Published by IOP Publishing for Sissa Medialab

Received: *March 27, 2025*
Accepted: *June 9, 2025*
Published: *July 16, 2025*

Classifying low-Z materials with cosmic-ray muon tomography for border security

N. Reed[a,*] **M. Mägi**[b] **K. Tark**[b] **M. Kiisk**[b] **E. Avots**[b] **V. Paštšuk**[b] **and N. Wardle**[a]

[a] *Department of Physics, Imperial College London, Prince Consort Road, London, United Kingdom*
[b] *GScan OÜ, Mäealuse tn., Tallinn, Estonia*

E-mail: nicholas.reed15@imperial.ac.uk

Abstract: Muon tomography utilizes natural cosmic-ray radiation to characterise a diverse range of objects that cannot be imaged by traditional techniques, exploiting the high penetration power and material atomic number sensitivity of muons. This paper describes a low-Z material classification exercise, using a prototype scanner system developed by GScan OÜ. and a novel tomographic reconstruction method that have been presented previously. Measurements were made of nine 6 cm side-length cubes composed of different low-Z materials, as well as a medium-Z steel, to demonstrate the feasibility of discriminating between these materials with this scanner system and technique. After tomographic reconstruction, machine learning methods were applied to classify the material measured. The final classification results achieved greater than 98% accuracy for 30 minute measurements with a limited scanner system, corresponding to just 4 minutes in an ideal scanner system. These results motivate further developing this technology and the corresponding methods for low-Z material classification as the next step towards commercialised low-Z muon tomography.

Keywords: Detection of contraband and drugs; Detection of explosives

*Corresponding author.

© 2025 The Author(s). Published by IOP Publishing Ltd on behalf of Sissa Medialab. Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

https://doi.org/10.1088/1748-0221/20/07/P07032

Contents

1 Introduction

Trafficking of illegal goods is a major challenge in modern society, necessitating effective border controls. Trafficking of firearms, drugs, and people all damage local economies and create risk to the public, while remaining very profitable for the criminals engaged in these activities [1]. There is also a particular need to prevent the illicit transport of Special Nuclear Materials (SNM) [2]. One of the most common techniques deployed for cargo scanning to detect prohibited materials is X-ray radiography. However, X-ray radiography systems have inherent limitations such as low penetration power restricting the scanning of large cargo objects, low sensitivity to material composition, and radiation safety restrictions, limiting its use on occupied vehicles.

First proposed in 2003 for the detection of SNM [3], muon scattering tomography (MST) is an emerging research field that has expanded rapidly in recent years, with applications in border security [4, 5], non-destructive testing and imaging [6], and monitoring of spent nuclear fuel [7–9]. This wide applicability is due to the high penetrating power of muons, which allows muon tomography techniques to overcome issues with shielding encountered by more typical X-ray scanners or passive radiation portal monitors. Furthermore as muon tomography relies solely on atmospheric rays, it does not carry the radiation safety concerns associated with X-ray systems or active radiation interrogation systems.

MST specifically relies on multiple Coulomb scattering of muons in materials. The distribution of the scattering angle θ after traversing a medium can be approximated by a Gaussian distribution for small scattering angles, with a mean of 0 and a standard deviation θ_0 given by,

$$\theta_0 = \frac{13.6\,\text{MeV / c}}{p\beta}\sqrt{\frac{\Delta z}{X_0}}\left(1 + 0.038\ln\frac{\Delta z}{X_0\beta^2}\right), \tag{1.1}$$

where p is the particle momentum, β is its speed as a fraction of the speed of light c, Δz is the thickness of the medium and X_0 denotes the radiation length of the medium [10]. This is illustrated in figure 1.

2025 JINST 20 P07032

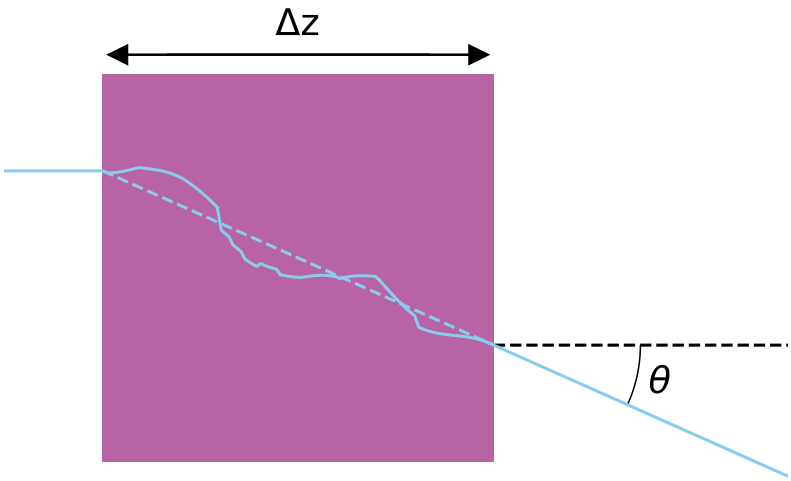

Figure 1. Schematic illustrating multiple scattering. The solid blue line indicates the true path of the particle through the medium. The total scattering angle θ is achieved after multiple scattering events through a medium of thickness Δz and is approximately distributed according to a Gaussian distribution with a mean of 0 and a standard deviation of θ_0.

The radiation length is a function of the medium density and atomic composition, and for a medium with a single type of nucleus it can be approximated as,

$$X_0 = \frac{716.4\ \mathrm{g\ cm^{-2}}}{\rho} \frac{A}{Z(Z+1)\ln\left(\frac{287}{\sqrt{Z}}\right)}, \tag{1.2}$$

where ρ, A, and Z denote the density, atomic weight, and atomic number of the medium respectively [10]. Equations 1.1 and 1.2 suggest that the radiation length is smaller for high-Z materials, resulting in a broader distribution of scattering angles. As a result, it is significantly easier to identify or discriminate between high-Z materials than between low-Z materials when using MST. The majority of applications since the inception of MST have been developed for detecting high-Z materials, such as uranium or plutonium, from amongst other high or medium-Z background materials. However, in recent years there have been efforts to extend this technique to low-Z materials.

