A binned clustering algorithm to detect high-Z material using cosmic muons

PUBLISHED BY IOP PUBLISHING FOR SISSA MEDIALAB

RECEIVED: *June 12, 2013*
REVISED: *September 17, 2013*
ACCEPTED: *September 23, 2013*
PUBLISHED: *October 11, 2013*

A binned clustering algorithm to detect high-Z material using cosmic muons

C. Thomay,[a,1] **J.J. Velthuis,**[a] **P. Baesso,**[a] **D. Cussans,**[a] **P.A.W. Morris,**[a] **C. Steer,**[b] **J. Burns,**[b] **S. Quillin**[b] **and M. Stapleton**[b]

[a] *University of Bristol, H.H. Wills Physics Laboratory, BS8 1TL, United Kingdom*
[b] *AWE, Aldermaston, Reading, RG7 4PR, United Kingdom*

E-mail: christian.thomay@bristol.ac.uk

ABSTRACT: We present a novel approach to the detection of special nuclear material using cosmic rays. Muon Scattering Tomography (MST) is a method for using cosmic muons to scan cargo containers and vehicles for special nuclear material. Cosmic muons are abundant, highly penetrating, not harmful for organic tissue, cannot be screened against, and can easily be detected, which makes them highly suited to the use of cargo scanning. Muons undergo multiple Coulomb scattering when passing through material, and the amount of scattering is roughly proportional to the square of the atomic number Z of the material. By reconstructing incoming and outgoing tracks, we can obtain variables to identify high-Z material. In a real life application, this has to happen on a timescale of 1 min and thus with small numbers of muons. We have built a detector system using resistive plate chambers (RPCs): 12 layers of RPCs allow for the readout of 6 x and 6 y positions, by which we can reconstruct incoming and outgoing tracks. In this work we detail the performance of an algorithm by which we separate high-Z targets from low-Z background, both for real data from our prototype setup and for MC simulation of a cargo container-sized setup.

KEYWORDS: Pattern recognition, cluster finding, calibration and fitting methods; Resistive-plate chambers; Large detector systems for particle and astroparticle physics; Analysis and statistical methods

[1] Corresponding author.

doi:10.1088/1748-0221/8/10/P10013

Contents

2013 JINST 8 P10013

1 Introduction

The threat of illicit trafficking of special nuclear material has brought about an increased interest in detection methods capable of identifying high-Z material in a non-invasive fashion on relatively short timescales. Muon scattering tomography (MST) is one method that uses the scattering behaviour of cosmic muons to estimate the material content of a volume of interest. Cosmic muons are attractive for this purpose due to being non-invasive, free and fairly plentiful (cosmic muon flux at sea level is $\sim 10000\,\mathrm{m}^{-2}\,\mathrm{min}^{-1}$, figure 1, with an angular distribution proportional to $\cos^2(\theta)$ at $E_\mu \sim 3\,\mathrm{GeV}$ [1]), easily detected with well-established and cost-effective detector technology, and virtually impossible to shield or trigger against.

Various groups world-wide have developed detectors for MST, notably starting with [4, 5]. More recently, results were shown from a large-scale system [6] built from prototype muon chambers of the CMS experiment at CERN. The standard reconstruction algorithm is an expectation-maximization (EM) method [7] which divides the volume of interest into voxels (small cubic volumes, usually with side lengths of $\sim 5\,\mathrm{cm}$) and reconstructs the most likely scattering length in each voxel.

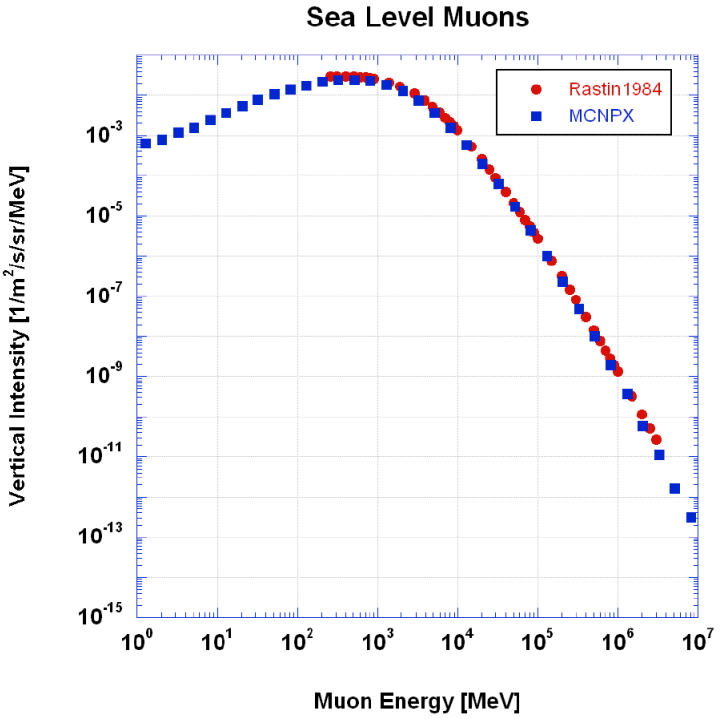

Figure 1. Muon energy spectrum at sea level (red/circles: measured data [2], blue/squares: MC simulation [3]).

The intended application of MST is scanning trucks at checkpoints or cargo containers in freight harbours. In order not to unduly disrupt cargo traffic operations, a decision between *ok* and *open and investigate* should optimally be reached in a timeframe of ~ 1 minute.

In this paper we detail a novel detection algorithm: it attempts to efficiently exploit the fact that a high-Z object will lead to strong scattering contained in a fairly small spatial volume, by estimating the 'clusteredness' within regions of the volume of interest, which can be relatively large (the current implementation uses cubes of 25 cm side length). The purpose of this algorithm is to provide a single decision value for the presence of a target and not to generate a detailed visual reconstruction.

2 Concept & system setup

We have built a prototype muon detection system using resistive plate chambers (RPCs) [8]. The system consists of 12 layers of RPCs with a fiducial region of $50\times50\,\mathrm{cm}^2$ each.

