Author:
Frazao, Leonor

Title:
Imaging and Material Identification of Nuclear Waste with Muon Scattering Tomography

University of BRISTOL

Imaging and Material Identification of Nuclear Waste with Muon Scattering Tomography

Maria Leonor Trigo Franco Frazão

A dissertation submitted to the University of Bristol in accordance with the requirements of the degree of PhD in the Faculty of Science, School of Physics, Bristol.

April 2018

26,271 words

Abstract

Nuclear waste currently created and placed in interim storage is expected to be fully known and characterised, with records kept and regularly checked, per IAEA regulations. However, there is also historical nuclear waste that was created at a time when these records were not required and presumably the problem of nuclear waste disposal was not such a concern as it is today. Not only the initial materials stored may not be known, but they may also have undergone changes, such as the oxidation of uranium that produces hydrogen gas. This brings a demand for techniques to characterise nuclear waste that both make sure that its records are current, and that historical waste can also be fully characterised and properly taken for final disposal or long term storage. Muon scattering tomography is a technique that can be used for this purpose. It consists of measuring individual cosmic-ray muons before and after they cross the volume of interest, and obtaining the angular distribution of their scatter and related variables. The width of the angular distribution is larger for materials with higher atomic numbers, so it allows for high-Z materials to be found in concrete. Several methods were developed that, when combined, can give a description of the contents of nuclear waste. This starts with an imaging algorithm that can first find lumps of high-Z materials and then detect the edges of these materials with a good precision. The same algorithm can also be used to determine the amount of gas present in the containers. This thesis shows that this algorithm can measure the length of uranium blocks in concrete with a resolution of 3.2 ± 0.6 mm when not using momentum information, for lengths down to 5 mm. A resolution of 0.98 ± 0.03 mm was obtained when including the muon momentum, for lengths down to 2 mm. In a following step, high-Z materials can be identified to verify if they come from nuclear fuel (uranium and plutonium), or if they are other materials such as lead and tungsten. It is shown that the distinction between uranium and lead or tungsten is possible for block sizes down to a cube of 2 cm side, with data taking times up to 70 hours. Some discrimination between uranium and plutonium was also obtained, for 3 cm side cubes and requiring more data, corresponding to 200 hours.

Acknowledgements

First and foremost I would like to thank my supervisor Jaap Velthuis for his guidance, support, time and effort through the development of this project and the writing of this thesis. I would also like to thank Chris Steer, my supervisor from AWE, for all his help and regular visits to Bristol. I would like to thank Christian Thomay for always answering when I asked any questions, as well as for his help with working with the existing software and implementing my own. Thanks to Sam and Magda for their collaboration in parts of this work. I want to acknowledge my colleagues from the Particle Physics Group at the University of Bristol for being a great group to work with, in particular Keith Clark and Paolo Baesso for making everything work in the lab, Chiara and Emma for their friendship, company and help in many discussions, and Martisse for her constant availability and cheerfulness. Thanks to my flatmates: Lana, who I looked up to all this time, for being the best company at home for three years and for being always there for anything work and non-work related, and Karina, for her company in the last few months. Lastly I would like to thank my family for always supporting me, in particular my parents for their endless encouragement and my brother Miguel for always bringing joy and laughs to our home.

Author's Declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

Signed

Date

Declaration

Part of the work described in this thesis has been developed in collaboration with other students. The metric distance algorithm described in Section 4.1 was developed by Christian Thomay. Part of the analysis described in Section 4.2 was developed in collaboration with Samuel Madrell-Mander. The analysis described in Section 4.3 was developed in collaboration with Magdalena Dobrowolska (Warsaw University of Technology). The analysis described in Chapter 5 was developed and implemented by the author.

Contents

List of Figures

List of Tables

1 Introduction

Nuclear waste needs to be characterised in order to be disposed of in the correct manner. It is required nowadays to keep records of the nuclear waste from its creation, and to monitor it regularly [1]. However, there is historical nuclear waste that lacks this information. This can be either legacy waste, or even current waste that has not been tracked according to the International Atomic Energy Agency (IAEA) regulations. Some other challenges to nuclear waste characterisation include the existence of shielded waste, the requirement to perform non-destructive assays in many cases, and the changes that the waste goes through, such as corrosion, breaches and overflow [2]. This characterisation is essential for the correct disposal or long term storage of the nuclear waste, where making the wrong classification can mean either creating dangerous waste, or needlessly increasing the cost by orders of magnitude.

For intermediate or low level waste, it is necessary to determine if there are any pieces of nuclear fuel, such as uranium or plutonium present. For high level waste, it is important to ascertain if there is any gas formed by the oxidation of uranium, and if the gas is in separate small bubbles or concentrated in big bubbles.

Some of this nuclear waste is encased in concrete, and this thesis proposes scanning this type of waste using muon scattering tomography. There are several other techniques that can be used to scan nuclear waste encased in concrete, which can be classified into passive and active scanning. Passive scanning consists mainly of measuring emissions (such as radiation or heat) from the materials present in nuclear waste. Uranium and plutonium emit radiation from nuclear decays, so neutrons and gammas can be measured in order to determine the presence of these materials. However, the concrete and other materials present in the containers can absorb both neutrons and gammas, in particular in large containers, preventing the identification of the radioactive materials. Uranium and plutonium also emit alpha particles, which are absorbed within a few millimetres of concrete and hence

cannot be detected outside. Another passive scanning technique is calorimetry, in which the total radiation rate can be estimated from how much it heats the sample. Active scanning consists of introducing external particle beams and measuring either the same beam on the opposite side of the sample, and thus inferring its scatter and absorption, or the radiation induced by the source. This can be achieved with X-rays, where high energy photons are introduced on one side and detected on the other side. Materials with a higher density will absorb more photons, and so can be distinguished against the lower density concrete. Neutron beams can also be used, by inducing nuclear fission of uranium and plutonium, which generates gammas with characteristic energies that can be used to identify the fissioned material and measure its amount. However, active scanning methods can be costly, since they need not only detectors but also a radiation source.

The method described in this thesis is muon scattering tomography, which is considered a passive scanning method, as it does not need an artificial source of particles. Instead, it uses cosmic-ray muons as probes. Muons are naturally occurring in the atmosphere, as a product of cosmic rays. They are highly penetrating particles, crossing dense materials and large volumes. Therefore very few will be absorbed by any materials in nuclear waste.

Muons undergo multiple Coulomb scattering in matter, which makes them change direction. The width of the angular distributions depends on the atomic number, so we can distinguish materials with higher or lower atomic number. More details on this technique will be given in Chapter 2.

It has been shown in several sources that it is possible to detect the presence of high-Z materials embedded in concrete, but it is more difficult to distinguish uranium and plutonium from other high-Z materials such as lead and tungsten, that could also be present in nuclear waste. This thesis will show a method that uses variables obtained from muon tomography and distinguishes pairs of these materials. It will then be possible to identify which high-Z material is present in a real container by comparing it to different cases and finding the most similar one.

This thesis is structured as follows: the theory behind and the state of the art of muon scattering tomography is discussed in Chapter 2. Chapter 3 describes the methods used in the simulations and analyses of the subsequent chapters. The imaging algorithm, previous results and new edge finding technique are given in Chapter 4. The material discrimination method and respective results are described in Chapter 5. Finally, conclusions are given in Chapter 6.

2 | Theory and state of the art

This chapter gives an overview of the origin and characteristics of cosmic-ray muons, as well as their interaction via multiple Coulomb scattering and absorption. It will then explain how these interactions can be exploited for passive scanning of large objects using muon scattering tomography. It will give some examples of potential applications of this technique, and address in more detail the motivation to use it for nuclear waste monitoring. Finally, it describes the muon detectors (resistive plate chambers) used in the University of Bristol particle physics laboratory, and it gives an overview of the different ways to measure muon momentum.

2.1 Cosmic rays

Cosmic rays are high energy particles produced by astrophysical sources and are for the most part protons, neutrinos, electrons, photons and alphas, although some heavier nuclei are also present.

When cosmic rays reach the Earth, they interact with the atmosphere, creating new particles, which interact with the atmospheric molecules, generating showers of particles. Muons are produced by the decay of charged pions, which decay predominantly into a muon and a neutrino ($\pi^+ \to \mu^+ + \nu_\mu$, $\pi^- \to \mu^- + \bar{\nu}_\mu$).

2.1.1 Cosmic-ray muons

Muons are leptons with a mass of $105.6583745 \pm 0.0000024$ MeV [3]. They are produced high in the atmosphere, at an altitude of around 15 km, and deposit around 2 GeV of their energy in the air before they reach the ground. Muons at ground level have a mean energy of ~ 4 GeV and their angular distribution is approximately proportional to $\cos^2\theta$, where θ is the zenith angle (angle between the muon trajectory and the vertical line) [3]. Low energy muons at higher angles are more likely to decay before reaching the surface, so there are fewer low energy muons at high zenith angles. This can be seen in Figure 2.1, which shows the muon spectrum at sea level for muons with zenith angles of 0° and 75°.

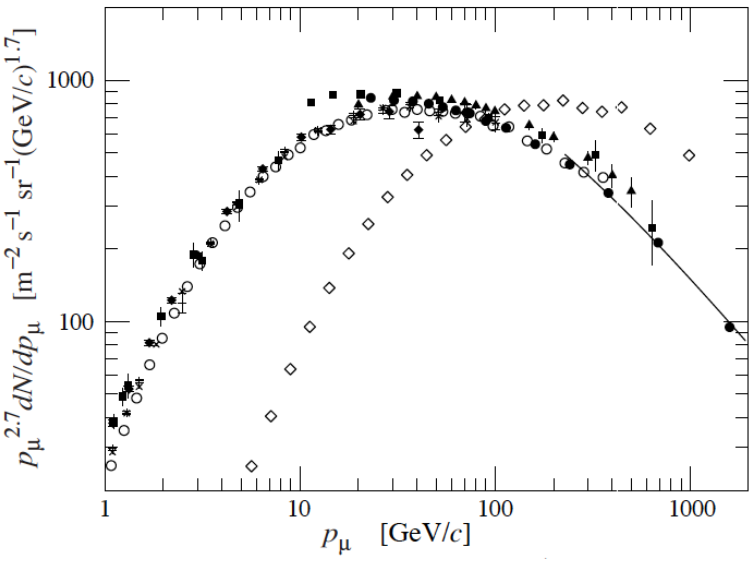

Figure 2.1: Muon spectrum, weighted by the momentum, at zenith angles of 0° and 75° [3]. The open diamonds are for $\theta = 75°$ and all the other points are for $\theta = 0°$.

2.1.2 Interaction of muons with matter

When high energy charged particles cross matter, they interact with the atoms present and transfer energy. The mean rate of energy loss in MeV g^{-1} cm^2, or stopping power, of muons and other relativistic charged heavy particles is described by the Bethe-Bloch equation as [3]

$$\left\langle -\frac{\mathrm{d}E}{\mathrm{d}x} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \tag{2.1}$$

where K is $4\pi N_A r_e^2 m_e c^2 = 0.307075$ MeV mol^{-1} cm^2 (N_A is Avogadro's number, r_e^2 is the electron radius and $m_e c^2$ is the electron mass in MeV), z is the charge number, Z is the atomic number, A is the mass number, I is the mean excitation energy in eV, $W_{\max}$ is the maximum energy transfer to an electron in a single collision and $\delta(\beta\gamma)$ is the density effect correction to ionization energy loss. This equation describes the mean rate of energy loss with an accuracy of a few percent, for particles in the range $0.1 \lesssim \beta\gamma \lesssim 1000$. Figure 2.2 shows the stopping power of muons in copper as a function of their momentum. In the momentum range of most cosmic-ray muons, ionisation is close to the minimum, so cosmic-ray muons can be considered minimum ionising particles. Since the mass of the muons is about

207 times larger than the electron's mass, the energy loss due to Bremsstrahlung by muons is small compared to the energy loss by electrons. This is because the Bremsstrahlung cross-section is proportional to the radius, which is inversely proportional to the square of the mass [3, 4]. For the electron the cross-section is

$$\sigma \propto r_e^2 = \left(\frac{e^2}{4\pi\epsilon_0 m_e c^2} \right)^2 . \tag{2.2}$$

Therefore, cosmic-ray muons lose much less energy in matter than electrons, so they can travel longer distances before losing all their energy and being absorbed. The mean life time of muons is $\tau \simeq 2.2 \times 10^{-6}$ s [3]. Due to their relativistic speed, the distance a muon would travel in this time can be considered to be about $c\tau \simeq 660$ m. But, also because they are relativistic, the distance the muons travel in the Earth's reference frame is even longer. In summary, cosmic-ray muons deposit little energy in matter and are able to travel long distances before decaying; because of these two properties, they are highly penetrating. This can be seen in Figure 2.3, which shows the vertical muon intensity as a function of the depth in rock, measured in km water equivalent, where 1 km.w.e. $= 10^5$ g cm^{-2} of standard rock (km.w.e. is the product between depth and density). The inset shows the vertical muon intensity for water and ice.

Muon scattering tomography (MST) exploits the fact that muons scatter in matter. More details on MST will be given in Section 2.2. Charged particles, like muons, interact with the nuclei present in matter via Coulomb's law, which results in elastic scattering. If the volume crossed is thick enough so that there are many independent scatterings (called multiple Coulomb scattering), this interaction can be treated statistically [4]. This results in a random walk, where the particle changes direction many times. A representation of this random walk is shown in Figure 2.4. It is shown therefore that when a muon crosses an object, it can exit with a different direction from that it entered the object. The scattering angle is the angle between the entry and exit vectors. The 2D projected scattering

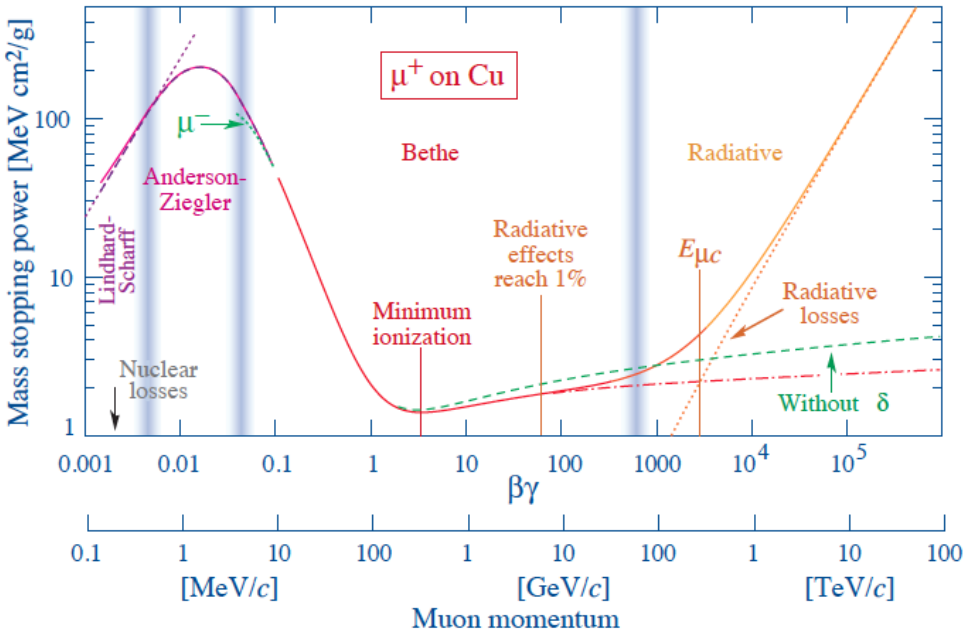

Figure 2.2: Mass stopping power for positive muons in copper as a function of $\beta\gamma = p/Mc$ [3].

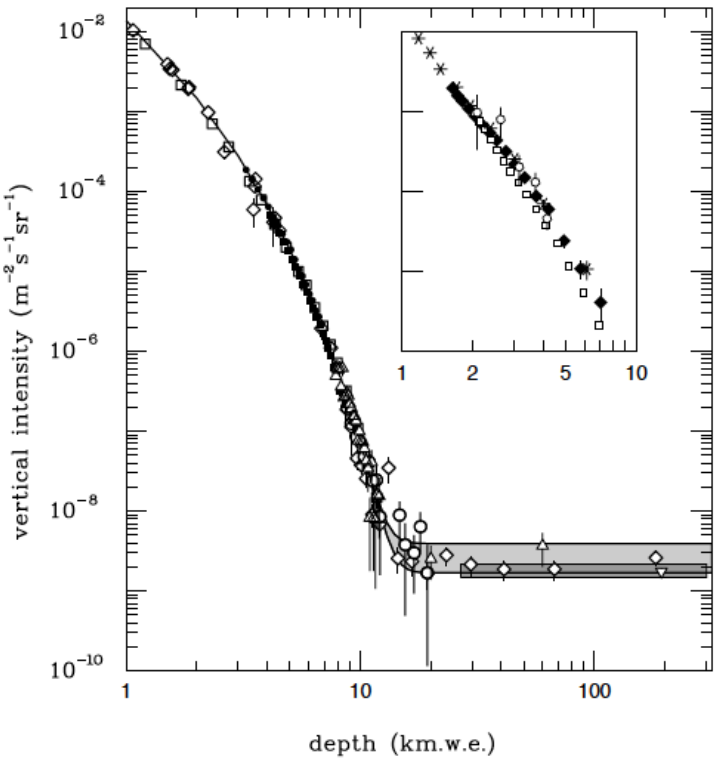

Figure 2.3: Muon flux versus depth in rock, in km water equivalent (1 km.w.e. = 10^5 g cm^{-2} of standard rock) [3]. The inset shows the vertical intensity curve versus depth for water and ice.

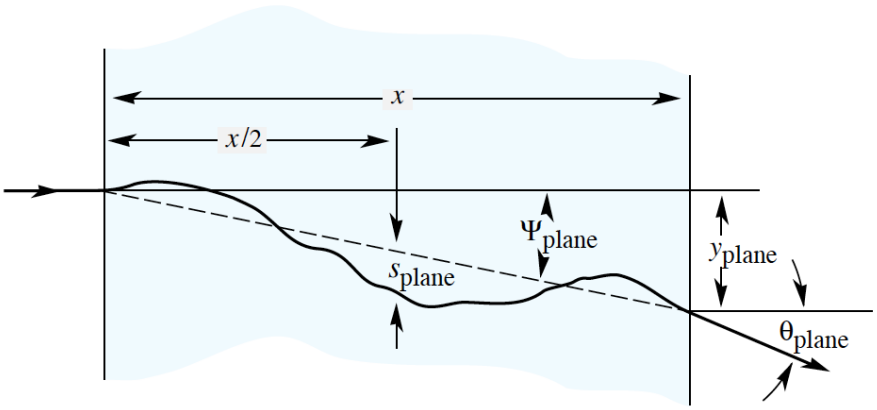

Figure 2.4: Two dimensional representation of the random walk due to multiple Coulomb scattering [3].

angle distribution is approximately Gaussian [5, 6], whose width σ depends on the radiation length of the material traversed,

$$\sigma \approx \frac{13.6 \text{ MeV}}{\beta cp} z\sqrt{X/X_0}(1 + 0.038 \ln(X/X_0)), \tag{2.3}$$

where p is the momentum of the muon, βc its velocity, z its charge number, X the thickness of the material, and X_0 the radiation length. Radiation length is the distance an electron can travel through a material, reducing its energy by a factor of $1/e$. It is a property of materials which depends on the atomic number and density as

$$X_0 \approx \frac{A \cdot 716.4 \text{ g/cm}^2}{\rho \cdot Z(Z+1) \ln(287/\sqrt{Z})}, \tag{2.4}$$

where A is the mass number, Z the atomic number and ρ the density [3]. Table 2.1 shows examples of radiation length values for different elements. It is clear that X_0, in g·cm^{-2}, decreases for an increasing atomic number. However, when weighted by the density to obtain X_0 in length (cm), its value does not always decrease monotonically with higher atomic number.

Element	Z	X_0 (g·cm^{-2})	X_0 (cm)	ρ (g·cm^{-3})
H_2 gas	1	63.05	7.528×10^{5}	8.376×10^{-5}
C	6	42.70	21.35	2.000
O_2 gas	8	34.24	2.571×10^{5}	1.332×10^{-3}
Si	14	21.82	9.370	2.329
Ca	20	16.14	10.42	1.550
Fe	26	13.84	1.757	7.874
Cu	29	12.86	1.436	8.960
W	74	6.76	0.3504	19.30
Pb	82	6.37	0.5612	11.35
U	92	6.00	0.3166	18.95
Pu	94	5.93	0.2989	19.84

Table 2.1: Values of atomic number, radiation lenght and density for different elements [3, 7]. Radiation length is given both in g·cm^{-2} and in cm, being weighed by the density in the latter case. Density of gases is given at 20°C and 1 atm.

2.2 Muon Scattering Tomography

As discussed in Section 2.1.2, the scattering angle distribution of muons depends on the atomic number of the material they traverse. This means that in principle, by measuring these distributions, it is possible to identify those materials. The muon scattering tomography (MST) concept is depicted in Figure 2.5. The incident muons are tracked before and after crossing the volume of interest, by using several detector layers. Then, there are different ways to obtain a scattering angle (and 2D projected scattering angles) from the track fit. For example, this can be done using the point of closest approach, which is the point where the incoming and outgoing tracks are closest to each other. It is possible then to obtain the angles projected in x and y. The algorithms that will be described in Sections 2.2.3.1, 2.2.3.2 and 2.2.3.3 use this method. Alternatively, a vertex can be calculated from the fit of the incoming and outgoing tracks, and both a 3D angle and projected angles obtained. A more detailed explanation of this second method will be given in Chapter 3.

2.2.1 Applications of Muon Tomography

The applications of muon tomography can be divided into two types: muon radiography, that exploits the absorption of muons; muon scattering tomography

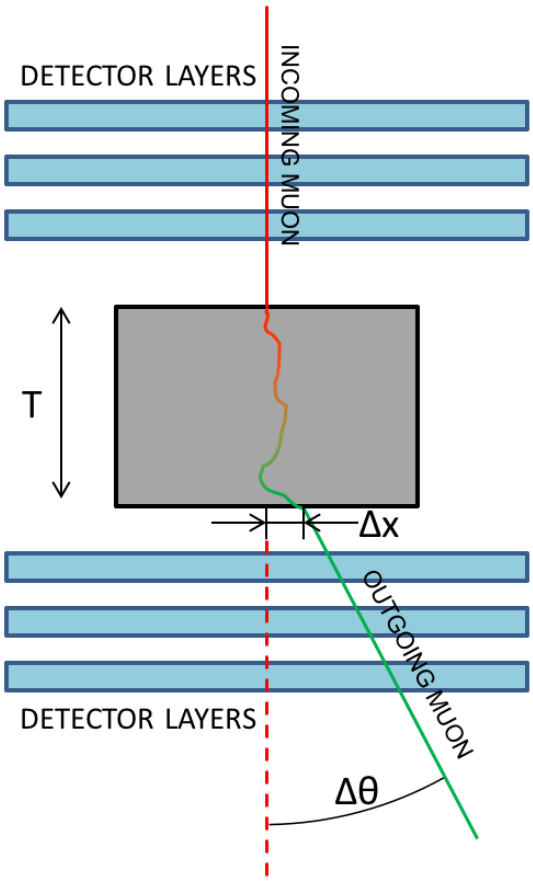

Figure 2.5: Illustration of the muon scattering principle. Each muon is tracked before and after crossing the volume to be scanned. As muons undergo a random walk in matter, a scattering angle $\Delta\theta$ and an offset Δx can be calculated.

(MST), that uses muon scattering. The results shown in this thesis use MST. A key difference between the two techniques is that radiography only requires one detector system to measure the muons that made it through the object under investigation, while in MST it is necessary to measure the incoming and outgoing trajectory of each muon.

2.2.1.1 Muon radiography

It is possible to use muon absorption instead of scattering to perform density maps of large structures, such as pyramids, volcanoes and nuclear reactors. In 1970 results from muon radiography on the Second Pyramid of Chephren, one of the three pyramids of Giza, were published in [8]. Simulations of the pyramid, which assumed no unknown chambers existed, were compared to measurements of muons crossing the pyramid. If there were chambers, more muons would be detected from that area, as fewer muons would be absorbed than if the whole volume were made of rock. No difference was observed, within uncertainties, between the number of muons detected and the expected number of muons from simulations, which proved the absence of unknown chambers. More recently, in November 2017, a study of the Great Pyramid of Cheops showed the presence of a new void which indicates the existence of unknown chambers in this pyramid [9]. Three different detector systems were used in this experiment, where two were placed in one of the known chambers inside the pyramid and another was placed outside the pyramid, close to its base. An example of the results obtained from the detectors at the pyramid base can be seen in Figure 2.6. The 2D figures in the **a** and **d** graphs show the number of muons, in logarithmic scale, measured in two detectors. It is possible to distinguish the shape of the pyramid, with fewer muons, against the sky, with more muons. Horizontal slices taken from the 2D graphs show the excess of muons both from the known (**c** and **f**) and the previously unknown chambers (**b** and **e**). In **g** and **h** the position of the detectors and the triagulation of the position of the chambers are shown. The fact that the position of the known chambers

was correctly identified confirms the hypothesis that the new void is an unknown chamber, or collection of chambers.

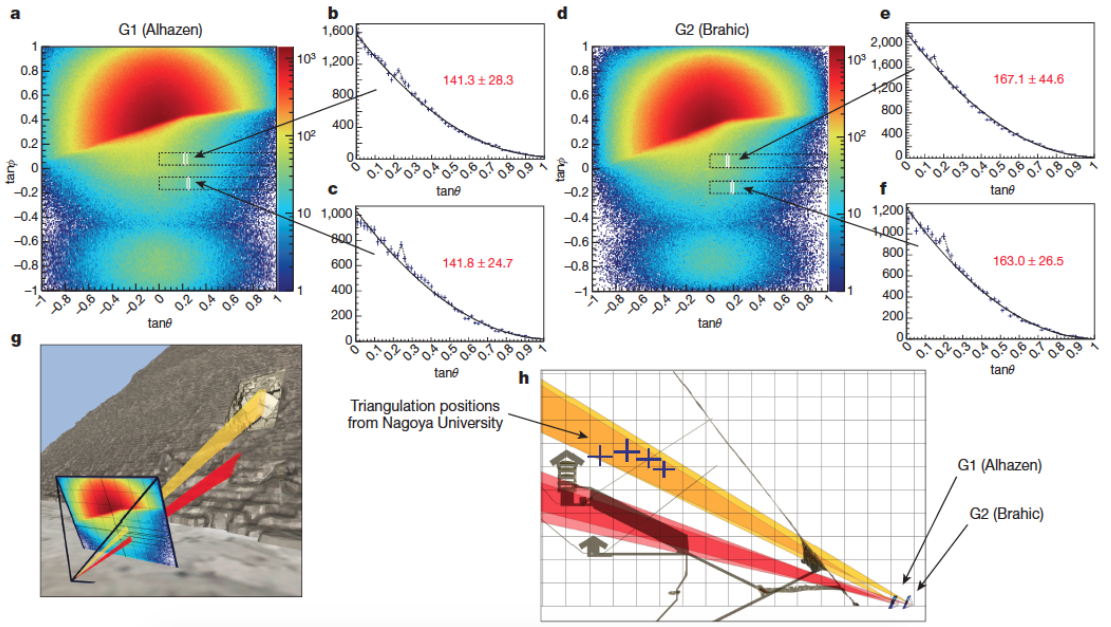

Figure 2.6: Results from muon radiography of the Great Pyramid of Cheops using gas detectors at the base of the pyramid [9]. **a** and **d** show the number of muons, in logarithmic scale, from the two detectors. From these 2D images, horizontal slices were obtain that show an excess of muons: in **c** and **f** it is possible to see the muon excess from the Grand Gallery, and in **b** and **e** from the new void. **g** and **h** illustrate the position of the detectors relative to the pyramid and the triagulation of the position of the chambers. In **h** the Queen's chamber can be seen at the bottom, the King's chamber in the middle, with the new void above it.

As it is possible to use muon radiography to determine the internal structure of pyramids, it is also possible to use it to monitor geological structures, such as volcanoes. The measurement of the density in a volcano can provide a survey of the changes that may occur, such as to the localisation of magma, water and gas, as well as the chemical alterations of the volcanic rock [10]. Figure 2.7 shows the muon radiography concept applied to scan a volcano. In this case, the muons that crossed the volcano are compared to the ones coming from the opposite direction, in order to calculate the attenuation provided by the volcano. In [11], a study is described that uses this technique to obtain a map of the average density distribution along the paths of muons that crossed the Satsuma-Iojima volcano (Mt. Iwodake). The average density projected in the plane parallel to the detectors, calculated from the muons measured in this experiment is shown in Figure 2.8. A region of lower

density is visible inside the volcano. In [10], results are shown from the scanning of the La Soufrière volcano in Guadeloupe over 5 weeks. Regions of lower and higher densities are observed, as shown in Figure 2.9. The low density region RF4 corresponds to a known cave, RF2 is a hydrothermal area and RS3 a dense rock separating these low density regions. Higher density regions of massive lava are also visible: RS2, RS5 and corresponding RF1, and RF5.

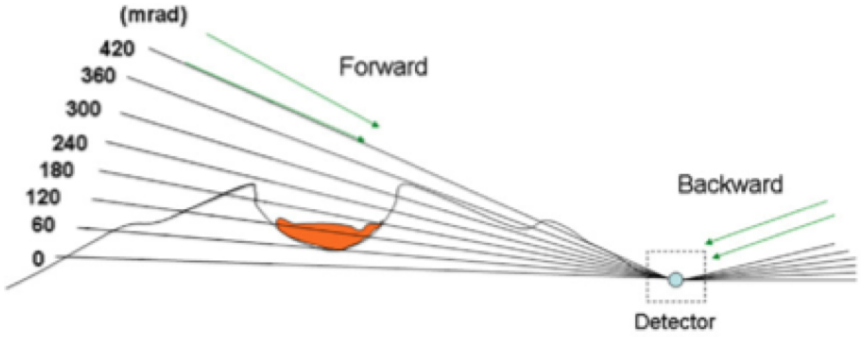

Figure 2.7: Illustration of the muon radiography concept for volcanos [12].

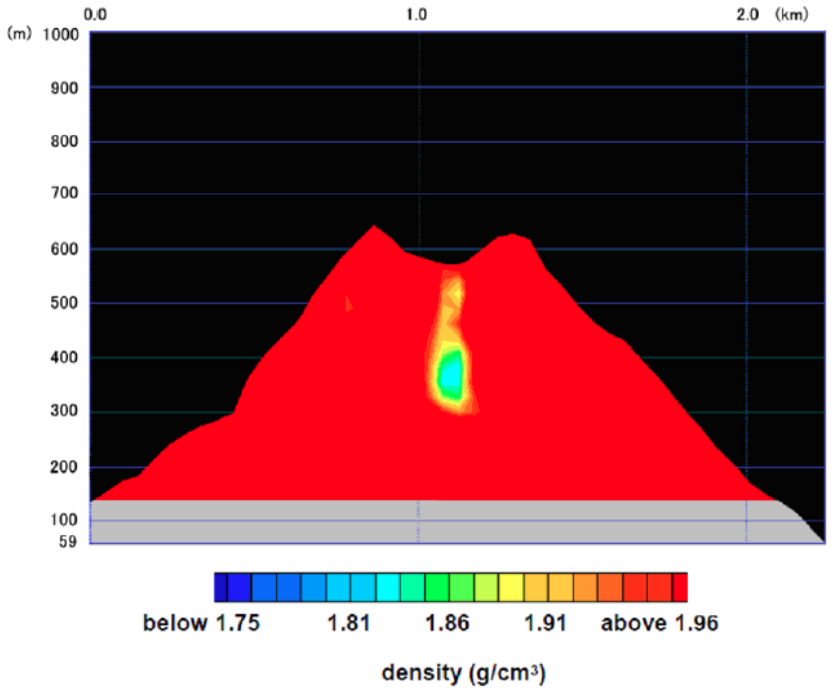

Figure 2.8: Average density distribution of the Satsuma-Iojima volcano with muon radiography [11].

Muon radiography has also been used to perform imaging of the nuclear fuel in the damaged reactors of Fukushima Daiichi. Current publications such as [13]

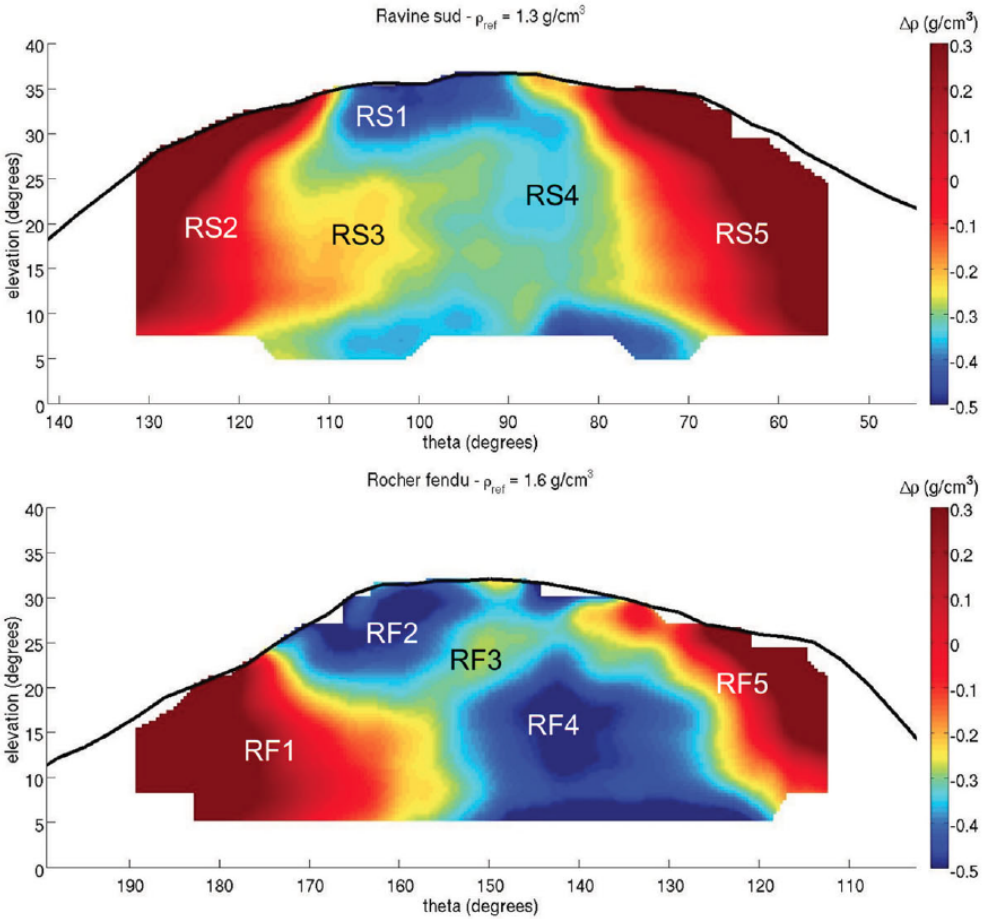

Figure 2.9: Average density relative to a reference density (in title) [10]. Ravine Sud is south of the mountain and Roche Fendue is at the East.

only describe the proposal and validation studies of this possibility, but results from muon detectors at the power plant can be found in handouts from TEPCO (Tokyo Electric Power Company), available online: results from muon radiography performed in Units 1, 2 and 3 reactors are shown in [14], [15] and [16] respectively. Imaging results form Unit 1 reactor indicate that there is very little fuel left in the reactor pressure vessel (RPV). It is assumed then that most of the fuel melted and fell into the primary containment vessel (PVC). Imaging results from Unit 2 reactor can be seen in Figure 2.10, where high density material, believed to be fuel debris, is located at the bottom of the RPV. In Unit 3 reactor, like for Unit 2, some fuel was found at the bottom of the RPV. Part of the fuel of Units 2 and 3 is expected to have fallen into the PVC (more for Unit 3 than 2).