Evaluating the feasibility of other physics signatures beyond muon scattering has proven an attractive field of research. Cuéllar et al. investigated the use of the stopping of low momentum cosmic rays as an additional signature of material composition for medium and low-Z materials using simulations of cosmic ray muon interactions, which was subsequently demonstrated experimentally in 2010 [11, 12]. Other muon-induced signatures were considered by Klimenko et al, including muonic X-rays and electromagnetic showers, although it was noted that additional detector components would be required to measure these [13]. Alternate designs of detectors with muon momentum measurement capabilities have also been proposed, including a multiple Coulomb scattering approach, Cherenkov detector-based designs, and time-of-flight measurements, and it has been demonstrated

2025 JINST 20 P07032

that accurate measurement of muom momentum significantly improves discriminating power between low-Z materials [14–17]. At the TUMUTY facility in China, machine learning techniques have been applied to discriminate low-Z materials from high-Z ones with both simulated and experimental data, with an error rate as low as 1% for a 10–30 minute measurement time [18, 19]. Measurements by Decision Sciences have led to success in identifying and discriminating low-Z materials, including locating marijuana hidden inside metallic rolls with a 4.5 minute measurement, and discriminating drug and explosives surrogates from other cargo such as bananas and bottled water [20]. While there has been success in identifying low-Z materials amongst materials of higher Z, further work is required to develop robust low-Z material classification.

In 2020, a lab-scale prototype tomography system was developed by GScan OÜ [21]. This system measures the muon trajectory above and below an interrogation volume using 1 mm diameter plastic scintillating fibres, with a resulting spatial and angular resolution in track reconstruction of 120 µm and 1 mrad respectively. A particle filtering approach to separate cosmic-ray muons and electrons was used to significantly increase the discriminating power for low-Z materials, which was demonstrated for a simulation of an RDX explosive cube, $(CH_2N_2O_2)_3$, inside a larger cube of flesh-like material.

The work presented in this paper focuses on low-Z material classification within a limited time frame, as anticipated for security applications. A set of low-Z material samples was collected, as well as a medium-Z steel sample. Each of these samples were measured with the prototype tomography system for a total exposure time of eight hours. These measurements were subsequently used to train classification algorithms to discriminate between different low-Z materials. For an exposure time of 30 minutes with a limited field-of-view scanner system, the best performing classification algorithm achieves an accuracy greater than 98%, indicating the potential of this technology for security applications.

2 Method

2.1 Prototype scanner system and experimental setup

Previously, a hodoscope design using scintillating fibres has been proposed for measuring incoming atmospheric ray tracks [21]. These hodoscopes are constructed of three detector plates, with a separation of 7.5 cm between each detector plate and a plate area of $25.0 \times 25.0\,\mathrm{cm}^2$. Each detector plate is composed of two orthogonal double-layered mats of 1 mm round Saint-Gobain BCF-12 single-cladded scintillating fibres, with a pitch of 1.1 mm in each layer. The top layer of an individual fibre-mat is offset by one half-pitch and aligned with the positions of fibres in the lower layer. Each detector plate has an active area of $24.7 \times 24.7\,\mathrm{cm}^2$. The fibre configuration results in a 120 µm spatial resolution and an angular resolution for in-hodoscope scattering of approximately 1 mrad. Further detail regarding the construction of individual detector plates and the derivation of the spatial and angular resolution can be found in [21]. A schematic of an individual double-layered fibre mat can be seen in figure 2.

For this exercise, measurements were made using a prototype scanner system with one such hodoscope above the volume of interest (VOI) and a single detector plate positioned 25.0 cm below, for a total volume of interest of $24.7 \times 24.7 \times 25.0\,\mathrm{cm}^3$. As only a single detector plane is placed below the VOI instead of a full hodoscope, the total scattering angle inside the VOI cannot be measured. For the purposes of demonstrating material classification the total scattering angle can be approximated by calculating the scattering angle between the bottom plane of the top hodoscope and the bottom detector plate.

2025 JINST 20 P07032

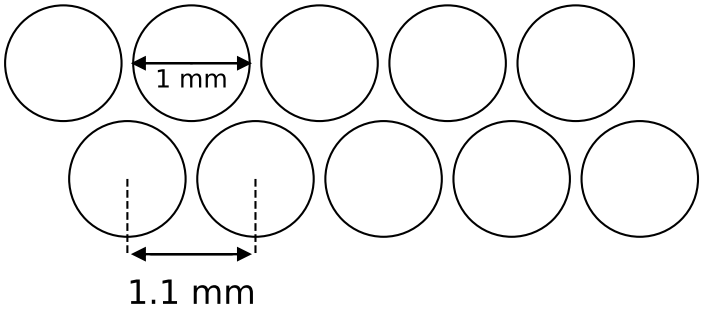

Figure 2. Schematic of a single double-layered fibre mat, viewed end-on. Each fibre has a diameter of 1 mm and the spacing between fibres in a single layer is 1.1 mm. The two layers are offset by a half-pitch, 0.55 mm.

The scintillating fibres are instrumented using Ketek PA3325-WB-0808 and Hamamatsu S13361-3050AE-08 silicon photomultiplier (SiPM) arrays, coupled to eight CAEN DT550W readout boards. Additionally, two $30 \times 30\,\mathrm{cm}^2$ plastic scintillator plates positioned above the upper hodoscope and below the lower detector plate respectively are used to trigger the data readout. These are instrumented with two $6 \times 6\,\mathrm{mm}^2$ Hamamatsu S13360-6050CS SiPMs with CAEN A1702 SiPM readout boards. When a charged particle passes through both plastic scintillator plates, the combined trigger system sends a signal to prompt a readout from all detector plate boards.