Muons undergo multiple Coulomb scattering when traversing matter (figure 2a); the angular distribution is approximately Gaussian, with width σ_0 and the radiation length X_0 given by [1]

$$\sigma_0^2 \approx \left(\frac{15\mathrm{MeV}}{pc\beta}\right)^2 \frac{T}{X_0}, \tag{2.1}$$

$$X_0 \approx \frac{A\cdot 716.4\mathrm{g/cm}^2}{\rho\cdot Z(Z+1)\ln(287/Z)}, \tag{2.2}$$

2013 JINST 8 P10013

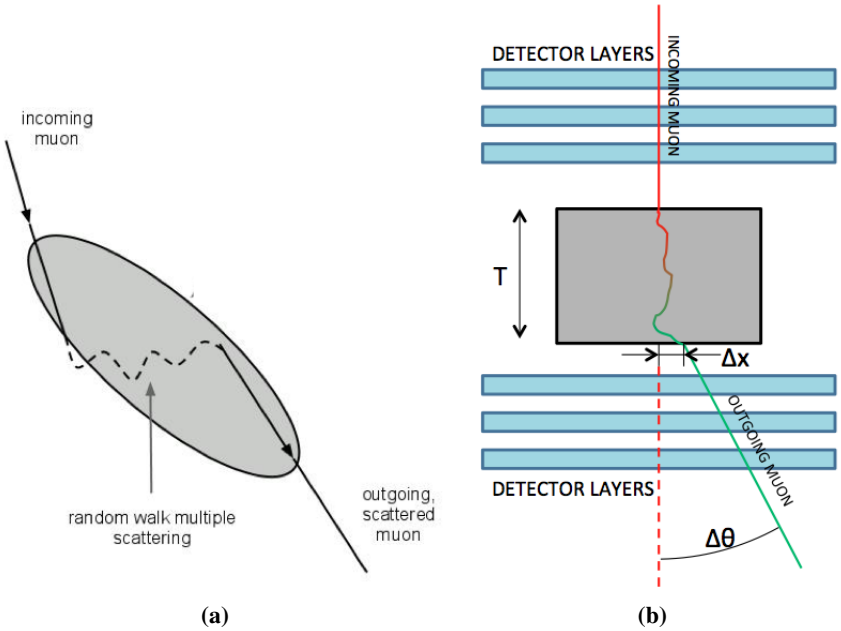

Figure 2. (a) Muon scattering principle. (b) Schematic of MST setup.

where p is the muon momentum, βc the muon velocity, T the material thickness, ρ the material density, A the mass number, and Z the atomic number.

From this it follows that the amount of scattering is approximately proportional to the Z^2 of the material traversed. By reconstructing incoming and outgoing muon tracks and studying track variables it is possible to discriminate between high-Z and low-Z material.

The analysis presented here is based on simulated data generated using a model of the prototype system. We will also present results from a simplified version of the algorithm on real data.

The muon tracks are measured using 12 layers of RPCs, paired in two for x- and y-readout (figure 2b). The muons pass the upper 6 RPCs (whereby we reconstruct the incoming track), traverse the volume under inspection, and pass the lower 6 RPCs from which we reconstruct the outgoing track. From the two reconstructed tracks, variables relating to the scattering behaviour can be calculated, and thereby material with large radiation length identified.

The setup is simulated in Geant4 [9], consisting of the 6 detector layers above and below a volume with sufficient space to accomodate a 20 ft cargo container ($5.7\times2.3\times2.4\,\mathrm{m}^3$, see figure 3). Cosmic muons were simulated with the library CRY [3]. The RPCs are sized $7\times4\,\mathrm{m}^2$ with 65 cm spacing between the layers. The area of the RPC layers was chosen to be larger than the area of the container so that the angular acceptance of the detector system doesn't affect muon statistics near the edges of the container. We have used data from the prototype system to tune the simulation: since our results indicate that our detectors have an intrinsic resolution of $< 500\,\mu\mathrm{m}$ [8], a resolution of 450 µm was chosen for the simulated detectors.

2013 JINST 8 P10013

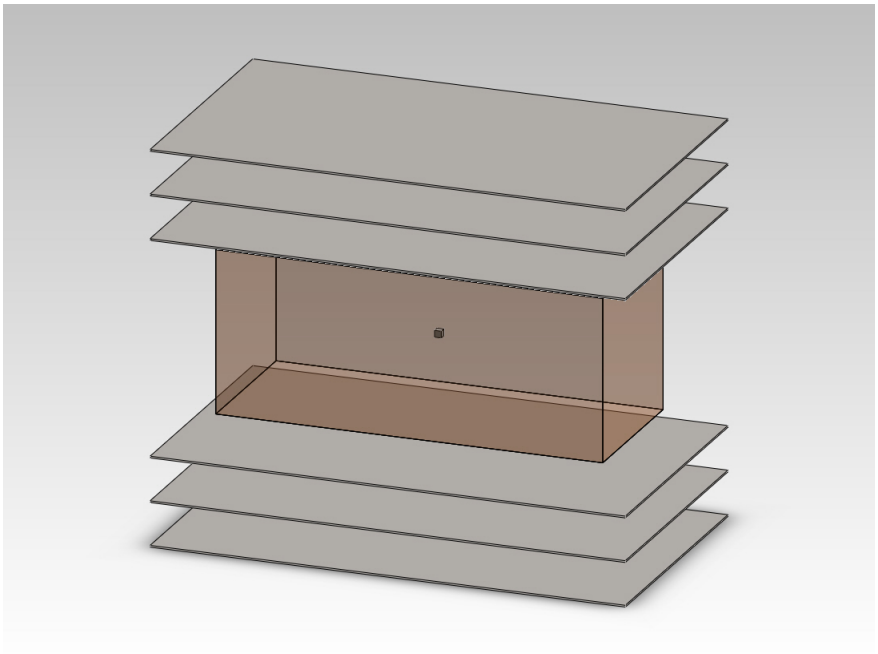

Figure 3. Large-scale simulation geometry. The RPCs above and below the cargo container volume measure $7\times4\,\mathrm{m}^2$. A $10\times10\times10\,\mathrm{cm}^3$ target block is indicated in the middle.

2.1 Event reconstruction

Track fitting is performed in two steps: first a straight line fit is performed through each triplet of hits (upper/lower, x/y) to ensure we have a correct muon track. The second step is a combined fit: it works under the assumption that the tracks share a common point (the vertex). Since muons undergo multiple Coulomb scattering in matter (figure 2a), the vertex assumption is not strictly correct. However, it is an useful approximation, since it nonetheless provides a roughly correct localization of the area of the muon scattering. A similar concept (the point of closest approach) is also used in other MST methods [5, 7]. We will investigate in section 4.4 how the uncertainty of the vertex position affects the results of the method.