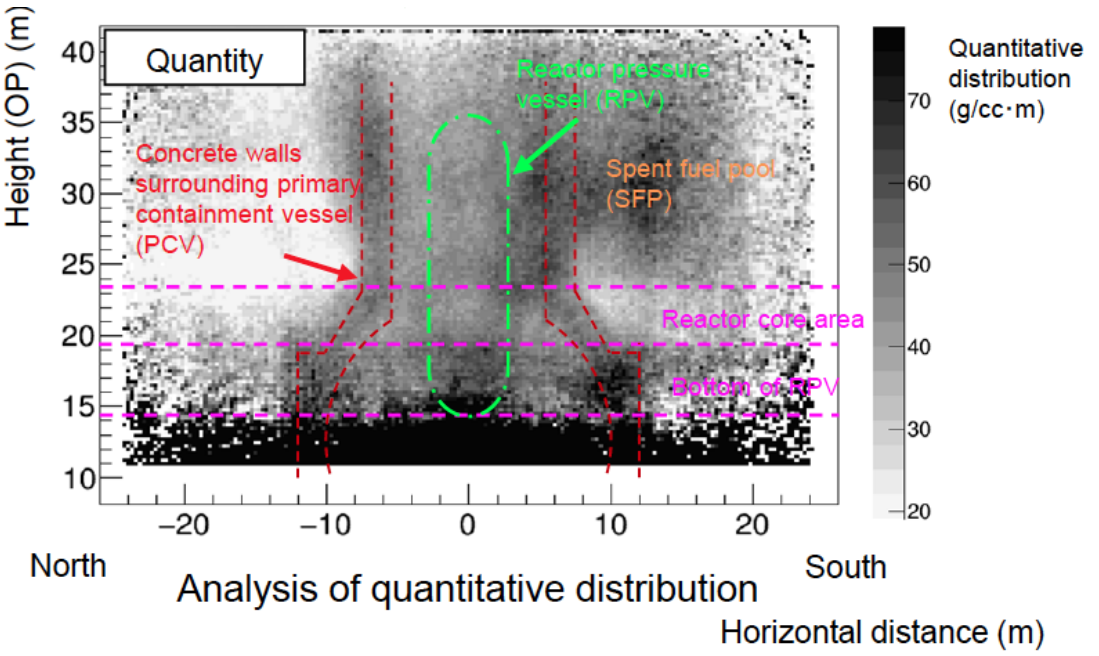

Figure 2.10: Quantitative distribution of material in Unit 2 reactor [15].

In summary, muon radiography is good to perform density maps of large structures, although precision is limited because of their size.

2.2.1.2 Muon scattering tomography

While muon radiography uses only one tracking detector system to measure the muons coming from the volume of interest, muon scattering tomography (MST) needs two tracking systems: one to track the muon before and the other after it

traverses the object to be scanned. This limits the size of the object to the space between both detector systems, and inside their angular acceptance, making MST more suited to smaller objects than muon radiography. Assuming it is possible to place the tracking systems below and above the object of interest, it is possible to use the muons coming from vertical or almost vertical directions.

MST can be used, for example, to scan shipping containers to find smuggled special nuclear material, and to scan nuclear waste and determine its contents.

The first proposal of using MST to scan shipping containers was in 2003 [17]. Several systems and algorithms have been developed, and can be found for example in [18–21].

In [18], simulations were performed of cargo vans, containing a 10 cm cube of tungsten to replicate a high-Z threat object. Different scenarios for the cargo were simulated, with and without the threat object. The scattering density was obtained using the Maximum Likelihood/Expected Maximisation algorithm (ML/EM), which will be described in Section 2.2.3.2. In order to compare the performance of different scanning times, ROC graphs (Receiver Operating Characteristics) can be used, which plot the rate of detection versus the false positive rate. More details on ROC graphs will be given in Chapter 5. The ROC graphs obtained from the average scattering density in the different scenarios, for different data taking times, are shown in Figure 2.11. The identification of the tungsten block is very good after 90 seconds.

In [20], lead and iron blocks with a volume of $10 \times 10 \times 20$ cm^3 were scanned, with a layout shown in Figure 2.12a. Figure 2.12b shows the reconstruction of the scattering density of the blocks, using an iterative tomographic algorithm (ML/EM), with 3 cm side voxels and 40 iterations. It is possible to distinguish the lead blocks, with a larger scattering density, from the iron blocks. It is also clear that the reconstruction has a limited space resolution, especially in the vertical direction.

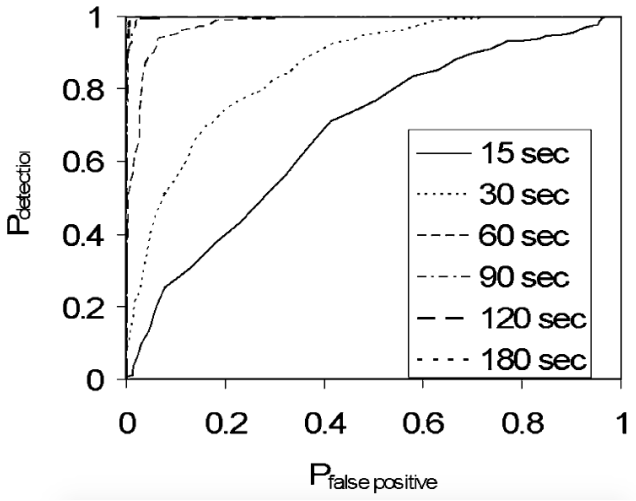

Figure 2.11: ROC graphs for the detection of a 10 cm cube of tungsten in a cargo van, average over different scenarios [18].

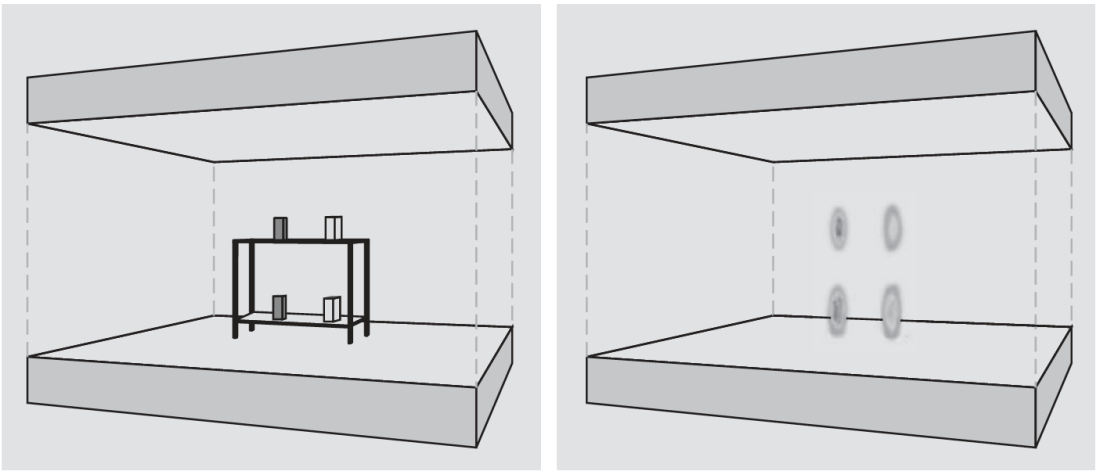

(a) (b)

Figure 2.12: Illustration of the placement of two blocks of lead (dark) and two blocks of iron (light) with a volume of $10 \times 10 \times 20$ cm^3 (a) and image reconstruction of the blocks (b) [20].

An algorithm has been developed at the University of Bristol to use MST for cargo screening [22]. Figure 2.13 shows the ROC graphs for the identification of a simulated $10 \times 10 \times 10$ cm^3 uranium block in a cargo container. In one scenario, it is shielded by rock (1.1 cm thick, above and below the block), where 2 minutes of data are needed to detect the uranium, and in another scenario, it is shielded by scrap iron (filling the container to the weight limit), where 3 minutes of data are needed for identification. These results include momentum information. If this is not available more data would be needed. Smaller uranium blocks will also require more data to be calculated. The area under the curve (AUC) of the ROC graphs from different uranium block sized and scanning times can be seen in Figure 2.14. The AUC is a measure of the classification performance, obtained from ROC graphs, which will be explained in more detail in Section 5.1.1.3.

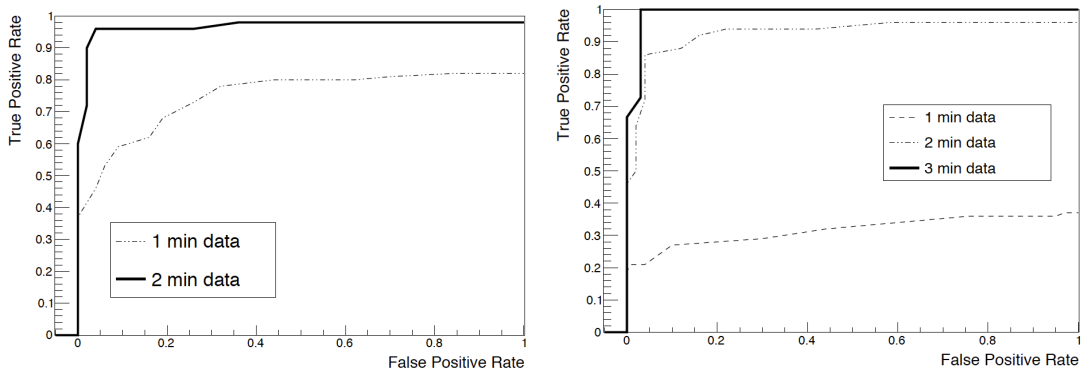

(a) (b)

Figure 2.13: ROC graphs for finding a uranium block in a cargo container, shielded by rock (a) and scrap iron (b) [22].

In a real application for border security, most containers going through are safe. It is important therefore to determine how much data is needed to assure that a container is safe, hence clearing most containers under a minute and taking more data only in the few cases where it is necessary. This can be done by defining the variable 'clear fraction' as the percentage of cases that can be clearly classified as "safe", therefore not containing any special nuclear materials. The clear fractions are shown in Figure 2.15, with a comparison between using the true momentum information, using 50% momentum resolution and not using the momentum [23].

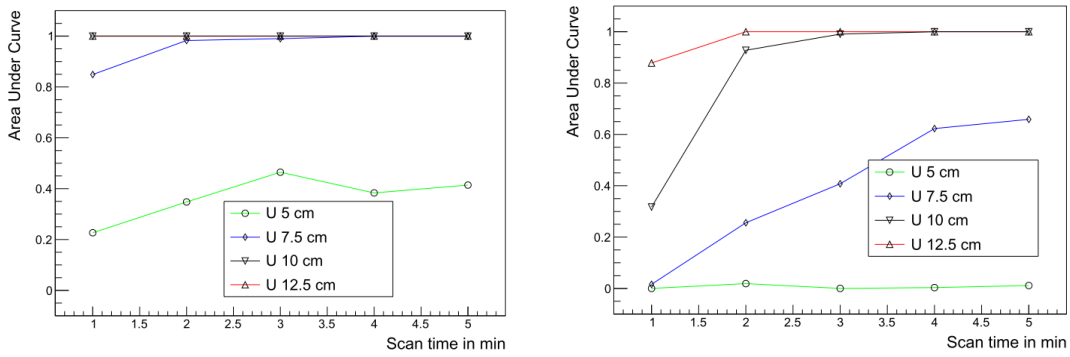

(a) (b)

Figure 2.14: Area under the curve (AUC) of ROC graphs for finding a uranium block in a cargo container, shielded by rock (a) and scrap iron (b) [22].

The value shown is the total sum of clear fractions for different weight categories of the containers, weighted by the frequency of containers in each weight category (up to 0.1 g/cm^3, 0.3 g/cm^3, 0.5 g/cm^3 and 0.7 g/cm^3). Two values are shown: the clear fractions for the highest weight in each category, labelled "clear fraction", and the "average clear fraction" for each category. The same algorithm was shown to allow for imaging. It will be described in Chapter 4, and results from the imaging of simulated nuclear waste containers will be shown.

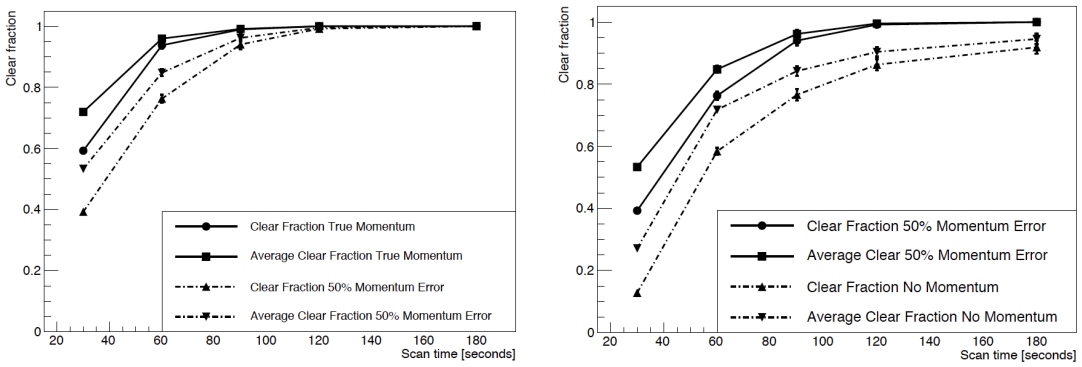

(a) (b)

Figure 2.15: Clear fractions overlapping 50% momentum error to perfect momentum (a) and to no momentum (b) [23].

A commercial system was developed and marketed by Decision Sciences, a spin-off company from the Los Alamos National Laboratory. This system is called Multi-Mode Passive Detection System (MMPDS) and it combines MST

with gamma detection. A test system was operated in Freeport, Bahamas, which can be seen in Figure 2.16.

Figure 2.16: Multi-Mode Passive Detection System from Decision Sciences, in Freeport, Bahamas [24].

It was also proposed that MST can be used for nuclear waste characterisation. This is the main application considered for the studies in this thesis, hence more details will be given in Section 2.2.2. In both cases, MST has advantages over X-ray or neutron scanning, such as the fact that it is a passive method and that muons are more penetrating than X-rays or neutrons, and it is also more easily scalable to scan large volumes.

2.2.2 MST for nuclear waste characterisation

There is a need to characterise nuclear waste in order to meet its disposal requirements. Nuclear waste can be stored in concrete drums, so it is necessary to develop methods to look through these drums from a safe distance and without opening them. This is of particular importance for legacy nuclear waste, that may have been produced several decades ago, potentially contains unknown materials and has possibly undergone changes, such as uranium oxidation, producing hydrogen bubbles. This section attempts to restrict the broad subject of nuclear waste characterisation to the cases where MST can be applied. The assessment of the applicability of a specific method depends on several factors, such as the

form and type of the waste, the existing regulations, the disposal concept and the characteristics to be measured.

Radioactive nuclear waste can be divided into different types, depending on its level of radioactivity. These types are high-level waste (HLW), intermediate level waste (ILW) and low level waste (LLW), very low level waste (VLLW), very short lived waste (VSLW) and exempt waste (EW). The exact definition of each category can differ slightly from country to country, but this division is generally the same. The definitions given here are from the IAEA [25]. An illustration of the classification of nuclear waste from the IAEA is show in Figure 2.17, which illustrates how this classification depends on both the activity and the half-life of the isotopes present in the waste.

HLW consists mainly of spent nuclear fuel and by-products from nuclear fuel reprocessing, and it generates a significant amount of heat. ILW can be filters and residues from treatment of the reactor cooling water, parts of the nuclear reactor and graphite from reactor cores. LLW comprises contaminated materials from the every-day operation or decommissioning of nuclear power plants, such as paper, plastics, protective clothing and scrap metal. VLLW has an even lower activity limit and can include for example soil and rubble. VSLW is composed of isotopes with very short half-lives, that can be stored for decay for up to a few years, after which it can be cleared from regulatory control. EW is waste that can be cleared from regulatory control and subsequently handled as normal waste.

Dry waste is generally stored in cylindrical steel containers, and use a shielding material such as concrete. They have different sizes and structures depending on the needs and regulations for the type of waste. Their sizes vary, and some examples are small containers for LLW and ILW with diameter of 40 cm and height of 50 cm, or diameter of 18 cm, and height of 38 cm. HLW may need a different structure in order to store safely fuel assemblies and allow for their transportation. Examples of containers for intermediate storage of spent fuel are the CASTOR® (cask for storage and transport of radioactive material) and the

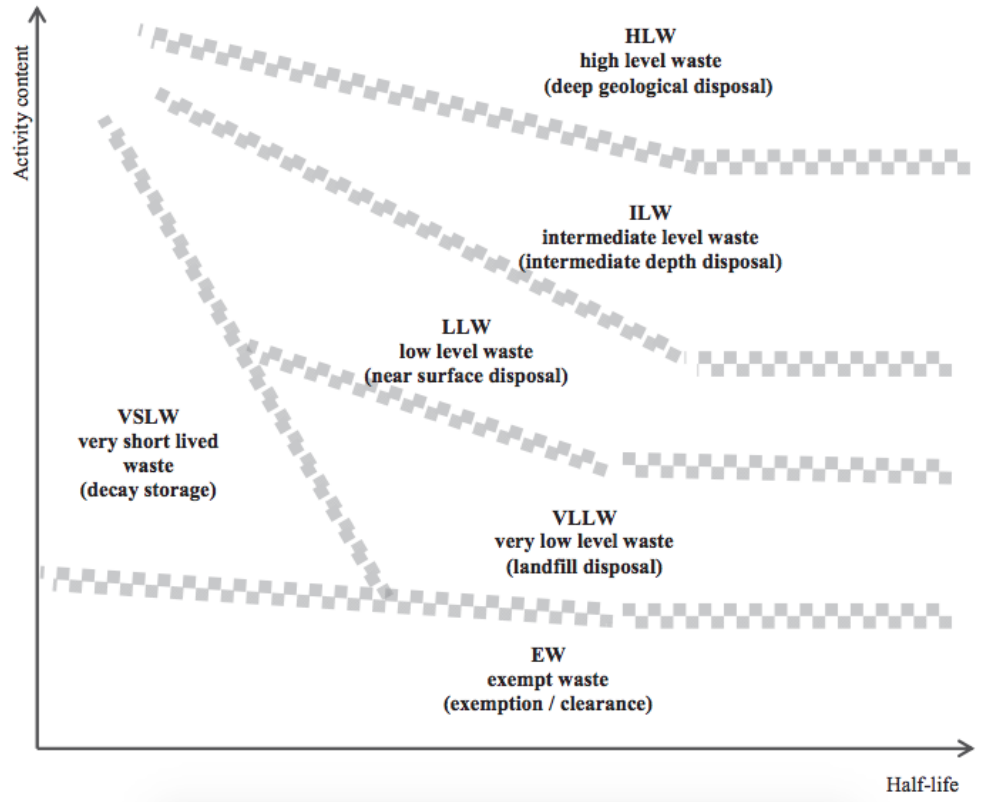

Figure 2.17: Concept of the waste classification scheme from the IAEA [25].

CONSTOR® (concrete storage cask), produced by Gesellschaft für Nuklear-Service (GNS), in Germany. They are 6 m long and about 2 m in diameter, and their purpose is to store nuclear fuel safely, providing radiation shielding, dissipation of heat and protecting it from impact [26]. Figure 2.18 shows some designs of CASTOR®, for different types of fuel assemblies.

MST is most applicable to solid waste, already conditioned, and stored in containers, for LLW, ILW or HLW. It is of particular interest to detect nuclear fuel and other high-Z materials. However, another application to HLW will also be shown, in which MST is used to measure the amount of gas present in a container, that can be formed by spent fuel in HLW.

The methodologies used for characterising nuclear waste can be non-destructive examination (NDE), which is the examination of physical and mechanical properties, non-destructive assay (NDA), which involves measuring radiation from the waste (either spontaneous or induced), and destructive assay (DA), which involves sampling the waste and chemically treating the samples in order to determine their

Figure 2.18: CASTOR® designs for different types of fuel assemblies and HLW: for fuel from pressurised water reactors (PWR) on the left, from boiling water reactors (BWR) in the middle, and high-active waste (HAW) generated from the reprocessing of spent fuel on the right [26].

composition and activity [1]. In this section, only NDA methods will be mentioned, since MST belongs to this category. In particular, NDA can be divided between passive and active scanning methods.

Active scanning methods probe the samples with artificially introduced particles using radiation sources or particle beams. X-ray radiography is an active scanning method, which uses a source of high-energy photons and a detector on the opposite side to measure the photons that are were able to cross the sample. This allows an image to be created, since the absorption of X-rays depends on the density and thickness of the material. So, for example, high-Z materials can be seen in concrete. In addition, neutrons can induce nuclear fission in uranium and plutonium. The resulting emitted gamma rays have specific energies depending on which nucleus fissioned. The main advantages of active neutron radiography over the other techniques are that it allows for the identification of the isotopes present in nuclear fuel and for the measurement of their amount. The general disadvantages of active scanning methods are their high cost compared to passive methods and the additional hazard created by the use of an external radiation source.

Passive scanning methods do not introduce artificial radiation sources. In most cases, passive methods use detectors to measure radiation or heat naturally emitted by the sample. Spontaneous decay of uranium and plutonium results in the emission of neutrons and gammas, which can then be detected and used to find these materials. The alpha particles emitted in these decays cannot be used in the same way, because they are absorbed within a few millimetres from where they are produced. The main disadvantage of this method is that some of these neutrons and gammas can be absorbed by the concrete, especially if the container is large. This would result in a signal dependent on the unknown location of the radiation sources, and thus introduce a larger uncertainty in the isotope identification. Calorimetry is an example of a passive scanning method, which measures the heat produced by the sample. The heat is due to the radiation originating from all the sources present in the sample, and it can be used to estimate the total radiation

rate. It is used in waste that produces high levels of radiation, such as HLW and, in some cases, ILW.

MST uses the naturally occurring cosmic-ray muons as scanning probes. However, as this is not an artificial source, MST can be considered a passive method. It offers a good solution for characterisation of solid nuclear waste, because it is sensitive to different atomic numbers, so it can both distinguish high-Z materials such as uranium, plutonium, lead and tungsten, and low-Z materials such as hydrogen pockets, from the concrete background, as will be shown in Chapters 4 and 5.

2.2.3 Algorithms for muon scattering tomography

Several algorithms have been developed to perform imaging of high-Z materials. The most common algorithms found in literature will be discussed here: the Point of Closest Approach, the Maximum Likelihood/Expected Maximisation and the Scattering Density Estimation.

2.2.3.1 PoCA algorithm

The PoCA (point of closest approach) algorithm was the first one to be developed for MST, and is described in detail in [27]. The principle of this algorithm is shown in Figure 2.19, which is in 2D for simplicity, although in practice it is applied to 3D. The volume is divided in voxels. As the scattering angle tends to be small, the path of the muon is at first considered to be a straight line connecting the entry and exit points, and the voxels through which the line passes are selected. It is assumed that there was only one scattering point, which is chosen to be the point of closest approach of the extrapolated incoming and outgoing tracks, and a signal S is assigned to the voxel that contains the point. This value is calculated as

$$S = \frac{\Delta\theta_x^2 + \Delta\theta_y^2}{2L(1+E_p^2)} \cdot \left(\frac{\hat{p}}{p_0}\right), \tag{2.5}$$

where $\Delta\theta_x$ is the scattering angle in x, L is the height of the voxel, $\hat{p}$ is the estimate of the muon momentum, E_p is the fractional momentum uncertainty and p_0 is

reference momentum. The momentum is estimated from the multiple scattering (this method is more detailed in Section 2.2.4). A value of 0 is assigned to the other selected voxels. The scattering strength λ of each voxel is calculated as the mean of S (including the ones of 0) from all the muons that crossed it.

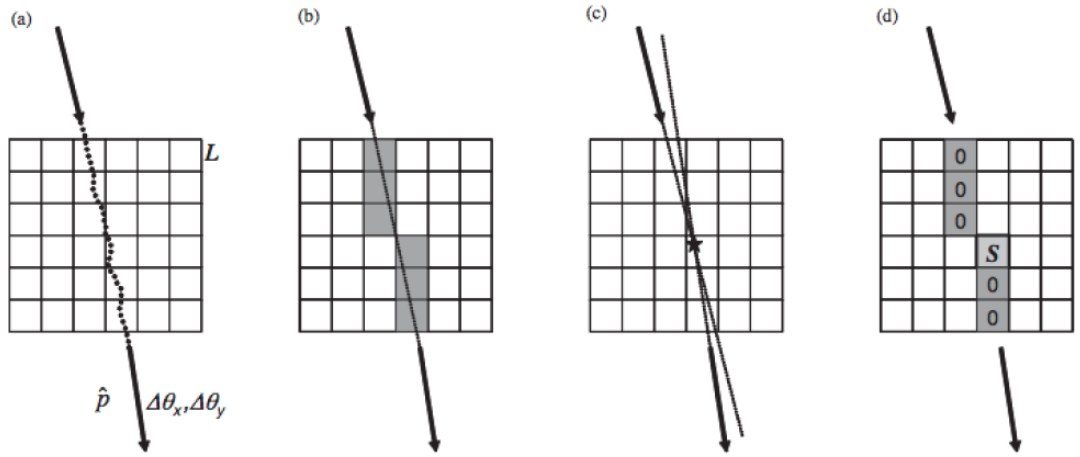

Figure 2.19: Representation of the PoCA method [27]. Stochastic path of a muon (a). Straight line apporoximation and selected voxels (b). Extrapolated incoming and outgoing trakcs, and respective point of closest approach (c). Signal values attributed to each selected voxel (d).

Results of the application of the PoCA algorithm to simulated materials are shown in Figure 2.20. Three levels are attributed to different ranges of λ: high-Z ($\lambda \geqslant 20$), medium-Z ($3 \leqslant \lambda \leqslant 20$) and low-Z materials ($0.25 \leqslant \lambda \leqslant 3$). Each of the three simulated materials (tungsten, iron and carbon) produced a λ in one of these ranges, and hence corresponds to one of these levels. Tungsten was measured as a high-Z material (dark in Figure 2.20), iron as medium-Z (lighter colour) and carbon as low-Z (the lightest colour). This results show that materials with very different atomic numbers can be distinguished with the PoCA algorithm. It does not show, however, if it is possible to distinguish between different materials with similar atomic numbers (such as different high-Z materials), or how good the imaging reconstruction can be (as there is some blurring for all blocks).

2.2.3.2 Expected Maximisation algorithm

The Maximum Likelihood/Expected Maximisation (ML/EM) [28], also applied by others such as in [20], [2] and [29], first calculates the point of closest approach,

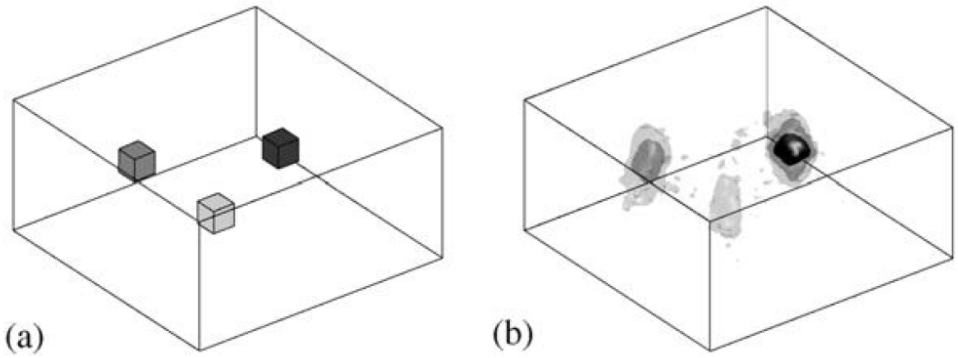

Figure 2.20: Simulated $5 \times 5 \times 5$ cm^3 cubes of different materials (a) and respective reconstruction with the PoCA algorithm (b) [27]. The darkest is tungsten, the medium is iron and the lightest is carbon.

to obtain the scattering angle and offset. Then, from the simplified formula for the scattering width

$$\sigma_\theta = \frac{15 \text{ MeV}}{\beta c p} \sqrt{\frac{H}{X_0}}, \tag{2.6}$$

where p is the muon momentum and X_0 is the radiation length of the material, it calculates a scattering density for a material as

$$\lambda(X_0) = \left(\frac{15}{p_0}\right)^2 \frac{1}{X_0}, \tag{2.7}$$

where p_0 is a nominal muon momentum and $\beta = 1$. The variance of the scattering angle can then be obtained as

$$\sigma_\theta^2 = \lambda H \left(\frac{p_0}{p}\right)^2, \tag{2.8}$$

For each muon, a scattering angle $\Delta\theta$ and an offset Δx are calculated and a data vector is defined as

$$D_i = \begin{bmatrix} \Delta\theta_i \\ \Delta x_i \end{bmatrix}, \tag{2.9}$$

and a covariance matrix as

$$\Sigma_i = \begin{bmatrix} \sigma^2_{\Delta\theta_i} & \sigma_{\Delta\theta_i \Delta x_i} \\ \sigma_{\Delta\theta_i \Delta x_i} & \sigma^2_{\Delta x_i} \end{bmatrix}, \tag{2.10}$$

The likelihood of the density profile λ is

$$P(D|\lambda) = \prod_i P(D_i|\lambda), \tag{2.11}$$

where

$$P(D_i|\lambda) = \frac{1}{2\pi|\Sigma_i|^{1/2}} \exp\left(-\frac{1}{2} D_i^T \Sigma_i^{-1} D_i\right). \tag{2.12}$$

The density profile to be calculated is the one that maximises the likelihood. This is an iterative algorithm, which may need substantial computing resources.

Some examples of results using this method to perform imaging of nuclear waste drums filled with concrete can be found in [29]. A steel drum was filled with concrete and three different objects were encased in it, as shown in Figure 2.21. The steel drum is 1.2 cm thick, with a height of 25.5 cm and a diameter of 17.5 cm. The samples are a lead cube with 4 cm side, a uranium cylinder with 3 cm in height and 2 cm of diameter, and a brass cylinder with 2 cm in height and 4 cm of diameter. The reconstruction of the scattering density λ for uranium and lead using the ML/EM algorithm is shown in Figures 2.22 and 2.23, using data collected over 30 weeks. Both materials can be clearly seen in these images. It is also possible to see that they have different sizes, but the shapes are blurred, due to the limitation of the voxel size ($5 \times 5 \times 10$ mm^3). The article claims that the amount of data used is not indicative of the real time needed to determine the presence of high-Z materials, and that there is indication of their presence within a "short timescale". However, this time is not given. Although good imaging results are shown, a quantitative study in terms of the imaging performance and the time required to obtain good images is essential to determine the efficiency of the method.

This method was also used to perform material discrimination, which is shown in Figure 2.24, where λ is shown for different materials, in air and in concrete. The two points furthest to the right for both cases correspond to lead and uranium. It is clear that discrimination between lead and uranium is better when they are in air than in concrete because the difference of their λ in air is larger than in

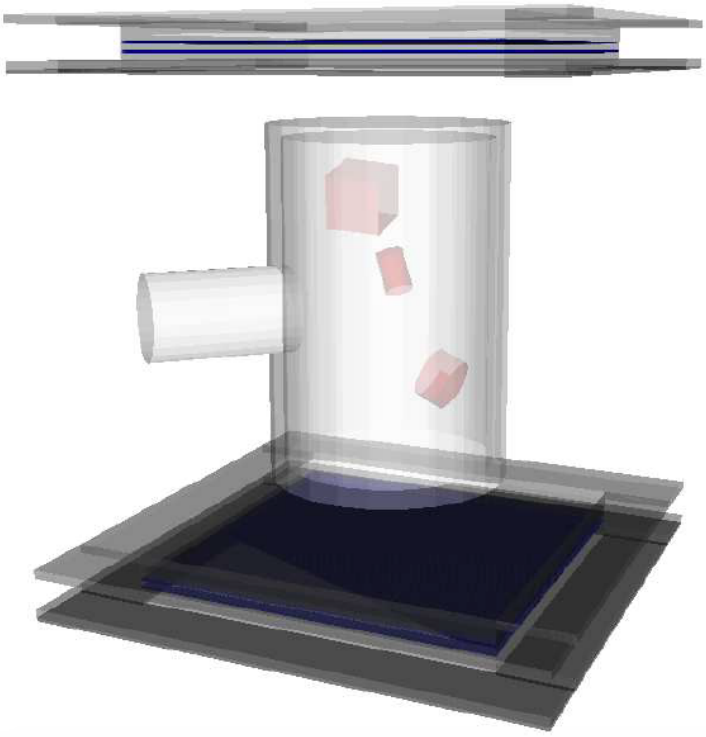

Figure 2.21: Geometry of the experiment performed in [29]. Three different materials are encased in concrete: lead (top), uranium (centre) and brass (bottom).

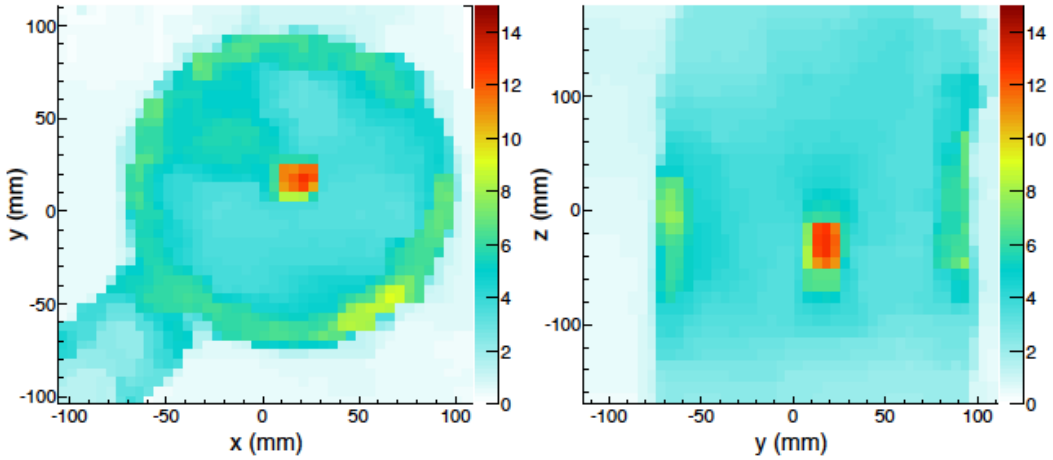

Figure 2.22: Section containing uranium from the image reconstruction of the drum in Figure 2.21. The colour scale is the most-likely λ value in each voxel [29].

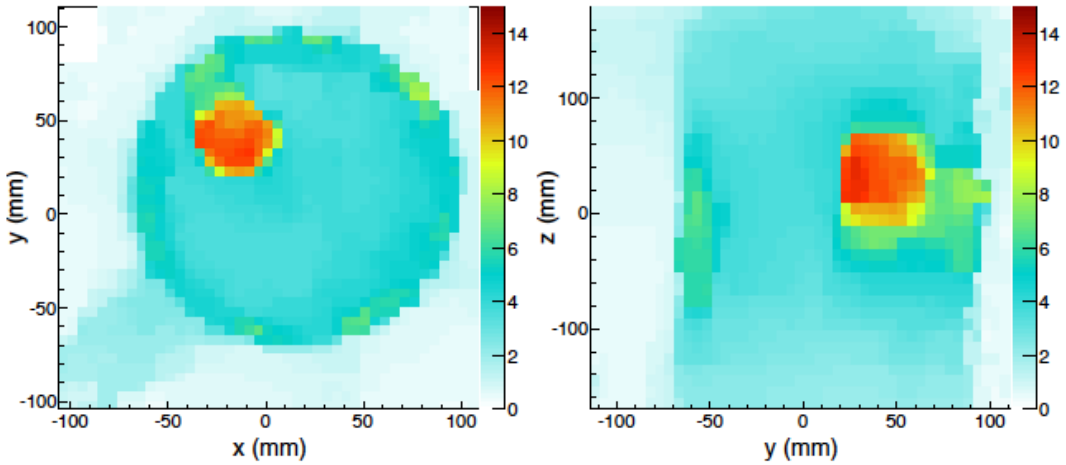

Figure 2.23: Section containing lead from the image reconstruction of the drum in Figure 2.21. The colour scale is the most-likely λ value in each voxel [29].

concrete. The different λ for materials in air and in concrete is understood as being due to the higher contribution of the scattering from the concrete along the muon path, which reduces the contribution of the scattering in other materials. The point for the external steel casing (effective Z of around 24) has a similar λ to the corresponding point for steel in air because, although the steel drum contains concrete, it is surrounded by air. The lines are exponential fits to the data, showing an exponential dependence between λ and Z. This demonstrates that material identification is possible, but a more detailed study is required, to quantify the performance of the material discrimination, and to show how good it is when using a shorter data taking time.

