. *Monotonicity restriction.* No "time loops" are allowed in the mapping: $\phi_t \geq \phi_{t-1}$ and $\psi_t \geq \psi_{t-1}$.
4. *Step-size condition* limits the warping path from long jumps (shifts in time) while aligning sequences.[34]

The start- and end-point constraints guarantee that the alignment does not partially consider one of the series. The monotonicity restriction and step-size condition ensure that the alignment path does not go backwards or forwards in time, which could omit potentially important features.

The optimal warping path, which minimizes the distance between points in the two series subject to the above constraints, is then calculated:

$$\boldsymbol{\Phi} = (\phi_t, \psi_t) = \arg\ \min_{\phi_t, \psi_t} \sum_{t=1}^{T} \frac{d(x_{\phi t},\ y_{\psi t}) m_{t,\Phi}}{M_\Phi}$$

Here $x_{\phi t}$ *and* $y_{\psi t}$ are the elements of the warped input time series, d is the local distance function, $m_{t,\Phi}$ is a local weighting coefficient, and $M_\Phi = \sum m_{t,\Phi}$.

A measure of dissimilarity between time series can be calculated by summing the Manhattan distance between the points, which have been matched by DTW. Represented by the following cost-function:

$$c_{\Phi}\left(T_A,\ T_B\right) = \sum\nolimits_{t=1}^{T} c(xAi,\ yBi)$$

To avoid unfairly favoring short series owing to the cumulative sum element of this cost-function, this is normalized to give the average per-step dissimilarity, known as the pairwise distance.

RPM alarming dataset – overview

The data used in this analysis was provided by the IAEA under a Coordinated Research Project (CRP).[35] It includes scanning records from multiple RPM lanes at maritime facilities in three countries in the form of "daily files."[36] The files contain conveyance information on all alarming and non-alarming containers that passed through a RPM over 24 hours. The officials at each site also provided additional information relating to the shipment manifests for alarming occupancies. The data included a description of the declared container contents and the corresponding six-digit Harmonized System (HS) code which was critical for this study. HS codes are an international nomenclature used to classify traded goods for customs purposes.[37] At the international level these are six-digit product codes where specific commodities are classified within a hierarchal structure composed of chapters, headings, and sub-headings. For example, all ceramics are within Chapter 69, within which bathroom ceramics come under heading 10, with all non-porcelain, non-china ceramics under sub-heading 90. Consequently, a shipment of non-porcelain, non-china bathroom ceramics will be given the HS code 691090.

A total of 720 usable alarming records triggered by NORM were identified in the dataset. These were separated into individual lanes before comparative calculations were performed to minimize the influence of local factors on the gamma radiation received. Accounting for variations in local factors is a key focus area of the IAEA CRP. However, the IAEA CRP research is ongoing so the authors decided to isolate the local factors and focus on the challenges associated with variable speed. Splitting the dataset into specific RPM lanes and analyzing the alarming occupancies in each lane reduced the size of the dataset, across which comparisons could be made. Nevertheless, in the most populous lane, there were still over 150 records. This was deemed promising enough to explore the potential utility of DTW and agglomerative hierarchical clustering in this proof of concept study.

Preliminary analysis

We compared and contrasted RPM spatial profiles from a random selection of alarming occupancies within specific lanes to evaluate the utility of

DTW. The DTW pairwise distance for these profiles was used as measure of their dissimilarity.[38] Results obtained from applying DTW to these small trial datasets was promising, with the relative pairwise distance proving a useful metric for distinguishing between commodities. An example from one of these initial tests, comparing five RPM alarming profiles is shown in Figure 3 and Table 1. According to the supplied HS-codes, three of these five alarms were triggered by shipments of fertilizer (HS-code 310490), shown in the red, green, and gold profiles. The other two alarms were triggered by ceramics (HS-code 691090) in blue, and clay (HS-code 240840) in purple. Examining the RPM spatial gamma profiles in Figure 3, the red and gold profiles are almost identical, while the green profile, despite having a similar shape, has a shorter time series as a result of being moved through the RPM at higher speed.