Both track parameters and vertex position are fitted in the same step, which means 12 data points are fitted with 7 parameters: 3 parameters for the vertex position, and 4 for track slopes (for upper/lower tracks in x and y). An energy function E is defined (in the x-plane):

$$E_x = \sum_{i=1}^{3} \frac{(h_i - (v_x + k_{\mathrm{x,upper}} \cdot t))^2}{\sigma_{h_i}^2} + \sum_{i=4}^{6} \frac{(h_i - (v_x + k_{\mathrm{x,lower}} \cdot t))^2}{\sigma_{h_i}^2}, \tag{2.3}$$

$$\text{with } t = z_i - v_z \tag{2.4}$$

and h_i the measured hit positions, z_i the positions in the z-plane (i.e. the known RPC positions), v_x, v_y, v_z the vertex positions, k the track parameters, and σ_{h_i} the measurement errors on the hit positions.

E_y is defined analogously, and the total energy function

$$E = E_x + E_y \tag{2.5}$$

2013 JINST 8 P10013

is then minimized using MINUIT [10] to obtain the vertex position and track parameters. Tracks are selected by a number of quality cuts:

- χ^2 of the 3-point straight line fits,
- χ^2 of the combined fit,
- vertex position.

In order to be accepted, the tracks must have χ^2 values below a cutoff, and the vertex has to be in the volume of interest. After event reconstruction and quality cuts, 1 minute of simulated cosmic muon flux results in between 100000 and 150000 useable tracks going through the volume of interest. The number of useable tracks decreases if more scattering material is present in the inspection volume, since then more muons are lost by absorption and scattering out of the detector area.

3 Algorithm

The decision-making algorithm is designed to exploit the spatial clustering of muons going through the target high-Z material. First, the volume of interest is subdivided into cubic sub-volumes large enough to contain targets that we wish to detect. The reference target we will look into in the subsequent sections is a uranium block of 10 x 10 x 10 cm^3; consequently, a sub-volume side length of 25 cm was chosen.

Momentum information is a key component in all MST methods, since without at least a rough estimate of the muon momentum it is hard to distinguish between compound scattering in a large volume of low- to medium-Z material and localized scattering in a small high-Z volume. The dependence of the algorithm's performance on the momentum information will be investigated in section 5. As a reference, and in order to allow for comparison with other MST methods (e.g. [7]), we use the true muon momentum in the following sections. It has been shown to be possible to obtain an estimate of the muon momentum from the multiple scattering within the detector layers [11]; the application of this in our prototype setup is currently under investigation.

Strong scattering of particles with high momentum is more indicative of short radiation lengths, and conversely, strong scattering of particles with low momentum is not necessarily caused by high-Z material, but could also be compound scattering in low- or medium-Z material. For that reason, for muons with a momentum less than p_{norm}, the scatter angle will be weighted down, and weighted up for muons with larger momentum; in the current implementation, $p_{\mathrm{norm}} = 1\ \mathrm{GeV}/c$, based on the distribution in figure 1. Relating the muon momentum to a reference momentum is a method used in other MST algorithms as well.

Within each of the sub-volumes, the algorithm calculates a 'weighted clusteredness' value. For every pair of muon tracks with scatter vertices $\mathbf{v}_i$, $\mathbf{v}_j$ within the respective sub-volume, the metric distance m_{ij} is calculated:

$$m_{ij} = \|\mathbf{v_i} - \mathbf{v_j}\|. \tag{3.1}$$

The metric distance is then weighted by the product of the scatter angles, normalized with the muon momenta:

$$\widetilde{m}_{ij} = \frac{m_{ij}}{(\theta_i \tilde{p}_i) \cdot (\theta_j \tilde{p}_j)}, \tag{3.2}$$

where θ_i is the scatter angle of muon i and $\tilde{p}_i$ the normalized momentum, $\tilde{p}_i = p_i / p_{\mathrm{norm}}$.

2013 JINST 8 P10013

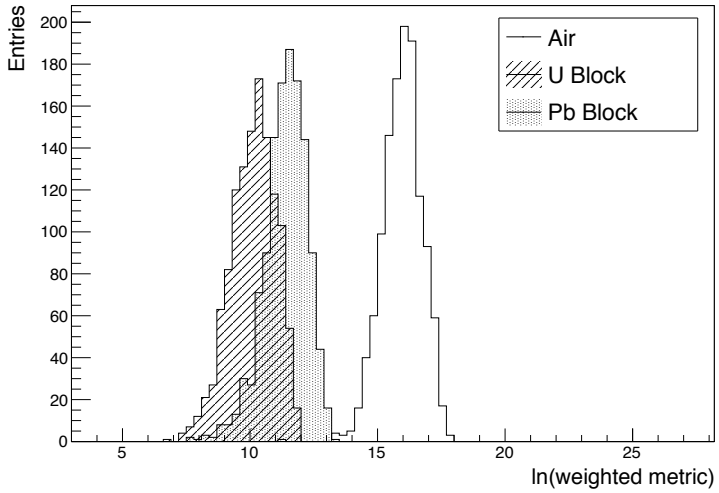

Figure 4. Natural logarithm of weighted metric values for one sub-volume, 1 minute of data. U and Pb blocks are sized 10x10x10 cm^3. The high-Z and low-Z distributions are clearly separated, so the scenarios can be clearly distinguished. Separation between lead and uranium is also possible (see section 5.5).

The shape of the distribution of all weighted metric distances is indicative of the material in the sub-volume. It can be seen in figure 4 that for larger amounts of scattering close together (as is the case with the uranium block present), the distribution has a lower median, since there are many pairs of vertices with both a short distance and strong scattering. In comparison, if only air is present, scattering on average is much lower and there is no area where the vertices would be clustered more closely.

The algorithm returns the median of the distribution of all weighted metric distances as the final discriminator.

The value of the discriminator is dependent on two things: the material content of the sub-volume, and the number of muon tracks used in the calculation; the same material yields a different discriminator if the number of muons going through the volume is different. For that reason, the muon tracks in each sub-volume are sorted by descending normalized scatter angle and only the n most strongly scattering tracks are considered. For a fixed n at fixed exposure time, the median is only dependent on the material content. For 1 minute of data, we expect between 50 and 70 tracks on average in each sub-volume, so a working point of $n = 50$ was chosen.