2.2.3.3 Scattering Density Estimation algorithm

The Scattering Density Estimation (SDE) algorithm, shown in [2, 31], starts by calculating the point of closest approach for each muon. It then calculates a scattering density for each voxel with a pitchfork method, which adds four 'satellite tracks' around the incoming and outgoing muon trajectories (the dashed lines shown in Figure 2.25). In order to obtain these tracks, a set of four perpendicular vectors is randomly generated on the plane perpendicular to the muon track. Their

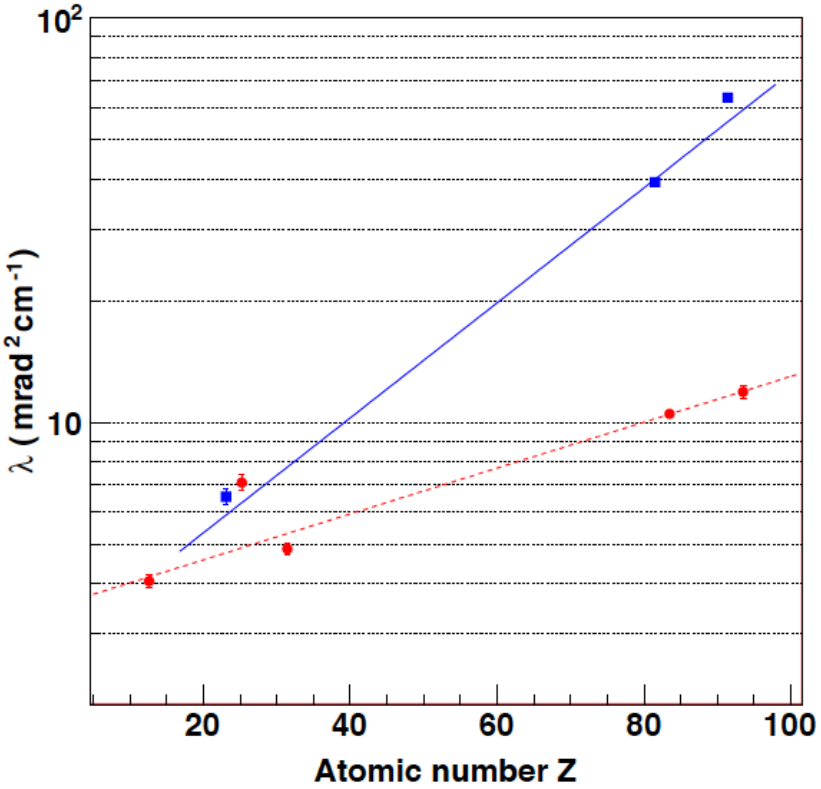

Figure 2.24: Material discrimination using the ML/EM algorithm, using experimental data [29]. The blue squares are the λ from different materials in air and the red circles are for materials embedded in concrete. The lines are exponential fits to the data.

magnitude corresponds to the detector resolution. The scattering density for each voxel is given by

$$\lambda = \frac{\left[\Sigma_i w_i |\hat{\boldsymbol{\mu}}_2^i - \hat{\boldsymbol{\mu}}_1^i| \cdot |\beta c \mathbf{q}_2^i|\right]^2 \Sigma_i w_i (|\mathbf{d}_i| \cdot |\beta c \mathbf{q}_2^i|)^2 |\hat{\boldsymbol{\mu}}_2^i - \hat{\boldsymbol{\mu}}_1^i| \cdot |\beta c \mathbf{q}_2^i|}{\Sigma_i w_i (|\mathbf{d}_i| \cdot |\beta c \mathbf{q}_2^i|)^3}, \tag{2.13}$$

where $\hat{\boldsymbol{\mu}}_1$ and $\hat{\boldsymbol{\mu}}_2$ are the incoming and outgoing muon trajectories respectively, w_i is the product of the weights given to the incoming and outgoing tracks (a weight of 1/3 is given to the real track and 1/6 to the satellite tracks), and $\mathbf{q}$ is the momentum vector. The sums are over the absolute value of the differences between each of the 25 possible combinations of a pair of vectors, for each of the 5 incoming and 5 outgoing vectors.

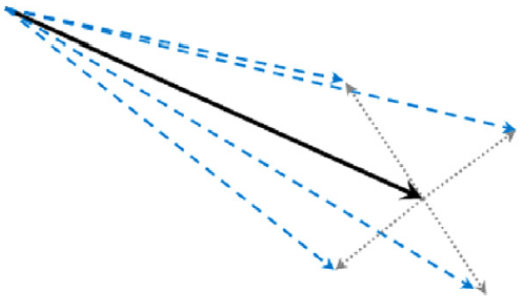

Figure 2.25: Representation of the pitchfork method [2]. The central solid line is the measured muon track and the four dashed lines are the randomly calculated satellite tracks.

Figure 2.26 shows an image reconstruction of a simulated steel drum, filled with paraffin wax and with a UO_2 fuel pin. A day of muon data was used for this simulation. The ML/EM algorithm, with 60 iterations, was used to produce Figure 2.26a and the SDE algorithm for Figure 2.26b. The fuel pin is clearly visible in both images. The performance of both methods for imaging this high-Z object is comparable but, since the SDE method is not iterative, it can be faster than the ML/EM, providing a good alternative.

2.2.3.4 Metric distance algorithm

A weighted metric distance algorithm, previously developed in the University of Bristol, is described in [22] for nuclear security applications, and in [32] for

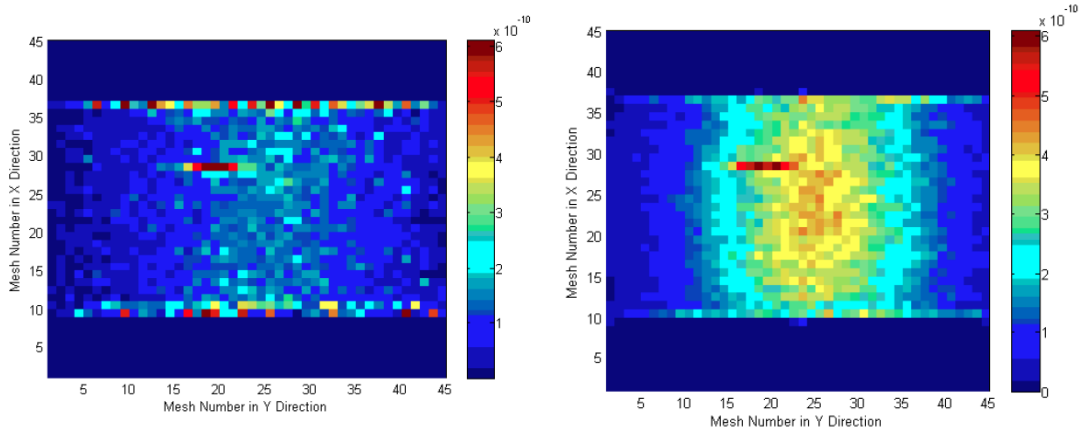

(a) ML/EM

(b) SDE

Figure 2.26: Scattering density per voxel, in arbitrary units, for a simulated UO_2 fuel pin inside a steel drum filled with paraffin wax [2]. The image reconstruction was done using the ML/EM algorithm with 60 iterations (a), and using the SDE algorithm (b).

nuclear waste imaging. In Chapter 4 this method will be described and applied to perform imaging of nuclear waste. It will also be shown how this algorithm can be used to determine the amount of gas in a concrete container. This is a fast algorithm, and the time that it takes to run is negligible compared to the data taking time needed for imaging. As an example, although there is variability in the run time, the algorithm takes about 25 minutes to analyse data corresponding to 6 days of muon exposure.

2.2.4 Momentum measurement

The muon momentum is used in all the algorithms described in Section 2.2.3. It is also possible to perform analyses without the momentum information, as will be seen in several sections of this thesis. However, some of the algorithms mentioned before require the muon momentum, or can be improved by using the momentum information. Therefore, it can be beneficial to add the momentum measurement to muon tomography systems.

In high energy physics experiments, muon momentum is measured by applying a magnetic field: the lower the momentum, the more bent the particle track will be. This can also be used to measure the momentum of cosmic-ray muons [33].

Due to their relativistic speeds, measuring the muon momentum using time of flight (TOF) requires detectors with very good timing resolution. An example of this is a fast scintillator coupled to a silicon photomultiplier. In [34], it was shown that fast scintillator crystals, such as LYSO:Ce (lutetium-yttrium oxyorthosilicate doped with cerium) and LSO:Ce,Ca (lutetium oxyorthosilicate doped with cerium or calcium), coupled to an SiPM (silicon photomultiplier) yield a time resolution under 20 ps for a 150 GeV muon beam. Another example are multigap RPCs (MRPCs), which are stacks of resistive plates separated by gas gaps, and where the voltage is applied to the outermost plates. This detector has a better time resolution than a single gap RPC because the variation on the position of the initial electron cluster introduces a time jitter [35]. Therefore, RPCs with smaller gaps (700 μm in this case) have better time resolution than larger gaps (for example, the RPCs for the Bristol system use a 2 mm gap). Having multiple small gaps allows for a detector with the same performance of one with a larger gap, as the signal is the sum of all gaps, but with a better time resolution [36]. Results for the TOF measurement of 150 GeV muons with a MRPC can be found in [37], which had 5 gaps of 220 μm. The time resolution obtained was 70 ps. Table 2.2 shows some values of the time a muon would take to cross a 3 m tall detector system, for different values of momentum. The reported time resolution of 20 ps would be good enough to measure the momentum for low-momentum muons using TOF, while the resolution of 70 ps would not be useful to measure the momentum of cosmic-ray muons.

momentum (GeV)	time (ps)
0.1	14557.82
1	10062.62
2	10020.88
3	10013.13
4	10010.41
5	10009.16
10	10007.48
50	10006.95

Table 2.2: Time that a muon takes to traverse 3 m, as a function of the momentum.

Another way to measure the muon momentum is to use the multiple Coulomb scattering. This is for example described in [38]. If the material of the detectors is known, the respective radiation length is also known. The scattering angle between each detector layer is measured, and the momentum can be calculated from the width of their scattering angle distribution, as seen in equation 2.3. However, when applied to a muon scattering tomography system, this method gives a poor momentum resolution of 100% to 200% (from preliminary results of a feasibility study). It is possible to improve the resolution by introducing a high-Z material, such as lead, between the layers to increase the scattering angles (for example conceptualised for muon tomography in [27]), increasing the number of layers, or increasing the spacing between layers. This concept is illustrated in Figure 2.27, which shows a MST system with two upper and two lower detector layers for muon tracking, and an added three more detector layers at the bottom, separated by a known material, in order to measure the momentum using the multiple Coulomb scattering information. The OPERA collaboration in [38] also shows this momentum measurement technique, using several layers separated by lead.

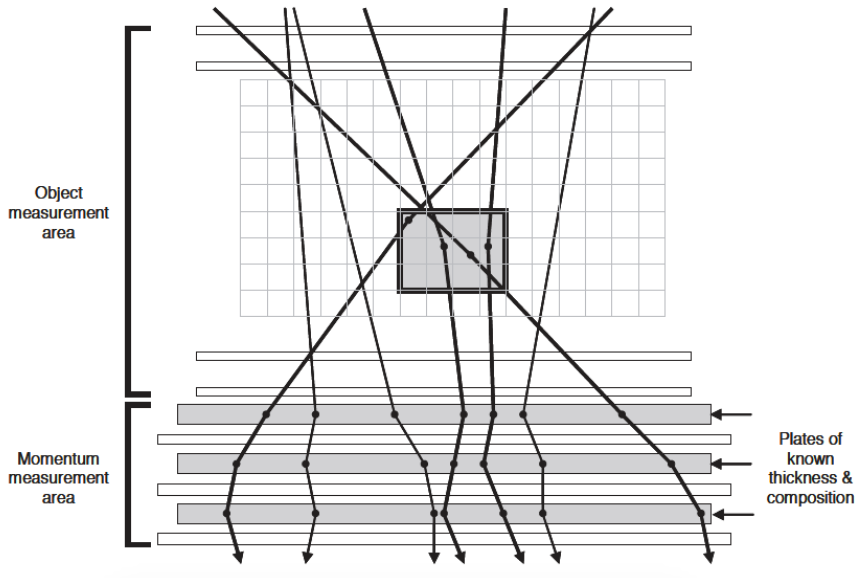

Figure 2.27: Illustration of the concept for momentum measurement using multiple Coulomb scattering integrated in a muon scattering tomography system [27].

3 | Simulations and data analysis

Since it would be difficult to have access to real nuclear waste to test the methods proposed in this thesis, Monte Carlo simulations were performed instead, with a detector system tuned to match the performance of a real detector. Testing novel methods using data from simulations is useful to show feasibility studies, before investing more effort in a muon tomography system.

The different results shown in Chapters 4 and 5 use similar simulation setups and analysis methods. This chapter describes the details of the simulations, track fitting, and multivariate analysis method that are common to those chapters. The application to simulate the muon tomography system includes muon detectors and a test nuclear waste drum. Simulations were performed of muons arriving from the atmosphere, producing hits in the detectors and scattering in the test drum, to reproduce what happens in a real system. The simulations are described in Section 3.1. The hits obtained in the simulated detectors are then fitted in order to reproduce the expected muon tracks through the system and measure the scattering they underwent in the sample. The method for track fitting is discussed in Section 3.2.

A multivariate analysis method, the Fisher linear discriminant, was used for different purposes in both Chapters 4 and 5. This method is explained in Section 3.3, while its applications to the muon tomography data will be detailed in the respective chapters.

3.1 Simulations

Simulations of a muon tomography system were performed in Geant4, which is a toolkit developed to simulate the passage of particles through matter [39–41]. The Geant4 version used was 9.4, including the muon processes of multiple

scattering, ionisation, Bremsstrahlung and pair production. The cosmic muons were generated by CRY, a library that generates muons according to the measured energy and angle distributions [42]. These distributions were obtained for the CRY software by simulating primary protons with energies between 1 GeV and 100 TeV through a model of the atmosphere, creating showers of secondary particles. These simulations were performed using MCNPX [43, 44] and benchmarked against other simulation software and against cosmic-ray measurements, as seen in Figure 3.1. As the particle distributions depends on the altitude, three altitudes are available in CRY: sea level, 2100 m, and 11300 m, from which sea level was chosen. It also depends on the effect of the magnetic at different latitudes, so the latitude of 51.46 deg was chosen, corresponding to the coordinate of the H. H. Wills Physics Laboratory at the University of Bristol.

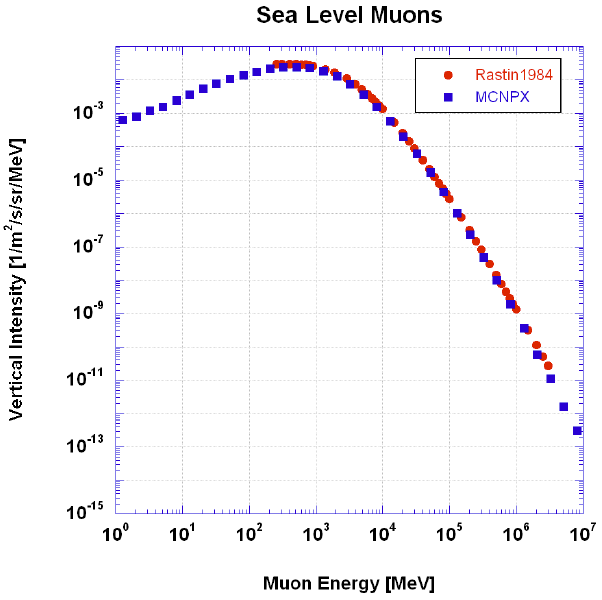

Figure 3.1: Muon spectrum at sea level [42]. The red circles are measurements from [33] and the blue squares are from Monte Carlo simulations in [42].

Each simulation event corresponds to a single muon. In reality, it is possible to detect multi-muon events, where two or more muons are detected after a single trigger event. However, the multiple tracks in these events can usually be distinguished, by performing all possible linear fits of the hits in the different detector

layers, and assigning the best fits as the most likely real muon tracks. Therefore, the simulations did not create multi-muon events.

The concept of muon tomography was explained in Section 2.2. Tomography requires the measurement of the incoming and the outgoing muon, hence detectors need to be placed on either side of the object under test. Our default geometry is shown in Figure 3.2, which was also used for the simulations. A 1×1 m^2 plane was used as the initial source of the muons. There were 6 pairs of RPCs in the simulation, with 3 pairs above and 3 under the sample, each with an area of 1×1 m^2. In each detector pair, one RPC was used to define the position in x and another in y. The RPCs were 6 mm thick, and made of glass. The spacing between each xy pair was 19 mm, and 58 mm between each of the pairs. The height between the upper and lower tracking systems was 548 mm. The drum was placed in the centre of this space, which means that there were 274 mm between the centre of the drum and both the upper and lower plates closest to the drum. The test sample was a small concrete drum, consisting of a cylinder 40 cm long and with 13 cm radius. This concrete cylinder was encased in a 1.5 cm thick steel container. The base and cap of the steel container (the ends of the cylinder) were 2 cm and 3.5 cm thick. In order to simulate realistic RPCs, the detectors had a resolution of 450 μm, which is close to the values measured in [45] and [46] (these values will be shown in Section 3.1.1.2).

The key parameters of the materials used were taken from the Geant4 material database [47]. The simulated concrete has a density of 2.3 g/cm^3, the air of 1.225×10^{-3} g/cm^3, the lead of 11.35 g/cm^3, the tungsten of 19.3 g/cm^3, the uranium of 18.95 g/cm^3, and the plutonium of 19.84 g/cm^3.

3.1.1 Simulation validation

The results shown in this thesis depend on simulations. Therefore, it is necessary to prove the validity of these simulations by using realistic detector performance and showing the comparison between results from simulations and ex-

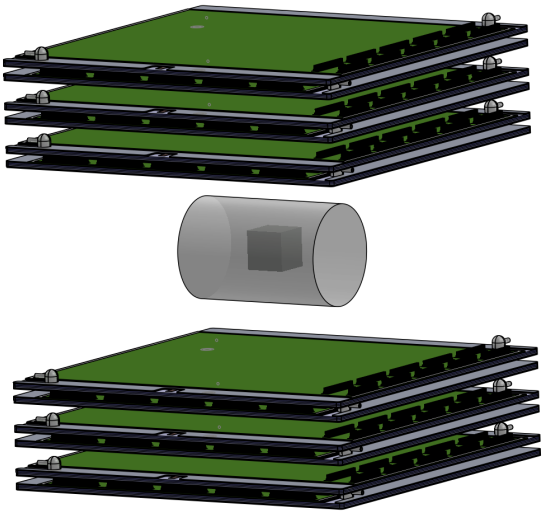

Figure 3.2: Illustration of the simulation geometry. The detectors have an area of 1×1 m^2 and the concrete cylinder has a length of 40 cm and a radius of 13 cm. The $10 \times 10 \times 10$ cm^3 block inside the concrete is an example of a material to be found and identified.

periments. All performance parameters like the hit resolution are consistent with measured with our prototype system [45]. This system will be briefly described.

3.1.1.1 Muon detectors: Bristol system

Since muons are charged particles, they can be detected with various different types of radiation detectors. Several types of detectors have been used for muon radiography and tomography, such as spark chambers [8], nuclear emulsion films [9], scintillator bars [9], and micro-pattern gaseous detectors (Micromegas) [9].

The detectors used in the University of Bristol particle physics laboratory are resistive plate chambers (RPCs), which are gaseous detectors. Additionally, plastic scintillators, coupled to photomultiplier tubes, are used for coincidence triggering, with one above and another below the system.

Resistive plate chambers have been developed since the 1980s [48]. They are a good candidate for muon tomography applications because of their good detection efficiency and low cost. They can be built with a good spatial resolution [46], and they can have large areas. A photograph of a 180×60 cm^2 RPC can be seen in Figure 3.3.

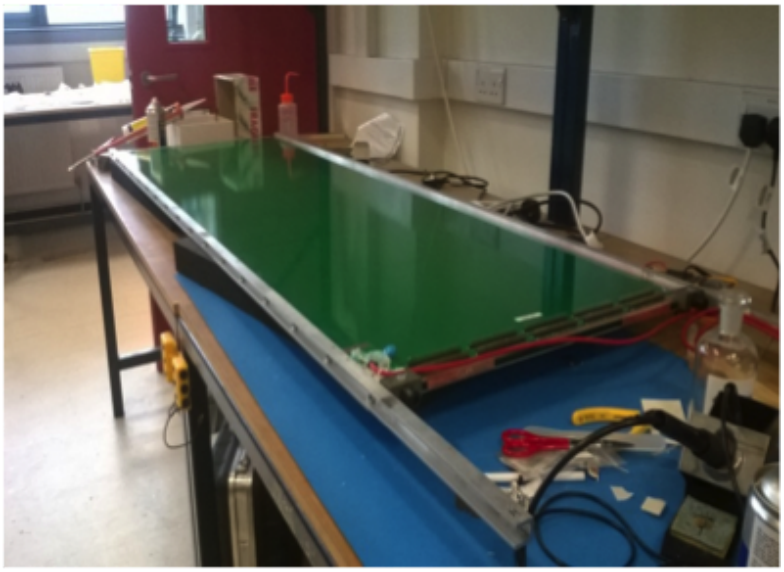

Figure 3.3: Resistive plate chamber (180×60 cm^2) [46].

An RPC consists of two parallel electrodes, to apply an electric field over a gas volume encased between the electrode plates. Charged particles ionise the gas. The resulting electrical charge drifts to an electrode, due the electric field, where it induces a signal on metal strips. Combining the signal on several strips allows to extract the position of the traversing particle. The RPCs are made of glass and have a single gap, which include glass spacers, which are lines of glass touching both side of the gap, in order to better maintain the gap size constant. They have an active area of 50×50 cm^2. A printed circuit board (PCB) with 320 strips at a pitch of 1.5 mm is glued on top of each chamber. This gives the readout position in one direction, and the other chamber in the same cassette has the readout strips perpendicular to the first. The gas mixture consists of about 95% of freon and 5% of isobutane. The components of an RPC detector are shown in an exploded diagram in Figure 3.4. An example of the full system can be seen in Figure 3.5, with a total of 16 RPCs: 8 cassettes (4 above and 4 below the scanning area), containing 2 RPCs each can be seen.

3.1.1.2 Previous experimental results

The efficiency obtained for the RPCs in was between 87% and 95%. The simulated detectors were 100% efficient in order to be able to use all the simulated

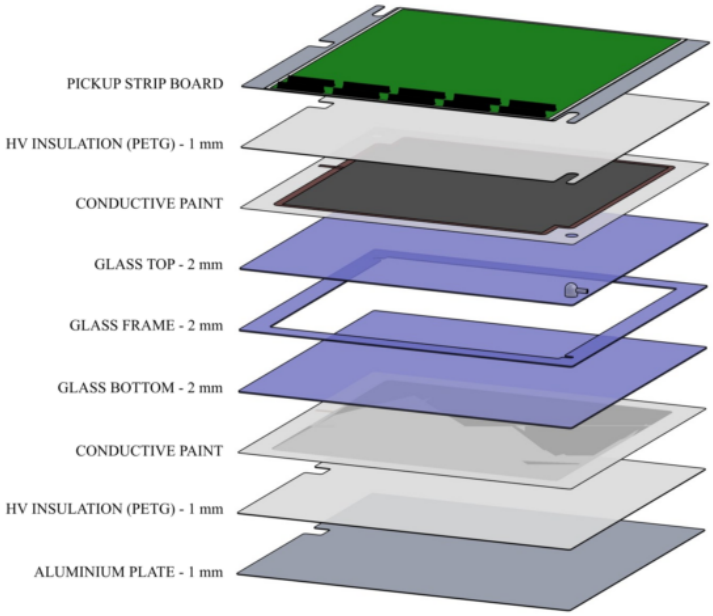

Figure 3.4: Exploded view of an RPC [46].

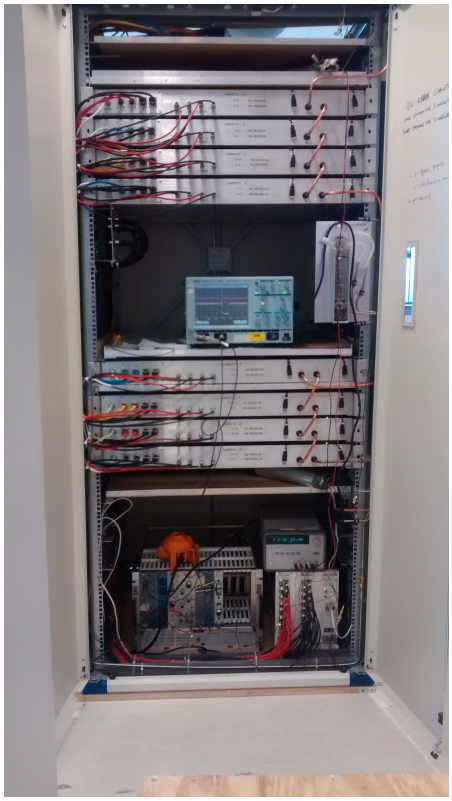

Figure 3.5: Muon telescope with 4 layers (each with a pair of xy RPCs).

muons that crossed the full height of the detector system, and because the real efficiency is also high. The only difference when comparing the results from simulation to real data would be a time factor (in reality it takes slightly longer to detect

Layer	Resolution (μm)
0	310
1	410
2	430
3	410
4	910
5	700
6	340
7	840
8	910
9	410
10	490
11	550

Table 3.1: Spatial resolution obtained for 12 50×50 cm^2 RPCs [45].

the same number of muons as in the simulation). The spatial resolution obtained for 12 50×50 cm^2 RPCs can be seen in Table 3.1 [45] and for 24 180×60 cm^2 RPCs in Figure 3.6 [46]. The resolution values are between 300 μm and 1 mm. However, most values are under 500 μm. A resolution value of 450 μm was used in the simulations, as a reasonable value that can be achieved with a real system.

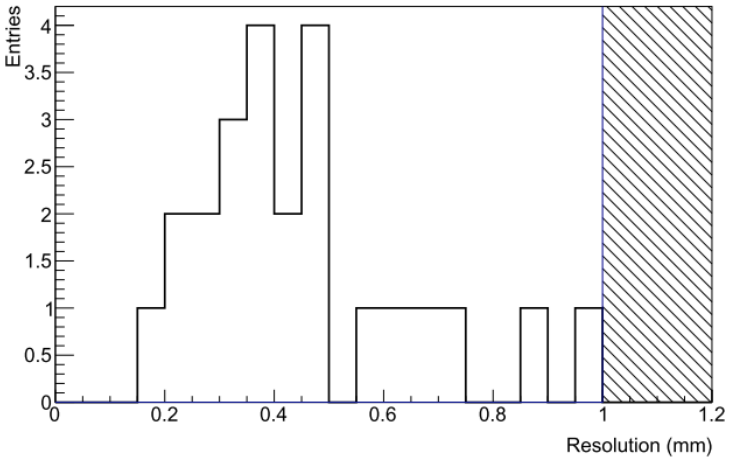

Figure 3.6: Spatial resolution obtained for 24 180×60 cm^2 RPCs [46].

The parameters of the simulated system used in this thesis were the same as described in [22], whose algorithm (that will be explained in Chapter 4) was first applied to simulations. These results were then compared to real data in [46], by showing that it was possible to find a $9 \times 9 \times 9$ cm^3 tungsten block. This block was used because the development of muon tomography systems and analysis methods

at the University of Bristol was at first with the purpose of being used to find special nuclear materials (SNM) in shipping containers. The tungsten block was used as an SNM surrogate in experiments. Additionally, a different algorithm (not used in this thesis) can be found in [49]. It is based on the same simulation work. That paper compared results from simulations with other published EM results. This showed that the simulations are consistent with the results of other groups using different algorithms.

3.2 Track fitting

Muons scatter very little when crossing the 6 mm detectors, so their tracks are approximated as a straight line through the upper detector layers and another through the lower detector layers. The incoming and outgoing muon tracks are therefore reconstructed by fitting a straight line through the hit positions measured by the upper and lower detectors in the simulation. The incoming and outgoing tracks allow for the calculation of the scattering angle and offset of the muons after crossing the sample. The method for the track fit will be explained in more detail.

The output from Geant4 of the hit positions on the detector plates was analysed with a ROOT [50] application that performed a fit of the upper and lower tracks, as described in [22]. For a system with 3 pairs (for x and y directions) of upper and 3 of lower detector layers, 12 hit points are obtained for each muon. A first fit is performed separately for each 3 points in upper and lower, x and y directions, to guarantee that a muon track is present, by accepting only fits with a χ^2 under a cutoff value. In the second step, all the 12 points obtained from an event are fitted at the same time. The fit has 7 parameters, which are the 4 track slopes (for upper and lower, x and y) and the 3 coordinates of the vertex position. This is done by minimising, using MINUIT [51], the function

$$E = E_x + E_y. \tag{3.1}$$

E_x (and, analogously, E_y) is given by

$$E_x = \sum_{i=1}^{3} \frac{(h_i - (v_x + k_{x,\text{upper}} \cdot (z_i - v_z)))^2}{\sigma_{h_i}^2} + \sum_{i=4}^{6} \frac{(h_i - (v_x + k_{x,\text{lower}} \cdot (z_i - v_z)))^2}{\sigma_{h_i}^2} \tag{3.2}$$

where h_i are the measured hit positions, z_i the vertical position of the detectors, v_x, v_y, v_z the vertex positions, k the track slopes, and σ_{h_i} the errors on the measurement of the hit positions. As explained in Section 2.2, the scatter vertex point calculated by the fit (Figure 3.7) is not a real point where the muon scattered, since muons scatter in multiple points, but it is a useful approximation, as will be seen in Chapters 4 and 5. The χ^2 values from the 3-point fits and from the combined fit are calculated and used as a cut: tracks with χ^2 values above the cutoff are discarded. Tracks with a reconstructed vertex outside the volume of interest are also rejected.

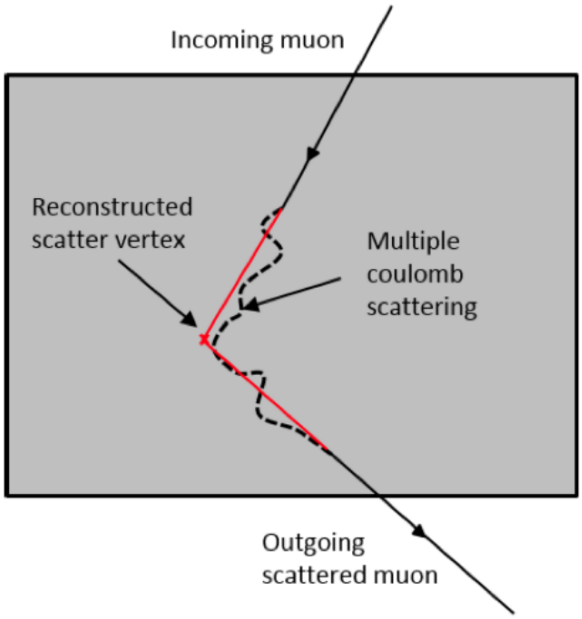

Figure 3.7: Illustration of the vertex reconstruction principle. The angle is exaggerated for better visualisation.

3.3 Fisher linear discriminant

The Fisher linear discriminant is a multivariate analysis (MVA) method that was used in this thesis. MVA is a broad subject that comprises techniques to combine several variables to classify data. In this case, it is used to separate the data into two classes (which we can call signal and background) and requires three steps. The first step is training, in which data samples, correctly identified as signal or background, are analysed in order to calculate weights for each variable, depending on how much they discriminated between the two classes. Finally, the weights calculated for each variable are used to classify data of an unknown category.

The Fisher linear discriminant was first developed by Ronald A. Fisher in [52] for taxonomy applications, to classify different species of plants based on data with several variables. In summary, this method finds the direction $\mathbf{w}$ in the L-dimensional space, for L variables measured, so that orthogonally projected samples, $\mathbf{y}_i = \mathbf{w}^T\mathbf{x}_i$, from the two classes onto $\mathbf{w}$ are better discriminated [53]. It is based on principal component analysis (PCA), but contrary to PCA, it is supervised (the training data includes label information on the respective class) and it also takes into account the spread of the data.

In the Fisher linear discriminant, the $\mathbf{w}$ direction maximises the difference between the mean values (projected in the $\mathbf{w}$ direction) of classes 1 and 2, $|\tilde{\mu}_1 - \tilde{\mu}_2|$, weighted by their scatter (the spread of the data points around the mean). This normalisation is done because if the points are projected on a direction that has a large $|\tilde{\mu}_1 - \tilde{\mu}_2|$ but also a large scatter, the classes can be more overlapped than if projected on a direction with a smaller $|\tilde{\mu}_1 - \tilde{\mu}_2|$ but a small scatter. So the function to be maximised is

$$J(\mathbf{w}) = \frac{(\tilde{\mu}_1 - \tilde{\mu}_2)^2}{\tilde{s}_1^2 + \tilde{s}_2^2} \tag{3.3}$$

where $\tilde{s}_c$ is the scatter for projected samples of class c, given by

$$\begin{aligned}\tilde{s}_c^2 &= \sum_{y_i \in \text{ class c}} (y_i - \tilde{\mu}_c)^2 \\ &= \sum_{y_i \in \text{ class c}} (\mathbf{w}^T x_i - \mathbf{w}^T \mu_c)^2 \\ &= \sum_{y_i \in \text{ class c}} \mathbf{w}^T (x_i - \mu_c)(x_i - \mu_c)^T \mathbf{w} \\ &= \mathbf{w}^T S_c \mathbf{w},\end{aligned} \tag{3.4}$$

where S_c is the scatter matrix for each class (before projection), defined as

$$S_c = \sum_{x_i \in \text{ class c}} (x_i - \mu_c)(x_i - \mu_c)^T. \tag{3.5}$$

We can define matrices S_B and S_W as the scatter matrices between and within classes, respectively,

$$S_B = \sum_c N_c (\mu_c - \bar{x})(\mu_c - \bar{x})^T \tag{3.6}$$

$$S_W = S_1 + S_2 = \sum_c \sum_{i \in c} N_c (x_i - \mu_c)(x_i - \mu_c)^T, \tag{3.7}$$

where μ_c and $\bar{x}$ are

$$\mu_c = \frac{1}{N_c} \sum_{i \in c} x_i \tag{3.8}$$

$$\bar{x} = \frac{1}{N} \sum_i x_i = \frac{1}{N} \sum_c N_c \mu_c, \tag{3.9}$$

N_c is the number of points in class c and N is the total number of points. The denominator of the ratio in equation 3.3 can be rewritten as

$$\begin{aligned}\tilde{s}_1^2 + \tilde{s}_2^2 &= \mathbf{w}^T S_1 \mathbf{w} + \mathbf{w}^T S_2 \mathbf{w} \\ &= \mathbf{w}^T S_W \mathbf{w}.\end{aligned} \tag{3.10}$$

The square difference between the projected means can also be rewritten as

$$\begin{aligned}(\tilde{\mu}_1 - \tilde{\mu}_2)^2 &= (\mathbf{w}^T\mu_1 - \mathbf{w}^T\mu_2)^2 \\ &= \mathbf{w}^T(\mu_1 - \mu_2)(\mu_1 - \mu_2)^T\mathbf{w} \\ &= \mathbf{w}^T S_B \mathbf{w}.\end{aligned} \tag{3.11}$$

Therefore, equation 3.3 can be expressed as [53, 54]

$$J(\mathbf{w}) = \frac{\mathbf{w}^T S_B \mathbf{w}}{\mathbf{w}^T S_W \mathbf{w}}. \tag{3.12}$$

Intuitively, the direction that maximises the ratio in equation 3.12 has the largest scatter between classes, relative to the scatter within classes, for samples projected in that direction [55]. To maximise $J(\mathbf{w})$, its derivative is set to 0 and solved for $\mathbf{w}$. Non-labelled data can then be classified, by being projected onto the calculated direction $\mathbf{w}$.

3.4 Summary

This section described the simulations performed of a muon tomography system and a nuclear waste container. These simulations will be used in Chapters 4 and 5 to evaluate the imaging and material discrimination methods described in those chapters. It was shown how the simulations were tuned using the resolution from the real detectors and benchmarked against real data.

The track fitting algorithm was explained, which fits an upper and lower track for each muon, as well as a scatter vertex. This vertex is not a real point where the muon scattered, but it is a useful approximation, as the results in subsequent chapters will show.

Finally, a description of the Fisher linear discriminant was given. This method is a multivariate analysis, which combines different variables from 2 categories, and is trained with known data in order to better classify unknown data into

one of those categories. The Fisher linear discriminant will be used for different applications in Chapters 4 and 5.