Simple Euclidean matching shows the green profile to be considerably different from the red and gold profiles, as their peaks and troughs occur at different points in time. DTW takes this variation into account

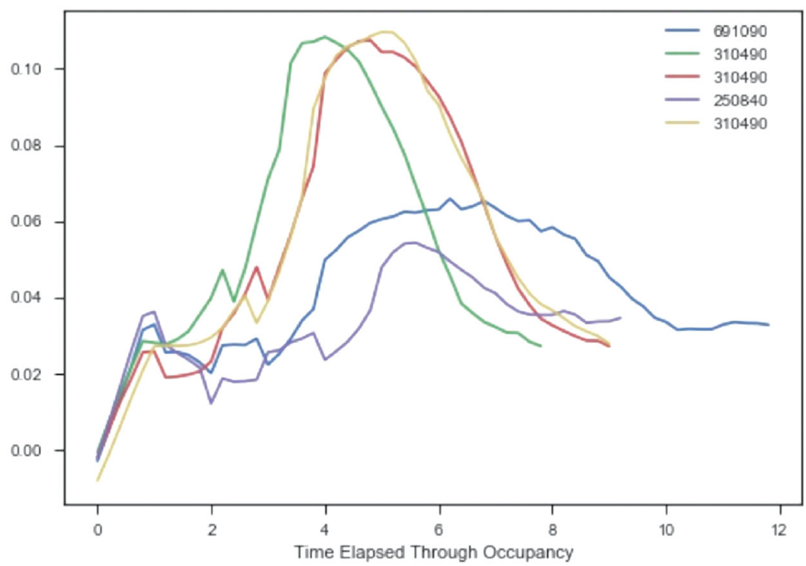

Figure 3. RPM spatial profiles for five alarming occupancies with corresponding HS codes.

Table 1. Pairwise distance dissimilarity matrix for the five RPM profiles illustrated in Figure 3.

	1 (Blue) (691090)	2 (Green) (310490)	3 (Red) (310490)	4 (Purple) (250840)	5 (Gold) (310490)
1 (Blue) (691090)	0	1.33	1.28	0.93	1.34
2 (Green) (310490)	1.33	0	0.53	1.50	0.55
3 (Red) (310490)	1.28	0.53	0	1.23	0.32
4 (Purple) (250840)	0.93	1.50	1.23	0	1.36
5 (Gold) (310490)	1.34	0.55	0.32	1.36	0

showing there is a comparable level of dissimilarity (a pairwise distance of 0.32, 0.53, and 0.55) between the three profiles. This is shown for all five profiles in the pairwise distance matrix in Table 1, with the level of dissimilarity between the shipments of fertilizer considerably less than their dissimilarity with respect to the shipments of ceramics and clay. It is also considerably less dissimilar than the pairwise distance (0.93) between the shipments of ceramics and clay. Consequently, for these five records the pairwise distance can be used to identify that three of shipments are of the same commodity while the other two shipments contain different commodities.

Clustering time warped RPM profiles

Following initial testing, DTW was then applied to larger datasets containing all the alarming records available from specific lanes. In assessing the ability of DTW to distinguish between similar and different commodities, the unsupervised data mining technique of clustering was employed. Clustering does not utilize any preexisting information or labeling of objects, but instead relies on establishing an effective criterion for defining "similar" data. Clustering attempts to "identify structure in an un-labelled data set by objectively organizing data into homogeneous groups where the within-group-object similarity is minimized and the between-group-object dissimilarity is maximized."[39] As was the case for the preliminary tests, the pairwise distance was used as the measure of dissimilarity between warped RPM spatial gamma profiles.

There are several different clustering options.[40] Here agglomerative hierarchical clustering was selected due to its ability to split a dataset into clusters naturally without having to specify their number. This is important in the context of this study as the number of different commodities passing through an RPM is unknown in advance. Specifying the number of commodities beforehand might have caused the datasets to split in an artificial manner. Agglomerative hierarchal clustering initially treats each record as its own one-record (singleton) cluster then utilizes a bottom-up stepwise iterative process to pair similar clusters until they are members of one cluster. This produces a tree-based representation of the records known as a dendrogram.