Measurements were made of 6 cm side-length cubes of eight different low-Z materials and a medium-Z steel. An additional measurement was made with no object present, which provides a sample for air for the purpose of material classification. The materials used are summarised in table 1. For each measurement, the cube of material was fixed 2 cm below the centre of the bottom detector plate of the upper hodoscope. A schematic of the experimental setup with an example true and reconstructed track can be seen in figure 3, and a photograph of the setup is shown in figure 4. The nine measurement samples are shown in figure 5. Each of the chosen samples was measured for 8 hours, divided into 30 minute segments such that there were 16 separate measurements for each material.

2.2 Data analysis

For object reconstruction, a voxel grid with a 5 mm spatial resolution is defined across the VOI. An approximation to the track of a detected particle through the VOI is found by finding the straight line between the hits in the lowest upper hodoscope plate and the bottom hodoscope plate.

To use a given particle measurement for object reconstruction and material classification, hits in all four detector planes are required. The hodoscope internal scattering angle is used to group tracks into three angle ranges, in which 0–0.2° is predominantly muons, 0.2–0.9° is a mixture of muons and electrons, and above 0.9° is solely electrons [21]. This pre-filtering is also described in further detail in the corresponding patent, and an application for an additional patent covering the method of material classification presented later using an instrument similar to the one presented here has been filed [22, 23].

To express the scattering angle in terms of the detector plane hits, each plane is numbered 1–4 counting from the highest plane down, such that planes 1–3 are the detector planes in the upper hodoscope and plane 4 is the lowest detector plane. The ray from a hit in plane 1 to a hit in plane