4 Imaging and edge finding

Performing imaging of nuclear waste drums can help identify the materials present, as well as their amounts. Some algorithms have already been developed to perform imaging of nuclear waste containers using muon tomography. This chapter describes one of these algorithms [22, 23], and shows previous imaging results, in Section 4.1. Section 4.2 shows an extension to this method to find edges of uranium blocks in concrete with a very good resolution. Results both with and without momentum information are shown. The existing imaging methods use the muon momentum, but by not requiring momentum information, the method can be also applied by systems without momentum measurement, which is important as it is not feasible to measure the muon momentum with great precision. These imaging methods were first developed and tested to find and image high-Z materials encased in concrete. However, the oxidation of uranium present in nuclear waste leads to the production of hydrogen. Therefore it is also necessary to determine the presence of gas in containers. Section 4.3 shows results using the imaging method to find gas bubbles in concrete and calculate the amount of gas in a container. Discussion and conclusions are given in Section 4.4.

4.1 Metric distance method for imaging

The method described in [22] and applied to imaging in [32] uses an algorithm based on the density of reconstructed vertices, that also takes into account the scatter angle and muon momentum. This algorithm divides the volume into cubic voxels and selects a fixed number of the most scattered tracks in each voxel. The method relies on the fact that there is a larger density of high angle scatters in high-Z materials, so the distance between two reconstructed vertices, divided by their scatter angle, will be smaller. A further advantage is that highly scattered tracks allow to extract a scattering vertex position with good precision, because a vertex reconstructed from a low-scattering track has a larger uncertainty, especially

in the vertical direction. The metric distance is also weighed by the momenta of both muons because a highly scattered muon with a large momentum is more likely to have scattered in a higher Z material, while a highly scattered muon with a low momentum could also have been scattered by lower Z nuclei. For each pair of selected vertices reconstructed in a voxel, a weighted metric distance is calculated as

$$m_{ij} = \frac{\|v_i - v_j\|}{(\theta_i p_i)(\theta_j p_j)}, \tag{4.1}$$

where v_i is the reconstructed vertex position of muon i, θ_i its scatter angle and p_i its momentum (if it is available). This is done for a fixed number of vertices in each voxel in order to make the resulting distribution equivalent for all voxels. The distributions of the natural logarithm of the weighted metric distance were used. Figure 4.1 shows these distributions for 4 different 1 cm^3 voxels: with only uranium or concrete, both far and close to the edge of a uranium block. The distributions for uranium and concrete far from the edge are well separated. However, the distributions of each material at the edge between concrete and uranium are partly overlapped.

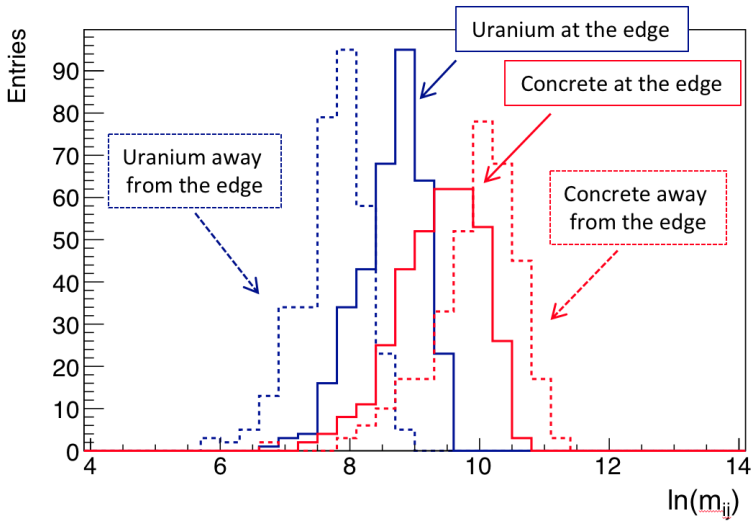

Figure 4.1: Distributions of $\ln(m)$ for 4 voxels: with only uranium or concrete, both far and close to the edge of a $10 \times 10 \times 10$ cm^3 uranium block.

The median of $\ln(m)$ from each voxel was used as a discriminator value: as it is clear in Figure 4.1 the median is lower for higher radiation length materials.

In Figure 4.2 the first imaging results of a simulated drum using this method are shown, where the colour scale corresponds to the discriminator value (red is a lower and blue a higher value) [23, 32]. It can be seen that this method allows to distinguish high-Z materials from the concrete background. It used data corresponding to 2 weeks of muon flux, divided in 4 different rotations (0°, 90°, 180°, 270°), and using true momentum information. The fixed number of most scattered tracks per voxel was chosen to be 70 because it guarantees that all the voxels in the volume of interested have enough tracks to be included. In this example, it is possible to see the steel flanges from the steel drum, a ($0.5 \times 10 \times 10$ cm^3 uranium sheet, a uranium rod with 1 cm radius and 10 cm height, and two 1 cm thick tungsten pennies with 2 cm and 4 cm diameter. It is also possible to distinguish a cylindrical air enclosure with 10 cm radius and 5 cm height. The only high-Z object that is not visible by eye in the image is the 1 cm diameter tungsten penny. An edge finding method was also developed in that article, resulting in a resolution of 1.2 ± 0.4 mm for uranium blocks. A new edge finding method, which will be described in Section 4.2, was later developed, showing an improvement on the resolution, as well as new results without momentum information.

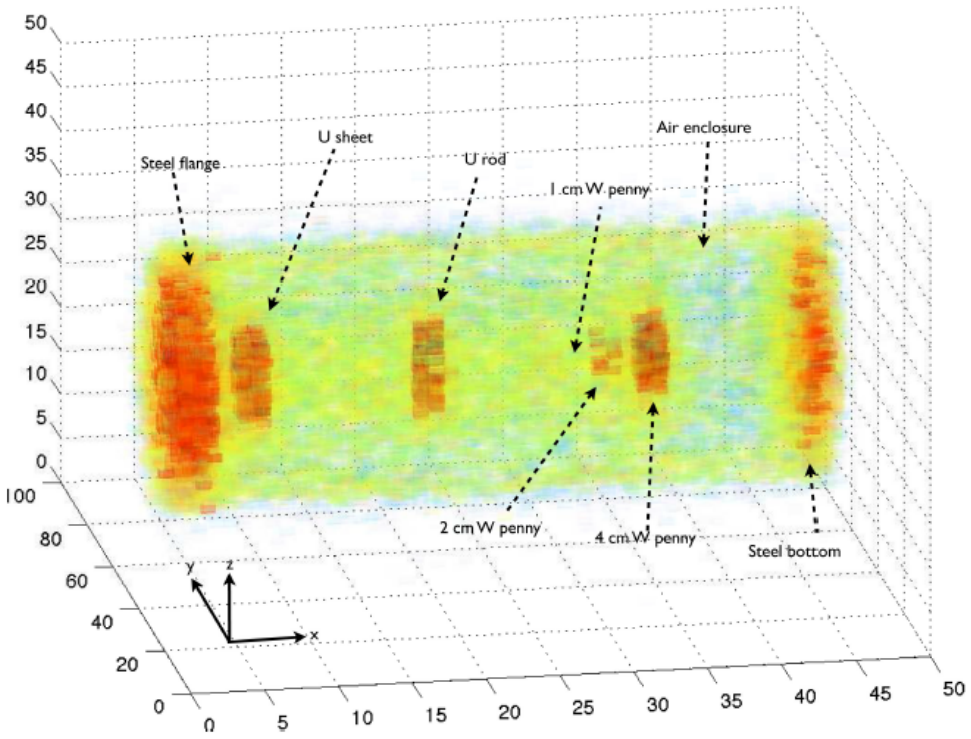

Figure 4.2: Imaging of high-Z and high density materials in a simulated concrete container using the metric distance method. The colour scale corresponds to the discriminator value, where red is lower and blue is a higher value. A region of lower density material (air) can also be seen [23, 32].

4.2 Edge finding

Figure 4.2 shows that lumps of material can be distinguished from concrete by eye. This raises the question of whether it is possible to measure their size. To this purpose, a method for determining edges of uranium blocks encased in concrete was developed and tested in simulations. Several simulations of the concrete cylinder described in Section 3.1 were performed, each containing a block of uranium with different sizes. The length of the blocks along the cylinder ranged from 0.2 cm to 10 cm. Blocks with a length of 3 cm or more were cubic, while smaller blocks had a width and height of 3 cm.

For the case of the edge finding, the fixed number of most scattered tracks was chosen to be 26. It is not obvious that using as many tracks as possible would yield the best results with this method, because the chosen tracks are the most scattered ones. Therefore, using more tracks will not necessarily add more information. The value of 26 was chosen because it results in enough data for the metric distance

distributions (the combination of pairs of 26 vertices results in 325 entries of m_{ij} in each voxel), and because, when using more tracks with the chosen muon exposure time of 5.8 days, some concrete voxels did not have enough tracks to be included.

4.2.1 Previous edge finding method

In the previous edge finding method [32], the distribution of $\ln(m_{ij})$ was plotted for each voxel and it was fitted with a Landau distribution, where the peak location was fixed to the value for the metric distance obtained when only using concrete or uranium. For each voxel the ratio of the amplitude of the fitted Landau for concrete and uranium was calculated. To obtain a good resolution determining the edges of the blocks, this was done for several grids, with the same voxel size but each shifted by 1 mm in the direction of the edge scanning (the x-axis, along the length of the concrete cylinder). When defining the start and finish of each block as the voxel where the amplitude fraction exceeded 99% for uranium, a resolution of 1.2±0.5 mm was obtained using 25 days of muon data.

4.2.2 New edge finding method

In the new method shown in this thesis, the distributions of $\ln(m_{ij})$ were fitted with a Landau distribution convoluted with a Gaussian. This was proven to be a better fit than only a Landau distribution. Figure 4.3 shows the distribution of $\ln(m_{ij})$ for a voxel inside uranium fitted with a Landau (Figure 4.3a) and a Landau convoluted with a Gaussian (Figure 4.3a). A goodness of fit χ^2 test was performed. Using Figure 4.3 as an example, the χ^2 test on the Landau only fit resulted in a p-value < 0.00001, while the Landau and Gaussian fit resulted p-value of 0.9999. Therefore, it the Landau convoluted with a Gaussian fits this metric distance distribution better than the Landau alone.

An example of the metric distribution for a voxel in a uranium block and a voxel in concrete, including the fit, can be seen in Figure 4.4. It is clear that the distributions have different positions and shapes. This indicates that the edge

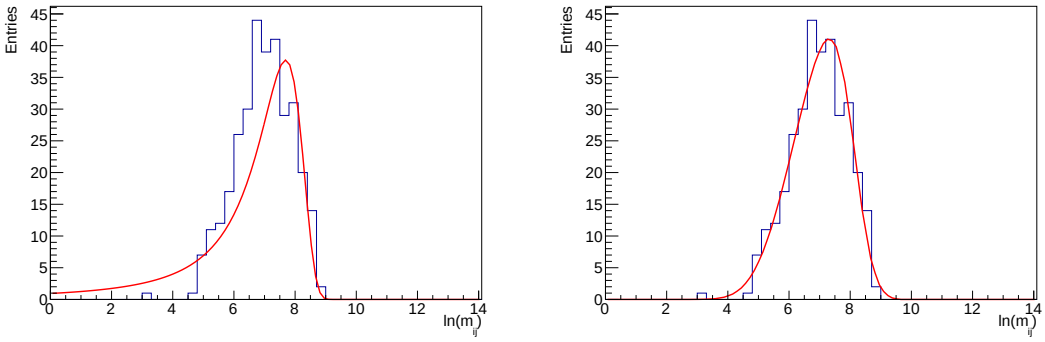

(a) (b)

Figure 4.3: Distribution of the logarithm of the metric distance for a voxel inside uranium, fitted with a Landau (a) and a Landau convoluted with a Gaussian (b).

between uranium and concrete could be better defined than with the previous method by combining different variables from these distributions.

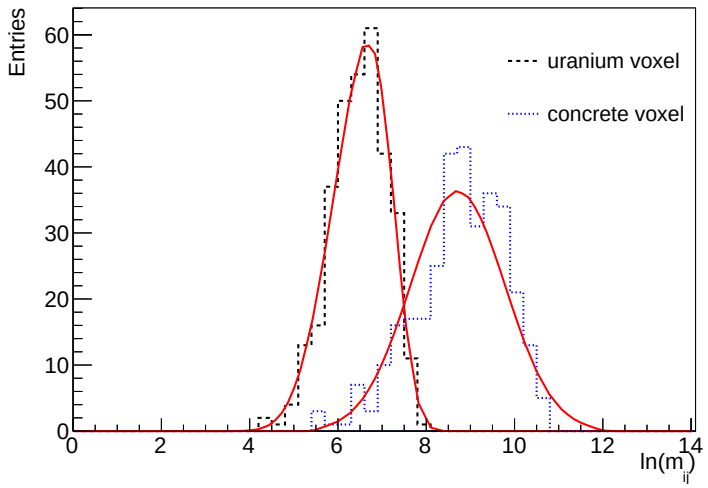

Figure 4.4: Distributions of the logarithm of the metric distance for a voxel inside uranium and a voxel in concrete, as well as respective fit with a Landau distribution convoluted with a Gaussian.

A multivariate analysis (MVA) was performed to classify each voxel as concrete or uranium. The MVA method used was the Fisher linear discriminant, and the value used was the Fisher probability output, which is between 0 and 1, where a voxel with probability of 0 is more likely to be in concrete and 1 to be in uranium. Details about the Fisher linear discriminant can be found in Section 3.3. Variables for the MVA were obtained from the metric distribution. These were the maximum value and respective bin position, and from the fit of a Landau and

Gaussian: their widths, peak positions and peak amplitudes. The training data were simulations with a sample of 1 cm^3 voxels with pure concrete (from a tube filled with concrete) and another of 1 cm^3 voxels with pure uranium (at the edges of $4 \times 4 \times 4$ cm^3 blocks encased in a concrete tube), and its output can be seen in Figure 4.5. It is clear that it is possible to separate concrete from uranium voxels, since most concrete voxels result in a Fisher probability value close 0, and uranium to 1. Additionally, it can be seen in Figure 4.5b, in a logarithmic scale, that there are very few entries which are neither close to 0 nor 1. As, in practice, the muon momentum is either not measured, or measured with a large uncertainty, the momentum was not used in these results. Instead, the fixed value of 3 GeV, which is the peak of the energy spectrum of cosmic muons at sea level, was used as the momentum for the calculation of the metric distance m_{ij}. Section 4.2.4 will show results for an ideal system that measures the muon momentum.

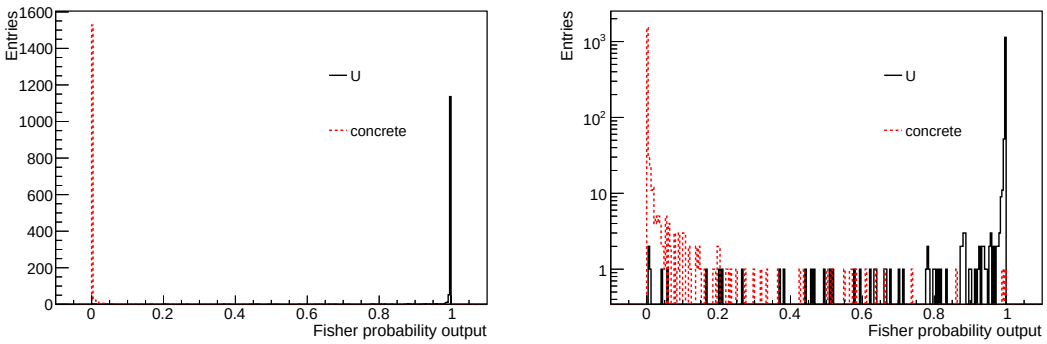

(a) (b)

Figure 4.5: Output from the training with the Fisher linear discriminant without momentum information, using samples of pure concrete and pure uranium. In (b) the same data are shown in a logarithmic scale.

The metric distributions were obtained, and MVA performed, for 1 cm^3 voxels. The grid that divides the volume in voxels was then shifted by 1 mm along the concrete cylinder, and the analysis was carried out for the new voxels. This was done for 9 grid shifts, until the shift matched the first grid. A Fisher probability value was therefore obtained for sections of 1 mm, which allows for a more precise measurement of the edge than using just the 1 cm side voxels. For each uranium block size, 5 simulations were performed, using the amount of muons corresponding

to 139h (or about 5.8 days), with 7000 muons per m^2 per minute. With this rate, the total data taking time for the 5 simulations corresponds to 29 days.

4.2.3 Results

In order to quantify its quality, the edge finding method was applied to simulated uranium blocks of different lengths, ranging from 0.3 cm to 5 cm, which were then measured from the method output. The Fisher probability output for reconstruction of the blocks is shown in Figure 4.6. Each point in these figures is the average of the values from the 5 simulations with the same block size for single voxels along x, at fixed y and z values. Unlike in the training, the Fisher probability is not exactly 1 for the uranium region. However, the Fisher probability for most points in the concrete background was zero, so it is possible to distinguish the uranium from the much lower concrete background, and consequently reconstruct the block length. These were measured by taking the region between the first and the last point whose ξ was above a given threshold, for a range where at least two consecutive points were above zero. Thus a single spike above the threshold is considered concrete, but more consecutive points above the threshold are considered uranium. The threshold value was chosen by testing different thresholds, as seen in Table 4.1. The chosen value was 0.1 for several reasons. Firstly, it resulted in the most accurate reconstructed lengths. Secondly, the decrease in the reconstructed lengths with the threshold increase was not linear for all lengths, therefore there would be varying accuracies for different lengths if the threshold were higher. Lastly, although the value of 0.1 seems low, it is high when compared to the points classified as concrete, which had values of the order of 10^{-12} to 10^{-10}. Uncertainties on the measurement of the length of each block were calculated by performing the same length measurement for 5 different simulations with data corresponding to about 5.8 days of muon flux each, and calculating the standard error on the mean. Blocks of 0.3 cm or smaller yielded an output that would be interpreted as only concrete, as Figure 4.6a shows for a 0.3 cm block.

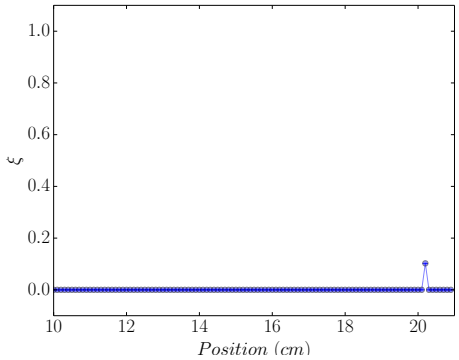

(a) 0.3 cm

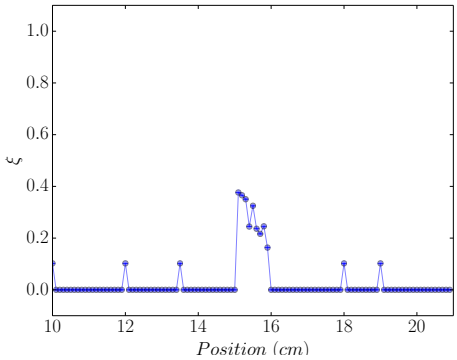

(b) 0.5 cm

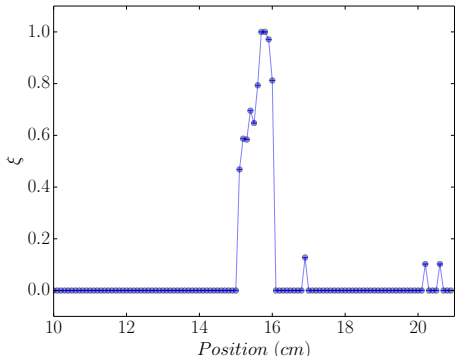

(c) 1 cm

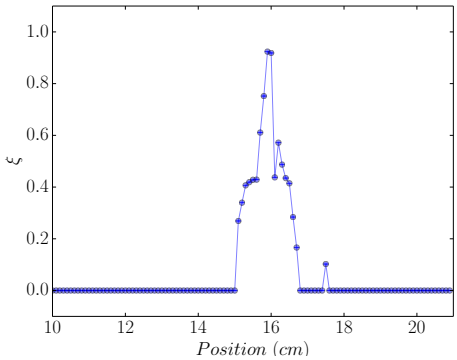

(d) 1.5 cm

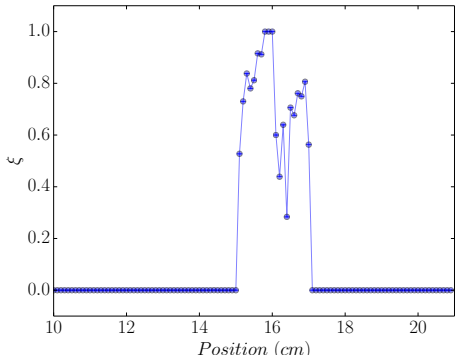

(e) 2.25 cm

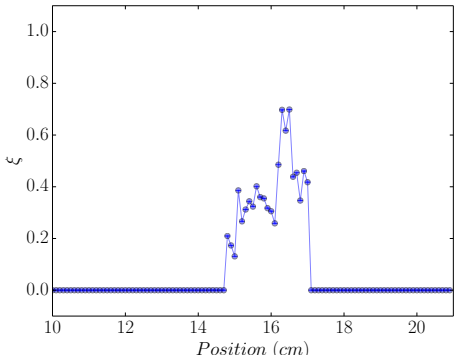

(f) 3 cm

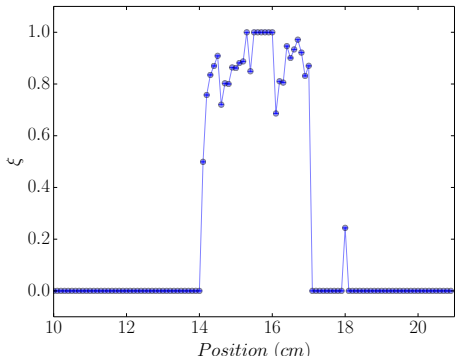

(g) 4 cm

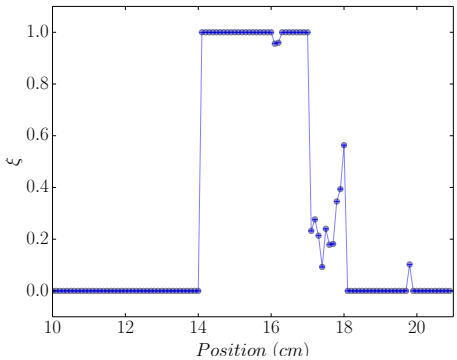

(h) 5 cm

Figure 4.6: Fisher probability outputs, ξ, as a function of the x coordinate for different uranium block sizes without momentum information.

Real length (cm)	Reconstructed length for each threshold (cm)								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.5	0.7	0.6	0.5	0	0	0	0	0	0
1	0.9	0.9	0.9	0.9	0.8	0.7	0.5	0.4	0
1.5	1.5	1.4	1.2	1.1	0.6	0.4	0.3	0.2	0.3
2.25	1.9	1.9	1.9	0.8	1.9	1.7	0.9	0.8	0.2
3	2.2	2.2	1.9	2.2	0.3	0.3	0	0	0
4	2.9	2.9	2.9	2.9	2.8	2.8	2.8	2.7	1.5
5	3.9	3.9	3.0	3.0	3.0	3.0	3.0	3.0	3.0

Table 4.1: Reconstructed lengths using different threshold values for the Fisher probability output.

The reconstructed sizes versus the real sizes are shown in Figure 4.7. It is clear that there is a monotonic relation between the reconstructed and simulated lengths, and the respective linear fit is given by

$$\text{Reconstructed Size} = (0.65 \pm 0.08) \times \text{Real Size} + (0.34 \pm 0.03)\ \text{cm}. \tag{4.2}$$

Although it was possible to perform a good linear fit, the linear function obtained has an offset from zero, and its slope is not 1, which implies that the simulated length cannot be reconstructed directly from the length measured with this method. This is likely due to the fixed size of the 1 cm voxels, and to the "mixing region" around the uranium blocks, where some vertices of muons that scattered in uranium are reconstructed in concrete and vice-versa: for blocks that are close to the voxel size or smaller, most voxels from the 1 mm shifts contain the mixing regions from both sides of the block, stretching their measured length, while for larger blocks the voxels at each edge only contain the mixing region of the respective edge. Therefore, it is necessary to correct the reconstructed values to obtain a 1:1 correspondence with the real sizes. Each point was corrected by performing the fit excluding that point and using the resulting equation to obtain the corresponding corrected length. Excluding the point to be corrected is necessary to guarantee that the correction is independent of the data from this particular simulation. In a real scenario, new data would be obtained, so the parameters from Equation 4.2 could be used instead, and the real size would be calculated as

Real Size = (Reconstructed Size − 0.34)/0.65. The corrected results can be seen in Figure 4.8. They were fitted with a line of equation

$$\text{Reconstructed Size} = (1.06743 \pm 3 \times 10^{-5}) \times \text{Real Size} - (0.11593 \pm 9 \times 10^{-5}) \text{ cm}. \tag{4.3}$$

Having a slope close to 1 and an intercept value close to 0 is good, as it shows a direct correspondence between the real and corrected values.

The uncertainty in the reconstructed length is independent on the real size. This can be seen from the size of the error bars, which indicate the error on the mean of 5 simulations for the same real size. Because of this and since the reconstructed size is linearly dependent on the real size, we can combine all data points in this graph to calculate the position resolution as the standard deviation of the reconstructed minus the real size, as

$$\sigma = \sqrt{\frac{\sum_i (l_{t,i} - l_{r,i})^2}{n-1}}, \tag{4.4}$$

where $l_{t,i}$ is the true simulated size, $l_{r,i}$ the reconstructed sizes, and n the number of data points used. The value obtained was $\sigma = 3.2 \pm 0.6$ mm. The smallest length that was possible to measure with this method was 0.5 cm.

4.2.4 Results using momentum information

The method above does not use the momentum information in the determination of the discriminator because, in reality, it is difficult to determine the momentum as, in a system such as the one simulated, the only information on the momentum can come from the multiple scattering in the detectors. However, it is of value to perform the same study using the true momentum to show how that information can improve the results, and to compare with a previous edge finding method that used true momentum information [32]. Therefore, the same method was applied using momentum information to uranium blocks with lengths from 0.2 to 10 cm. The training for the MVA was also done with momentum since this

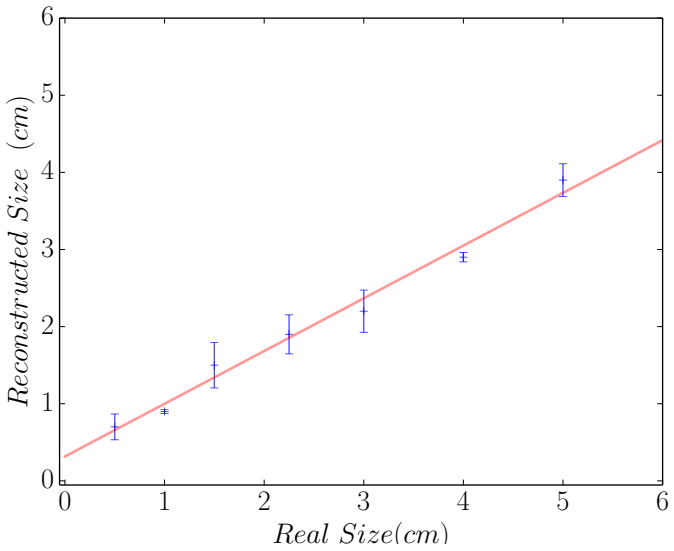

Figure 4.7: Reconstructed lengths as a function of the real lengths without momentum information before correction, including the linear function fitted to the data.

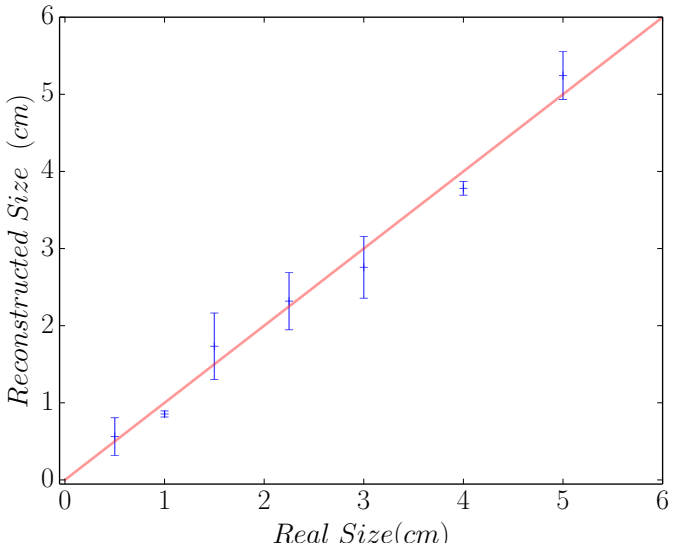

Figure 4.8: Reconstructed lengths as a function of the real lengths without momentum information, after correction.

changes the distribution of the metric distance. The concrete voxels used to obtain the training data were from a tube filled with concrete and the uranium voxels were the ones from the edges of a $18 \times 18 \times 18$ cm^3 block encased in a concrete tube. The Fisher probability output from the training data with pure concrete and pure uranium voxels can be seen in Figure 4.9. The difference in geometry for the training with and without momentum information may account for the apparently better classification after the training without momentum (Figure 4.5) than with momentum (Figure 4.9). However, when applied to new data, including the momentum improves the results.

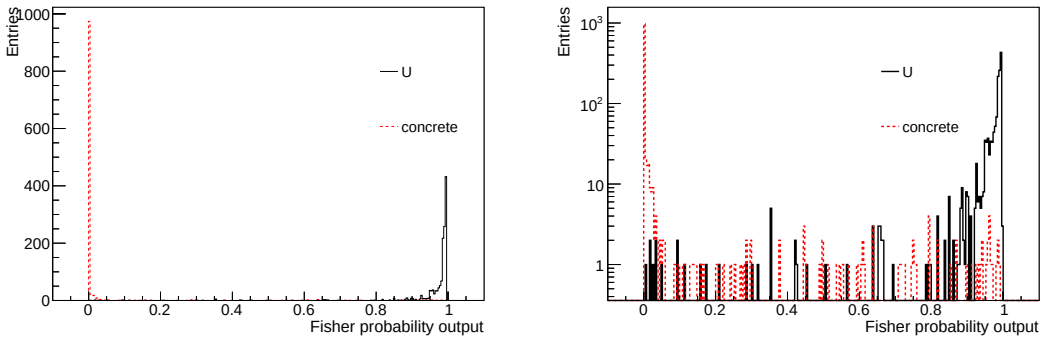

(a) (b)

Figure 4.9: Output from the training with the Fisher linear discriminant using momentum information, with samples of pure concrete and pure uranium. In (b) the same data are shown in a logarithmic scale.

Figure 4.10 shows examples of the Fisher probability output for the reconstruction of 0.5 cm and 3 cm block lengths. For the study using momentum information, the reconstructed length of the blocks was considered to be from the first voxel which had a probability value over 0.9 to the last voxel over this value. The uncertainties were calculated as in the study without momentum. Figure 4.11 shows the reconstructed sizes as a function of the simulated sizes. The linear relation obtained was

$$\text{Reconstructed Size} = (0.76 \pm 0.01) \times \text{Real Size} + (0.44 \pm 0.02)\ \text{cm}. \qquad (4.5)$$

With this equation, the reconstructed sizes were corrected, as shown in Figure 4.12. Figure 4.13 is the same graph, but zoomed in on the lengths under 2 cm. The correction was performed in the same way as for the previous points without momentum information: a fit was done with a point omitted, and that point was corrected using the resulting equation. The linear fit performed over the corrected values follows equation

$$\text{Reconstructed Size} = (0.998716 \pm 7 \times 10^{-6}) \times \text{Real Size} + (0.03117 \pm 3 \times 10^{-5}) \text{ cm}. \tag{4.6}$$

The resolution obtained using momentum information was $\sigma = 0.98 \pm 0.03$ mm, which shows an improvement over the previous result in [32]. As expected, the resolution is better than the one without momentum information ($\sigma = 3.2 \pm 0.6$ mm). It was also possible to measure smaller lengths when using the momentum, down to a 0.2 cm length.

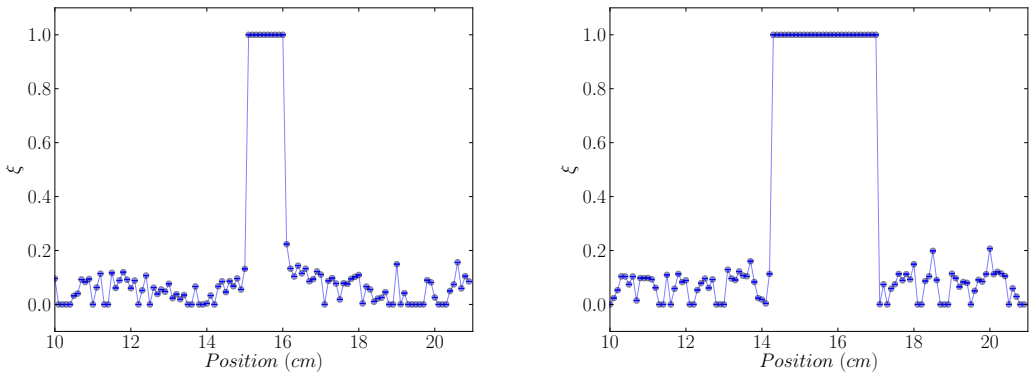

(a) 0.5 cm (b) 3 cm

Figure 4.10: Fisher probability output, ξ, as a function of the x coordinate for a 0.5 cm (a) and a 3 cm block (b) of uranium inside concrete.

The new edge finding method was shown to improve the resolution over the previous method, and to allow the measurement of the edges of uranium blocks in a more realistic case without momentum information.

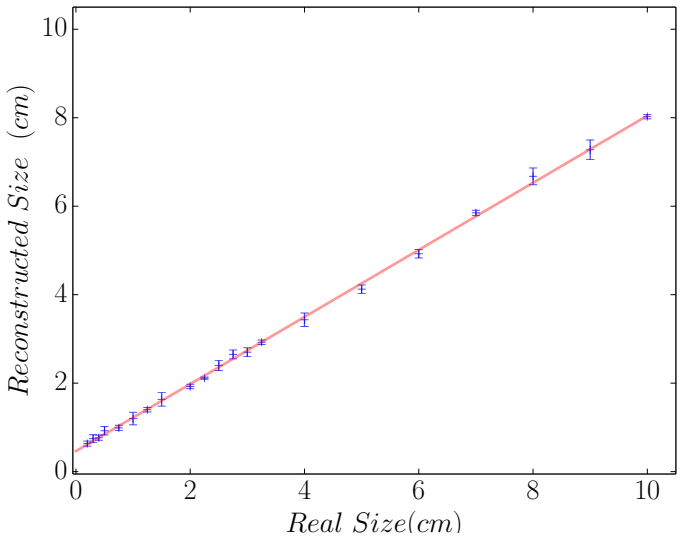

Figure 4.11: Reconstructed lengths as a function of the real lengths before correction using momentum information, including the linear function fitted to the data (equation 4.5).

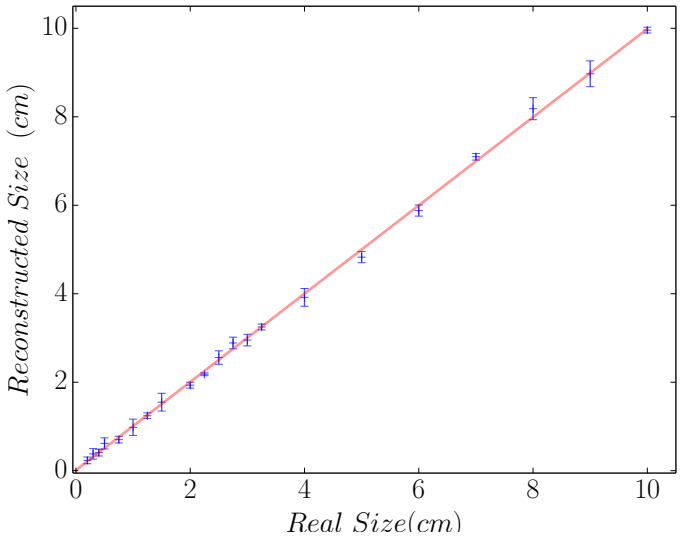

Figure 4.12: Reconstructed length as a function of the real length using momentum information, after correction.