The process of agglomerative hierarchal clustering is illustrated in Figure 4 for an example dataset of six records labeled 0 through 5. The order of the six records from left to right is not sequential according to their labels but arranged to show the incremental clustering approach. Here each record starts as its own singleton cluster before being paired and merged into larger clusters as the measure of dissimilarity is increased, moving up the vertical axis of the dendrogram. In the first step, records 2

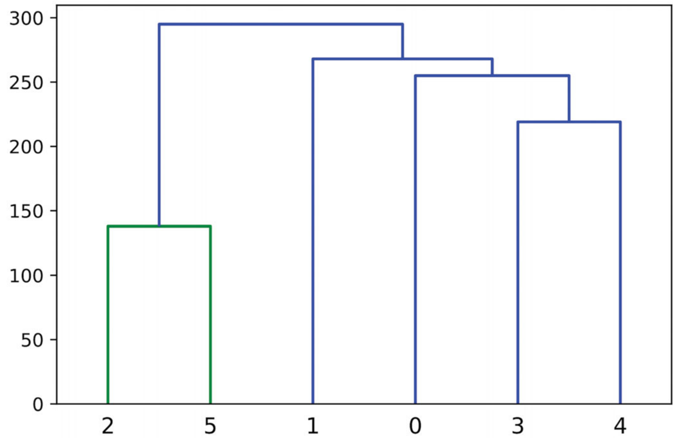

Figure 4. Dendrogram illustrating the process of agglomerative hierarchical clustering for a small dataset of six records, labeled 0 to 5. The vertical axis is the measure of dissimilarity. The branch linking records 2 and 5 is green to illustrate that they are like one another but considerably different to any other record.

and 5 are paired, in step two records 3 and 4 are paired, in step three record 0 is merged with records 3 and 4, in step four record 1 is merged with records 3, 4 and 0, and finally in step five all records are merged into a single cluster. This demonstrates that records 2 and 5 are like one another but considerably different than any other record in the dataset, clustering only with the other records at high levels of dissimilarity.

In determining how clusters are merged at each step there are several linkage methods that can be used to weigh their relative difference based on their constituent records.[41] The utility of the following commonly used methods was explored:

- *Single link*. Mergers are decided by the minimum pairwise distance value between any of the records within two clusters (i.e. the similarity of their most similar members).
- *Average link*. Mergers are decided by the average pairwise distance value between all the individual records within each of the two clusters.
- *Complete link*. Mergers are decided by the maximum pairwise distance value between any of the records within two clusters (i.e. the similarity of their most dissimilar members).

The results from the most populous RPM lane containing 153 alarming occupancies are illustrated as dendrograms in Figures 5–7 for single, average, and complete link cluster merging. In each figure the y-axis provides a measure of the dissimilarity between commodities, with individual alarming occupancies represented under the x-axis. The three most common commodities extracted from the HS codes on the shipment manifests were glazed ceramics (HS-code 690890), ceramic bathroom fixtures (HS-code 691090), and fertilizer (HS-code 310490). These alarming occupancies are represented in Figures 5–7 as three vertical red dots (glazed ceramics), two vertical yellow dots (ceramic bathroom fixtures) and four vertical blue dots

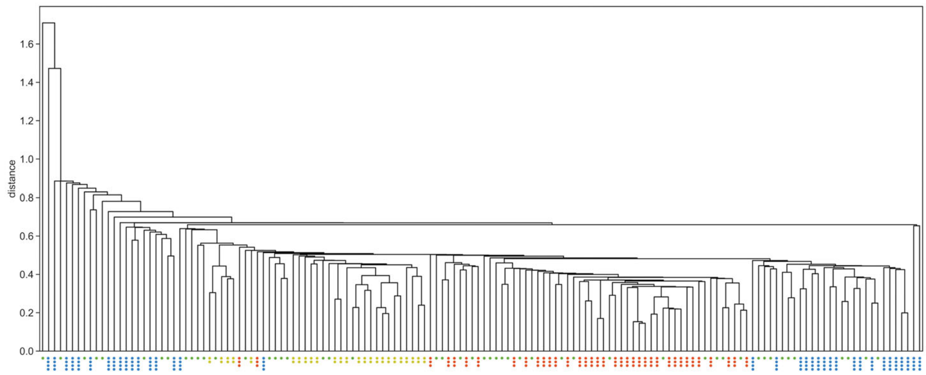

Figure 5. Single link clustering for DTW RPM spatial profiles.