2025 JINST 20 P07032

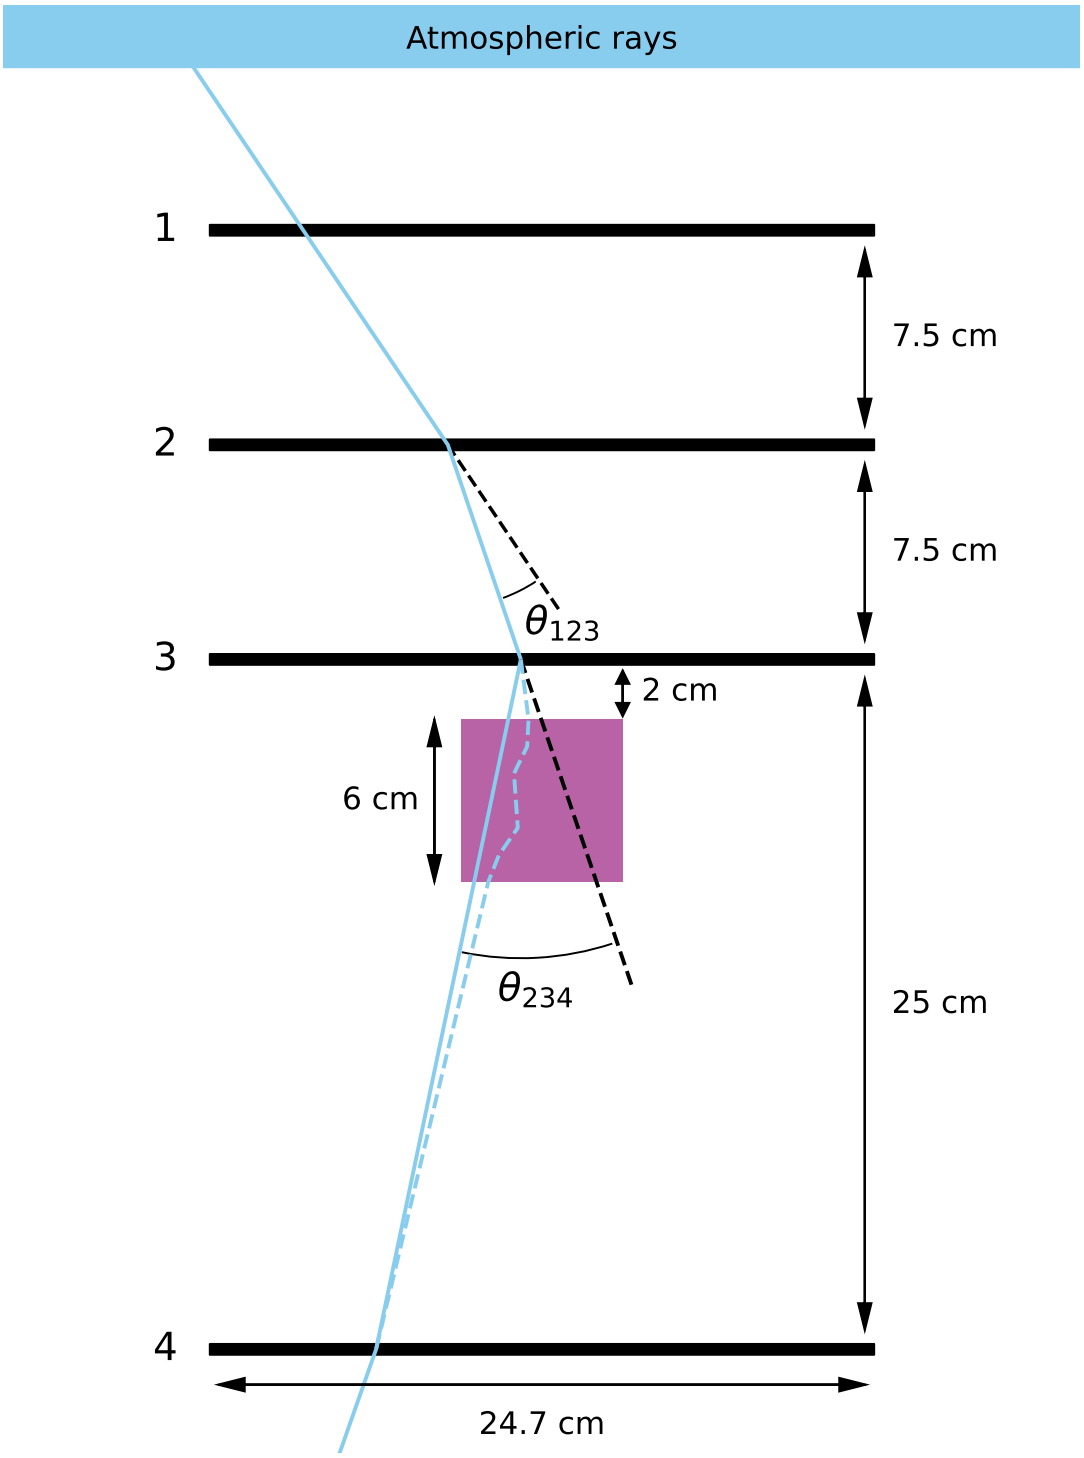

Figure 3. A schematic of the prototype scanner system with four horizontal detector plates, in a 3-1 configuration. The measurement object is aligned with the centre of the detector plates and positioned 2 cm below the third plate. The solid and dashed blue lines indicate the reconstructed and true tracks of an incident muon respectively. The hodoscope and VOI scattering angles are labelled as θ_{123} and θ_{234} respectively.

2025 JINST 20 P07032

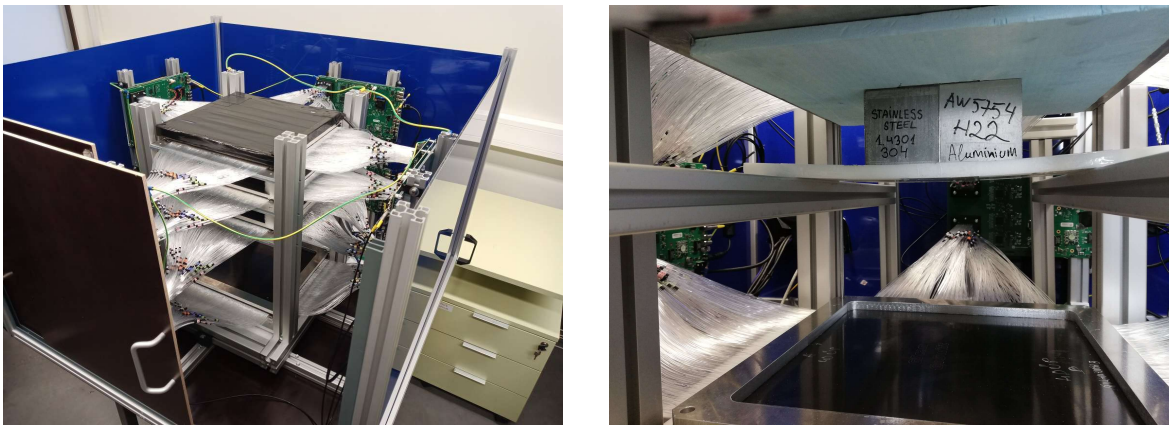

(a) A photograph of the whole scanner system.

(b) Measurement samples inside the scanner system.

Figure 4. Photographs of the prototype scanner system. The 3-1 detector plate configuration can be seen, as well as the scintillating fibres and the test sample.

Figure 5. The test samples used for this exercise. Starting from the top left, moving left-to-right and top-to-bottom: ammonium nitrate with sulphur pellets, PENO plastic explosive, plywood, water, glass, plexiglass, concrete, steel, and aluminium.

2025 JINST 20 P07032

Table 1. Measured materials for the material classification exercise. With the exception of steel, all measured materials are comprised of low atomic number constituents.

Material	Weight / g	Notes
Air	-	Measurement with no sample present
AW5754 Aluminium	574.2	-
C30/37 Concrete	512.0	-
Glass	514.8	-
Water	256.9	2 mm thick plexiglass container
Ammonium nitrate with sulphur pellets (NH_4NO_3)	199.6	2 mm thick plexiglass container
PENO plastic explosive	342.0	-
Plexiglass	251.6	-
15 mm FF Birch Plywood	147.6	-
S355MC Steel	1684.9	-

2 is denoted as ray 12, and then the scattering angle between ray 12 and ray 23 is denoted as θ_{123}, and is referred to as the hodoscope scattering angle. The same convention is used to derive the VOI scattering angle θ_{234}. These angles are illustrated in figure 3. Typically, all three particle filter groups are used separately for object reconstruction as they each have different sensitivity to materials in the VOI. Measurements of θ_{234} encode information about the VOI object scattering but are position sensitive and as such require all prototype measurements be of objects of the same size and in the same location, as described in the experimental setup.

Every track that is propagated through the VOI is registered in each voxel that it passes through. VOI maps for each desired quantity are constructed by assigning the relevant quantity from each track to every voxel it traverses. Separate VOI maps are calculated for each particle filter group. A proprietary object detection algorithm using a series of iterative edge detection and noise suppression kernels is used to segment the VOI into individual objects. However, due to the limited tomographic setup for this prototype, the vertical distribution of the detected objects was not sufficiently accurate, which can be seen in further detail in [21]. As a result instead of using reconstructed objects, the object position, size, and shape in the VOI were defined in software for the purposes of material classification.