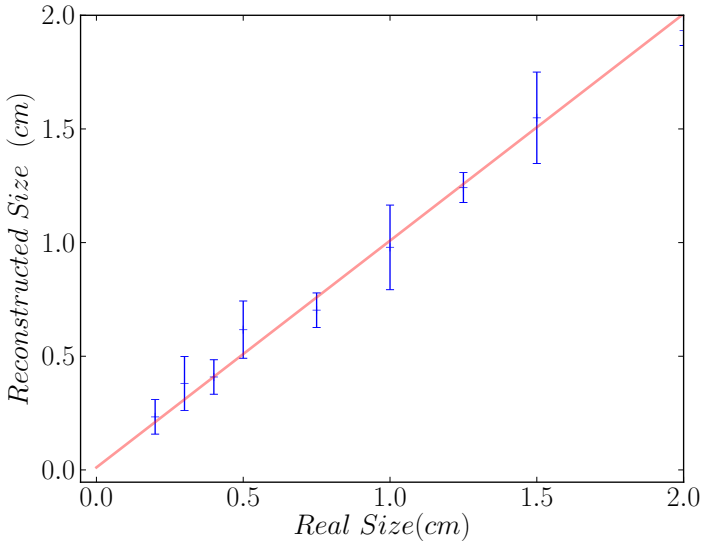

Figure 4.13: Reconstructed length as a function of the real length for the smallest lengths using momentum information, after correction (Figure 4.12 zoomed in the small values).

4.3 Low-Z materials

It has already been shown in Sections 4.1 and 4.2 that it is possible to perform imaging of high-Z materials in concrete drums, but it is also important to find gas pockets in nuclear waste, to predict any overflowing. In several countries, nuclear waste can be stored in bitumen instead of concrete. Uranium and other metals present in nuclear waste can oxidise, producing bubbles of hydrogen gas, which can move in the bitumen and possibly accumulate in bigger bubbles. It has already been observed that a large amount of gas can be trapped in real containers: sealed 220 L drums of bitumenised waste that were filled to 180 L have been found to overflow, which shows that the accumulation of gas can make the content of the containers expand by 20% of their volume.

Since the concrete has higher atomic number and density than gas, the muon scatter angle distribution will be different for both materials. This was already seen for example in Figure 4.2, where an air enclosure can be seen, but a method to measure the volume of gas had not been developed.

The metric method was also used to measure the amount of gas embedded in concrete. Just like the imaging method from [32] and summarised in Section 4.1,

the discriminator value chosen was the median of the metric distance distributions from every voxel in the volume. This value is plotted in Figure 4.14, for steel drums full of air and concrete. Several variables from this distribution were studied: the

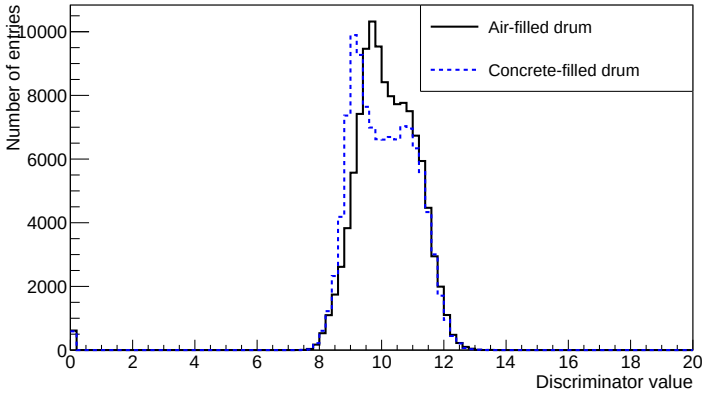

Figure 4.14: Distribution of the discriminator for air and concrete. Each entry corresponds to a 1 cm^3 voxel.

mean, the skewness, the kurtosis and the RMS. The dependence between each of these variables and the bubble volume can be seen in Figure 4.15. The first point of each graph corresponds to a drum full of gas and the last point to a drum full of concrete. RMS was rejected as a discriminator because it had the lowest slope. Kurtosis was also rejected because it was not monotonic for low volumes, unlike the other variables. Therefore only the performance of the mean and the skewness were compared. For a preliminary comparison between the performance of the mean and the skewness, and also between using 72 million muons and 159 million muons, three points were used, for the volumes of 7,125 cm^3, 4,423 cm^3 and 2,001 cm^3. The resolutions were calculated as $r = \frac{\Delta V}{V}$ and are shown in Table 4.2. The best resolutions can be seen in the last column of the table, which corresponds to using the mean, with 159 million muons. Therefore, the mean was found to be the most sensitive variable to the amount of gas in a concrete container. By calculating the mean of several containers with different gas bubble volumes, their volumes were measured, as will be described in the next section. The chosen number of muons corresponds to an exposure time of about 15.8 days at the rate of 7000 muons per m^2 per minute.

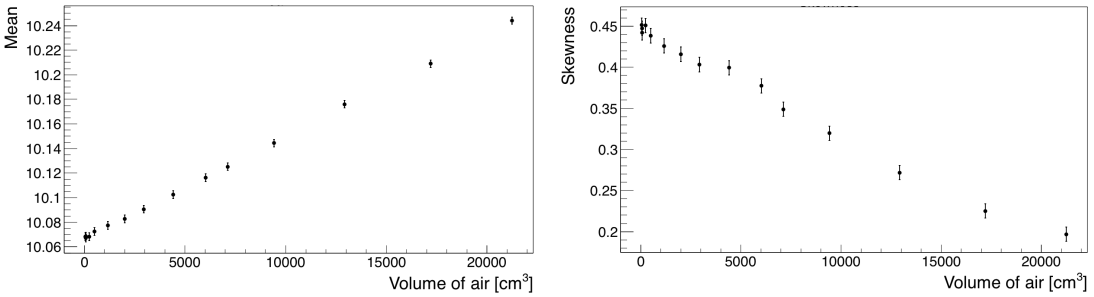

(a) Mean.

(b) Skewness.

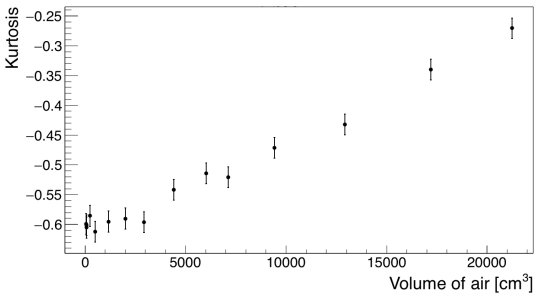

(c) Kurtosis.

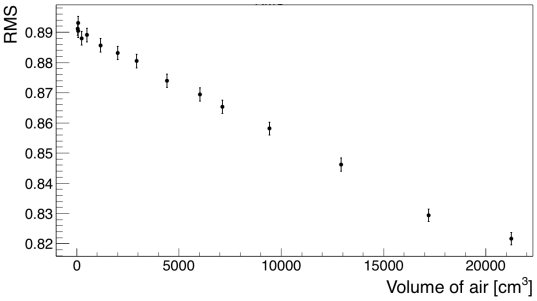

(d) RMS.

Figure 4.15: Variables obtained from the discriminator distribution. In each graph, the first point was obtained from a drum full of gas and the last point from a drum full of concrete.

Gas volume	Resolution (using the skewness)		Resolution (using the mean)	
(cm^3)	72M muons	159M muons	72M muons	159M muons
7,125	4.8%	30%	3.9%	0.9%
4,423	4.4%	88.2%	6.1%	3.8%
2,001	34.7%	103.8%	20.1%	3.5%

Table 4.2: Resolution of the bubble volume measurement obtained using the mean or the skewness, and using 72 million muons or 159 million muons.

A study was performed to determine the best fixed number of tracks for each voxel. Figure 4.16 shows the mean of the discriminator for 2 cm slices along the concrete drum, for a drum with one gas bubble and a drum with two bubbles, using different fixed numbers of tracks (15, 25, 35 and 45 tracks). It can be seen that using 35 tracks results in a better distinction between the drum with one bubble and the drum with two bubbles, than the other chosen values. This is a larger number of tracks than the one used for the edge finding method, which was expected because the chosen exposure time is longer.

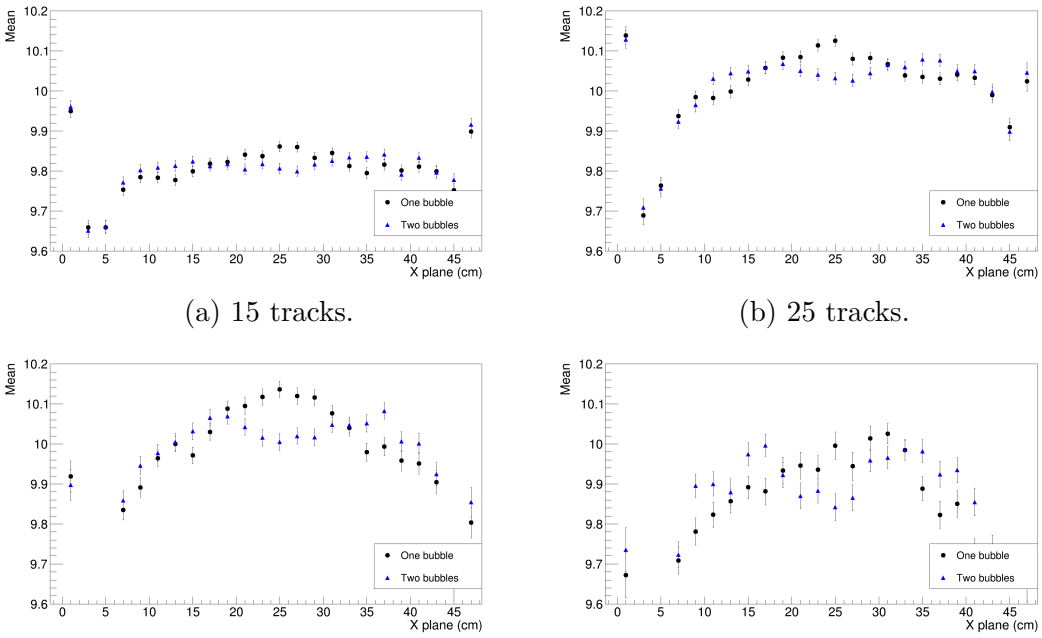

(a) 15 tracks.

(b) 25 tracks.

(c) 35 tracks.

(d) 45 tracks.

Figure 4.16: Mean of the discriminator for 2 cm slices of the volume in the x direction, for a drum with a 4.4 L gas bubble and a drum with two 2.2 L bubbles, for different fixed number of tracks.

4.3.1 Total amount of gas

To evaluate this method when applied to measure the gas volume present in a contained, several concrete containers were simulated, with air bubbles of different sizes. The container shape and size was the same as detailed in Section 3.1.

The simulated bubbles were cylindrical, with different volumes, with varying radius and length. Their dimensions are shown in Table 4.3. For each configuration, the metric distance distributions were calculated (using the 35 most scattered

tracks) and the discriminator value was obtained from those distributions. Finally, the mean of the discriminator value, μ_{discr} was calculated. Figure 4.17 shows a monotonic dependence of μ_{discr} on the volume of the gas bubble. It is clear that μ_{discr} depends only on the bubble volumes and not their shapes (which are shown by the different radii and lengths in Table 4.3). This dependence was parametrised and fitted with a linear function in a range starting at 1 L (because for smaller volumes there is no linear dependence). The line equation obtained for the mean as a function of the volume V was

$$\mu_{\text{discr}} = (6.0 \pm 0.2) \times 10^{-6}\ \text{V} + (10.010 \pm 0.002). \tag{4.7}$$

Each bubble volume was reconstructed by performing a linear fit excluding the corresponding point, and using the parametrisation obtained to calculate the volume for that point. Figure 4.18 shows the reconstructed bubble volumes as a function of the simulated volumes. These values are also shown in Table 4.3. Figure 4.19 shows the relative uncertainties on the volume reconstruction of bubbles of 1 L or more. It can be seen that for 1 L bubbles the relative uncertainty obtained was approximately 20%. The resolution for the volume reconstruction for bubbles of 2 L or more was $1.6 \pm 0.8\%$, which is the standard deviation obtained from the Gaussian fit to these uncertainties (Figure 4.19). This shows that it is possible to reconstruct these gas volumes with a good resolution.

The measurement of the total gas volume in a concrete container was performed using a single cylindrical gas bubble, placed in the centre of the drum, for each volume measured. However, it is important to determine if the results would be influenced by the position and shape of the bubble. The different shapes of these cylinders, shown in Table 4.3, suggest that the method does not depend on the shape of the bubble. To confirm this, the study was also performed using spherical bubbles, comparing them to the cylindrical bubbles. In order to determine if the location of the bubble in the drum would alter the results, the study was also done with a spherical bubble shifted away from the centre of the concrete drum.

Radius (cm)	Length (cm)	Volume (L)	Reconstructed volume (L)
2	4	0.050	0.2 ± 0.2
2	5	0.063	0.3 ± 0.2
2	6	0.075	0.1 ± 0.2
3	9	0.254	0.2 ± 0.2
4	10	0.503	0.7 ± 0.2
5	15	1.178	1.3 ± 0.2
7	13	2.001	2.0 ± 0.2
7	19	2.925	3.0 ± 0.2
8	22	4.423	4.3 ± 0.2
8	30	6.032	6.0 ± 0.2
9	28	7.125	7.1 ± 0.2
10	30	9.425	9.4 ± 0.3
11	34	1.2925	13.1 ± 0.3
12	38	17.191	17.1 ± 0.4
13	40	21.237	21.3 ± 0.4

Table 4.3: Radius, length and volume of the simulated cylindrical air bubbles and respective reconstructed volumes.

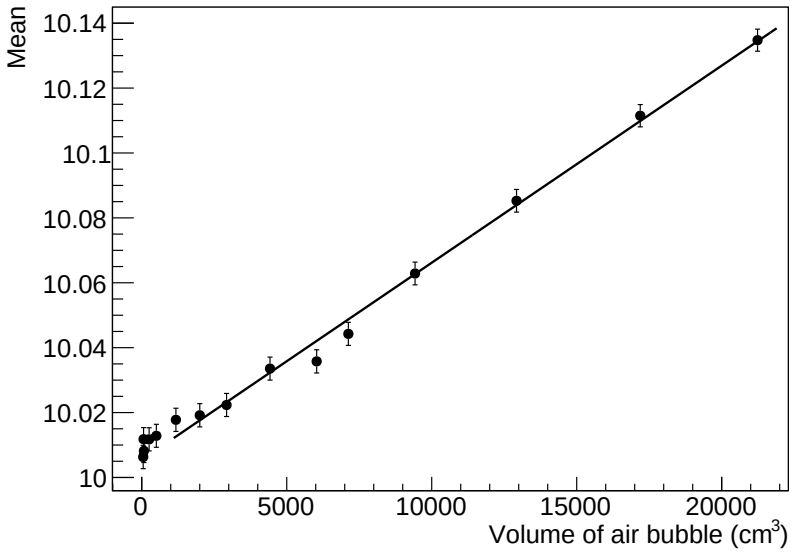

Figure 4.17: Mean of the discriminator distribution as a function of the volume of the air bubble.

Finally, the method was applied to simulations with two spherical bubbles with the same volume, to show if the total gas volume can be measured when it is split in two bubbles. Figure 4.20 shows the results from these studies. The values of μ_{discr} obtained for these four different geometries are the same within errors, demonstrates that the method does not depend on the shape or position of the

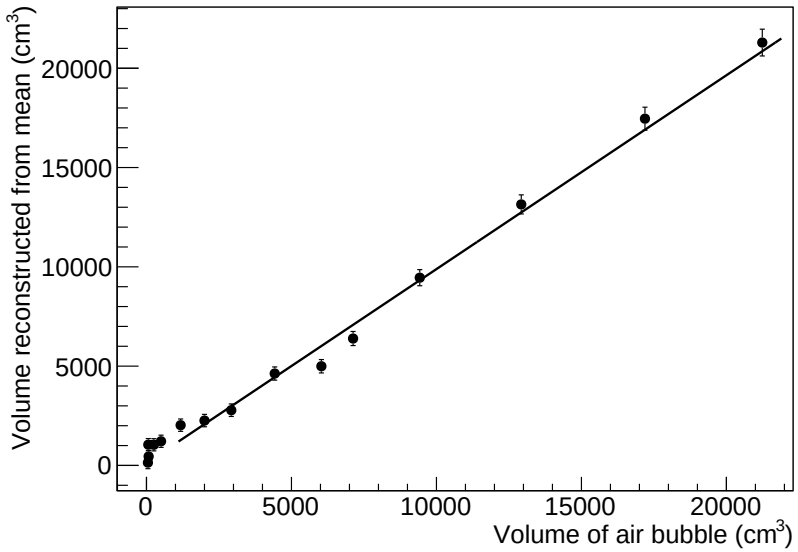

Figure 4.18: Reconstructed volume as a function of the generated volume of the bubble using the mean of the discriminator distribution.

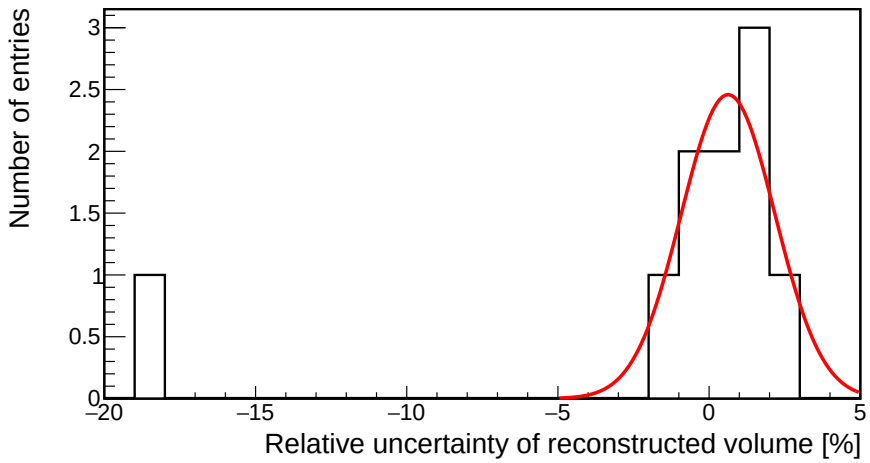

Figure 4.19: Relative uncertainties on the reconstructed bubble volume, for bubbles of 1 L or more. The parameters obtained from the Gaussian fit (shown in red) were $\mu = 0.6 \pm 0.7\%$ and $\sigma = 1.6 \pm 0.8\%$. The entry on the left of the graph is the relative uncertainty on the volume reconstruction of a 1 L bubble, which was excluded from the calculation of this resolution.

bubbles. This shows the robustness of the method and its applicability to real waste drums.

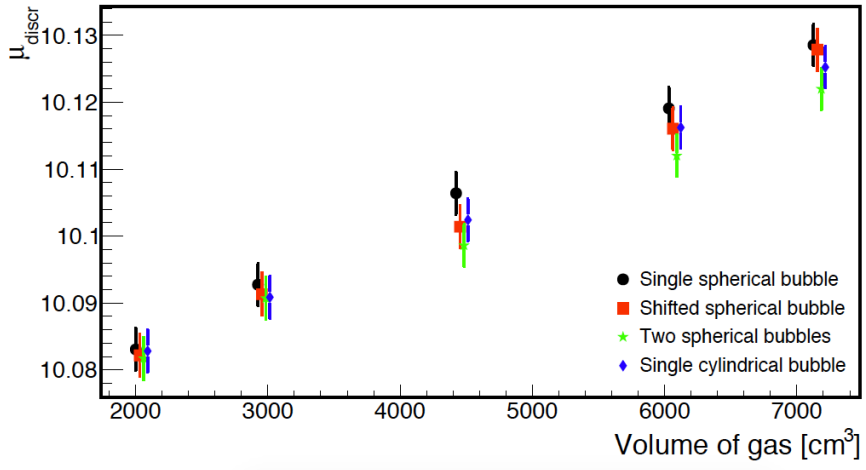

Figure 4.20: Mean of the discriminator distribution as a function of the volume of the air bubble for different simulation geometries: a single spherical bubble, a spherical bubble shifted from the centre, two spherical bubbles (each with half of the total gas volume) and a single cylindrical bubble. The actual volume of gas corresponding to each point is the one from the single spherical bubble series, and a small horizontal offset was applied to the other points to make the results visible.

4.3.2 Localisation of gas bubbles

In the previous results the amount of gas in a single bubble in a container was measured. It was also shown that the method gives the same results when measuring the total gas volume when it is split into two bubbles. Therefore, a different technique is required if one wants to determine if the gas is in only one bubble, or split in more than one smaller bubbles. This is important for the risk assessment of each container, since a large bubble poses a higher risk than several separated smaller bubbles. To determine the position of the bubbles, the mean of the discriminator was plotted for slices of the container. The simulations performed for this study used spherical bubbles. Figure 4.21 shows the mean for 2 cm slices along X for a concrete container with one 3 L bubble (which is seen in the middle) and for a container with two separate bubbles of 1.5 L each, as well as for a container with only concrete. Figure 4.22 shows the same thing, but for a 4.4 L bubble compared to two 2.2 L bubbles. The bracket lengths correspond to

the diameter of each bubble. The edges of the graphs have a higher μ_{discr} because these sections are located in the air outside the container, and the points with lower μ_{discr} just inside the edges correspond to the steel caps of the drum. Slices in the concrete have a higher μ_{discr} value, and it can be seen that the μ_{discr} is even higher where bubbles are located. In the slices outside the bubbles, and at the edges of the bubbles (which correspond to the edges of the brackets), it can be seen that the μ_{discr} values are the same within errors for the drum with only concrete and the ones with bubbles. However, in the centre of the bubbles, the μ_{discr} value is higher than the value obtained for the concrete.

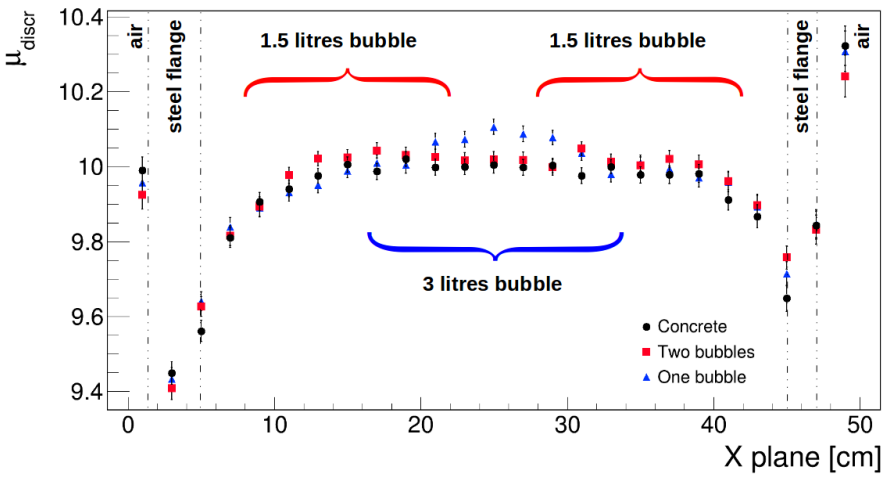

Figure 4.21: Mean of the discriminator for 2 cm slices of the volume in the x direction, for a drum with a 3 L gas bubble, and a drum with two separate bubbles of 1.5 L bubbles. Each bracket length represents the diameter of the corresponding bubble.

To determine the location of the bubbles, the significance of the difference between μ_{discr} from the container with bubbles and from the container with only concrete was calculated. These values can be seen in Figures 4.23 and 4.24, for a total amount of gas of 3 L and 4.4 L, respectively. A value of 3σ was chosen as a threshold to detect bubbles. In Figure 4.23a, the bubble is found between the 25 and the 27 cm slices, and in Figure 4.24a, it is found between the 21 and the 31 cm slices). In Figures 4.23b and 4.24b, where the same volume of gas is divided in two separate bubbles, the calculated differences are lower and, although two regions of higher differences can be seen, one for each bubble, these values are not over 3σ

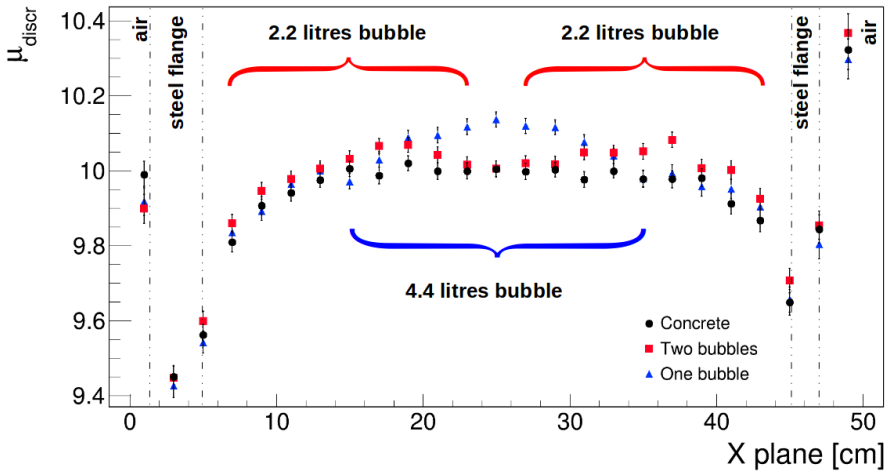

Figure 4.22: Mean of the discriminator for 2 cm slices of the volume in the x direction, for a drum with a 4.4 L gas bubble, and a drum with two separate bubbles of 2.2 L bubbles. Each bracket length represents the diameter of the corresponding bubble.

(except for one of the slices from the right-side bubble in Figure 4.24b). It has been shown in the previous section, in particular with Figure 4.20, that the total amount of gas can be measured even if it is divided in two bubbles. Therefore, if the total volume of gas from the containers in Figures 4.23b and 4.24b is measured, it is possible to conclude from these graphs that the gas is not located in one single bubble. This proves that this technique can distinguish between a container with one single bubble (or a big conglomerate of bubbles) and the same gas volume split in separate regions of the container, which is important for the risk assessment of the drums.

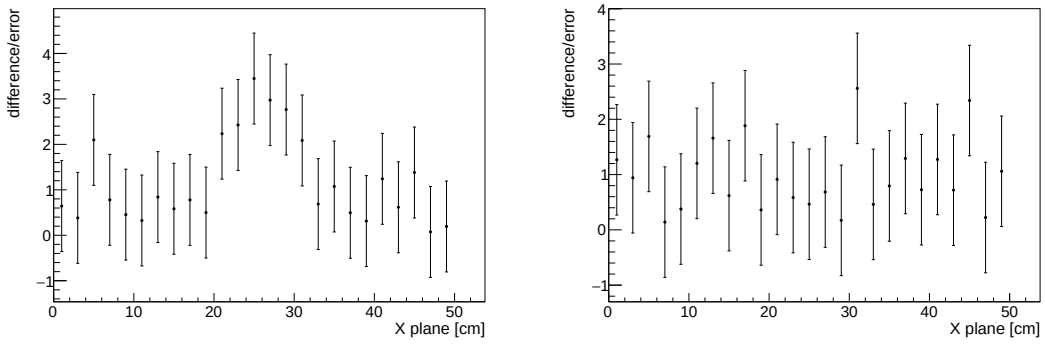

(a) (b)

Figure 4.23: Difference of μ_{discr} in each slice, between a 3 L bubble and a concrete container (a) and between two 1.5 L bubbles and a concrete container (b), divided by its error.

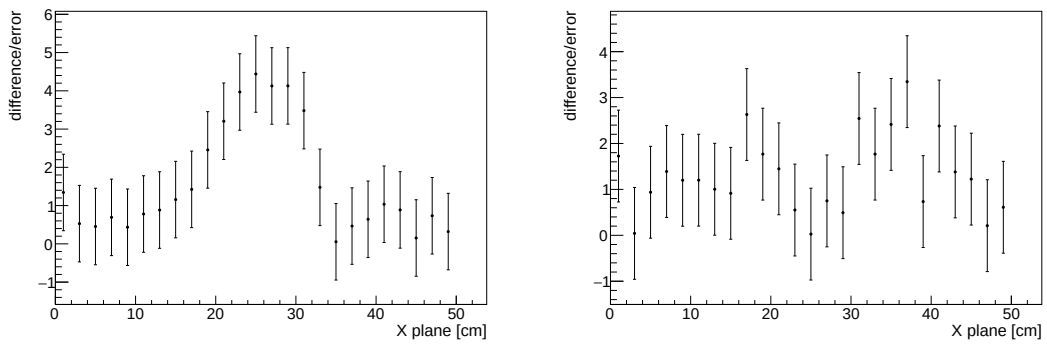

(a) (b)

Figure 4.24: Difference of μ_{discr} in each slice, between a 4.4 L bubble and a concrete container (a) and between two 2.2 L bubbles and a concrete container (b), divided by its error.

Additionally, the bubble shown in Figure 4.24a was also simulated in different positions in the drum, shifted from the centre by 2, 4, 6, 8 and 10 cm. This can be seen in Figure **??**. It is clear that a shift from the centre makes the bubbles harder to detect, shown both by having fewer points over 3σ and by the larger spread of values in the bubble region, when compared to Figure 4.24a. This is due to the smaller angular resolution and the lower acceptance closer to the edges of the scanned volume than at the centre.

In order to make the waste drum more realistic and show that it is still possible to locate a gas bubble next to uranium, a $3 \times 3 \times 3$ cm^3 uranium block was added to the geometry, 1 cm away from the edge of the bubble. Figure 4.26 shows the results obtained for this volume, compared to a drum containing only concrete. The uranium block is found where μ_{discr} is lower than the concrete baseline (slices at 13-15 cm) and the bubble where it is above (slices at 23-27 cm).

4.3.3 Time needed to detect a gas bubble

In order to determine how long a container needs to be scanned for a gas bubble to be found, plots were obtained for different scanning times. These can be seen in Figure 4.27. The difference of μ_{discr} between each slice of the waste drum containing the 4.4 L bubble and the drum containing only concrete (in Figure 4.22)

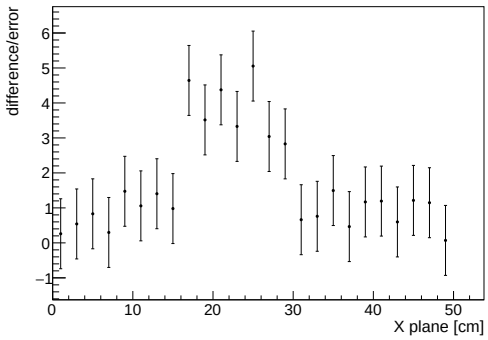

(a) 2 cm shift.

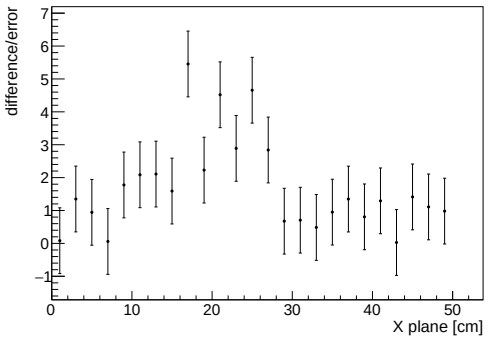

(b) 4 cm shift.

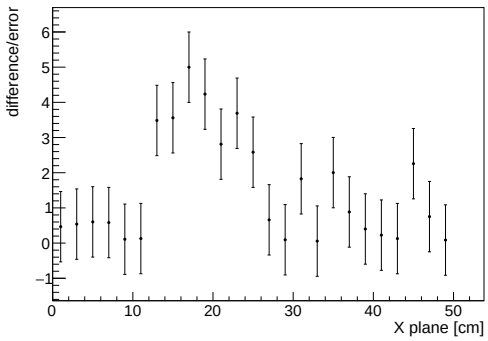

(c) 6 cm shift.

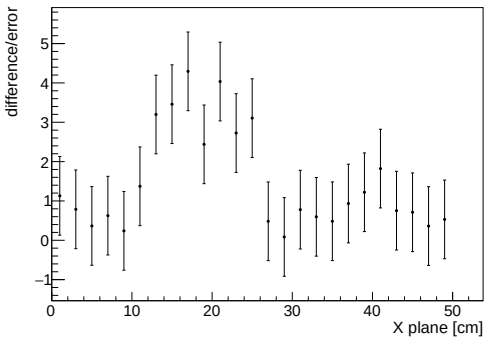

(d) 8 cm shift.

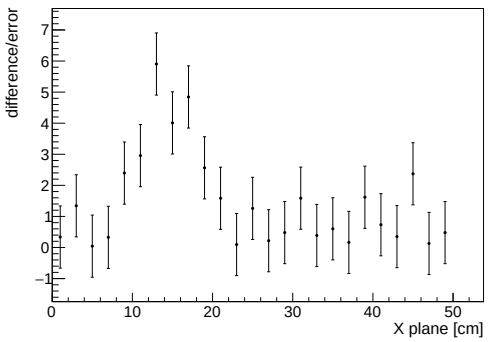

(e) 10 cm shift.

Figure 4.25: Difference of μ_{discr} in each slice, between a 4.4 L bubble and a concrete container, for different positions of the bubble in the drum: shifted from the centre by 2, 4, 6, 8 and 10 cm.

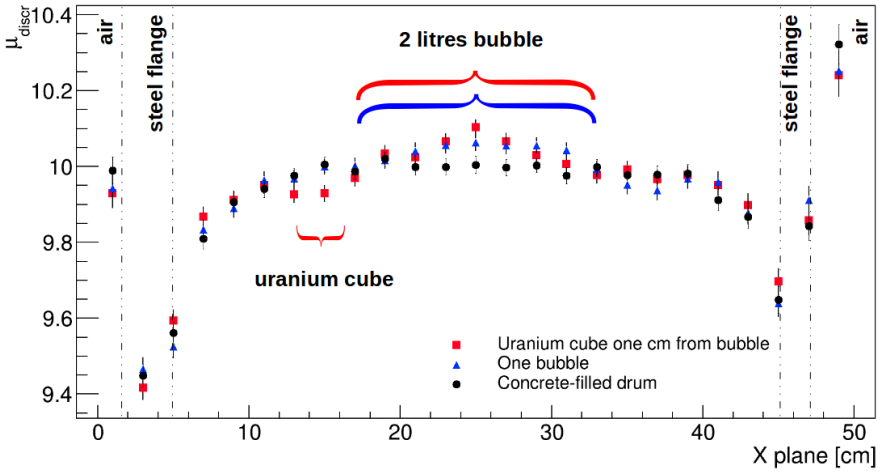

Figure 4.26: Mean of the discriminator for 2 cm slices of the volume in the x direction, for a drum with a 2 L gas bubble and a uranium cube 1 cm away from the edge of the bubble, and a drum with only concrete. Each bracket length represents the diameter of the corresponding bubble or uranium block length respectively.

was calculated (only for the slices inside the drum). This difference was calculated for each slice, then integrated and divided by its error, as

$$\text{weighted difference} = \frac{\sum\limits_{\text{slices}} \left(\mu_{\text{bubble}} - \mu_{\text{concrete}}\right)}{\sqrt{\sum\limits_{\text{slices}} \left(\sigma^2_{\mu_{\text{bubble}}} + \sigma^2_{\mu_{\text{concrete}}}\right)}}, \tag{4.8}$$

where μ_{bubble} and μ_{concrete} are the μ_{discr} from a slice in the drum with a gas bubble or in the drum with only concrete, respectively. This weighted difference is shown in Figure 4.28 as a function of the scanning time. If the minimum value to detect a bubble is chosen to be 3σ, the graphs shows that it takes about 9 days to find a 4.4 L bubble. It can be seen that for longer scanning times than 11 days the ratio of the integrated difference remains approximately constant, as the difference measurement itself does not improve anymore. The point at 5.4 days shows a negative integrated difference, which is due to statistical fluctuations in the simulations. This is proven by Figure 4.27b, which shows several points corresponding to the concrete to be below the uranium, but most of which within error bars.

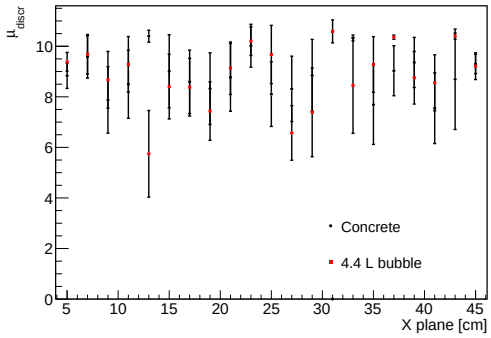

(a) 3.5 days.