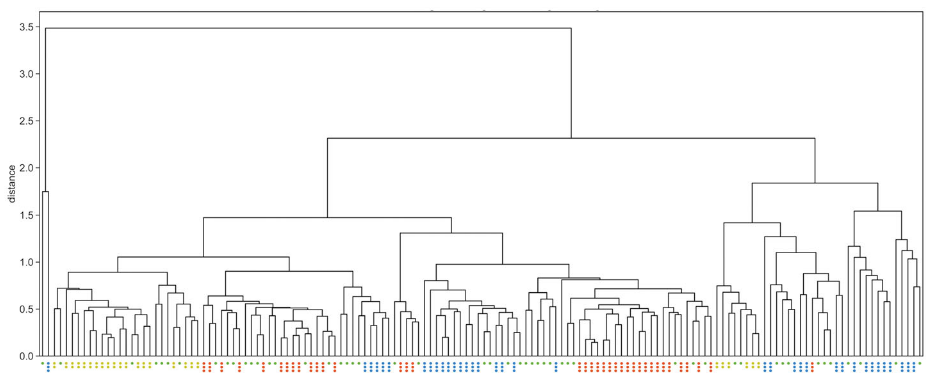

Figure 6. Average link clustering for DTW RPM spatial profiles.

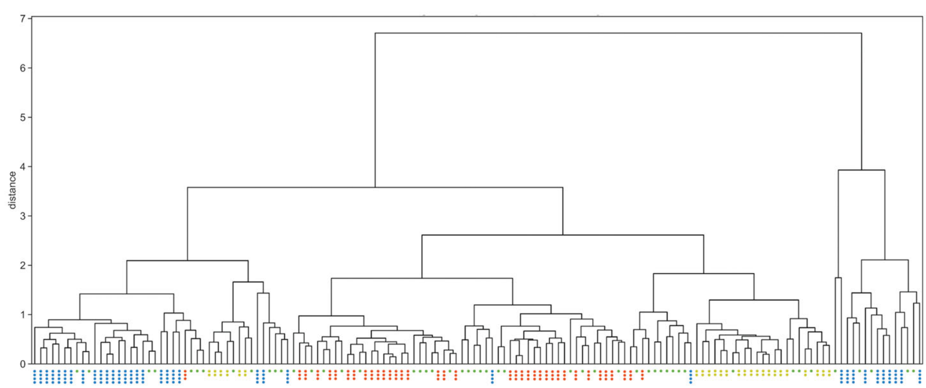

Figure 7. Complete link clustering for DTW RPM spatial profiles.

(fertilizer), with other commodities which were far less prevalent within the dataset represented as a single green dot.

The clusters of red, yellow, and blue records underneath the x-axis illustrate that for each of the three linkage methods tested there is a preferential grouping of similar commodities. Although imperfect, this is to be expected with real world data.

Integrating DTW and clustering into RPM alarm assessment

In field conditions the warped spatial gamma profile for a specific shipment against a database of warped spatial gamma profiles for common NORM commodities would be compared. Agglomerative hierarchical clustering would be used to associate the alarming shipment with a particular commodity type. The strength of this association would be calculated based on level of dissimilarity between the warped spatial gamma profile and other members of the commodity group, i.e., where the alarming shipment fits within a particular cluster. This would then be checked against the shipment manifest for any inconsistencies.

If this method grouped the cargo with a commodity that was different from what was listed on the shipment manifest then a secondary inspection would be triggered. However, if there was a match then a decision would have to made by the on-site official, based on the strength of association, whether the container could be returned to the supply chain or if a secondary inspection was warranted. For example, referring to Figure 5, if the warped spatial profile of a declared shipment of bathroom ceramics fell within the strong cluster of yellow records (that can be observed toward the left of center) then it could be reasonably assumed that the information recorded on shipment manifest was accurate. The container could then continue its journey. However, if the warped spatial profile was to fall within the red or blue clusters toward the left or right of Figure 5, then it would be highly likely that a commodity other than bathroom ceramics was contained within the shipment and a secondary inspection would be initiated.