Histograms were derived from quantities related to both scattering and stopping power measurements. For each event, the track length and average distance traversed per voxel were calculated and recorded in every voxel traversed by the track. All voxels occupied by the measurement samples were selected, as well as all events traversing these voxels, according to the specified size, shape, and position of the measurement samples. The VOI map quantities calculated included the total number of events in each voxel and the total average distance traversed per voxel. Scattering angle quantities were also calculated, including the mean scattering angle in each voxel, scattering angles from all events traversing the sample volume, the total scattering angle per unit track length in each voxel, and the average scattering angle per unit track length in each voxel. Histograms for each of these quantities were then built, with 101 bins for each of the four scattering angle histograms and 51 bins each for the remaining two histograms. This process was repeated for each particle selection group and possible combination of detector layers. For each measurement with two detector layer combinations and three

2025 JINST 20 P07032

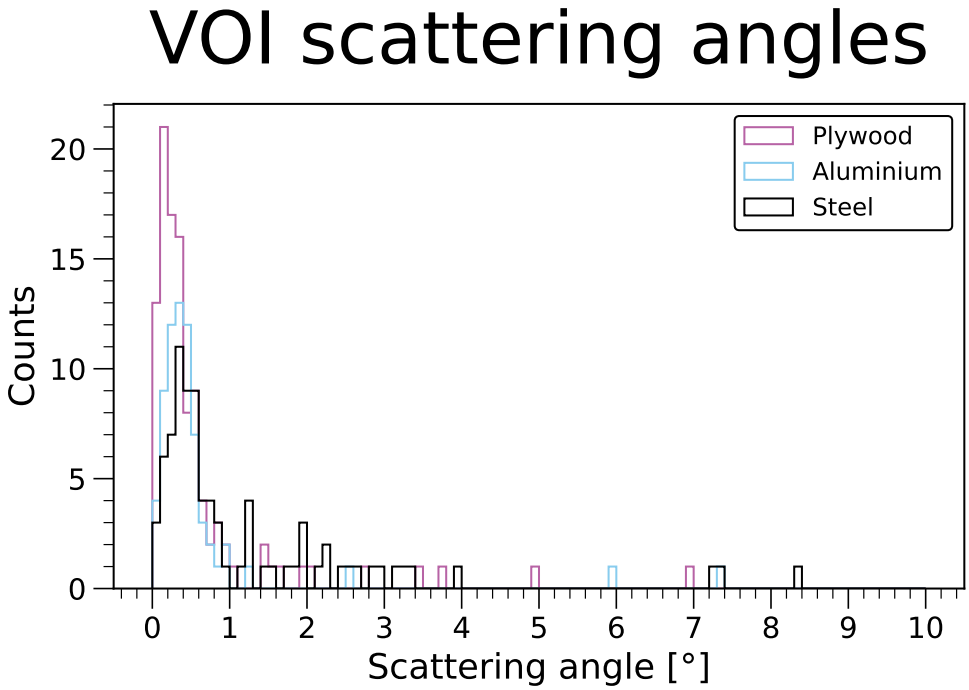

Figure 6. Example histograms of VOI scattering angles for 30 minute measurements of the plywood, aluminium, and steel samples.

particle selection groups, this resulted in a total of 36 histograms. The final amount of training data created for a given material sample is therefore 36 histograms × 16 measurements for a total of 576 histograms. Example histograms of angle θ_{234} from all events traversing the sample volume can be seen for measurements of plywood, aluminium, and steel in figure 6.

Each individual histogram was normalised to unity so that the resulting weights in each bin approximate a probability density for each feature. After all the histograms have been generated, the collective information was saved as feature vectors. For every bin in a histogram, the corresponding entry in the feature vector is the frequency in that bin. The total length of the feature vector is therefore equal to the sum of the number of bins in all histograms, for a total of 3036 features. Because each individual histogram is normalised separately the complete feature vector sums to 36, the number of histograms. Every feature was standardised according to the mean and variance of that feature across the entire dataset before application of any classification algorithm.

2.3 Feature selection

Before classification, a feature selection procedure was applied to determine which features have the greatest predictive power. This allows for removal of any features that have minimal or even detrimental effect on the accuracy of material classification. To this end, the ReliefF feature importance algorithm was applied to order the feature vectors by the calculated feature importance [24]. This algorithm computes a score for each feature based on how important that feature is for discriminating between different classes. Features that reduce class discriminating power have a negative feature score, and features with no variance across the entire dataset are not assigned a ReliefF score and

2025 JINST 20 P07032

instead are removed from the feature vector, kept in the same order, and appended to the end of the ordered feature vector. In the case of this dataset, features with zero variance have a value of zero for all data points in all classes and thus are not informative for material classification.

After ordering, an N-fold cross-validation was used to determine the optimal number of features for classification. For each fold, the total dataset was divided into three subsets: one for training, one for optimisation of feature count, and one for final performance evaluation. These are referred to as the training, validation, and test sets respectively. The top-ranked feature was selected for all datasets and an N-fold cross-validation was performed, training on the training dataset and evaluating performance on the validation dataset. The performance of each classification algorithm was quantified by calculating the average accuracy of classification across all folds, defined as the ratio of the number of correct predictions and the total number of predictions. Subsequently the procedure was repeated but including the next most highly ranked feature. This iteration continued until the accuracy had been calculated for all feature counts. The optimal number of features for a given classification algorithm was then determined according to the highest average classification accuracy. For multiple feature counts with the same average classification accuracy, higher feature counts are favoured to minimise the risk of overfitting based on a small number of features.

2.4 Material classification

The classification method adopted for this exercise was linear discriminant analysis (LDA), using the SCIKIT-LEARN Python package [25]. For this algorithm, feature values are assumed to be distributed according to an independent multivariate Gaussian distribution for each class. At training time, the parameters of each class Gaussian are estimated from the corresponding samples to minimise the in-class variance and to maximise the separation between the mean of different classes within the feature space [26, 27]. Prediction data points are then projected into the LDA feature space and assigned to the closest class according to the covariance-weighted distance from each class mean [28]. For this method, a 15-fold cross validation fold was used. The number of top-ReliefF-ranked features was selected according to the best validation set performance across all folds, and the final classification accuracy was evaluated on the test dataset for all folds.

Additionally, a one-vs-one LDA approach was investigated. For this algorithm instead of modelling each class with a separate Gaussian distribution simultaneously, an independent binary LDA classifier was trained for each unique pair of materials for a total of 45 individual classifiers. Each individual binary classifier is expected to capture different discriminating features between pairs of classes when compared to the multiclass classifier. To make predictions each classifier is used to predict the sample class and the final output is the most commonly predicted class across all of the individual classifiers. The number of training features used was optimised separately for each one-vs-one classifier with a 15-fold cross validation procedure, using the same data split as for the multiclass LDA method.