Dividing up the volume of interest into sub-volumes comes with the inherent risk of the target being on the edges or corners of multiple sub-volumes. For that reason the algorithm is run on multiple grids, where each sub-volume is shifted in respect to the original. This way, any target smaller than half the side length of the sub-volumes is ensured to be fully contained in at least one sub-volume of one grid.

In its final form including the multiple grids, the algorithm runs for ~ 10 sec on a desktop computer for this geometry and 1 min of cosmic muons. The most significant contribution to this comes from the sort operation that is performed on the set of tracks in each sub-volume, which has a worst-case time complexity of $O(n \log n)$. However, in a real-life scanning application, allocating tracks to sub-volumes and performing the sort operation can be done while the system is scanning.

2013 JINST 8 P10013

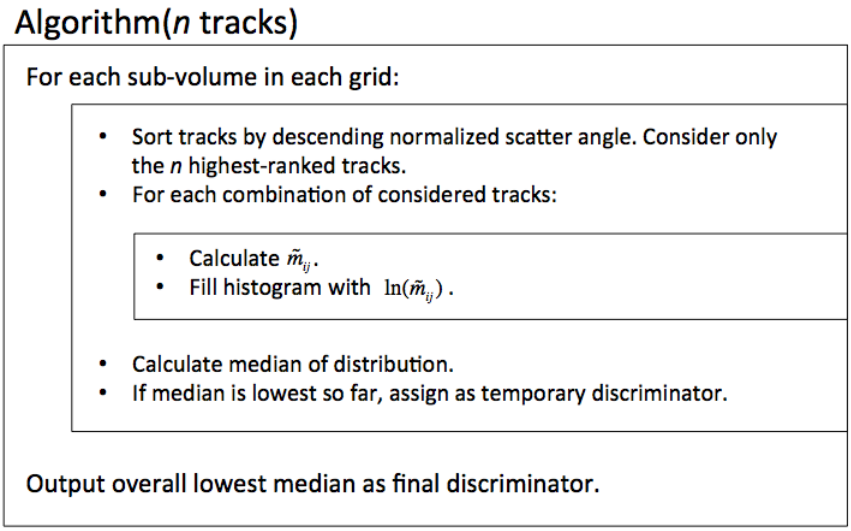

Figure 5. Algorithm workflow in pseudocode.

The runtime could also be further improved by parallelization, since after the initial assignment of tracks to the sub-volumes, each sub-volume can be treated independently.

The algorithm returns the smallest median value of all sub-volumes of all grids as the final discriminator value. The workflow of the algorithm is illustrated in figure 5. The final discriminator serves as a classifier: from simulations (or real data), a threshold value needs to be chosen. Based on where the observed value is in relation to the threshold value, the decision between *investigate* and *ok* can be made. This will be investigated in the performance section in a ROC analysis.

4 Performance

This section shows the performance on different simulated volumes. The main scenario of interest is that of a cargo container, in which the algorithm needs to be able to identify a block of high-Z material like uranium. This is complicated by the presence of shielding material, i.e. the uranium being hidden within material commonly found in cargo containers. In order to estimate how much shielding material one can reasonably expect, the geometry and weight limitations of a cargo container need to be taken into consideration. For MST in general, the most problematic geometry for any background material is vertical pile-up: the more material every muon traverses above and below the target one wishes to identify, and the lower the radiation length of that material is, the less accurate the vertex position becomes, and the harder it becomes to distinguish between the scattering that occured within the target, and the scattering that occured within the background material. Consequently, the most 'efficient' shielding of the uranium block would be to put it in the center of a cubical volume of dense shielding material.

Intermodal containers come in various sizes; the two most commonly used are the 20 ft and 40 ft container. All intermodal containers have a similar net weight limit (28200 kg for the 20 ft and 26580 kg for the 40 ft container), since the maximum weight of an intermodal container is mainly limited by the lifting capacity of container cranes. We limit ourselves here to the 20 ft container,

2013 JINST 8 P10013

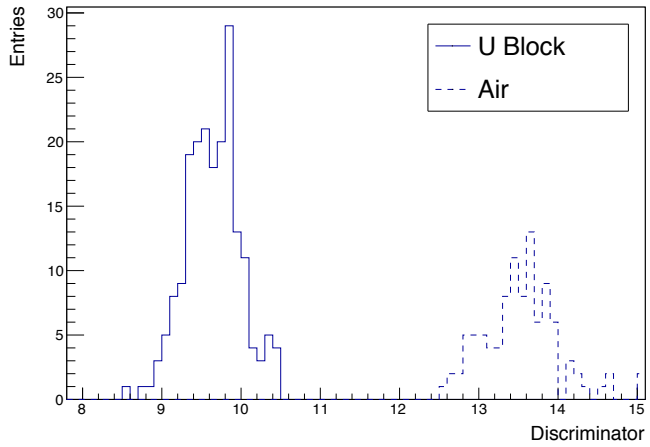

Figure 6. Discriminators for solitary U block. (1 min).

since due to the lower net weight limit and the geometry the 40 ft container will essentially always be less problematic in terms of background scattering.

The effectiveness of the discrimination is investigated in a ROC analysis (cf. [12]). ROC curves show the true positive rate (in our case, how often the uranium block was found at the correct location) versus the false positive rate (how often a uranium-less scenario was falsely classified as containing uranium), with each entry in the ROC curve corresponding to a threshold value of the discriminator. A result only counts as a true positive if the discriminator is below the threshold and the correct sub-volume was chosen, so it's possible that the ROC curve does not reach a true positive rate of 1. This is somewhat of a 'worst-case' evaluation, since it can happen that a sub-volume containing the target yields a discriminator below the threshold, and another yields an even lower discriminator; this case would not count as a true positive.

For a more exclusive choice of the threshold value (lower left region of the ROC curve), fewer data sets will be classified as 'uranium-like', but at the same time fewer data sets with the uranium block will be identified as such. A more inclusive threshold value (upper right region of the ROC curve) identifies the uranium more often, but will generally misidentify background more often as well. Optimally, there would be a threshold value at the (0,1) point, meaning correct classification in every case.

4.1 Solitary U block

This section shows the ability of the algorithm to identify a $10\times10\times10\,\mathrm{cm}^3$ block of uranium in an otherwise empty container. Figure 6 shows the distribution of the discriminator for 100 sets of 1 minute data each, for a container with and without the block of uranium. It can be seen from the fact that the two distributions are clearly separated that these two scenarios can be distinguished with 100% efficiency and 100% purity.