For this approach to be operationalized, a large and comprehensive database of warped RPM occupancies against which comparisons could be made is required, ideally containing hundreds, if not thousands, of occupancies for common commodity types. Databases could be constructed for individual lanes, or at the facility, national, or international level if it proved possible to account for the influence of local factors. Ideally databases would be populated with records from shipments that had their NORM cargo unambiguously confirmed through secondary or tertiary inspections. As highlighted previously, hundreds of thousands of alarming records that could potentially be drawn upon for this purpose and new confirmed alarming records already exist and could be added over time thus increasing the size of the reference database.[42]

Conclusions and future work

Identifying key nuclear threat materials during border monitoring is a challenging task. This challenge is exacerbated in environments where significant shipments of NORM and commercial radioactive sources trigger nuisance alarms as they move through radiation detection systems. Currently, nuisance alarms are resolved using time-intensive secondary and tertiary inspections. This paper has explored an alternative approach to assessment, involving the characterization of nuisance alarms, radioactive shipments, and possibly threat materials at the primary scanning stage. Rather than suggesting the development of a new type of detector, such as a SPM, it uses data science tools to extract valuable insights from information generated by existing RPM technology. Focusing on the RPM spatial gamma profile, DTW was used to account for differences in container speed so different alarming occupancies could be compared. The application of agglomerative hierarchical clustering to these warped profiles, showed a preferential grouping of similar NORM commodities. Following the development of a reference database, this approach could in theory be used to identify the likely commodity within a container and hence the cause of an alarm without the need for additional inspections. Given the relatively small dataset used in the analysis, it is difficult to precisely estimate the operational gains from this technique. However, the strong groupings of similar commodities suggest that it could resolve a significant fraction of alarms caused by NORM.

It should be emphasized that this study represents exploratory data analysis due to the relatively small dataset across which comparative analyses could be run. Its results should be confirmed using larger datasets, ideally with thousands of alarming NORM occupancies which would also enable a quantitative measure of the quality of clustering for a record to be determined. For example, if a record is a relative outlier within a particular cluster or if it has high levels of similarity to other records. This is an essential step in the development of a practical tool based on this approach, as the quality of clustering would likely determine whether on-site officials will allow a container back into the supply chain or launch a further investigation.

Future research could explore whether the approach outlined in this paper could be used to classify commercial radioactive shipments or possibly threat materials. Given its comparative approach, in theory there is no reason it could not be extended to non-NORM alarming occupancies, as long as an appropriate dataset against which comparisons could be made could be constructed. However, this may pose a significant challenge to identify threat materials given the virtually unlimited number of potential smuggling scenarios involving unknown amounts of material and shielding. Here a potentially fertile line of research would be to explore how the insertion of different quantities of threat materials into a container of NORM would affect the quality of clustering within specific commodity types. It would also be interesting to probe the limits of the approach advocated in this paper with regards to container speed i.e. explore whether there is a maximum speed at which comparisons between gamma spatial profiles utilizing DTW become ineffective. This could be accomplished by passing the same container of NORM through an RPM multiple times, increasing the speed until it was no longer possible for it to be identified as the same shipment.

Notes

1. United States Central Intelligence Agency, "The Clandestine Introduction of Nuclear Weapons into the U.S.," TS190512, July 1970, 4, https://www.cia.gov/library/readingroom/docs/DOC_0001211144.pdf.; Richelson, Jeffrey, *Defusing Armageddon: Inside NEST, America's Secret Nuclear Bomb Squad* (New York: W.W. Norton, 2009), 1.
2. The smuggling of materials across borders is seen by many policy makers and analysts as a key enabling step towards a successful act of nuclear terrorism. For example, see Kouzes, Richard T, "Detecting Illicit Nuclear Materials: The Installation of Radiological Monitoring Equipment in the United States and Overseas Is Helping Thwart Nuclear Terrorism," *American Scientist* 93 (September–October 2005): 422–427, https://www.jstor.org/stable/27858641.
3. Nuclear Threat Initiative, "Radiological and Nuclear Detection Devices," last modified April 2019, http://www.nti.org/analysis/articles/radiological-nuclear-detection-devices/.
4. United States Government Accountability Office, "Combating Nuclear Smuggling: Megaports Initiative Faces Funding and Sustainability Challenges: Report to Congressional Requesters," 2012, http://purl.fdlp.gov/GPO/gpo33533; Oak Ridge National Laboratory, "Nuclear Smuggling Detection and Deterrence FY 2018: Data Analysis Annual Report," ORNL/SPR-2019/1087, 2019, 6, https://info.ornl.gov/sites/publications/Files/Pub123138.pdf.
5. Flynn, Stephen, "Port Security is Still a House of Cards," *Far Eastern Economic Review* 169 (Jan/Feb 2006): 5–11; Gaukler, Gary M., Chenhua Li, Yu Ding, and Sunil S. Chirayath, "Detecting Nuclear Materials Smuggling: Performance Evaluation of Container Inspection Policies," *Risk Analysis* 32 (2011): 531–554, doi:10.1111/j.1539-6924.2011.01696.x.
6. For analysis of nuisance alarms caused by NORM at seaports, relative to airports and vehicle crossings, please see Oakridge National Laboratory, "Nuclear Smuggling Detection and Deterrence FY 2018," 6.
7. World Nuclear Association, "Naturally-occurring Radioactive Materials (NORM)," last modified March 2019, www.world-nuclear.org/information-library/safety-and-security/