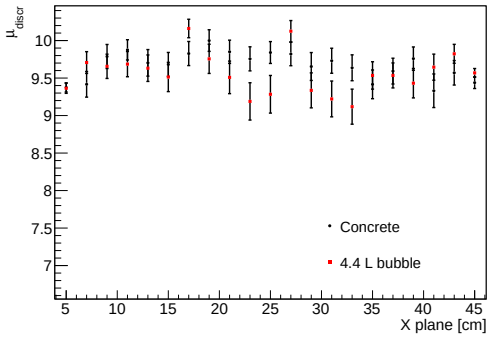

(b) 5.4 days.

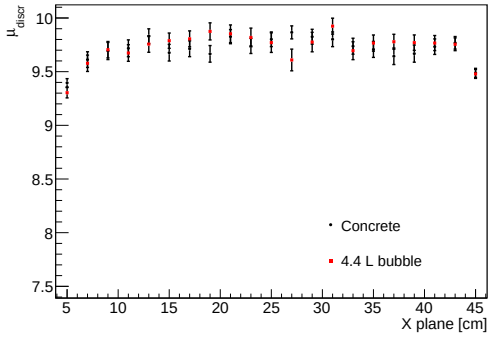

(c) 7.1 days.

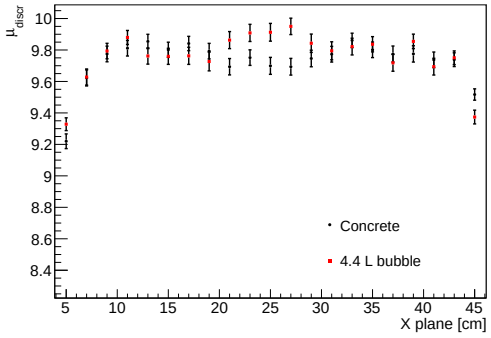

(d) 8.9 days.

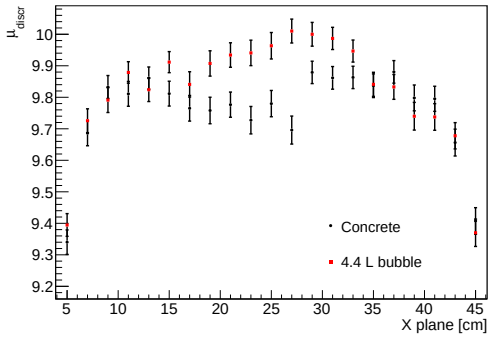

(e) 10.7 days.

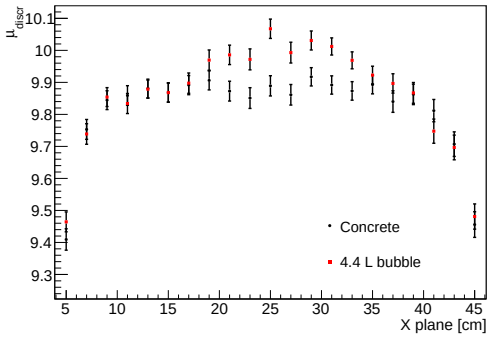

(f) 12.5 days.

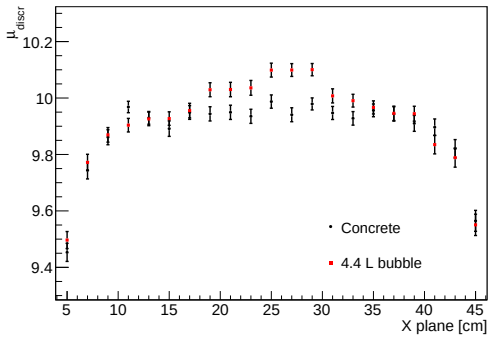

(g) 14.3 days.

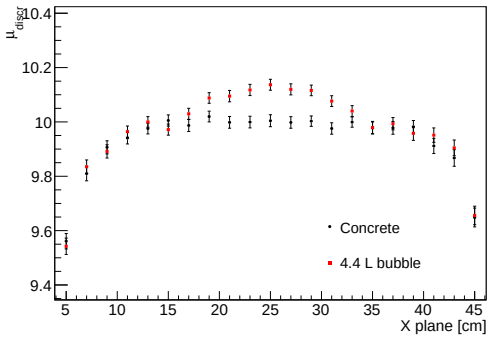

(h) 15.8 days.

Figure 4.27: Mean of the discriminator for 2 cm slices of the volume in the x direction, for a drum with a 4.4 L gas bubble and a drum with only concrete for different scanning times. Only points for the slices inside the drum are shown, because these were the ones used to calculate the integrated difference.

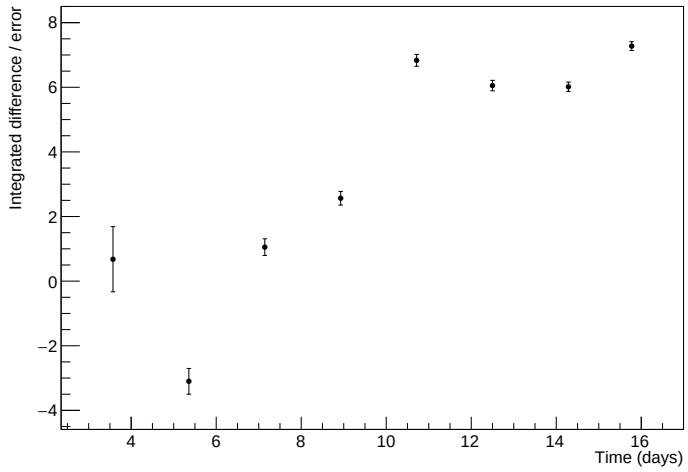

Figure 4.28: Integrated difference of μ_{discr} divided by the error in each slice, between a 4.4 L bubble and a concrete container.

4.4 Discussion and Conclusions

It has been shown in this chapter that it is possible to use MST to perform imaging of concrete nuclear waste containers, in particular to find lumps of high-Z materials, defining edges of uranium blocks, measuring the amount of gas present in a waste drum, and finding large gas bubbles.

The imaging results using the metric distance method, showing lumps of high-Z materials, had been previously developed, and published in [32], as well as a first edge finding method, with a resolution of $\sigma = 1.2 \pm 0.4$ mm, and with a smallest length measured of 5 mm. However, these results used the true momentum information, which is not a realistic assumption since most muon tomography systems cannot measure the momentum, or do not have a good resolution. For example, a momentum resolution of 30% to 200% is expected if multiple Coulomb scattering were used to measure the momentum. Therefore, a new edge finding method was developed, and tested both with and without momentum information. The resolution obtained without momentum was $\sigma = 3.2 \pm 0.6$ mm, and the smallest length was also 5 mm. The resolution obtained when using momentum information was $\sigma = 0.98 \pm 0.03$ mm, which is an improvement over the previous

method. The smallest uranium block length was also improved from 5 mm down to 2 mm.

Another possible future study is to apply the method to irregular shapes, which could be done by performing the measurements along x but for several different y and z values, resulting in different lengths and allowing to better define different shapes.

A possible future development is the combination of the 5 data points in a way that does not only take their average. For example, if one of the points has a value of 1, and the others a value closer to 0, it is still likely that it corresponds to a uranium voxel, since concrete voxels have a very low probability of obtaining a value of 1 (which can be seen in the graphs from the MVA training). Other future developments include performing the same analysis method using other materials instead of uranium and reduce the data taking time.

It had been shown that the metric distance method was also sensitive to a gas volume enclosed in concrete [32], but a more detailed study on its application to low-Z materials had not been developed. Therefore, a study was also shown in this chapter on measuring the gas content of concrete containers. This is an important application because the oxidation of uranium forms hydrogen gas, which accumulates in the drums and causes them to overflow. Results from simulations of gas bubbles show that it is possible to use the metric distance method to measure the volume of gas in concrete, for bubbles of 2 L or more, achieving a resolution of $1.6 \pm 0.8\%$ on the volume. Furthermore, it is possible to localise a single gas bubble, for example with 3 L, distinguishing it from a scenario where the same volume of gas was separated in two smaller 1.5 L bubbles. This is particularly important because the same gas volume can pose a higher risk when it is located in a larger bubble than when separated in smaller ones.

In summary, it has been shown that a muon tomography clustering algorithm, based on the density of high scatter angles, can be used not only to perform

imaging of high-Z materials, but also to find gas bubbles in nuclear waste. These are both important applications for the characterisation of nuclear waste.

5 | Discrimination of high-Z materials

It was shown in Chapter 4 that it is possible to find high-Z materials in concrete. The next step for nuclear waste characterisation is to identify these materials. Distinguishing between different high-Z materials, such as uranium, plutonium, lead and tungsten, is challenging, since they have similar radiation lengths. There have been other attempts to identify materials, such as in [20], but for high-Z materials this distinction had not been shown yet.

In this chapter a method is presented to distinguish between two high-Z materials. The majority of this research has been published in [56]. This method is based on the premise that the position of the high-Z materials is known. Chapter 4 shows an imaging method that allows for the position information to be obtained with a good resolution. The discrimination between two materials is the first step towards material identification. The principle is that an unknown material can be compared to known data from standard high-Z materials. Any material from which it is discriminated can be excluded, and if there is only one material from which the unknown object is not discriminated, it can be determined that this is the material present in the unknown sample.

Section 5.1.1 describes the material discrimination method using multivariate analysis, and explains how the performance of this method is assessed. Section 5.1.2 shows the results of discrimination between uranium and lead, uranium and tungsten, and uranium and plutonium for different block sizes and data taking times. Section 5.1.3 shows how the material discrimination is improved by adding muon tracks reconstructed in the concrete region around the target block. It also determines the ideal volume of this extra region that maximises the discrimination. Section 5.1.4 discusses how the material discrimination results were compared between the case with perfect momentum information, adding an error

to the momentum, and removing the momentum information altogether. It shows how including the momentum information does not improve the material discrimination. Section 5.1.5 shows the comparison between using only the 3D scatter angle and using the MVA. Discussion and conclusions are given in Section 5.2.

5.1 Material discrimination simulations

The simulations were performed using Geant4, with the same settings, as well as detectors, concrete tube, cosmic muon generator and fit of muon hits as described in Section 3.1. Only the sample materials inside the concrete were different. The materials used were lead, tungsten, uranium and plutonium. These were chosen to be high-Z materials that could potentially be found in nuclear waste. In particular, it is important to determine the presence or absence of uranium and plutonium. The block shapes were all cubic, with edges of 1 cm, 2 cm, 3 cm, 5 cm and 10 cm. Different data taking times were compared as well, in order to determine for how long a drum with a high-Z object needs to be scanned, depending on the size of the object. A multivariate analysis method, presented in Chapter 3, was used to combine information from several variables obtained from the track fit, in order to distinguish pairs of materials.

The momentum information was used because most MST analysis methods use it (as it was shown in the description of some of these methods in Chapter 2). The Monte-Carlo truth momentum was smeared by adding a random value from a Gaussian distribution centred around zero and with 50% of the true momentum as the standard deviation (as in [18]). All the results used this smearing, except where stated otherwise. Additionally, these results were compared to the ones using the true momentum and not using any momentum information. As will be shown in Section 5.1.4, the discrimination performance is similar with and without momentum information.

5.1.1 Method for material discrimination

Several variables are obtained from the fitting of muon tracks, such as scatter angles, offsets and χ^2 from the fit, that can be sensitive to different materials. A multivariate analysis (MVA) was performed in order to combine all these variables to distinguish between two materials, namely uranium and lead, uranium and tungsten, and uranium and plutonium. The MVA method used was the Fisher linear discriminant. Details on MVAs and Fisher linear discriminants can be found in Section 3.3. The training was done using the same geometry as the application.

5.1.1.1 Variables for MVA

To choose the variables to use, simple simulations of different materials (uranium, tungsten and lead) irradiated by mono-energetic 1 GeV, 2 GeV and 10 GeV vertical muons were performed. The geometry and resolution of the detectors were the ones described in Chapter 3. Two geometries were used for each sample: a $10 \times 10 \times 10$ cm^3 block, and a plate with thickness of 6 mm and the same area as the detector plates, in order to guarantee that all detected muons crossed the material in question. Several variables obtained from the track fitting of those simulations were compared in order to choose the ones most sensitive to different materials. These variables were the projected and 3D scatter angles, offsets, vertex positions, errors on the vertex position, χ^2 from different fits (upper and lower track fits, 6 point fit with both upper and lower points, and combined fit of the upper and lower hits and the vertex), and distance between the reconstructed vertex and the lower and upper tracks (labelled track distance). Graphs of all these variables can be found in Appendix A. This first comparison used all the variables that can be obtained from the track fitting without any additional analysis. A significance value was defined as

$$\text{Significance} = \frac{1}{N} \sum_{\text{i=1}}^{N} \frac{S_i}{\sqrt{B_i^2 + S_i^2}}, \tag{5.1}$$

where S and B are signal and background, and N is the number of bins in the distribution. Examples of the values obtained, for the case where uranium was considered "signal" and lead "background", can be seen in Table 5.1, which were obtained using the 6 mm thick plate. Both tables show the significance values for each variable, but the table on the left is in descent order of the significance values. This makes clear the gap in significance between the variables with a value over 0.4 and the ones with a value under 0.2. Therefore, variables with significance values over 0.4 were chosen for the MVA. Although the χ^2 from combined fit shows a lower value in Table 5.1, it had a significance above 0.4 for the simulations with a $10 \times 10 \times 10$ cm^3 block, so it was also included in the MVA.

Variable	Significance
Angle (x)	0.565775
Angle (y)	0.531012
3D Angle	0.484635
Upper offset (x)	0.439282
Upper offset (y)	0.443182
Lower offset (x)	0.536945
Lower offset (y)	0.587794
Longitudinal error	-0.0627532
Transversal error (x)	0.132827
Transversal error (y)	0.13854
Vertex position (x)	-0.0462523
Vertex position (y)	-0.0240665
Vertex position (z)	-0.176026
χ^2 from upper fit (x)	0.122719
χ^2 from upper fit (y)	0.137903
χ^2 from lower fit (x)	0.191506
χ^2 from lower fit (y)	0.123077
χ^2 from six point fit (x)	0.610759
χ^2 from six point fit (y)	0.639472
χ^2 from combined fit	0.13668
Upper track distance (x)	0.0568133
Upper track distance (y)	0.124223
Lower track distance (x)	0.121605
Lower track distance (y)	0.13497

Variable	Significance
χ^2 from six point fit (y)	0.639472
χ^2 from six point fit (x)	0.610759
Lower offset (y)	0.587794
Angle (x)	0.565775
Lower offset (x)	0.536945
Angle (y)	0.531012
3D Angle	0.484635
Upper offset (y)	0.443182
Upper offset (x)	0.439282
χ^2 from lower fit (x)	0.191506
Vertex position (z)	-0.176026
Transversal error (y)	0.13854
χ^2 from upper fit (y)	0.137903
χ^2 from combined fit	0.13668
Lower track distance (y)	0.13497
Transversal error (x)	0.132827
Upper track distance (y)	0.124223
χ^2 from lower fit (y)	0.123077
χ^2 from upper fit (x)	0.122719
Lower track distance (x)	0.121605
Longitudinal error	-0.0627532
Upper track distance (x)	0.0568133
Vertex position (x)	-0.0462523
Vertex position (y)	-0.0240665

Table 5.1: Significance values (equation 5.1) obtained from the distributions of each variable, comparing simulations of 2 GeV muons crossing a plate of uranium and a plate of lead. Both tables show the same values, but the table on the left is in descent order of the significance value.

The chosen variables were then used as inputs into a multivariate analysis to compare pairs of materials. There were 11 chosen variables, which were the scatter angles (the 3D angle and the 2D projections in x and y), the lower and upper offsets in x and y (distance between the extrapolated track and the real hit, as shown in Figure 2.5) the χ^2 of a linear fit to all six hit points (in x and y), and the χ^2 of the combined fit of incoming and outgoing tracks, including the reconstructed vertex. The muon momentum was initially also included because it was expected to enhance sensitivity to the material, as will be explained in Section 5.1.4. However, the momentum cannot be easily obtained without additional analysis apart from the track fitting.

An example of each type of variables used in the MVA is shown in Figures 5.1-5.6, normalised, and both in linear and logarithmic scales. It can be seen that the all distributions for uranium (except the momentum) are shifted to higher values than the lead. The 3D angle is the variable whose distributions for uranium and lead are more clearly different. The muon momentum was not expected to show any discrimination by itself, since it does not depend on the material the muon crosses.

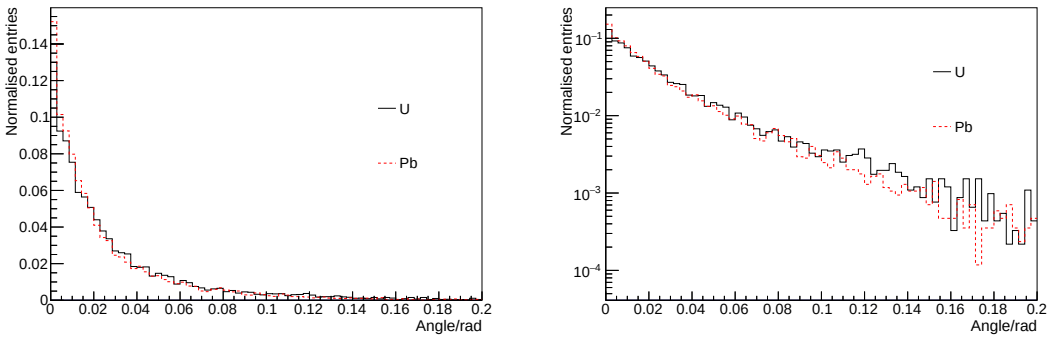

(a) (b)

Figure 5.1: Projected angle in x for uranium and lead, for a $10 \times 10 \times 10$ cm^3 block, in linear (a) and log (b) scales. The histograms are normalised.

The 3D scatter angle is linearly correlated to all the other variables used, which can be seen in Figure 5.7, which shows the correlation coefficients for all variables, from the training using a $10 \times 10 \times 10$ cm^3 uranium block as signal and a lead

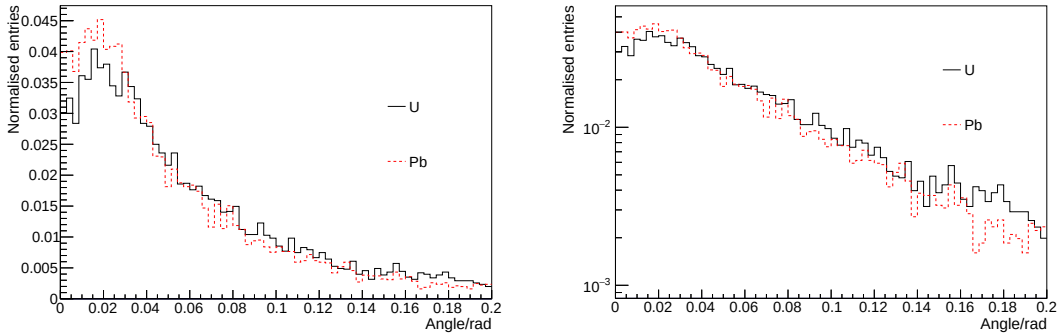

(a) (b)

Figure 5.2: 3D angle for uranium and lead, for a $10 \times 10 \times 10$ cm^3 block, in linear (a) and log (b) scales. The histograms are normalised.

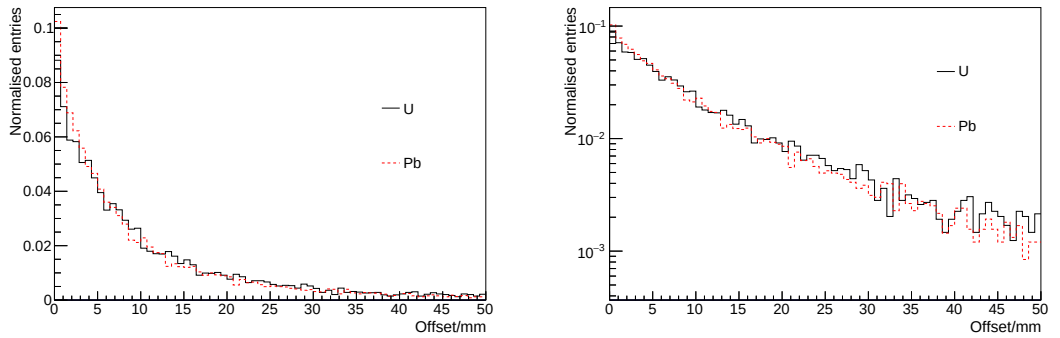

(a) (b)

Figure 5.3: Upper offset in x for uranium and lead, for a $10 \times 10 \times 10$ cm^3 block, in linear (a) and log (b) scales. The histograms are normalised.

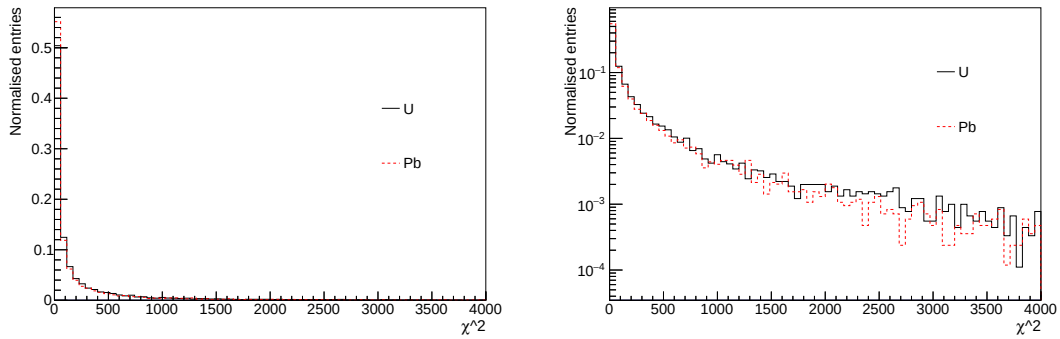

(a) (b)

Figure 5.4: Six points fit χ^2 (in x) for uranium and lead, for a $10 \times 10 \times 10$ cm^3 block, in linear (a) and log (b) scales. The histograms are normalised.

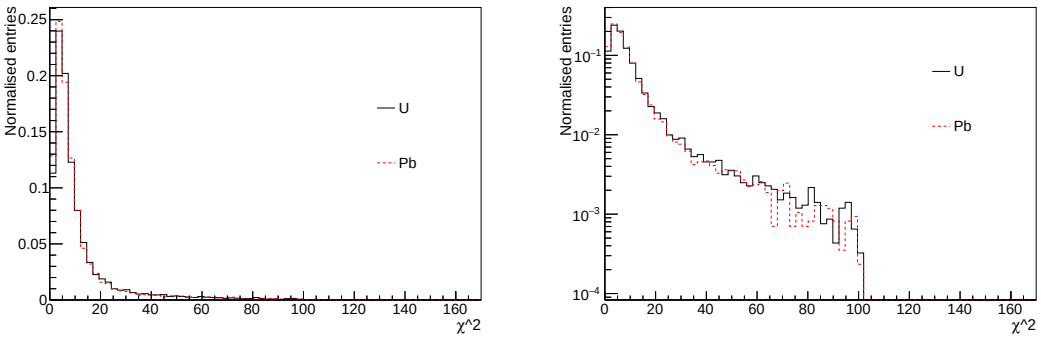

(a) (b)

Figure 5.5: Combined fit χ^2 for uranium and lead, for a $10 \times 10 \times 10$ cm^3 block, in linear (a) and log (b) scales. The histograms are normalised.

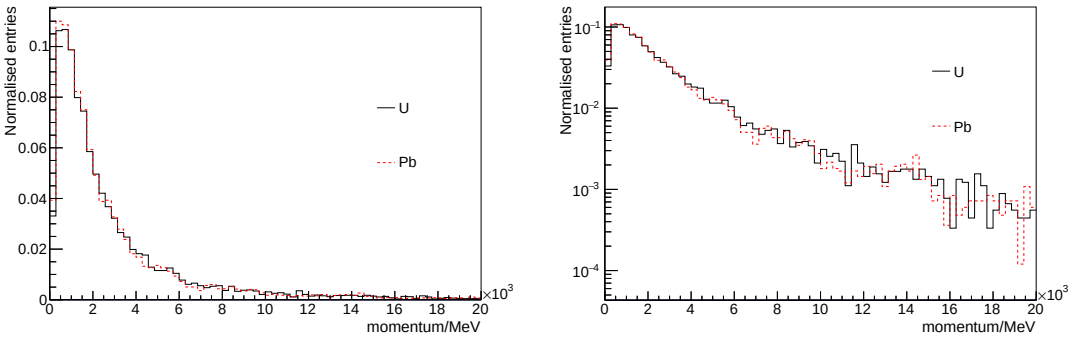

(a) (b)

Figure 5.6: Muon momentum for uranium and lead, for a $10 \times 10 \times 10$ cm^3 block, in linear (a) and log (b) scales. The histograms are normalised.

block the same size as background. Its linear correlation coefficients with the 2D angles and the offsets are over 80%. The 3D angle has correlation coefficients with the different χ^2 of between 50% and 70%, and with the momentum of 22% (in absolute value). Variables that are highly correlated (such as 3D angle, 2D angles and offsets) may be redundant for the linear discriminant, and may even diminish the discrimination power. The comparison done in Section 5.1.5 between using only the 3D angle and the multivariate analysis shows that, in fact, using the 3D angle would be enough to achieve material discrimination.

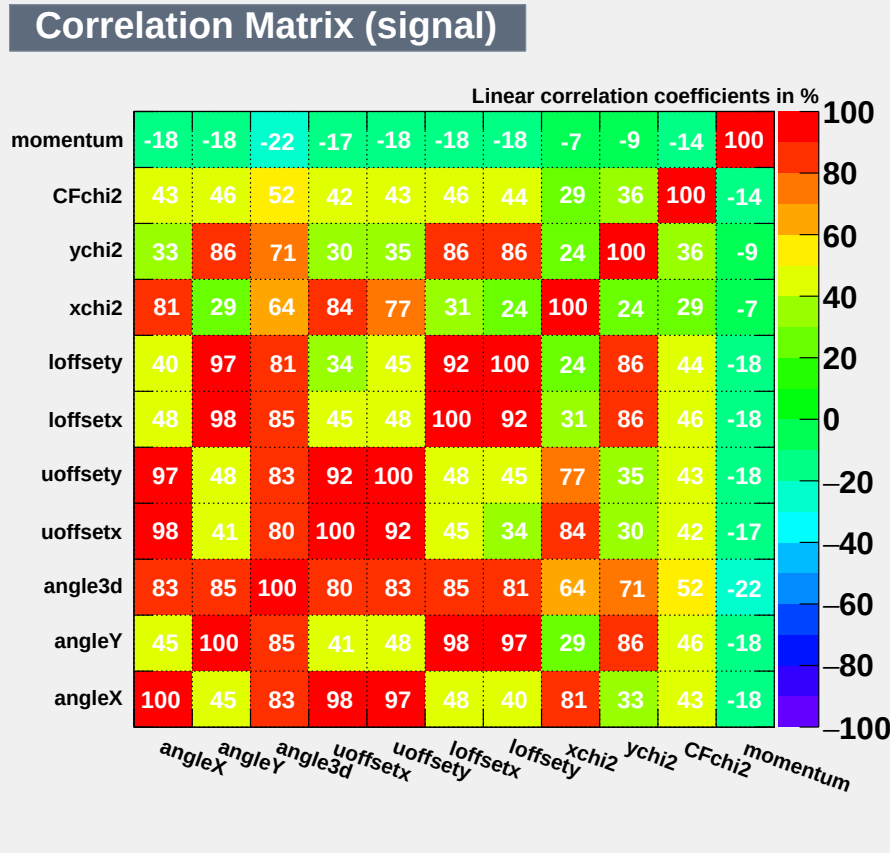

Figure 5.7: Correlation coefficients from the training for the comparision between a $10 \times 10 \times 10$ cm^3 uranium block and a lead block.

5.1.1.2 Application of MVA

As shown in Figure 4.2, the metric method yields an approximate position of the material and the method presented in Chapter 4 allows precise determination of the position of the material. Therefore, only the tracks that passed through the block were used for the training and testing. They were selected by registering, for

each simulated muon, if it entered the target block or not. For the application of MVA, the tracks were selected by choosing the ones whose reconstructed vertex was inside the target block (for 5 cm and 10 cm side cubes) or inside a small volume that included the whole block (for 2 cm and 3 cm side cubes). Additionally, for the 2 cm and 3 cm side cubes, different volume sizes were studied, as shown in Section 5.1.3, starting from the exact size of the block in question and then increasing the edge length in all sides by 1, 2, 3 and 4 cm. Due to the results obtained in that study, the volume in which the tracks were selected for the 2 cm and 3 cm side cubes included an extra 1.5 cm and 2 cm, respectively, from each edge of the cubes. The discrimination method can be executed using exactly the same muon sample as the imaging method, hence no new data needs to be acquired.

The number of muons used for training was different for each block size: for $10 \times 10 \times 10$ cm^3 and $5 \times 5 \times 5$ cm^3 blocks, 10h of muon exposure were used, which yielded about 13000 and 3300 selected muons respectively, 50h for the $3 \times 3 \times 3$ cm^3 and $2 \times 2 \times 2$ cm^3 blocks yielded about 6000 and 2600 muons, and 100h for the $1 \times 1 \times 1$ cm^3 block yielded about 1400 muons. To see how the amount of data affected the results, some training was performed with less data.

An example of the output of the training for the Fisher linear discriminant comparing $10 \times 10 \times 10$ cm^3 blocks uranium and lead is shown in Figure 5.8. Several values from this distribution were tested as potential discriminators between the two materials, such as the mean, the integral over a cut, the maximum value, the integral weighted more for higher values, and the weighted integrals over and under a cut. It was observed that the mean offered the best discrimination, which will be shown later in Section 5.1.1.3.

5.1.1.3 Analysis of discrimination performance

After training, the Fisher linear discriminant was applied to sets of events, where an event denotes one track of a simulated muon and a set of events is a collection of muon tracks, which corresponds to a fixed time of muon exposure.

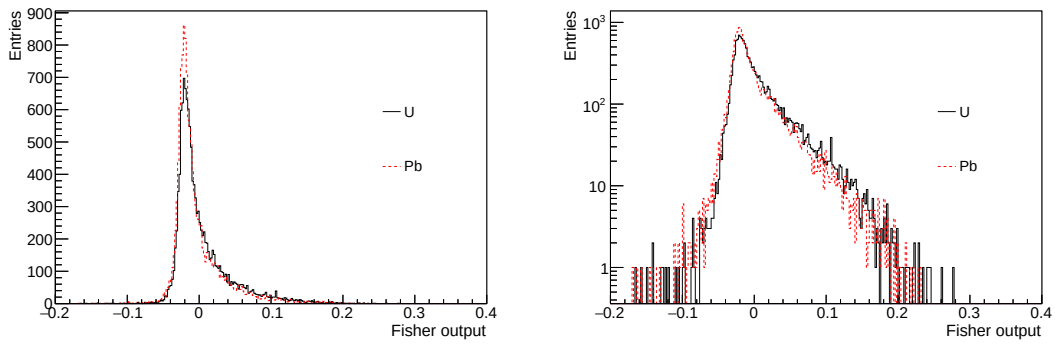

(a) (b)

Figure 5.8: Output of the Fisher MVA method, comparing uranium and lead, on linear (a) and logarithmic (b) scales for a $10 \times 10 \times 10$ cm^3 block.

Several sets with the same time of muon exposure were created. A distribution of the Fisher output was obtained for each set, and the respective mean was calculated and plotted as an entry in a histogram like in Figure 5.9a, where each entry corresponds to 1h of muon exposure. This process was also done for sets with other exposure times. These results will be discussed in Section 5.1.2.

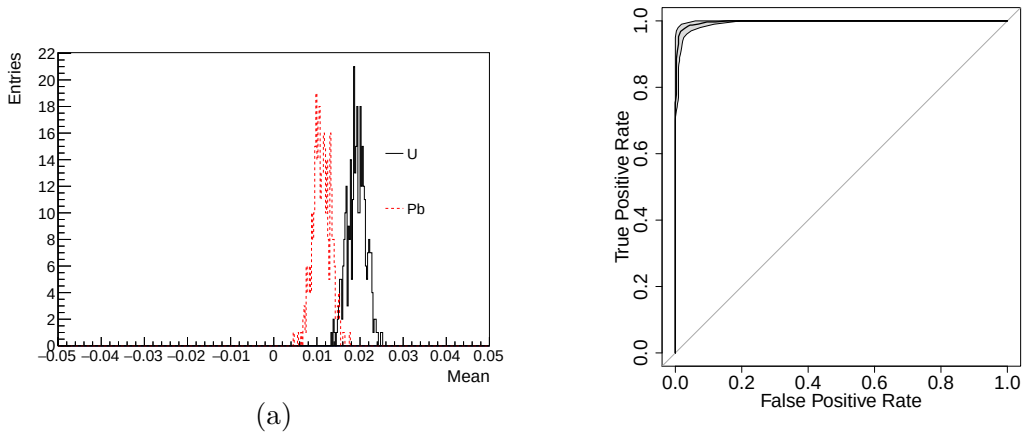

(a) (b)

Figure 5.9: Mean value (a) and ROC curve (b) from the Fisher discriminant method, distinguishing uranium from lead, for a $10 \times 10 \times 10$ cm^3 cube and 1h of muon exposure.

In order to assess the performance of this discriminator, ROC curves (Receiver Operating Characteristics) were made. These are graphs of the true positive rate (in the case of Figure 5.9b, correctly identifying uranium) as a function of the false positive rate (incorrectly identifying lead as uranium). ROC curves are useful to visualise the performance of classifiers [57]. A possible value for the quality of

the discrimination performance is the area under the curve (AUC) of the ROC curves. The AUC is shown here as a percentage, where 50% is the worst possible discrimination, corresponding to a random classifier, and 100% is an ideal classifier with total separation between the two materials. A classifier may perform better than another with a higher AUC in a given region of the ROC curve [57]. An example of this can be seen in Figure 5.10, where ROC curve B has a higher AUC than A. However, A has a higher true positive rate when the false positive rate is over 0.6, and hence performs better in that region. However, the ROC curves obtained do not show this behaviour (some examples of ROC curves will be shown in Section 5.1.2): ROC curves with lower AUC have not been found to perform better than curves with higher AUC in any region of the graphs.

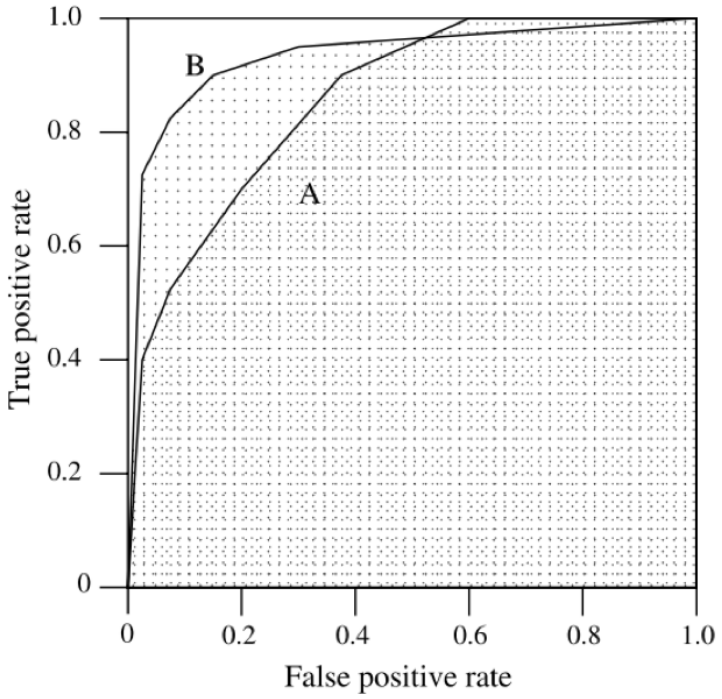

Figure 5.10: Example of a ROC curve A with a lower AUC than B, but that performs better at false positive rate over 0.6 [57].

A $10 \times 10 \times 10$ cm^3 uranium cube is distinguished from a lead cube, for 1h of muon exposure, in the example of the ROC curve shown in Figure 5.9b. To produce the ROC curve the true positive and the false positive rate were calculated for all possible cut values using Figure 5.9a, and plotted against each other. The

95% confidence regions of the ROC curves and respective confidence of the AUC were calculated using the R package pROC [58], which uses the DeLong method, described in [59].