4.1.1 Prototype system data

A similar scenario was evaluated on real data from our prototype system (small scale, fiducial region of $50\times50\,\mathrm{cm}^2$). We used a simplified version of the algorithm: it considered only one sub-

2013 JINST 8 P10013

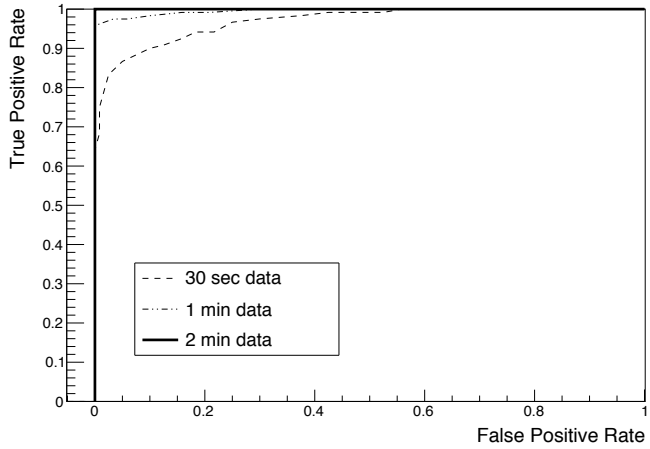

Figure 7. ROC curves for W block vs. air on prototype system data.

volume, and since momentum measurement is currently still under investigation, the momentum normalization was omitted. As a stand-in for uranium, a $9\times9\times9\,\mathrm{cm}^3$ block of tungsten was used. Figure 7 shows ROC curves for different scan times; for the purpose of this analysis, 1 minute of real data was defined to be as many reconstructed tracks as there are cosmic muons passing through the entire detector in 1 minute, so the detector was assumed to be 100% efficient. It can be seen that separation is close to optimal for 1 minute of data, and optimal for 2 minutes.

4.2 Shielded U block

In these scenarios a target block of uranium, sized $10\times10\times10\,\mathrm{cm}^3$, is shielded by different materials. In each scenario the cargo container volume is filled with as much of the material as the weight limit allows.

4.2.1 Rock

The first scenario contains rock, approximated as silicon dioxide with a density of $2.32\,\mathrm{g\,cm}^{-3}$. The maximum amount of rock that can be loaded into the container is arranged in two geometries, flat (filling the entire space in x and y, rising 90 cm high in z), and in a cubic shape (side length of 2.3 m). In both cases, the uranium block is placed in the middle of the rock volume. The impact of vertical pile-up can be seen here very clearly: figure 8a shows the ROC curves for the flat arrangement, and figure 8b for the cubic arrangement. Since every muon going through the uranium block has to traverse ~ 2.5 times more rock in the cubic case, it is much harder for the algorithm to cleanly identify the block: 1 minute is completely sufficient to clear the flat arrangement, the cubic arrangement takes twice as long.

4.2.2 Scrap iron

The second scenario contains scrap iron, which can be approximated as pure steel with 1/10 density ($0.7874\,\mathrm{g\,cm}^{-3}$). This is one of the most problematic scenarios commonly encountered in real life, since it combines a uniformly medium-scale atomic number ($Z_{\mathrm{Fe}} = 26$) with a density that allows for a 20 ft container to be filled almost completely; scrap iron is frequently shipped in this particular

2013 JINST 8 P10013

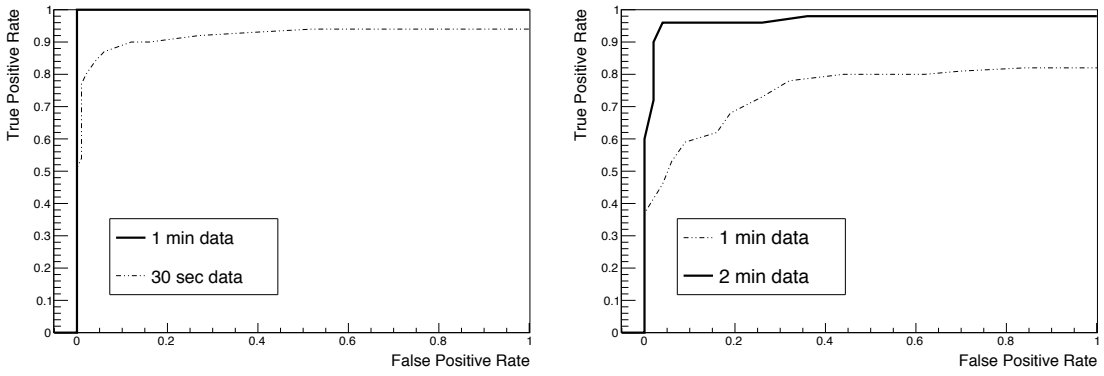

(a) **(b)**

Figure 8. ROC curves for (a) flat and (b) cubic rock scenarios.

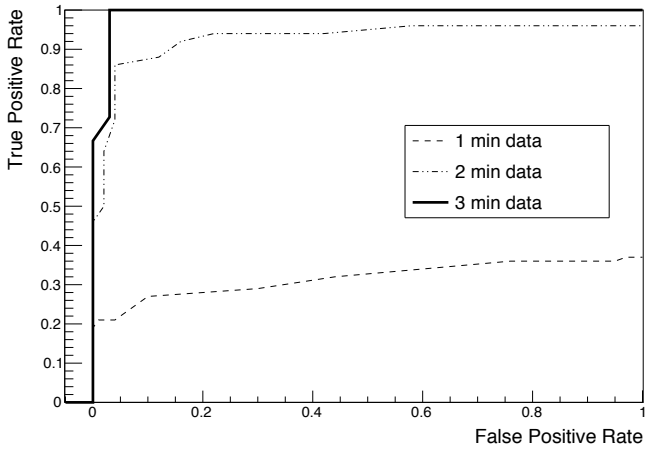

Figure 9. ROC curves for scrap iron scenario.

configuration. It can be seen from the ROC curves in figure 9 that the algorithm requires 3 min of data to reliably identify the uranium block.

4.3 Shielded U spheres

A more realistic way of hiding uranium could be dividing the block into smaller volumes. We simulated this scenario by dividing the 10 cm side length cube into 36 separate spheres with 1.9 cm radius. The spheres were positioned in a $4\times3\times3$ grid configuration with $\sim$ 1 cm spacing and embedded in flat rock (figure 10a) and scrap iron (figure 10b) as in the above scenarios. Similarly, the rock scenario is cleared in a minute, while 3 minutes of data are required for the scrap iron scenario.