The final performance for both of these algorithms with a selected optimal feature count was determined by evaluating the performance on the test dataset for each fold and calculating a final average classification accuracy. Confusion matrices were also calculated for the final trained algorithms. These are tables that display what percentages of samples with a specific true class are predicted as each possible class, for every specific class, and therefore allow for identification of any specific classification errors between materials.

2025 JINST 20 P07032

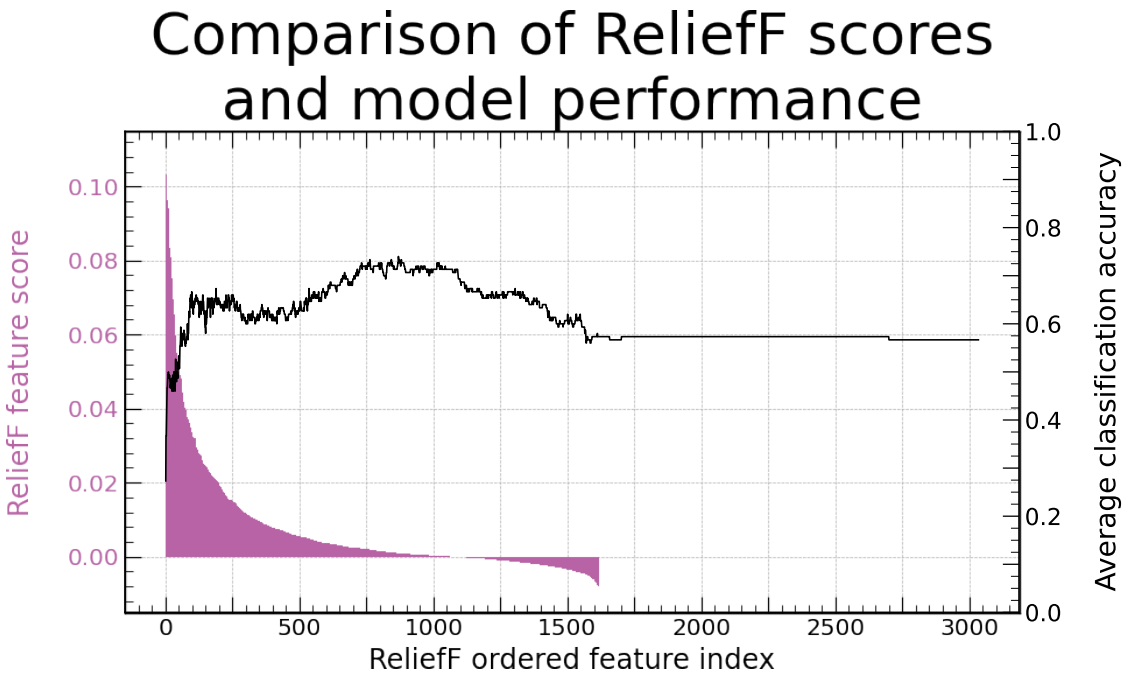

Figure 7. ReliefF feature scores and average model classification accuracy for multiclass LDA, as a function of the ReliefF ordered feature index used. Each classifier training uses the top-ranked n_{features} according to their ReliefF scores. Features with a negative ReliefF score reduce the performance of the classification and features with no variance across the dataset do not impact the performance in a significant way.

3 Results and discussion

The ReliefF feature scores ordered by score and the average classification accuracy for multiclass LDA as a function of the number of ReliefF-ordered features used can be see in figure 7. The maximum classification accuracy is found for training using the top-ranked 870 features, with an average accuracy of 74% across all 15 folds. The significant drop-off in efficiency observed after approximately 1100 features used corresponds to when the ReliefF feature scores become negative. These feature scores indicates that these features actively reduce material discrimination power, which is reflected in the classification accuracy results. After approximately 1600 features are used, adding any more features makes no significant difference to the classification performance, as these features have zero variance across the training dataset and thus provide no additional discriminating power.

The confusion matrix for the best-performing multiclass LDA, trained using the 870 top-ranked features, is shown in figure 8. The average accuracy for this classifier is 74.0%. While good performance is observed for some materials, namely Air, PENO, and Steel, the classification accuracy for more than half of the materials is less than 75%. Similarly, figure 9 shows the confusion matrix for the best-results 1v1 LDA approach. The performance is significantly improved relative to the multiclass approach, with only two misclassifications and a final total accuracy on unseen data of 98.7%. All individual materials are classified with an accuracy greater than 90%. The improved performance is attributed to the ability of individual binary classifiers to better capture distinctions between pairs of classes than a single multiclass classifier is capable of. The results shown here are believed to be the first published results showing high accuracy material discrimination between