radiation-and-health/naturally-occurring-radioactive-materials-norm.aspx; National Research Council, *Evaluating Testing, Costs, and Benefits of Advanced Spectroscopic Portals for Screening Cargo at Ports of Entry: Interim Report (Abbreviated Version)*. (Washington D.C: National Academies Press, 2009), 3.

8. Geelhood, B.D., J.H. Ely, R.R. Hansen, R.T. Kouzes, J.E. Schweppe, and R.A. Warner, "Overview of portal monitoring at border crossings" *Nuclear Science Symposium Conference Record* 1 (2003): 514, https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1352095.
9. Oak Ridge National Laboratory, "Nuclear Smuggling Detection and Deterrence FY 2016: Data Analysis Annual Report," ORNL/SR-2017/91, 2017, https://info.ornl.gov/sites/publications/files/Pub72426.pdf.
10. Caldwell, Stephen L. "The 100 Percent Solution," *The Maritime Executive* (5 December 2016), https://www.maritime-executive.com/magazine/the-100-percent-solution.
11. "Shanghai Port Sets World Box Volume Record," World Maritime News, 3 January 2018, https://worldmaritimenews.com/archives/239541/shanghai-port-sets-new-world-box-volume-record/.
12. Oak Ridge National Laboratory, "Nuclear Smuggling Detection and Deterrence FY 2016."
13. Ibid.
14. A more detailed overview of this dataset is provided in the Appendix.
15. International Atomic Energy Agency. "Nuclear Security Systems and Measures for the Detection of Nuclear and Other Radioactive Material out of Regulatory Control: Implementing Guide," 2013, https://www-pub.iaea.org/MTCD/Publications/PDF/Pub1613_web.pdf.
16. Hamrashdi, Hajir Al, Stephen D. Monk, and David Cheneler, "Passive Gamma-Ray and Neutron Imaging Systems for National Security and Nuclear Non-Proliferation in Controlled and Uncontrolled Detection Areas: Review of Past and Current Status," *Sensors* 19 (2019): 6, https://doi.org/10.3390/s19112638.
17. Ely, James, Richard Kouzes, John Schweppe, Edward Siciliano, Denis Strachan, and Dennis Weier, "The use of energy windowing to discriminate SNM from NORM in radiation portal monitors," *Nuclear Instruments and Methods in Physics Research A* 560 (2006): 375, https://doi.org/10.1016/j.nima.2006.01.053.
18. National Research Council, *Evaluating Testing, Costs, and Benefits of Advanced Spectroscopic Portals for Screening Cargo at Ports of Entry*, 3.
19. Bakshi, Nitin, Stephen E. Flynn, and Noah Gans, "Estimating the Operational Impact of Container Inspections at International Ports," *Management Science* 57 (2011): 10, https://doi.org/10.1287/mnsc.1100.1252.
20. International Atomic Energy Agency, "Coordinated Research project (J02005): Improved Assessment of Initial Alarms from Radiation Detection Instruments Annual Report 2016 and Results from the First Research Coordination Meeting," (2016), 9.
21. Burr, Tom, James R. Gattiker, Kary Myers, and George Tompkins, "Alarm criteria in radiation portal monitoring," *Applied Radiation and Isotopes* 65 (2007): 580, https://doi.org/10.1016/j.apradiso.2006.11.010; Ely, James, Richard Kouzes, John Schweppe, Edward Siciliano, Denis Strachan, and Dennis Weier, "The Use of Energy Windowing to Discriminate SNM from NORM in Radiation Portal Monitors," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 560 (2006): 373–387, doi:10.1016/j.nima.2006.01.053.