4.4 Size comparison

In this section, we look at uranium blocks of different sizes in the same shielding material, namely the flat rock and scrap iron geometries described above. The uranium blocks are cubes of 5, 7.5, 10, and 12.5 cm side length.

2013 JINST 8 P10013

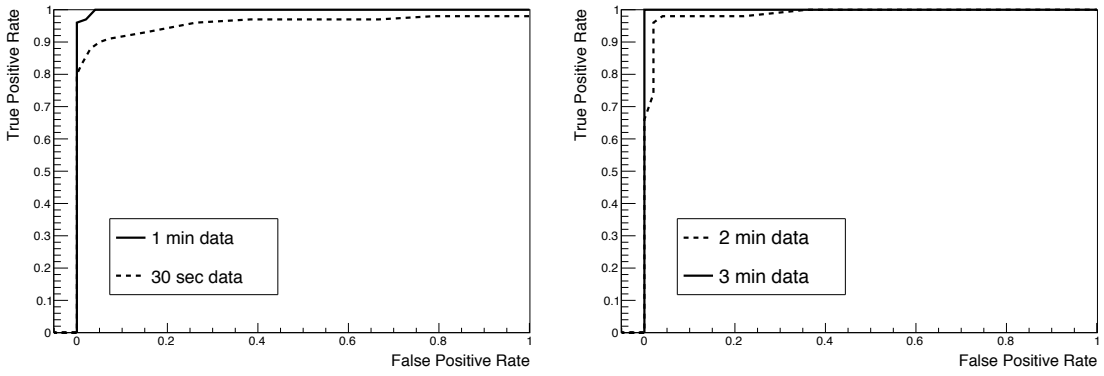

(a) **(b)**

Figure 10. ROC curves for U spheres in (a) flat rock and (b) scrap iron scenarios.

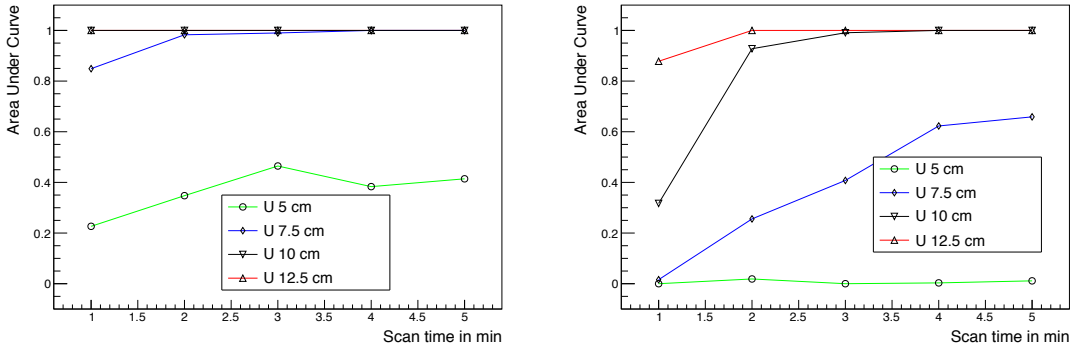

(a) **(b)**

Figure 11. AUC for various U block sizes in (a) rock and (b) scrap iron. Note that in (a) the red (12.5 cm) and black (10 cm) lines overlap.

A common metric used to compare different ROC curves is the AUC (Area Under Curve), i.e. the area enclosed by the ROC curve. The AUC ranges between 0 and 1, where 1 means perfect separation. Figure 11a shows the AUC values for the flat rock scenario, and figure 11b for the scrap iron scenario.

The smallest target size that can still be resolved within a few minutes depends on the number of muons actually passing through the target. To illustrate, figures 12a and 12b show the vertices from 'signal' tracks (muons which we know passed through the target from the MC truth) for 1 minute of data, for the small (5 cm side length) and large (12.5 cm) U targets in scrap iron. It can be seen that a) there are only a couple of muons actually traversing the small U block and b) most of the vertices are reconstructed outside of the true position of the block, due to the scattering occuring in the scrap iron.

For the analysis, this means that the weighted metric distribution is strongly impacted by the scattering in the large block (figure 13a), whereas the contribution to the metric distribution of the small block is basically negligible (figure 13b). On the other hand, it can be seen from figure 12a

2013 JINST 8 P10013

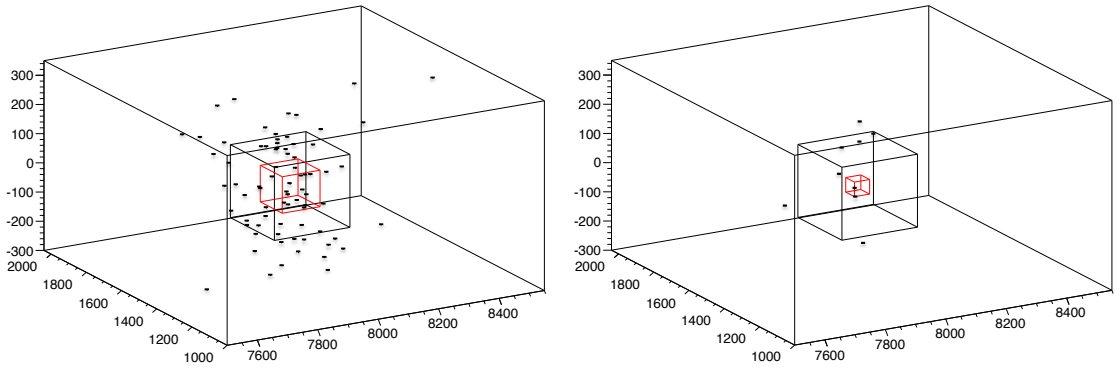

(a) **(b)**

Figure 12. Vertices for 'signal' muon tracks for the (a) large and (b) small U block in scrap iron (1 min of data). The outer black box indicates the size of one grid sub-volume, and the inner red box the size and position of the U target.