A perfect discrimination, corresponding to a discriminator value with no overlap between the two categories, requires an AUC of 100%. But a good discrimination can also be obtained with a slightly lower AUC. Figure 5.9 can be used as an example: there is a small overlap, so the operator can choose as the cut value (above which the block is identified as uranium instead of lead) a discriminator value which is lower than all the values obtained for uranium. Because of the overlap, in some cases a lead block will be incorrectly identified as uranium, but all uranium blocks would be correctly identified. Different cut values will result in different probabilities of incorrect identification. Furthermore, the operator can decide that if a discriminator value is in the overlapping region, more data need to be taken in order to obtain a perfect discrimination. But, as most of the entries in Figure 5.9a are outside the overlapping region, in most cases more data will not be needed. It is also up to the operator to choose what is the minimum AUC value needed to obtain the required discrimination performance. In this thesis discrimination with an AUC of at least 95% will be referred to as good discrimination.

Examples of ROC curves using these discriminators to distinguish between $5 \times 5 \times 5$ cm^3 blocks of uranium and lead, for 1h of muon exposure, are shown in Figure 5.11. These were chosen to be values that can distinguish between the two overlapped distributions: the mean, the integral over a cut (summing the bin entries only from the bins over the value above which the distribution for the higher Z material is larger), the maximum value, a weighted integral (giving more weight to higher values by multiplying the entries in each bin by the bin number), a weighted integral over the cut and a weighted integral under the cut (the same method of weighting the integral, but using only bins over or under the cut value, respectively). The value used for the cut was 0, which is the observed point above which the entries obtained from the uranium simulations are higher than the ones

from lead and tungsten (this can be seen in Figure 5.8b). The mean has the highest AUC, of $87.0 \pm 2.7\%$, although it is within errors of several other discriminator candidates (for example, the weighted integral yielded an AUC of $84.9 \pm 2.8\%$). However, all other block sizes and data taking times consistently showed a higher AUC for the mean. Another example can be seen in Table 5.2, for a $3 \times 3 \times 3$ cm^3 block of lead and using 10h of muon exposure. The mean was therefore chosen as the discriminator.

Discriminator variable	AUC for 5 cm, 1h	AUC for 3 cm, 10h
mean	$87.0 \pm 2.7\%$	$90.2 \pm 3.0\%$
integral over a cut	$84.2 \pm 3.0\%$	$83.2 \pm 4.0\%$
maximum value	$76.7 \pm 3.6\%$	$64.4 \pm 5.6\%$
weighed integral	$84.9 \pm 2.8\%$	$78.9 \pm 3.6\%$
weighted integral over the cut	$84.4 \pm 3.0\%$	$84.5 \pm 3.9\%$
weighted integral under the cut	$83.7 \pm 3.1\%$	$83.7 \pm 4.0\%$

Table 5.2: AUC of the ROC curves obtained from different discriminator variables, for a 5 cm side cube with 1h, and for a 3 cm side cube with 10h.

For a certain value of muon exposure time, the number of muons selected was smaller for smaller blocks, because there were fewer muons crossing the block. The number of muons that traverse each block size per hour is shown in Table 5.3. It is clear that fewer muons cross smaller volumes for the same scanning time, but it can also be seen that the number of muons per unit volume decreases with the total volume. This is due to the fact that most of the muons detected are vertical, or almost vertical, so the number of muons crossing the block depends more on its horizontal area than on its volume. This can also be seen in the last column of Table 5.3, which shows that the number of muons per unit area of the face of the blocks is similar for all blocks (approximately 13 muons per cm^2).

Block size	# muons h^{-1}	# muons h^{-1} cm^{-3}	# muons h^{-1} cm^{-2}
$1 \times 1 \times 1$ cm^3	14	13.8	13.8
$2 \times 2 \times 2$ cm^3	53	6.7	13.3
$3 \times 3 \times 3$ cm^3	124	4.6	13.8
$5 \times 5 \times 5$ cm^3	332	2.7	13.3
$10 \times 10 \times 10$ cm^3	1314	1.3	13.1

Table 5.3: Number of muons that cross each volume per hour, per unit volume, and per unit area of the upper face of the cubes.

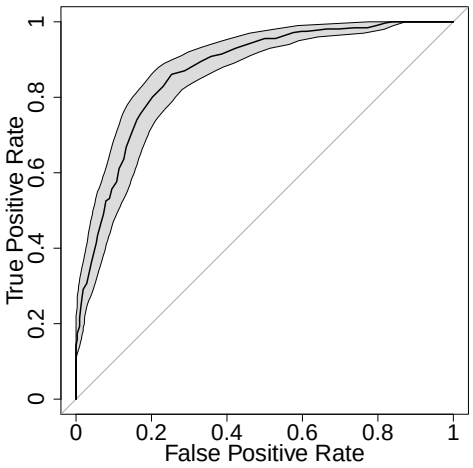

(a) ROC using the mean.

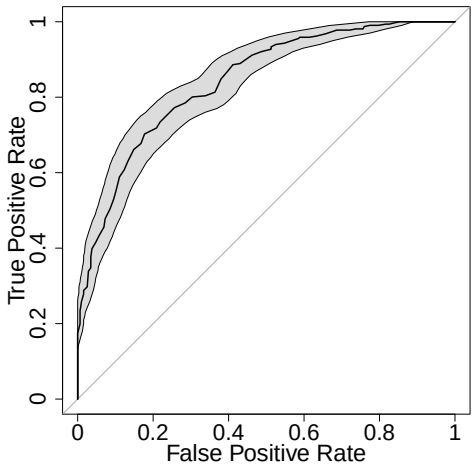

(b) ROC using the integral over a cut at 0.

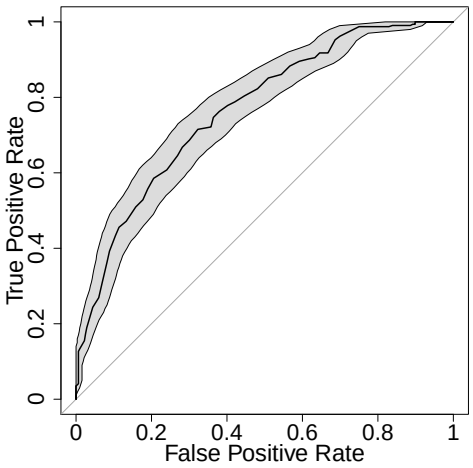

(c) ROC using the maximum.

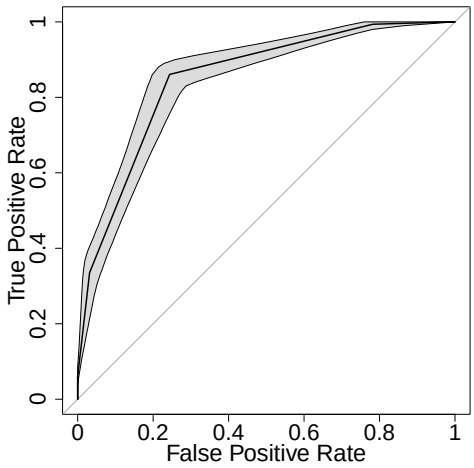

(d) ROC using the weighted integral.

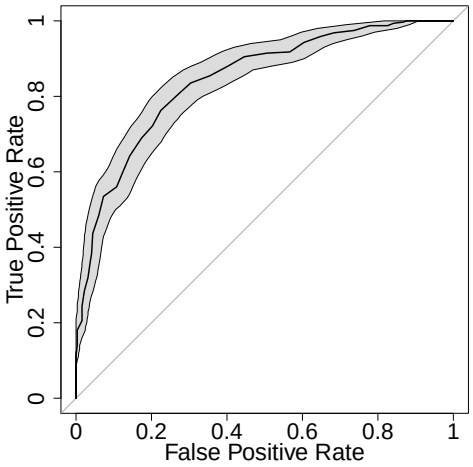

(e) ROC using the weighted integral over a cut at 0.

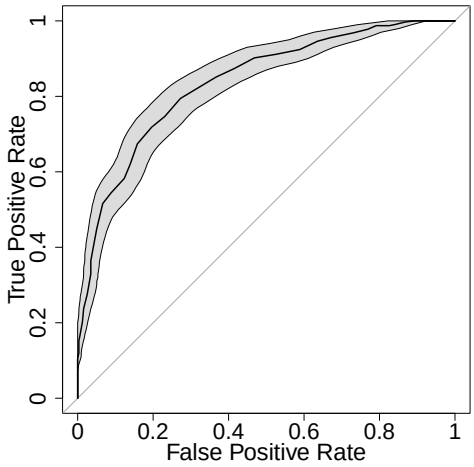

(f) ROC using the weighted integral under a cut at 0.

Figure 5.11: ROC curves obtained from the comparison between uranium and lead, for a $5 \times 5 \times 5$ cm^3 block and 1h of muon exposure. The AUC values obtained are (a) $87.0 \pm 2.7\%$, (b) $84.2 \pm 3.0\%$, (c) $76.7 \pm 3.6\%$, (d) $84.9 \pm 2.8\%$, (e) $84.4 \pm 3.0\%$ and (f) $83.7 \pm 3.1\%$.

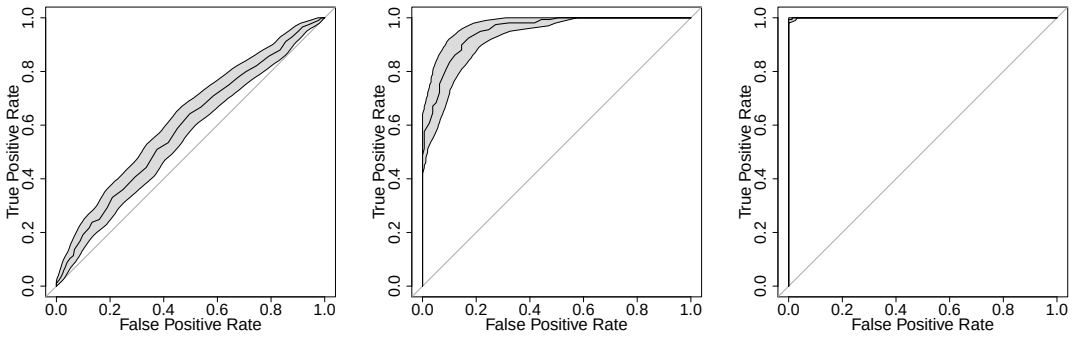

(a) $2 \times 2 \times 2$ cm^3 (b) $5 \times 5 \times 5$ cm^3 (c) $10 \times 10 \times 10$ cm^3

Figure 5.12: ROC curves for the mean of the Fisher determinant distinguishing uranium from lead for different block sizes, with a muon exposure time of 2h. The cubic blocks have a side of $2 \times 2 \times 2$ cm^3, with AUC $= 59.1 \pm 3.6\%$ (a), $5 \times 5 \times 5$ cm^3, with AUC $= 95.1 \pm 2.0\%$ (b) and $10 \times 10 \times 10$ cm^3, with AUC $= 99.9 \pm 0.1\%$ (c).

5.1.2 Results

As it was already shown in Figure 5.9b, a good discrimination (AUC $= 99.7 \pm 0.2\%$) can be obtained for comparing a $10 \times 10 \times 10$ cm^3 uranium block to lead with 1h of data. The AUC calculated for the same cube with half an hour of data was $97.3 \pm 0.7\%$, as seen in Figure 5.14a, which already shows a good discrimination. For a fixed time, fewer muons cross small blocks than large blocks. Because only muons tracks that crossed the block are used, the smaller the target volume, the less data is obtained. This results in a worse discrimination, with lower AUC, as shown in Figure 5.12, which consists of ROC curves from different block sizes of uranium and lead for 2h of data acquisition. While the AUC obtained for $2 \times 2 \times 2$ cm^3 blocks was low ($59.1 \pm 3.6\%$), the $10 \times 10 \times 10$ cm^3 yielded a high AUC, of $99.9 \pm 0.1\%$. The results also show that increasing the exposure time improves the discrimination power, so a longer data taking time can be used to achieve a good discrimination. This can be seen in 5.13, where the ROC curves for different exposure times of a $5 \times 5 \times 5$ cm^3 block can be seen.

Figures 5.14 and 5.15 show the AUC values obtained from comparing uranium to lead and uranium to tungsten for different block sizes, as a function of the muon exposure time. It is clear that the time needed to obtain discrimination depends on the block size. If, for example, the minimum AUC to accept discrimination is

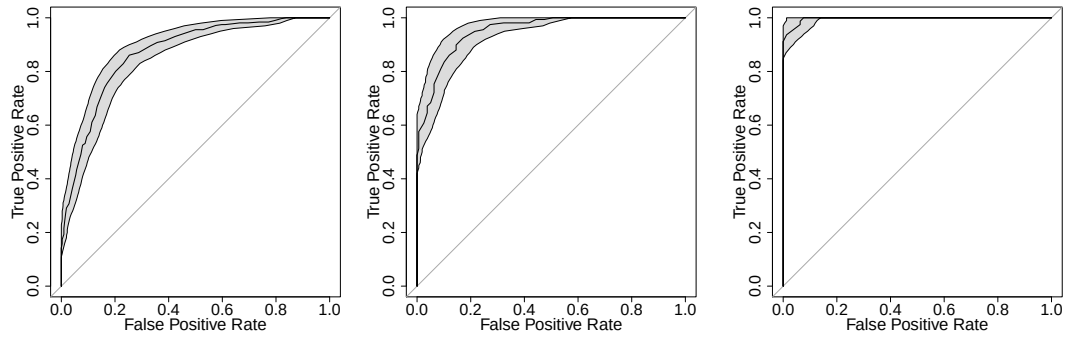

(a) 1h (b) 2h (c) 4h

Figure 5.13: ROC curves for the mean of the Fisher determinant distinguishing uranium from lead for different muon exposure times, for a cubic block with 5 cm side. The muon exposure times are 1h (a), 2h (b) and 4h (c).

chosen as 95%, the 10 cm side cube is separated with 30 min of muon data, while the 5 cm side cube needs 2h for lead or 3.1h for tungsten (extrapolating between the points of 2h and 4h in Figure 5.14). Figure 5.15 shows the AUC for longer muon exposure times, up to 70h, because the 3 cm and 2 cm side cubes require longer scanning times than the larger volumes. It can be seen that material discrimination would not be possible, with this system and the scanning times considered, for 1 cm side cubes, because the AUC obtained is close to 50%. This is due to the fact that a small amount of muons cross this block compared to the larger ones. Using the values from Table 5.3, about 967 muons cross the 1 cm side cube in 1 hour, while 3735 cross the 2 cm side cube during the same time. The slow increase of their AUC suggests that muons may not undergo enough scattering in the 1 cm^3 blocks to provide enough discrimination between high-Z materials. An additional source of error in the discrimination of 1 cm^3 blocks is the vertex approximation: if some of the few muons that scattered in the block are reconstructed outside the selected area, more information is lost. These arguments, along with the data from 1 cm^3 blocks seen in Figure 5.15, indicate that it would not be possible to obtain a high AUC value, even with more data.

Figures 5.16 and 5.17 show the same AUC values, but as a function of the block sizes, instead of the scanning time. These graphs were created in order to better illustrate the dependence of the discrimination performance on the volume.

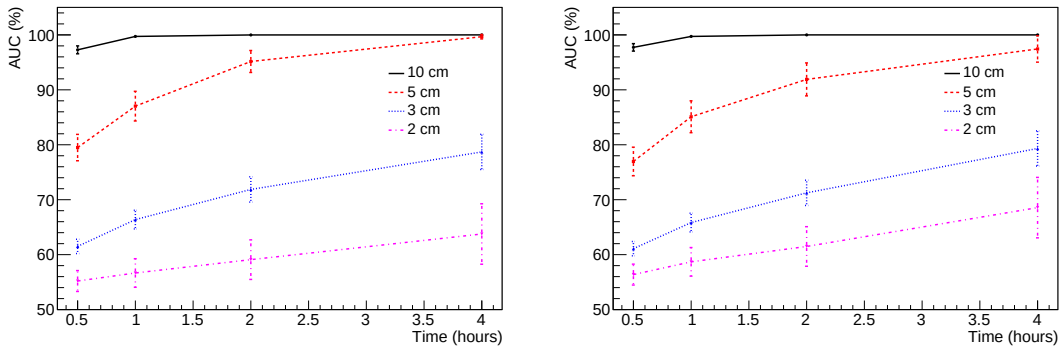

(a) Uranium vs Lead

(b) Uranium vs Tungsten

Figure 5.14: Area under the curve of ROC curves distinguishing uranium from lead (a) and uranium from tungsten (b), as a function of the scanning time, for different block sizes. The error bars are the 95% confidence intervals of the AUC.

It is important to note, however, that they do not take into account the shape of the object: for example, a cuboid with the same volume but a larger horizontal area and a smaller height would have more muons crossing it, but fewer multiple scattering interactions.

Table 5.4 shows the measurement times required for each block size and material to achieve a 95% AUC. The values were obtained by linearly extrapolating between the calculated points before and after 95%. Although the trend is not linear, this is a useful estimate, since the shape of the trends is not known. The time needed to reach the desired AUC value depends, in part, on the amount of data obtained for each case, which in turn depends on the number of muons that crossed the block. Table 5.5 shows the number of muons that crossed each block in the corresponding time from Table 5.4. These numbers were calculated using the number of muons that cross each block per hour from Table 5.3. It is interesting to note that, as the block size decreases, the number of muons required to obtain an AUC of 95% increases. This is due to the fact that muons can interact more in larger blocks, increasing the chance of having a high scattering angle.

After showing separation of uranium from lead and tungsten, whose atomic number is lower than that of uranium, the method was applied to the discrimination between uranium and plutonium. Plutonium is denser and has a higher Z

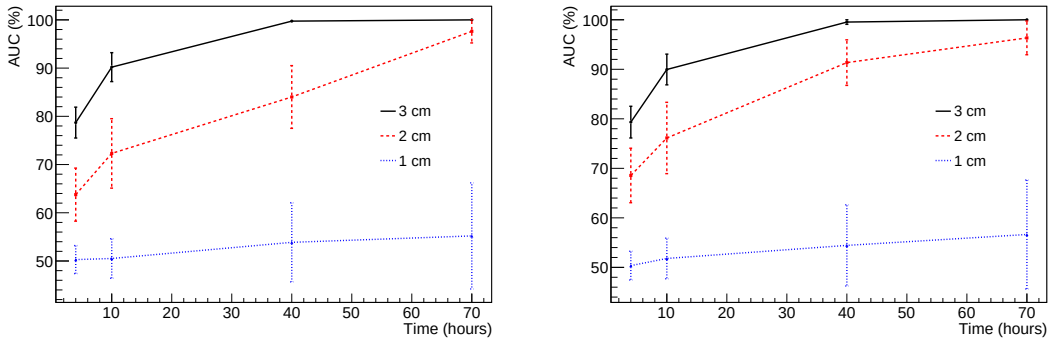

(a) Uranium vs Lead (b) Uranium vs Tungsten

Figure 5.15: Area under the curve of ROC curves distinguishing uranium from lead (a) and uranium with tungsten (b), as a function of the scanning time, for different block sizes. The error bars are the 95% confidence intervals of the AUC.

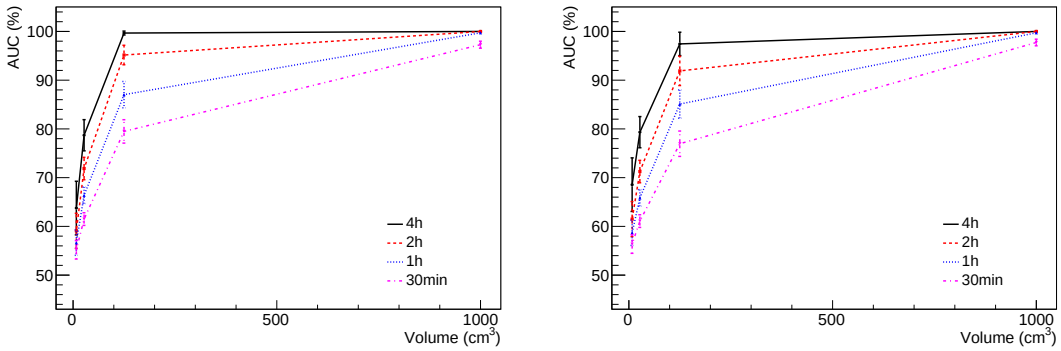

(a) Uranium vs Lead (b) Uranium vs Tungsten

Figure 5.16: Area under the curve of ROC curves distinguishing uranium from lead (a) and uranium from tungsten (b), as a function of the block size, for different scanning times. The error bars are the 95% confidence intervals of the AUC.

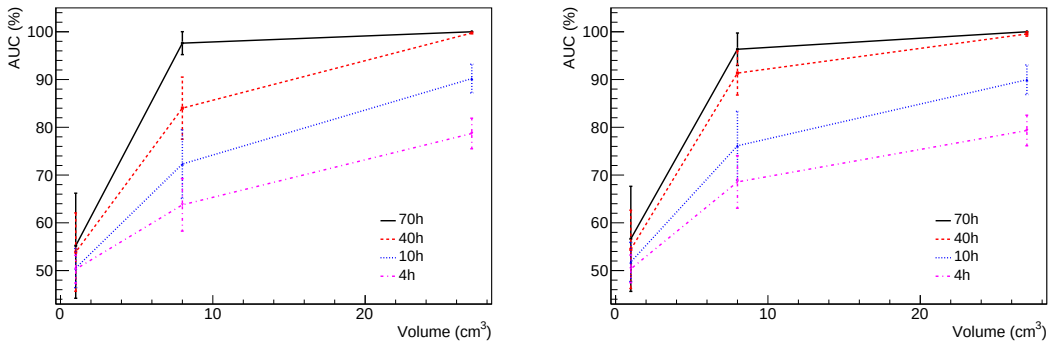

(a) Uranium vs Lead (b) Uranium vs Tungsten

Figure 5.17: Area under the curve of ROC curves distinguishing uranium from lead (a) and uranium with tungsten (b), as a function of the block size, for different scanning times. The error bars are the 95% confidence intervals of the AUC.

Block size	Time for U vs Pb (h)	Time for U vs W (h)
$2 \times 2 \times 2$ cm^3	64.2	61.9
$3 \times 3 \times 3$ cm^3	25.1	25.8
$5 \times 5 \times 5$ cm^3	2.0	3.1
$10 \times 10 \times 10$ cm^3	0.5	0.5

Table 5.4: Time to reach 95% AUC of several ROC curves for different block sizes.

Block size	# muons for U vs Pb (h)	# muons for U vs W (h)
$2 \times 2 \times 2$ cm^3	657	657
$3 \times 3 \times 3$ cm^3	665	1031
$5 \times 5 \times 5$ cm^3	3107	3193
$10 \times 10 \times 10$ cm^3	3425	3302

Table 5.5: Number of muons that crossed each block for the corresponding time in Table 5.4, for different block sizes.

(94) than uranium (92), and hence has a lower radiation length (0.2989 cm and 0.3166 cm respectively, from Table 2.1). Thus, because the radiation length of lead and tungsten is even higher than the one of uranium, the discrimination between those materials and plutonium can be assumed to be at least as good as that between uranium and lead or tungsten. Discriminating uranium and plutonium is more difficult because of their similar atomic number, which results in a similar radiation length, as equation 2.4 shows. Therefore, it requires longer scanning times to obtain discrimination. This can be seen in Figure 5.19, which is a graph of the AUC as a function of the scanning time, for $3 \times 3 \times 3$ cm^3 blocks. It was found that, for this case, the integral over a cut of the Fisher output was a better discriminator than its mean. For example, the AUC obtained for 200h of data with the integral over a cut was 80%, while with the mean a lower value of 69% was obtained. This is because the means of the Fisher output obtained from the uranium and the plutonium simulations were very similar, while the tail in the Fisher output distribution for plutonium is higher than for uranium, introducing some discrimination power. The value for the cut on the integral was chosen to be 0.01 which is the observed point from where the entries obtained from the plutonium simulations started being higher than the ones from uranium. This can be

seen in Figure 5.18, which is the Fisher output (in linear and logarithmic scale) from the training with uranium and plutonium.

The best AUC obtained was 80%, which still allows for a significant discrimination, even if not a perfect one. For a smaller block, of $2 \times 2 \times 2$ cm^3, it was not possible to obtain discrimination between uranium and plutonium, with times up to 200h, as Figure 5.20.

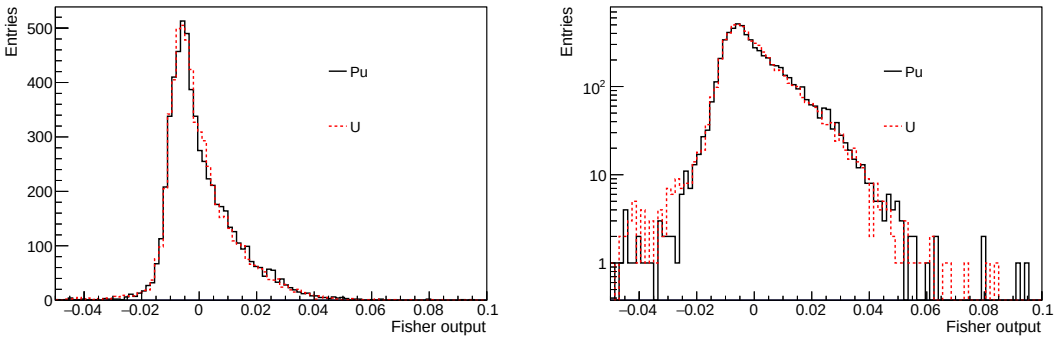

(a) (b)

Figure 5.18: Output of the Fisher MVA method, comparing plutonium and uranium, on linear (a) and logarithmic (b) scales for a $3 \times 3 \times 3$ cm^3 block.

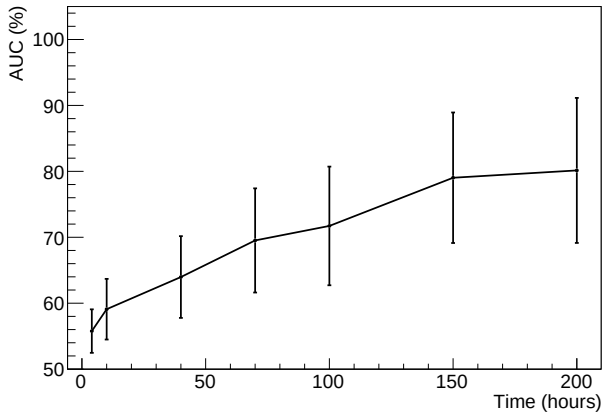

Figure 5.19: Area under the curve of ROC curves distinguishing uranium from plutonium for a $3 \times 3 \times 3$ cm^3 block.

5.1.3 Results for different volume sizes

The material discrimination method relies on the fact that the position of the high-Z material is known, so only muon tracks that were reconstructed in the material block are selected. However, it is possible that some muons that crossed

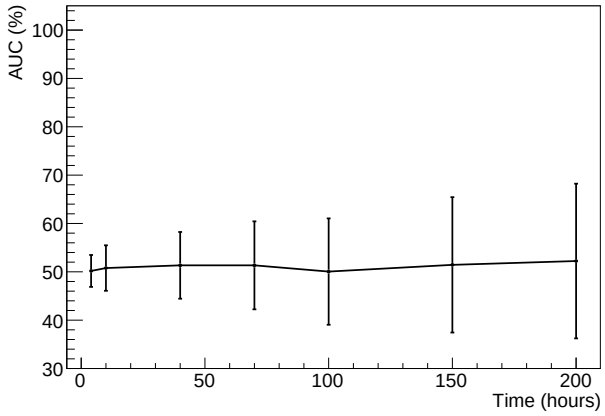

Figure 5.20: Area under the curve of ROC curves distinguishing uranium from plutonium for a $2 \times 2 \times 2$ cm^3 block.

and scattered in the target block, especially next to the edges, are actually reconstructed in the concrete. The information from these tracks is therefore lost when only exactly the target volume is selected. For this reason, the study was extended to add an extra volume around the edges of the blocks for the track selection. This is not only important in order to show if adding some tracks reconstructed next to the blocks can improve the results or not, but also to understand how precise the position information needs to be to obtain a good discrimination.

The method was applied to the $2 \times 2 \times 2$ cm^3 and $3 \times 3 \times 3$ cm^3 blocks, adding tracks that were reconstructed in concrete, but inside a fiducial region centred around the blocks, increasing the volume on all 6 sides by 1 to 4 cm. The AUCs for each added length are shown in Figure 5.21. It can be seen that an extra 1 to 2 cm improves the performance of the material discrimination method for $2 \times 2 \times 2$ cm^3 blocks. This is why 1.5 cm was added to each side of the $2 \times 2 \times 2$ cm^3 blocks in the previous section. Adding 3 cm results in a lower AUC, (although still within errors of the AUC for 2 cm extra). Adding 4 cm clearly worsens the performance compared to the 1 and 2 cm, resulting in a lower AUC, as seen in Figure 5.21d for the 70h scan. For the $3 \times 3 \times 3$ cm^3 blocks, 2 cm were added to each side, as the discrimination does not seem worse for larger lengths added (up to 4 cm).

These results, in particular the fact that the performance for 1 cm and 2 cm added is similar, are relevant to determine the resolution required for the measure-

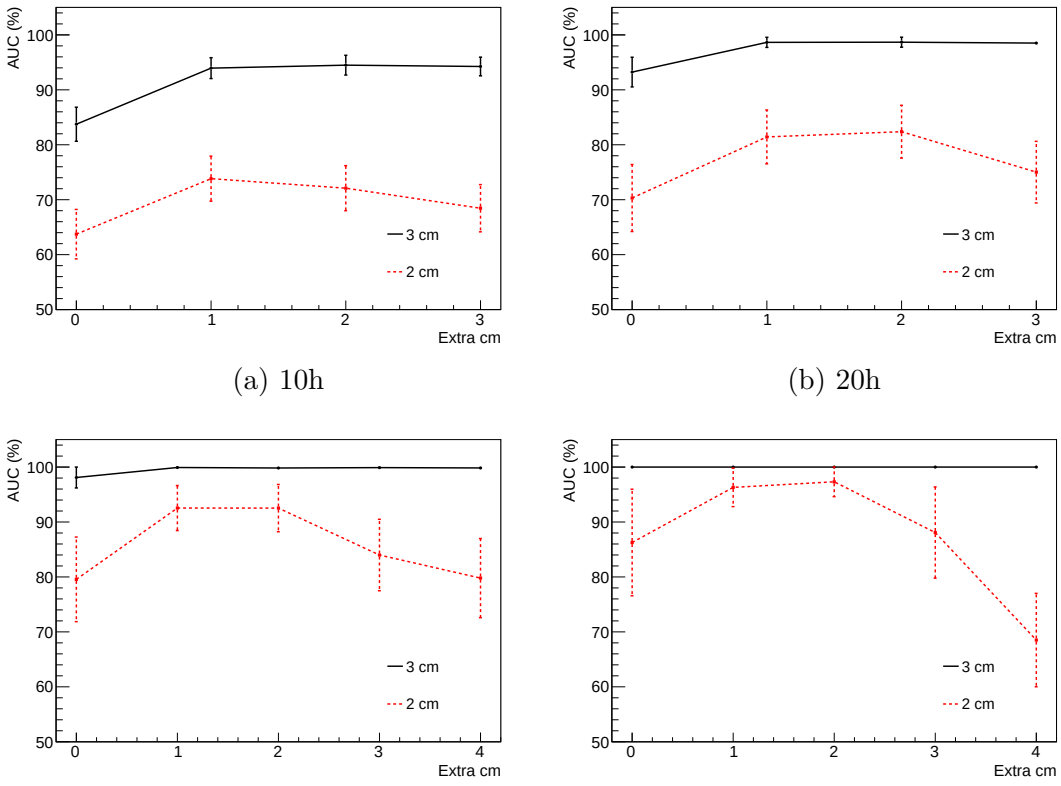

(a) 10h (b) 20h

(c) 40h (d) 70h

Figure 5.21: Area under the curve of ROC curves distinguishing uranium from tungsten, for different block sizes and different muon exposure times. The error bars are the confidence intervals of the AUC.

ment of the blocks. For example, if an imaging algorithm with 1 cm resolution is used, this method can be applied adding 2 cm on each side, and if the edge measurements are 1 cm off, the discrimination performance is not affected.

5.1.4 Results without momentum information and with momentum smearing

The muon momentum was one of the variables used in the MVA. Combined with the angle, the momentum was expected to improve the material discrimination because a large scatter angle from a high momentum muon strongly indicates localised scattering in a high-Z material, while the same large angle from a low momentum muon can also be a common scattering event due to multiple scattering in lower-Z material. To test this assumption, the AUCs for the comparison of $5 \times 5 \times 5$ cm^3 uranium and lead blocks with true momentum, with 50% momentum

smearing and without momentum information were compared. These are shown in Table 5.6. It can be seen that the results did not alter significantly when the muon momentum was smeared by 50% or even when the momentum information was removed, since the AUC remains the same for the three different cases, within errors.

Time	with true momentum	with momentum smearing	without momentum
1h	$87.0 \pm 2.7\%$	$87.0 \pm 2.7\%$	$87.1 \pm 2.7\%$
2h	$95.2 \pm 2.0\%$	$95.2 \pm 2.0\%$	$95.3 \pm 2.0\%$
4h	$99.5 \pm 0.5\%$	$99.7 \pm 0.3\%$	$99.6 \pm 0.4\%$

Table 5.6: AUC of several ROC curves from the mean of the Fisher discriminant for the $5 \times 5 \times 5$ cm^3 block, distinguishing uranium from lead, with true momentum information, smearing the momentum by 50% and without any momentum information.

It is clear that the momentum information does not help high-Z material discrimination. In Chapter 4, it was shown that not using the momentum resulted in a worse resolution in the measurement of the of length uranium blocks in concrete, effectively suggesting that the muon momentum combined with the angle improves the discrimination between uranium and concrete. This is due to the fact that the algorithm of Chapter 4 used the momentum directly in the calculation of the metric distance. It also selected a fixed number of the most scattered tracks, which may have a different momentum distribution depending on the material they crossed, while in the material discrimination method all the muons that crossed the target block were selected, so their momentum distribution will be the same regardless of the material. Additionally, in this chapter the materials to be distinguished all have a high atomic number, unlike in Chapter 4, where the materials compared were uranium and concrete, so it is possible that the correlation between the scattering angle of a muon in high-Z materials and its momentum is not enough to provide discrimination.

It has been shown that the momentum information is not necessary for material discrimination. Although momentum measurement would result in a better imaging resolution, the results from Section 5.1.3 proved that the resolution without

momentum (which was $\sigma = 2.9 \pm 0.5$ mm) would be good enough to identify the volume required for the selection of tracks for material discrimination. Therefore, the system and analysis do not need to measure or estimate the muon momentum in order to perform imaging and material discrimination.

5.1.5 Results using only the scatter angle

It was seen that the most significant variable used in the MVA was the 3D scatter angle, and all the other variables (except the muon momentum) were highly correlated to it. For this reason, the material discrimination method was applied using the mean of the 3D scatter angle as a discriminator instead of the mean of the Fisher output. Table 5.7 shows some AUC values to compare the two approaches. It shows a small improvement for MVA, which is however not statistically significant. Therefore, it is proven that using only the 3D scatter is enough to perform material discrimination.

Time	scatter angle	Fisher
1h	$86.3 \pm 2.8\%$	$87.0 \pm 2.7\%$
2h	$94.6 \pm 2.2\%$	$95.2 \pm 2.0\%$
4h	$99.1 \pm 0.9\%$	$99.7 \pm 0.3\%$

Table 5.7: AUC of several ROC curves from the mean of the scatter angle and the Fisher discriminant for the $5 \times 5 \times 5$ cm^3 block, distinguishing uranium from lead.

5.2 Discussion and Conclusions

It was shown in this chapter that it is possible to distinguish between uranium and lead or tungsten using MST, for blocks with a volume of $2 \times 2 \times 2$ cm^3 or bigger. Taking a minimum AUC of 95% to achieve discrimination between two materials, the $10 \times 10 \times 10$ cm^3 cubes only needed 30 minutes of data, the $5 \times 5 \times 5$ cm^3 blocks needed 2h (U vs Pb) and 3.1h (U vs W), the $3 \times 3 \times 3$ cm^3 blocks needed 25.1 (U vs Pb) and 25.8 (U vs W), and the $2 \times 2 \times 2$ cm^3 blocks needed 64.2 (U vs Pb) and 61.9 (U vs W). It was not possible to discriminate $1 \times 1 \times 1$ cm^3 blocks with scanning times up to 70h. However, it is not expectable that a longer scanning time

would improve much the discrimination of these blocks. This is because the vertex reconstruction, assuming that the muon scattered in a single point, implies an intrinsic uncertainty of this method: the tracks used are the ones reconstructed in the high-Z volume, but this reconstruction is only an approximation and not a real guarantee that all muons that scattered in the block are correctly reconstructed in the respective position (and conversely that all muons reconstructed in the block have actually scattered in it).