22. Li, Wenkai, *Planning and Scheduling for Maritime Container Yards* (Cham: Springer International PU, 2016), 15–16.
23. Robert Downes, Christopher Hobbs, and Daniel Salisbury, "Combating nuclear smuggling? Exploring drivers and challenges to detecting nuclear and radiological materials at maritime facilities," *The Nonproliferation Review* 26 (2019): 97, doi:10.1080/10736700.2019.1610256.
24. Ibid, 95–96.
25. U.S. Congress, House. Committee on Appropriations, *Homeland security: Hearings before the Subcommittee of the Committee on Appropriations,* 110th Congress, 2nd sess. (2008), 553
26. U.S. Government Accountability Office, "Combating Nuclear Smuggling: Lessons Learned from Cancelled Radiation Portal Monitor Program Could Help Future Acquisitions," GAO-13-256, 2013.
27. United States Government Accountability Office, "Testimony Before the Committee on Homeland Security and Governmental Affairs, U.S. Senate, Maritime Security: Progress and Challenges with Selected Port Security Programs," GAO-14-636T, 2014, 14–15, https://www.hsgac.senate.gov/imo/media/doc/Testimony-Caldwell-2014-06-04.pdf.
28. International Atomic Energy Agency, "IAEA Launches Mobile Application Tool for Radiation Alarm and Commodity Evaluation," last modified 9 June 2017, https://www.iaea.org/newscenter/news/iaea-launches-mobile-application-tool-for-radiation-alarm-and-commodity-evaluation.
29. Times correspond to a minimum of 26 and a maximum of 169 entries in the daily files for the smallest and largest record. Readings are taken every 0.2 seconds corresponding to occupancy periods of 5.2 seconds and 33.8 seconds. There were several records that had much longer time series, but these were judged to be erroneous (see Appendix) and consequently were removed from the analysis dataset.
30. International Atomic Energy Agency, "Coordinated Research project (J02005)," 13.
31. Myers, Cory, Lawrence R. Rabiner and Aaron E. Rosenberg, "Performance trade-offs in dynamic time warping algorithms for isolated word recognition," *IEEE Transactions on Acoustics, Speech, and Signal Processing* 28 (1980): 623–635, https://doi.org/10.1109/TASSP.1980.1163491; Muller, Meinard, Henning Mattes, and Frank Kurth, "An efficient multiscale approach to audio synchronization," *Proceedings of the 6th International Conference on Music Information Retrieval,* London, 11–15, September 2005, 192–197, https://archives.ismir.net/ismir2006/paper/000015.pdf; Rath, Toni M., and R. Manmatha, "Word Image Matching Using Dynamic Time Warping," *Proceedings of the 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition,* 18–20 June 2003, Madison, Wisconsin, 218–222; Viak, Jérôme, et al., "Combination of dynamic time warping and multivariate analysis for the comparison of comprehensive two-dimensional gas chromatograms: Application to plant extracts," *Journal of Chromatography A* 1216 (2009): 2866–2872, https://doi.org/10.1016/j.chroma.2008.09.027; Tormene, Paolo, Toni Giorgino, Silvana Quaglini, and Mario Stefanelli, "Matching incomplete time series with dynamic time warping: an algorithm and an application to post-stroke rehabilitation," *Artificial Intelligence in Medicine* 45 (2009): 11–34, https://doi.org/10.1016/j.artmed.2008.11.007.
32. Folgado, Duarte, Marília Barandas, Ricardo Matias, Rodrigo Martins, Miguel Carvalho, and Hugo Gamboa, "Time Alignment Measurement for Time Series," *Pattern Recognition* 81 (2018): 268–279.