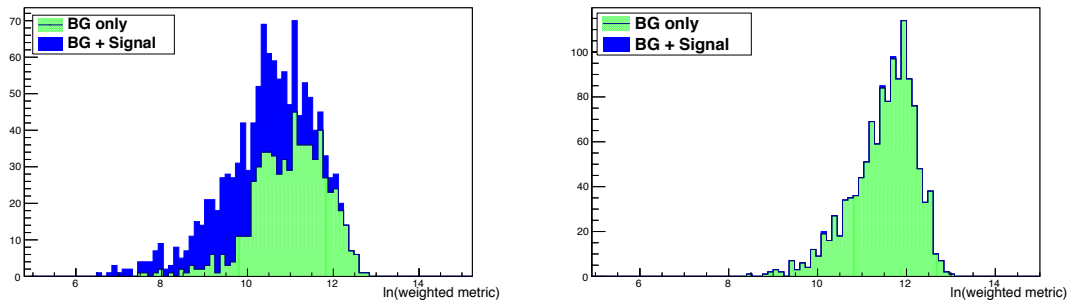

(a) **(b)**

Figure 13. Weighted metric distribution for the grid sub-volume containing the (a) large and (b) small U block in scrap iron (1 min of data). The green/light curve shows the contribution of track combinations where neither track went through the U block, while the blue/dark curve contains background and signal tracks. It can be seen in (b) that the contribution of the signal tracks is virtually invisible.

that even though the vertices of many muon tracks going through the target block are reconstructed outside the uranium volume, they still contribute to the metric distribution in the respective sub-volume, which shows the validity of the assumption of a single scattering vertex when used in conjunction with relatively large sub-volumes.

4.5 Material identification

This section illustrates the capacity of the algorithm to discriminate between materials of similar radiation length. A fundamental issue in MST is the similarity, in terms of muon scattering, of lead and tungsten to uranium. Consequently, MST methods run the risk, especially with low statistics, of being misled by lead or tungsten in the volume of interest. The algorithm was used on volumes containing one of 3 blocks, each $10\times10\times10\,\mathrm{cm}^3$ in size, of lead, tungsten, and uranium.

2013 JINST 8 P10013

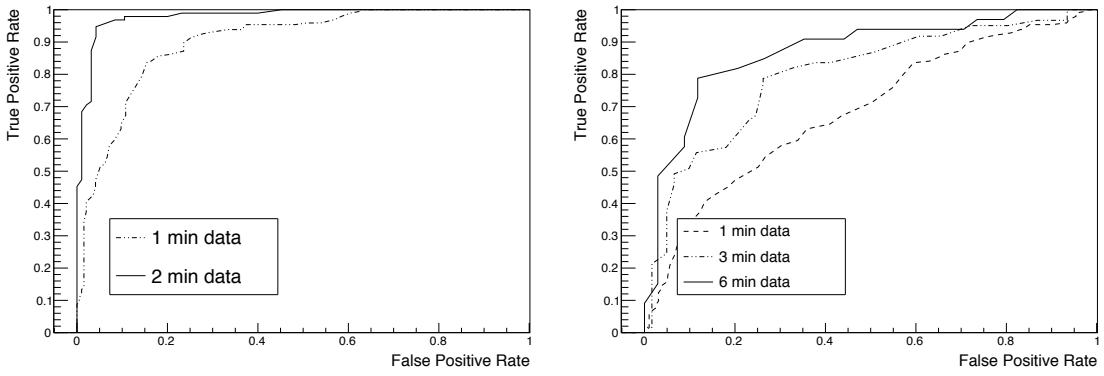

(a) (b)

Figure 14. ROC curves for (a) separation of uranium and lead and (b) separation of uranium and tungsten.

The discriminator was used to estimate whether uranium was present or not, and the separation performance again evaluated with ROC curves.

The separation between uranium ($X_0 \approx 0.32\,\text{cm}$) and lead ($X_0 \approx 0.56\,\text{cm}$) is close to optimal after 2 minutes (figure 14a). However, tungsten ($X_0 \approx 0.35\,\text{cm}$) due to its radiation length being very similar to uranium's, is likely to yield some false positives, even with more than 5 minutes of data-taking (figure 14b).

5 Momentum precision

In this section we investigate the effect of the muon momentum precision on the performance of the algorithm. To that end, we compared results using the MC true momentum, a momentum measurement with 50% uncertainty (similar to [5]), and no momentum information at all.

Figure 15a shows the flat rock scenario with the 10 cm U block, with 1 minute of data (as in figure 8a). It can be seen that for such simple background scenarios the effect of momentum uncertainty is minor, as the ROC curves for the true and 50% momentum are very close together. In fact, even entirely without momentum information, performance is very good after 2 minutes (figure 15b).

Figure 16a shows the scrap iron scenario with the 10 cm U block, with 3 minutes of data (as in figure 9). In this case there is a noticeable difference in performance between the true and 50% momentum information; with two more minutes of data-taking, the 50% scenario regains perfect separation (figure 16b). Without momentum information, it would take longer than 5 minutes to clear this scenario.

6 Discussion

The main limitation of the algorithm, and in fact MST methods in general, is the fact that it tends to be relatively insensitive to sparse distributions of high-Z material. As seen in section 4.4, for uranium targets below a certain size, there is insufficient physical interaction between each individual muon and the target to cause sufficient scattering, and the number of muons hitting the target

2013 JINST 8 P10013

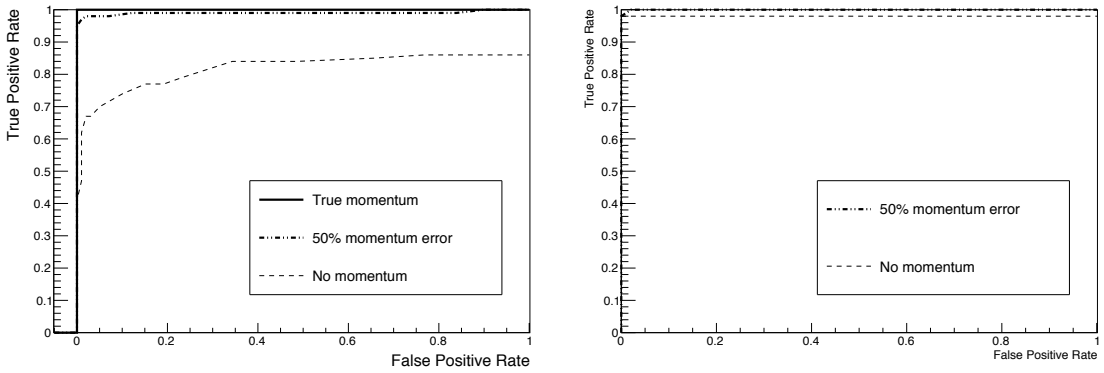

(a) **(b)**

Figure 15. ROC curves for flat rock scenario with varying momentum precision with (a) 1 min of data and (b) 2 min of data.