2025 JINST 20 P07032

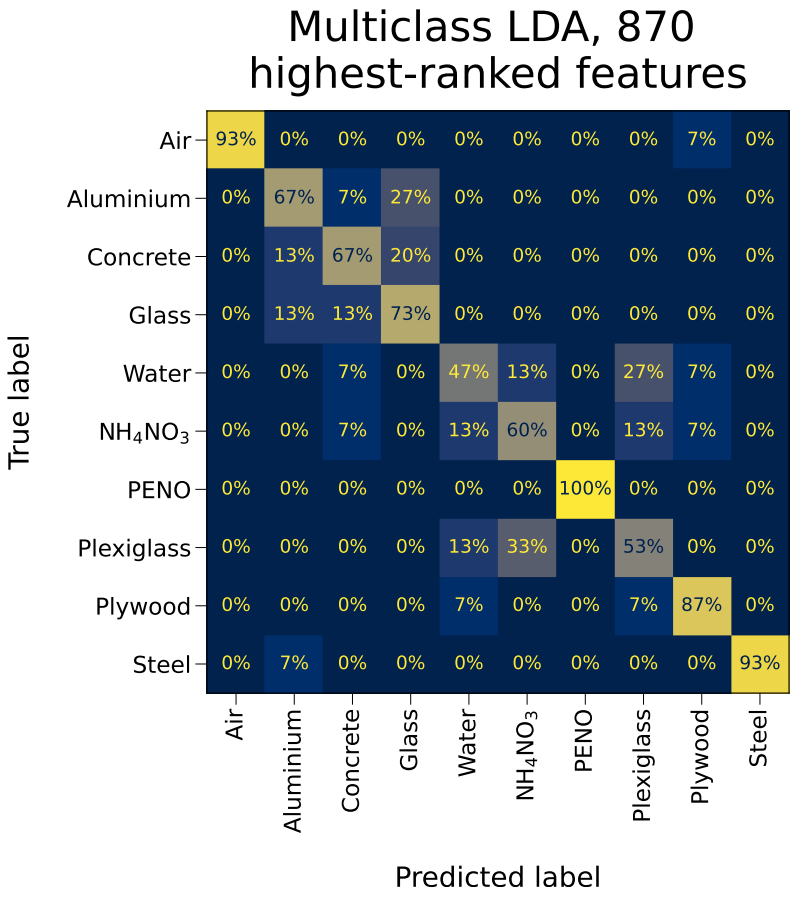

Figure 8. Confusion matrix for multiclass LDA, trained using 870 top-ranked features. The average classification accuracy for this model was 74.0%.

low-Z materials. While this experiment has been conducted on a limited lab-prototype tomograph and on a limited set of materials with only 8 hours of measurement time for each sample, the material classification accuracy is greater than 98% for a number of different low-Z materials.

One key limitation of this exercise is the simplified prototype, with only a single detector plate below the volume of interest. A complete tomography system with complete hodoscopes both above and below the volume of interest would be capable of more accurately measuring the total muon scattering angle. Additional hodoscopes at the sides of the detector would further improve performance by increasing the solid angle coverage of the scanner system, reducing the required measurement times for accurate classification. The mean number of muon tracks detected in the measurements of this exercise was 2637, which when compared to the expected incident count of 18300 muons for a 30 minute measurement indicates that this prototype only successfully detects approximately 14% of the incident muon flux, which is approximately 1 $\mathrm{cm}^{-2}\,\mathrm{min}^{-1}$ [10]. A system with 100% efficiency of detection for incident muons would therefore be capable of detecting the same number of muons in just 4 minutes, as opposed to the 30 minute measurements with the prototype system. An added benefit of this change would improve the object reconstruction performance and thus ensure the same method could be applied to objects of arbitrary shape and size.

2025 JINST 20 P07032

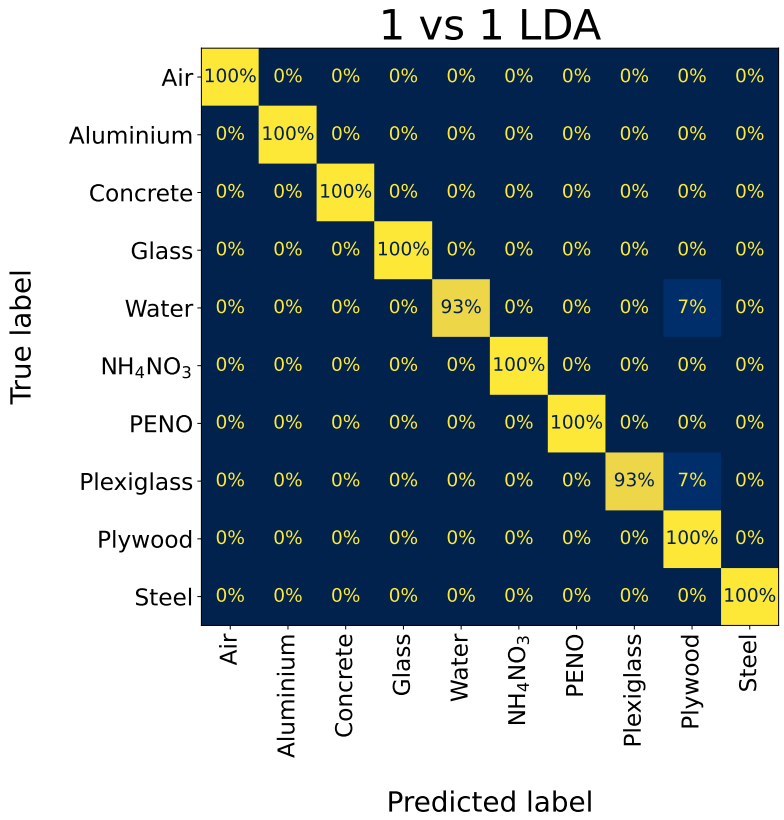

Figure 9. Confusion matrix for the best performing 1 vs 1 LDA model. The only misclassifications were predicting one sample of water and one of plexiglass as Plywood. The total accuracy of this classifier was 98.7%.

To develop classification further, one of the main requirements is to increase the material library available to the system. Measurements of more low-Z materials and objects common to security applications would be added to the training data for classification algorithms, increasing the range of objects which this instrument would be capable of detecting. At this stage, it would likely be necessary to move to other machine learning methods, which would need benchmarking to determine the most reliable way to classify object materials.

4 Conclusion

Muon scattering tomography is a promising technology for many applications as a non-invasive, radiation-free imaging technique. Most existing work has shown accurate results for high and medium-Z materials, but discriminating low-Z materials remains a challenging task. This paper has presented a prototype scanner system and tomographic method capable of discriminating between several low-Z materials common in security applications with better than 98% accuracy for 30 minute measurement times. The accuracy of tomographic reconstruction can be improved with an updated scanner system with two complete hodoscopes, and measurements of other materials and objects with such a scanner system would allow for classification with more complex configurations. Developments to improve this scanner system for more accurate reconstruction and classification are already underway.