33. Akila, A., and E. Chandra, "Slope Finder – A Distance Measure for DTW based Isolated Word Speech Recognition," *International Journal of Engineering and Computer Science* 2 (2013): 3413.
34. Adapted from Tormene, Paolo, Toni Giorgino, Silvana Quaglini, and Mario Stefanelli, "Matching Incomplete Time Series with Dynamic Time Warping: An Algorithm and an Application to Post-stroke Rehabilitation," *Artificial Intelligence in Medicine* 45 (2009): 11–34, doi:10.1016/j.artmed.2008.11.007.
35. International Atomic Energy Agency, "Improved Assessment of Initial Alarms from Radiation Detection Instruments," CRP-J02005 (2014), https://www.iaea.org/projects/crp/j02005.
36. To see how information is typically presented in an RPM daily file, see Lousteau, Angela, Georgia Adams, and Alexander Enders, "Benefits of Daily File Analysis Using DART," Oak Ridge National Laboratory, ORNL/TM-2013/56, May 2013, 3, https://info.ornl.gov/sites/publications/files/Pub41358.pdf.
37. United Nations, International Trade Statistics Database, Harmonized Commodity Descriptions and Coding Systems (HS), https://unstats.un.org/unsd/tradekb/Knowledgebase/50018/Harmonized-Commodity-Description-and-Coding-Systems-HS.
38. Calculations were carried out using the Fast DTW library, a widely used open-source package that performs efficient Dynamic Time Warping. A summary of the package can be found at Python, https://pypi.org/project/fastdtw/.
39. Lio, T. Warren, "Clustering of Time Series Data—A Survey," *Pattern Recognition* 38 (2005): 1857, https://doi.org/10.1016/j.patcog.2005.01.025.
40. Rodriguez, Mayra Z., et al., "Clustering algorithms: A comparative approach," *PLoS One* 14 (2019): 1–34, https://doi.org/10.1371/journal.pone.0210236.
41. Marie Lisandra Zepeda-Mendoza, and Osbaldo Resendis-Antonio, "Hierarchical Agglomerative Clustering," in *Encyclopaedia of Systems Biology*, ed. Werner Dubitzky, Olaf Wolkenhauer, Kwang-Hyun Cho, Hiroki Yokota (New York: Springer, 2013), 886–887.
42. Oak Ridge National Laboratory, "Nuclear Smuggling Detection and Deterrence FY 2016," 8.

Acknowledgment

The authors are also grateful to Charles Massey and Chonlagarn Iamsumang at the IAEA for many useful discussions in support of this research and the reviewers and editor for their helpful comments and suggestions.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was funded under IAEA Coordinated Research Project, "Improved Assessment of Initial Alarms from Radiation Detection Instruments," (CRP-J02005).

Appendix

Alarming dataset and pre-processing

The data for this research was provided to King's College London by the IAEA, it was collected from maritime facilities in three countries, from road, as opposed to rail-based RPMs. Information provided included RPM "daily-files" containing 24 hours of data on the count rates observed by RPM gamma and neutron detectors. Also included was information on the key RPM settings such as alarm thresholds, detector type, and fluctuating background radiation levels. "Daily-files" are relatively large, containing 10,000 s of lines of data, these were cleaned and pre-processed to reveal a total of 720 alarming occupancies, where the gamma radiation received exceeded the pre-set threshold. For these alarming records, spatial gamma profiles were created from the gamma readings captured every 0.2 s as the container passed through the RPM.

Information related to the commodity for alarming records was also made available in a Microsoft® Access file originally created by the on-site official, drawing on information from the shipment manifest. This included the HS-code of the commodity and its weight. Unfortunately, not all alarming occupancies could be associated with an HS-code as this had not been inputted correctly for every record. Because this study relies on ultimately associating alarms with specific commodities, records that could not be assigned an HS-code had to be removed. This work also requires gamma readings to be normalized by weight so that shipments of different weights could be compared. This was done by subtracting the background radiation before dividing by weight. Accurate weight information is therefore essential, however, here it was clear that for certain cases this had been inaccurately recorded and so these records were also disregarded. Finally, records were also removed if the gamma count readings produced were obviously erroneous. For example, there was a series of records taken in succession which had profiles with an extended tail, far longer than for other records. This is likely to have occurred due to the container not being completely evacuated from the RPM after scanning, possibly due to the formation of a traffic jam, where trucks were unable to drive away.

This pre-processing stage reduced the number of records from over 1000 to 720. Although not the focus of this study, improvements in local processes to capture RPM data would support future research.