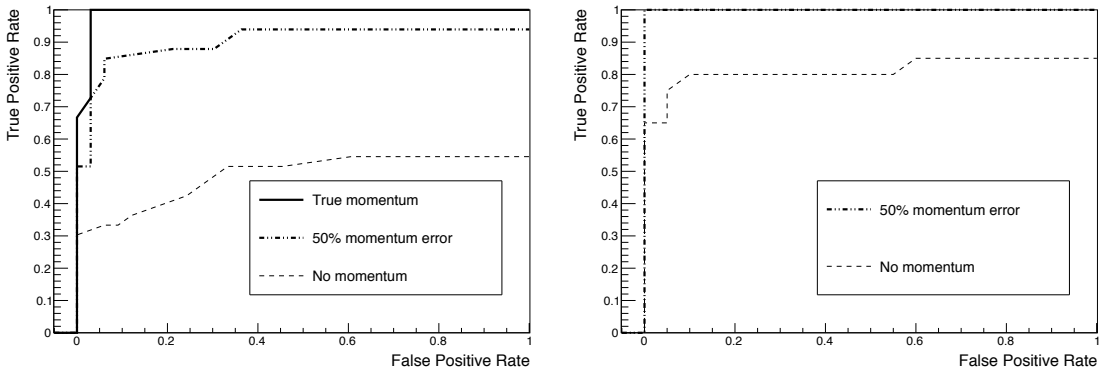

(a) **(b)**

Figure 16. ROC curves for scrap iron scenario with varying momentum precision with (a) 3 min of data and (b) 5 min of data.

in the desired timeframe is too low. The good performance of the method on the uranium spheres scenario (section 4.3) comes from the fact that, since there are multiple layers of spheres in z, there is a good chance that one muon traverses multiple spheres. The same limitation also applies to the potential scenario of a thin uranium sheet. However, in all of these scenarios shielding material (like lead) may be required on the part of the malicious agent to avoid radiation scanners, which would result in more scattering around the uranium, which in turn would cause a stronger signal.

In practical usage at e.g. a freight harbour, one would want to clear as many containers as possible within a minute, so as not to hold up the unloading queue. With the proposed method a decision between 'clear' and 'need more data' can always be made within a minute: due to the algorithm's short runtime, almost no time would be lost after data-taking. Furthermore, after a minute one would either have a conclusive negative, or for more densely loaded containers, multiple sub-volumes would show relatively low discriminators, which would make the necessity for a longer scan apparent. This is an important aspect, since this way a significant number of containers

2013 JINST 8 P10013

would require only the set minute without risking false negatives (even a single of which would evidently not be acceptable). With multiple scanning stations, the total time loss by containers that need to be scanned longer than a minute could be kept at a minimum.

7 Conclusions

We have presented the performance of an algorithm capable of identifying the presence of a block of high-Z material, such as uranium, in various large-scale configurations. The algorithm was shown to be fast (runtime of ~ 10 sec on a desktop computer for 1 min of cosmic muons) and efficient in a number of approximations of real-life situations. A $10\times10\times10\,\mathrm{cm}^3$ uranium block was shown to be identified correctly in every case within the desired timeframe of 1 minute, both on its own and shielded by commonly encountered materials, such as rock. In cases of more problematic shielding material (rock in a vertical pile-up configuration, scrap iron), longer scan times were required, but the algorithm was shown to regain optimal separation after a maximum of 2 additional minutes. An adapted version of the algorithm was also shown to be able to identify a solitary $9\times9\times9\,\mathrm{cm}^3$ block of tungsten in ~ 1 minute of real data from a small-scale prototype RPC system. The algorithm was also shown to be capable of distinguishing an equivalent amount of lead from uranium in ~ 2 minutes, and offered some separation for tungsten and uranium. These results were obtained from simulations using RPCs with a realistic intrinsic resolution of 450 µm. We have also investigated the effect of momentum uncertainty and found that for a 50% momentum error the algorithm required additional scanning time only for the more problematic shielding scenarios. We are looking forward to testing the method on real data obtained from a full-size prototype that is currently under construction.

References

[1] PARTICLE DATA GROUP collaboration, J. Beringer et al., *Review of Particle Physics (RPP)*, *Phys. Rev.* **D 86** (2012) 010001.

[2] B. Rastin, *An accurate measurement of the sea level muon spectrum within the range 4-GeV/C TO 3000-GeV/C*, *J. Phys.* **G 10** (1984) 1609 [INSPIRE].

[3] C. Hagmann, D. Lange and D. Wright, *Cosmic-ray shower generator (CRY) for Monte Carlo transport codes*, IEEE Nucl. Sci. Conf. Ser. **2** (2007) 1143.

[4] K. Borozdin et al., *Cosmic-ray muon tomography and its application to the detection of high-Z materials*, in proceedings of *46th Annual Meeting*, Institute of Nuclear Materials Management, Phoenix, AZ, U.S.A. (2005).

[5] C.L. Morris et al., *Tomographic imaging with cosmic ray muons*, *Science & Global Security: the Technical Basis for Arms Control, Disarmament, and Nonproliferation Initiatives* **16** (2008) 37.

[6] S. Pesente et al., *First results on material identification and imaging with a large-volume muon tomography prototype*, in proceedings of *First International Conference on Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA)*, Marseille, France (2009), pp.1-4.

[7] L.J. Schultz et al., *Statistical reconstruction for cosmic ray muon tomography*, *IEEE T. Image Process.* **16** (2007) 1985.

2013 JINST 8 P10013

[8] P. Baesso et al., *A high resolution resistive plate chamber tracking system developed for cosmic ray muon tomography*, 2013 *JINST* **8** P08006.

[9] GEANT4 collaboration, S. Agostinelli et al., *GEANT4: A Simulation toolkit*, *Nucl. Instrum. Meth.* **A 506** (2003) 250.

[10] F. James and M. Roos, *Minuit: A system for function minimization and analysis of the parameter errors and correlations*, CERN-DD-75-20 (1975) [*Comput. Phys. Commun.* **10** (1975) 343].

[11] L. Schultz et al., *Image reconstruction and material z discrimination via cosmic ray muon radiography*, *Nucl. Instrum. Meth.* **A 519** (2004) 687.

[12] T. Fawcett, *An introduction to ROC analysis*, *Pattern Recogn. Lett.* **27** (2006) 861.

2013 JINST 8 P10013