2025 JINST 20 P07032

Acknowledgments

This work is a part of the CosmoPort project and has received funding from the European Union's Horizon Europe programme, grant number 101121379. Imperial College London contributions were funded by U.K. Research and Innovation (UKRI) under the U.K. government's Horizon Europe funding Guarantee, grant number 10102097.

References

[1] C. May, *Transnational crime and the developing world*, https://gfintegrity.org/report/transnational-crime-and-the-developing-world/.

[2] United States Nuclear Regulatory Commission, *Safeguard Categories of SNM*, https://www.nrc.gov/security/domestic/mca/snm.html.

[3] K.N. Borozdin et al., *Radiographic imaging with cosmic-ray muons*, *Nature* **422** (2003) 277.

[4] W.C. Priedhorsky et al., *Detection of high-Z objects using multiple scattering of cosmic ray muons*, *Rev. Sci. Instrum.* **74** (2003) 4294.

[5] P. Baesso, D. Cussans, C. Thomay and J. Velthuis, *Toward a RPC-based muon tomography system for cargo containers*, 2014 *JINST* **9** C10041.

[6] P.M.R.-D. Arbol, P.G. Garcia, C.D. Gonzalez and A. OrioAlonso, *Non-destructive testing of industrial equipment using muon radiography*, *Phil. Trans. A. Math. Phys. Eng. Sci.* **377** (2018) 20180054.

[7] C. Thomay et al., *Passive 3D imaging of nuclear waste containers with Muon Scattering Tomography*, 2016 *JINST* **11** P03008.

[8] D. Poulson et al., *Cosmic ray muon computed tomography of spent nuclear fuel in dry storage casks*, *Nucl. Instrum. Meth. A* **842** (2017) 48 [arXiv:1604.08938].

[9] D. Mahon et al., *First-of-a-kind muography for nuclear waste characterization*, *Phil. Trans. A. Math. Phys. Eng. Sci.* **377** (2018) 20180048.

[10] Particle Data Group collaboration, *Review of Particle Physics*, *PTEP* **2022** (2022) 083C01.

[11] L. Cuéllar et al., *Soft cosmic ray tomography for detection of explosives*, in the proceedings of the *IEEE Nuclear Science Symposium*, Orlando, FL, U.S.A., October 24–November 1 (2009), pp. 968–970 [DOI:10.1109/nssmic.2009.5402469].

[12] L. Cuéllar et al., *A range muon tomography performance study*, in the proceedings of the *IEEE Nuclear Science Symposium, Medical Imaging Conference, and 17th Room Temperature Semiconductor Detectors Workshop*, Knoxville, TN, U.S.A., October 30–November 6 (2010), pp. 67–69 [DOI:10.1109/NSSMIC.2010.5873718].

[13] A.V. Klimenko et al., *Fusing Signatures of Different Physical Processes in Muon Tomography*, in the proceedings of the *IEEE Nuclear Science Symposium*, Fajardo, PR, U.S.A., October 23–29 (2005), pp. 307–311 [DOI:10.1109/nssmic.2005.1596259].

[14] V. Anghel et al., *A plastic scintillator-based muon tomography system with an integrated muon spectrometer*, *Nucl. Instrum. Meth. A* **798** (2015) 12.

[15] J. Bae and S. Chatzidakis, *Momentum-Dependent Cosmic Ray Muon Computed Tomography Using a Fieldable Muon Spectrometer*, *Energies* **15** (2022) 2666.

[16] J. Chen, H. Li, Y. Li and P. Liu, *Towards a muon scattering tomography system for both low-Z and high-Z materials*, 2023 *JINST* **18** P08008 [arXiv:2304.09489].

2025 JINST 20 P07032

[17] J. Marteau et al., *Implementation of sub-nanosecond time-to-digital convertor in field-programmable gate array: applications to time-of-flight analysis in muon radiography*, *Measur. Sci. Tech.* **25** (2014) 035101 [arXiv:1310.4281].

[18] Z. Yifan et al., *Discrimination of drugs and explosives in cargo inspections by applying machine learning in muon tomography*, *High Power Laser Part. Beams* **30** (2018) 086002.

[19] X.-Y. Pan et al., *Experimental validation of material discrimination ability of muon scattering tomography at the TUMUTY facility*, *Nucl. Sci. Tech.* **30** (2019) 120.

[20] International Atomic Energy Agency, *Muon Imaging: Present status and emerging applications*, Tech. Rep. IAEA-TECDOC-2012, IAEA, Vienna (2022), https://www.iaea.org/publications/15182/muon-imaging.

[21] G. Anbarjafari et al., *Atmospheric ray tomography for low-Z materials: implementing new methods on a proof-of-concept tomograph*, arXiv:2102.12542.

[22] A. Georgadze, M. Kiisk, M. Mägi, E. Avots and G. Anbarjafari, *Method and Apparatus for Detection and/or Identification of Materials and of Articles Using Charged Particles*, U.S. Patent US6107628A, 2019.

[23] M. Mägi, E. Avots and K. Aktas, *System and method for classifying material*, European Patent EP4411424A1, 2023.

[24] I. Kononenko, *Estimating attributes: Analysis and extensions of RELIEF*, in *Machine Learning: ECML-94*, F. Bergadano and L. De Raedt, eds., Springer Berlin Heidelberg (1994), p. 171–182 [DOI:10.1007/3-540-57868-4_57].

[25] F. Pedregosa et al., *Scikit-learn: Machine Learning in Python*, *J. Mach. Learn. Res.* **12** (2011) 2825 [arXiv:1201.0490].

[26] J. Cohen et al., *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*, Routledge (2013) [DOI:10.4324/9780203774441].

[27] R.A. Fisher, *The use of multiple measurements in taxonomic problems*, *Annals Eugen.* **7** (1936) 179.

[28] R.O. Duda, P.E. Hart and D.G. Stork, *Pattern Classification*, 2nd edition, Wiley-Interscience (2000).

2025 JINST 20 P07032