In addition, there is some discrimination between $3\times3\times3$ cm^3 cubes of uranium and plutonium, reaching an AUC of 80%. The same study was done for $2 \times 2 \times 2$ cm^3, but it was observed that discrimination was not possible, as the AUC obtained was always close to 50%.

Being able to distinguish between two different high-Z materials is not in itself a method to identify materials. However, material identification can be achieved by comparing an unknown material to several known ones, from materials that can be usually found in nuclear waste, and by determining from which ones it is discriminated, and finally which one is the most similar.

It has also been shown that using only the 3D scatter angle gives similar AUCs as when combining several variables with the Fisher linear discriminant. This result allows for future analyses to use the 3D scatter angle without the need of performing training for an MVA.

For this method to work, the position of the materials has to be known, which was shown to be possible in Chapter 4. The image reconstruction requires a reasonable amount of data, that can go up to 2 weeks of muon exposure, such as in Figure 4.2, or even more if edge finding is needed. However, it was also shown in this chapter that using the muon tracks from a volume slightly bigger than the target material can even improve the material discrimination, because it allows for the use of muons that scattered in the high-Z material but whose vertex was reconstructed outside it. This means that it is not essential to have a very accurate image of the materials in the concrete, so the amount of data, and consequently the

exposure time, can be lower: for example instead of the reported 2 weeks, only 70h could be used (almost 3 days), which allow for material discrimination of blocks as small as $2 \times 2 \times 2$ cm^3. Furthermore, the same muon scattering data can be used for imaging and material discrimination, so the muon exposure times shown in this chapter are not added to the imaging time. The scanning times required will differ for different detector sizes and spacing, as a smaller angular acceptance will detect fewer muons. Increasing the detector area would increase the amount of muons detected, not only because more muons will cross a larger area, but also because it would result in a larger angular acceptance, allowing the detection of muons with higher angles that are unavailable for detector systems with smaller areas. Using a smaller spacing between tracking planes would also result in a higher acceptance angle, increasing the number of muons detected. Another factor that will affect the required time is the position of the high-Z objects relative to the detectors. But, as it is assumed that their position is known, the high-Z objects of interest should be placed in the centre of the detector system in order to minimise the scanning time.

Another important result is the fact that the muon momentum does not improve the discrimination power of high-Z materials. Using the momentum improves the imaging resolution, as Chapter 4 showed, but if the resolution required for discrimination is, for example, 1 cm, as shown in Section 5.1.3, it is possible to perform imaging and discrimination of high-Z materials using a system which does not provide momentum information.

6 Conclusion

There is a need to characterise nuclear waste before it can be taken for final disposal or long term storage. This especially important in the case of legacy waste whose records may have not been kept, and that may have undergone changes in its composition. There are several techniques that can be used to characterise the nuclear waste, both destructive and non-destructive. This thesis proposed using muon scattering tomography as a non-destructive technique to scan nuclear waste containers and showed how good this application can be to characterise waste stored in concrete containers. Muon scattering tomography (MST) is a passive scanning technique that uses cosmic-ray muons to image and identify different materials. It exploits the fact that muons undergo multiple Coulomb scattering, where the scatter angle distributions are wider when crossing materials with higher atomic number. In MST, an incoming muon is tracked before and after scattering in the volume to be scanned, in order to measure its scatter angle, the horizontal offset and other related variables. This thesis explained and showed the evaluation of different algorithms that combine these variables in order to determine the content of nuclear waste. These methods were tested using simulation data. The simulation geometry and parameters, as well as the track fitting algorithm were described in Chapter 3.

6.1 Imaging

An imaging method was previously developed [32]. It applied the data obtained from a muon tomography system to find lumps of high-Z materials encased in concrete. It had also been extended to determine the edges of uranium blocks embedded in concrete. An improved edge finding method was developed (see Chapter 4), in order to obtain a better resolution and to include the possibility of not using momentum information. The existing imaging methods with muon tomography require momentum information. However, some systems may not be able

to measure it, or have a poor momentum resolution. For example, measuring the muon momentum using the multiple Coulomb scattering, the expected resolution for a system similar to ours ranges from 30% to 200%. It is therefore important to test the performance of methods without using the momentum. A very good position resolution was obtained of 3.2 ± 0.6 mm, with the smallest length which was possible to measure of 0.5 cm, without using momentum information. Applying the same method but including the true momentum information resulted in a resolution of 0.98 ± 0.03 mm, which is an improvement over the 1.2 ± 0.4 mm resolution obtained with the previous method. Another improvement over the previous method was the smallest length measured of 0.2 cm, which was 0.5 cm with the previous method.

6.2 Gas bubbles

An interesting application shown in this thesis is the measurement of gas bubbles in nuclear waste containers. Uranium present in nuclear waste can oxidise, leading to the production of hydrogen. It has been observed in certain nuclear waste containers that the accumulation of gas makes them overflow. Therefore, a scanning technique for nuclear waste assay will be more versatile if it allows not only for the imaging of high-Z materials, but also for finding gas bubbles. It was shown that it is possible to determine the total gas volume encased in concrete, with a resolution of $1.6 \pm 0.8\%$ for volumes over 2 L. Additionally, it was shown that it is possible to determine if the gas is located in one single bubble, which could pose a higher risk, or separated into smaller bubbles. This shows that muon tomography can also be used to find gas bubbles in concrete containers.

6.3 Material discrimination

After showing possibilities for imaging nuclear waste with muon tomography, a method for distinguishing high-Z materials was presented in Chapter 5. Since these materials have very similar radiation lengths, methods that perform imaging

of high-Z materials against the concrete background with a much lower Z, may not be able to identify which high-Z materials are present. It was shown that it is possible to distinguish between two high-Z materials when comparing data from simulations with the same geometry but different materials. A good separation was shown between uranium and lead, and uranium and tungsten, for different sizes of blocks, down to $2 \times 2 \times 2$ cm^3, needing data taking times of up to 70 hours. The comparison between uranium and plutonium shows some separation for $3 \times 3 \times 3$ cm^3 blocks, although not reaching a perfect discrimination with the times used (up to 200h).

We can summarise and combine the imaging and material discrimination studies. The full process to characterise the content of a nuclear waste drum would start by performing imaging to find lumps of different materials. Then the muon tracks reconstructed in the high-Z material regions would be selected, and compared to existing data or simulations from standard high-Z materials such as uranium, plutonium, lead or tungsten, being then identified as the most similar material. Finally, the edge finding method would be applied to each material found, using the MVA training previously performed with known data, if a good imaging resolution were required.

6.4 General overview

Results in this thesis have shown that muon tomography can be used to perform imaging of objects encased in concrete, of both high-Z and low-Z materials. The main limitation of muon tomography, in particular for imaging, is the fact that muons undergo multiple scattering along their path, and not in a single point, which introduces an inherent uncertainty in the technique, when using a reconstructed vertex, or point of closest approach. However, using the method described in Section 4.2 which uses voxels of a fixed size for imaging, but shifts them by 1 mm, it is possible to measure uranium lengths as small as 0.5 cm (without momentum information) and 0.2 cm (including momentum information). For low-Z materials,

it was shown that it is possible to measure the gas content of a concrete container, for gas volumes down to 2 L. It was also shown that it is possible to distinguish between different high-Z materials, namely between to distinguish uranium from lead and tungsten, for blocks down to 2 cm side cubes. These results were obtained from simulations where the objects of interest were placed in the centre of a muon tomography system. If these objects are shifted from the centre, more data may be required in order to obtain the same results, as the shifted bubbles results in Section 4.3.2 indicate.

6.5 Future outlook

Some of the future work required to develop a system for scanning real nuclear waste to be achieved includes the study of imaging and material discrimination in larger containers. The concrete container simulated for this thesis was chosen to be around 21 L because this was the size of a test drum built for a previous experiment. It is a good size to compare the results from simulations and experiments without having to change the existing experimental setup to fit a larger container. However, many real nuclear waste drums can be much bigger, for example 220 L. So studies on how the size affects these results are necessary, because not only larger drums contain more concrete, increasing the muon scatter, their height also affects the angular acceptance of the system, as well as the position and size of the detectors.

It is also interesting to apply the edge finding method to different shapes. This work showed edge finding results for regular cuboid shapes, but it is not expected that the objects found in nuclear waste have such regular shapes. This could be done by performing the measurements along x but for several different y and Z values, resulting in different lengths and allowing to better define different shapes. Another future development is the application of this method to different materials. The resolution obtained for measuring edges of materials with atomic number lower than uranium is expected to be worse, since their radiation length will be closer to the one of concrete.

The gas measurement study should also be performed replacing the concrete by bitumen (which has a lower density than concrete). This is because, unlike concrete, bitumen is impermeable, which can lead to the hydrogen accumulating in large bubbles. Another interesting extension to this study is to show if it is possible to improve the technique in order to use it for a range of smaller volumes. It is also important to determine how removing the momentum information will affect the results on gas measurement.

6.6 Summary

As a general conclusion, this thesis has shown the potential and limits of using muon scattering tomography for characterisation of dry nuclear waste encased in concrete. This technique shows good results and some more study and experiments should be performed in order to finally develop a system that can be used in nuclear waste storage facilities. Therefore it is recommended that the study continues, in particular following these suggestions for future work, and testing these methods with real data from prototype systems.

Appendix A | Variables considered for material discrimination

The variables used to calculate the significance value mentioned in Chapter 5 are shown in the current appendix. These distributions were obtained by performing simulations of monoenergetic vertical muons, with 1, 2 and 10 GeV, separately incident on a plate of lead and a plate of uranium (A.1) and on a $10 \times 10 \times 10$ cm^3 block (A.2).

A.1 Simulation of a plate

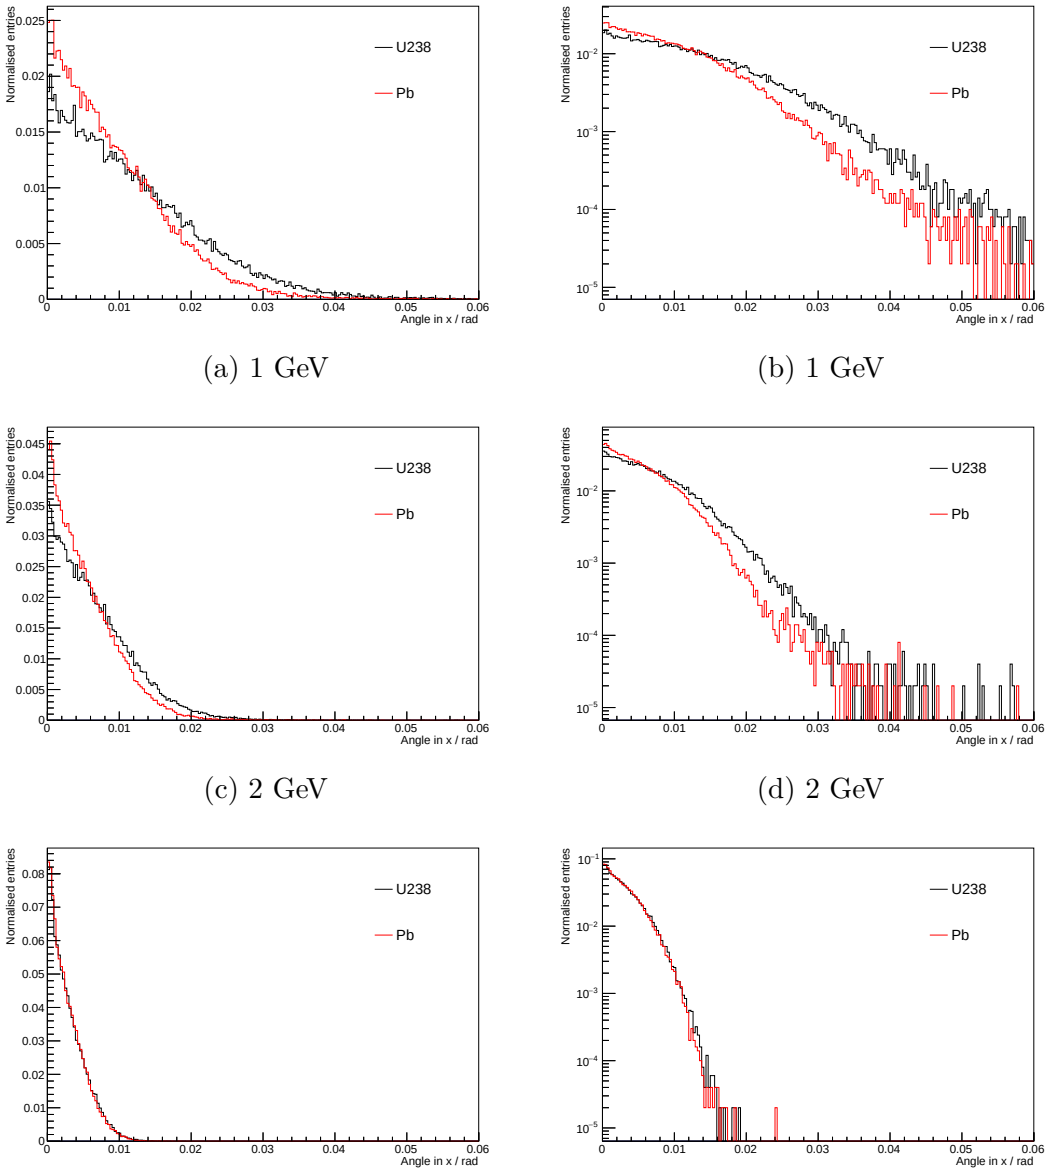

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.1: Projected angle in x for uranium and lead plates, in linear (left) and logarithmic (right) scales, for different muon energies. The histograms are normalised.

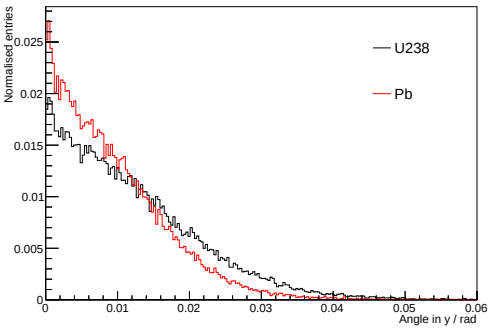

(a) 1 GeV

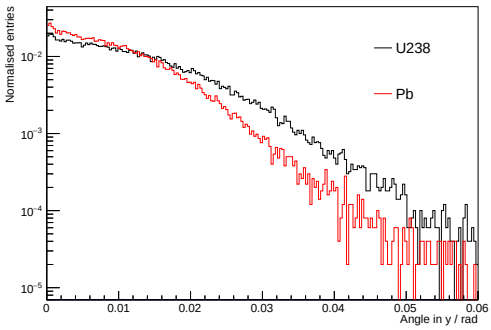

(b) 1 GeV

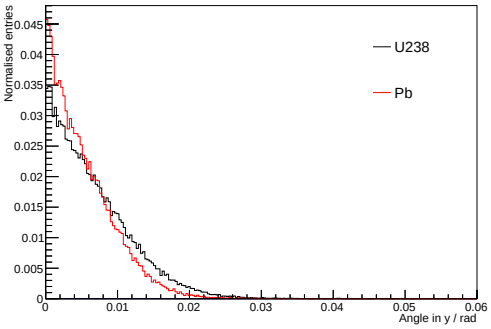

(c) 2 GeV

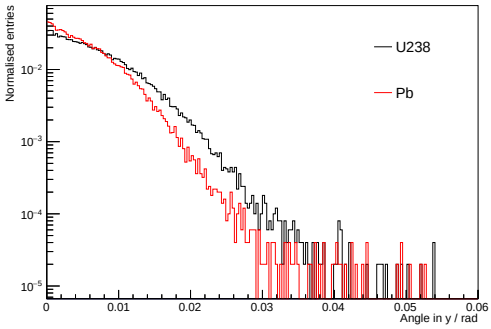

(d) 2 GeV

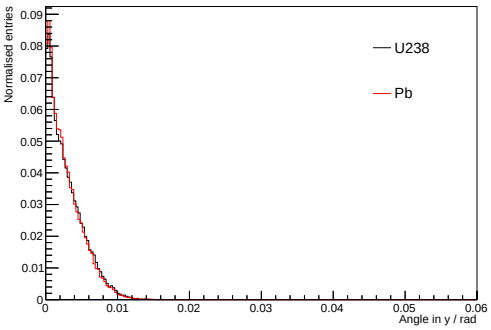

(e) 10 GeV

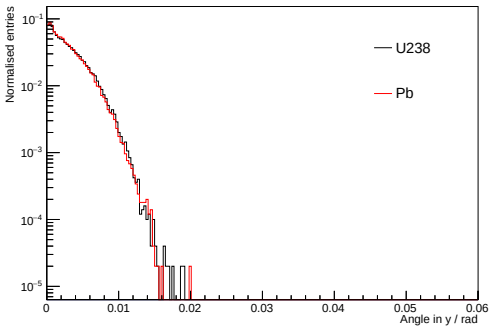

(f) 10 GeV

Figure A.2: Projected angle in y for uranium and lead plates, in linear (left) and logarithmic (right) scales, for different muon energies. The histograms are normalised.

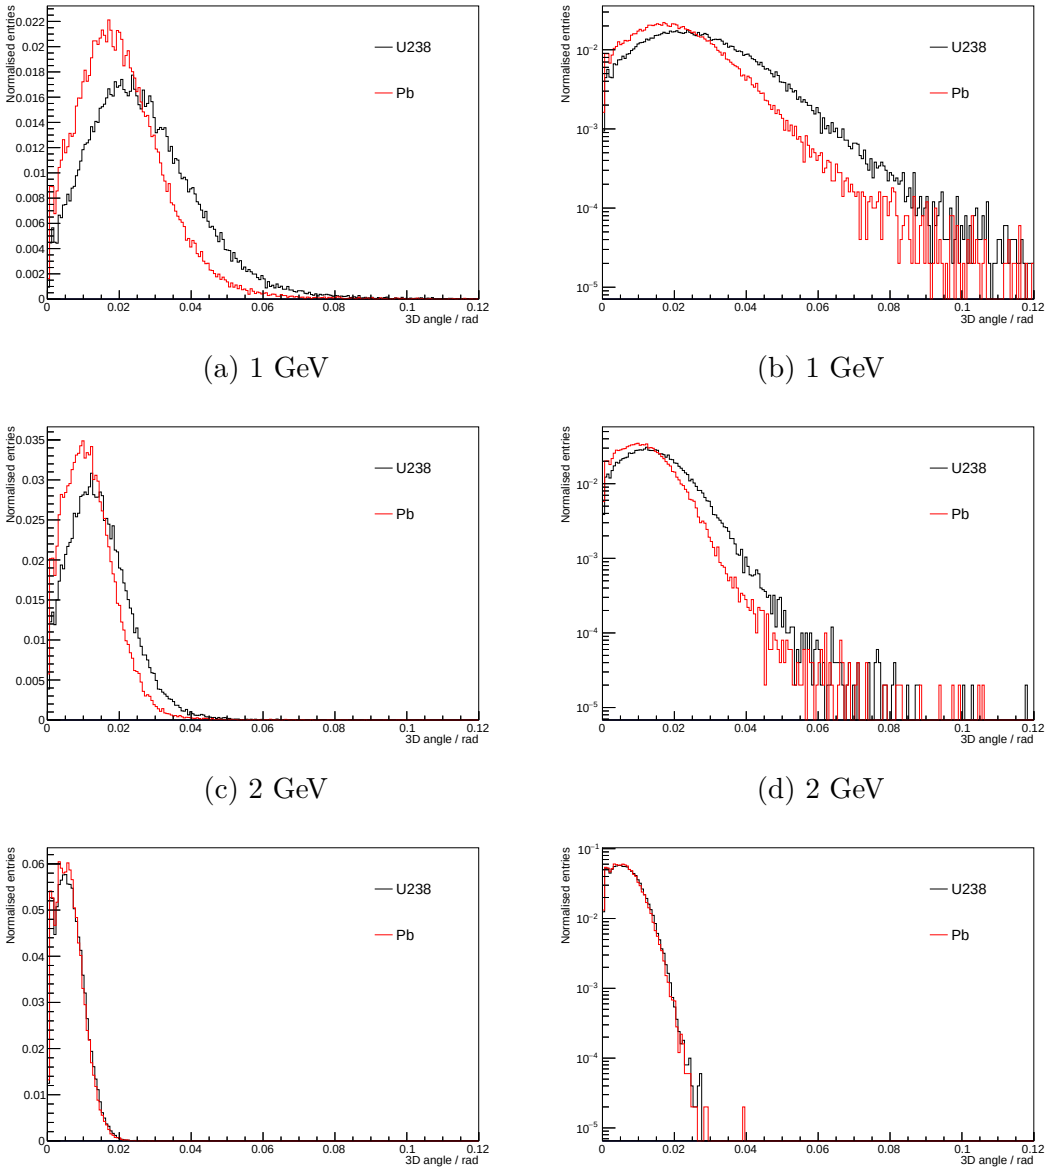

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.3: 3D angle for uranium and lead plates, in linear (left) and logarithmic (right) scales, for different muon energies. The histograms are normalised.

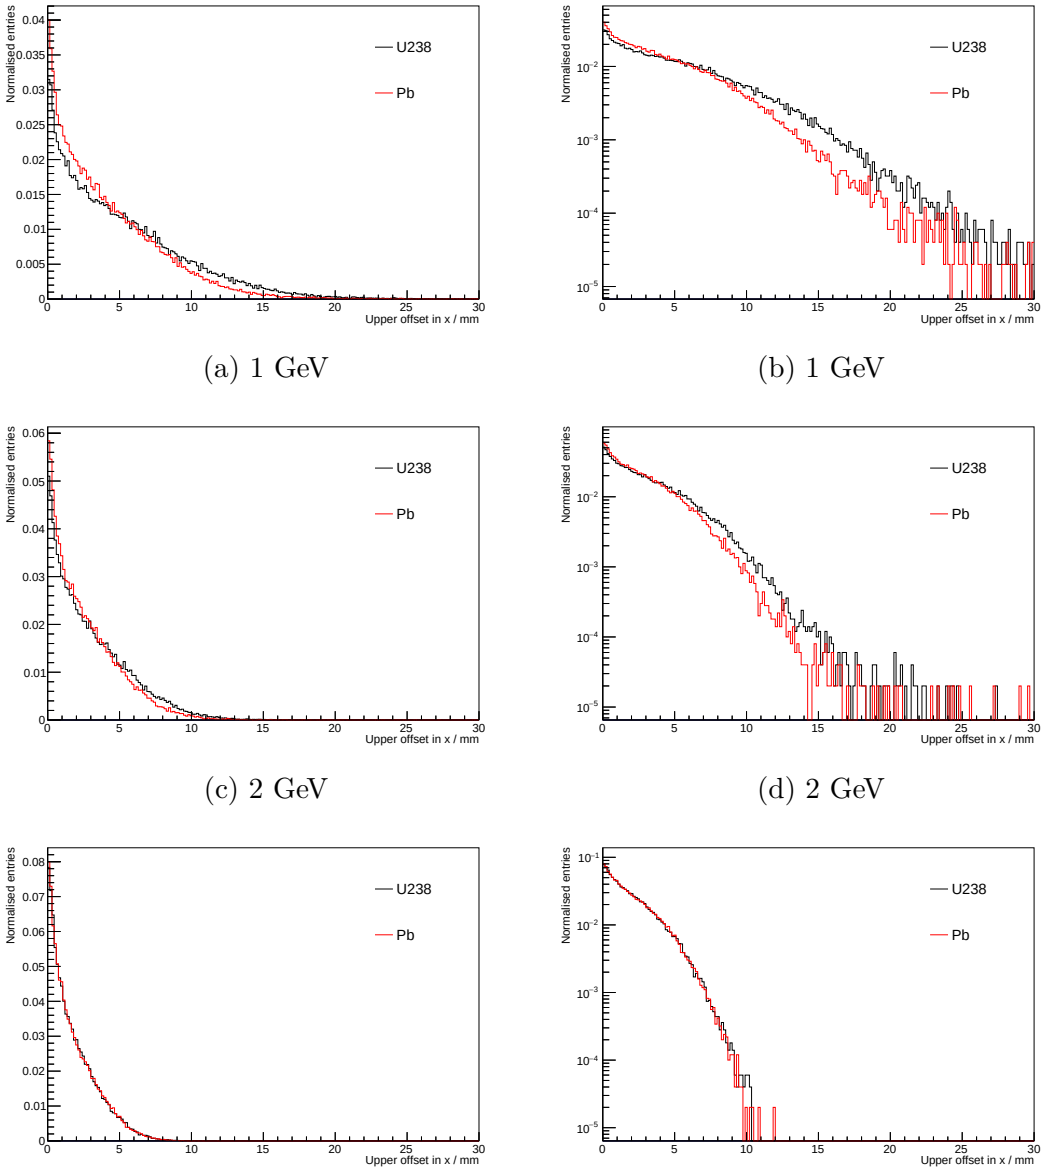

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.4: Upper offset in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

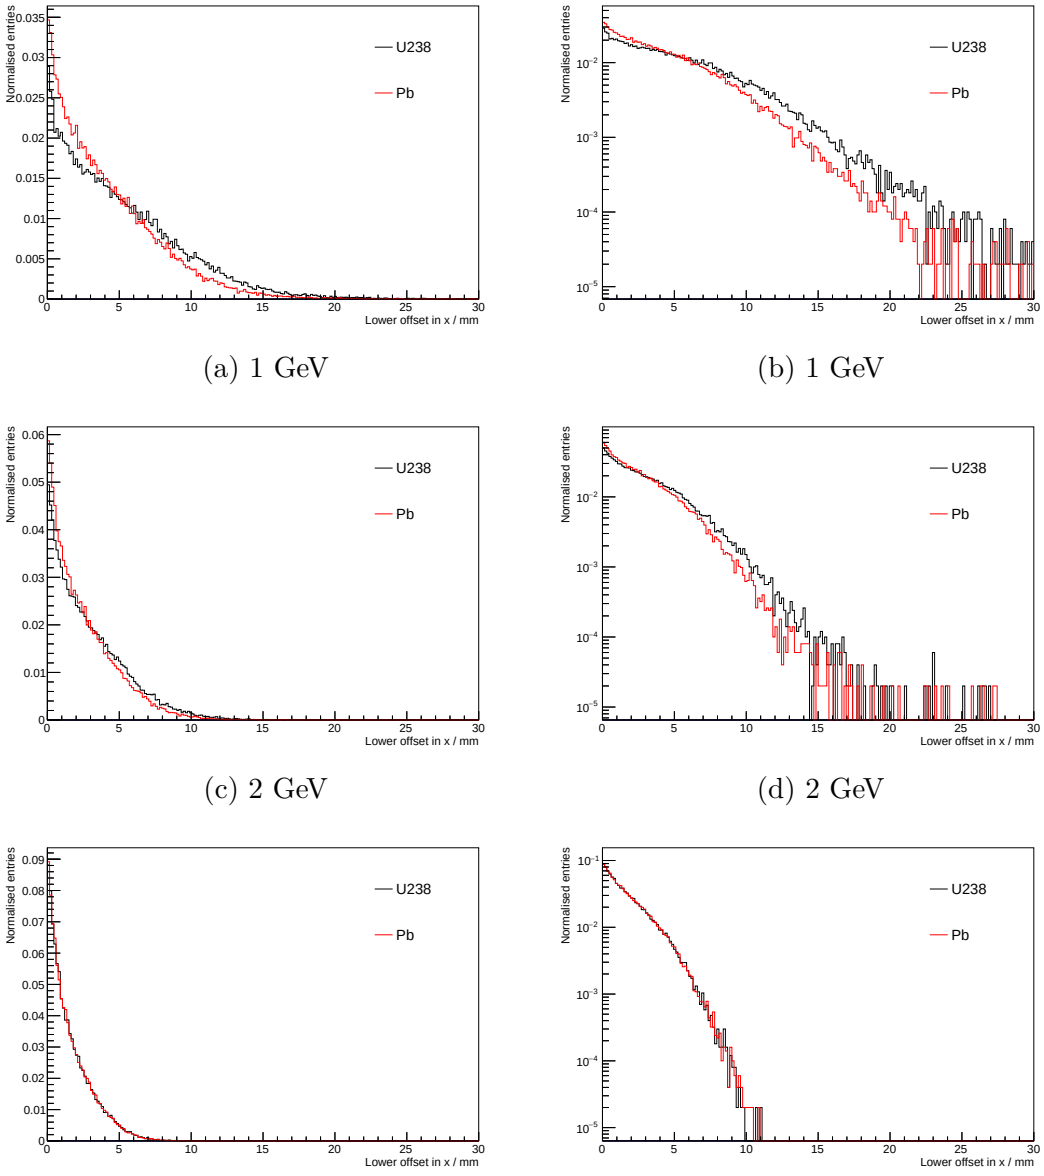

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.5: Upper offset in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

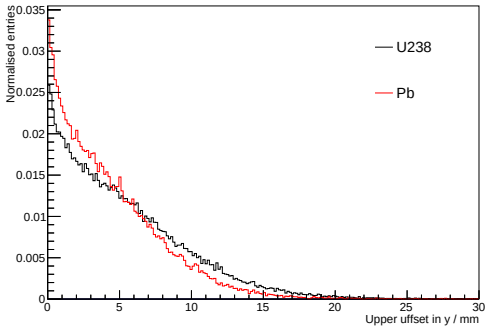

(a) 1 GeV

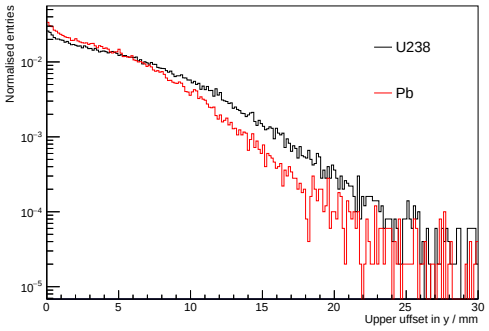

(b) 1 GeV

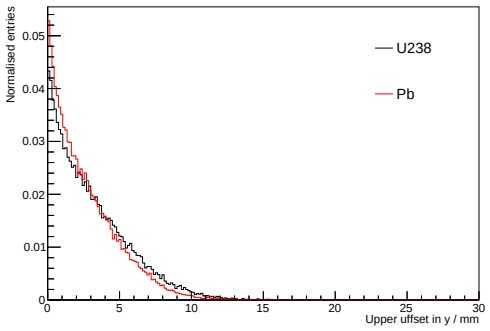

(c) 2 GeV

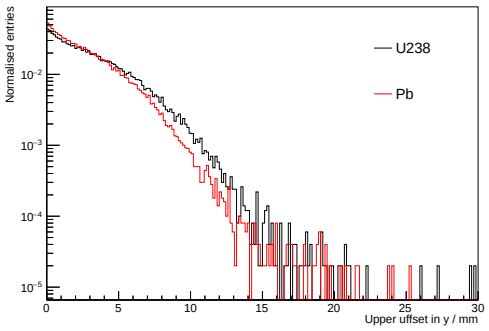

(d) 2 GeV

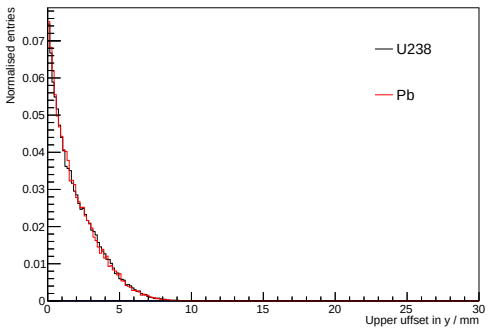

(e) 10 GeV

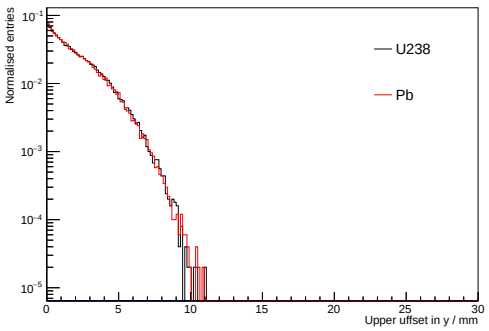

(f) 10 GeV

Figure A.6: Lower offset in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

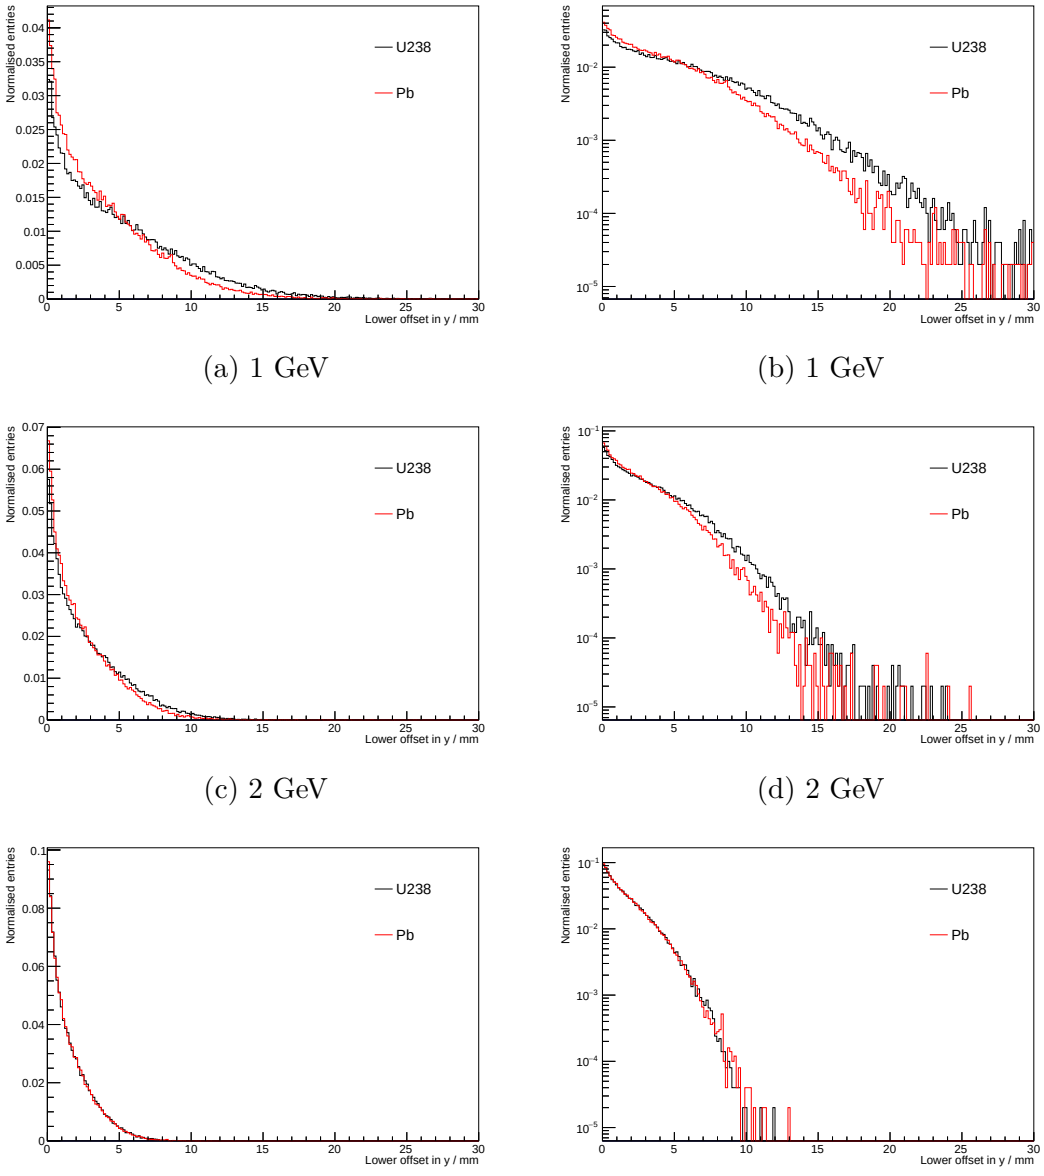

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.7: Lower offset in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

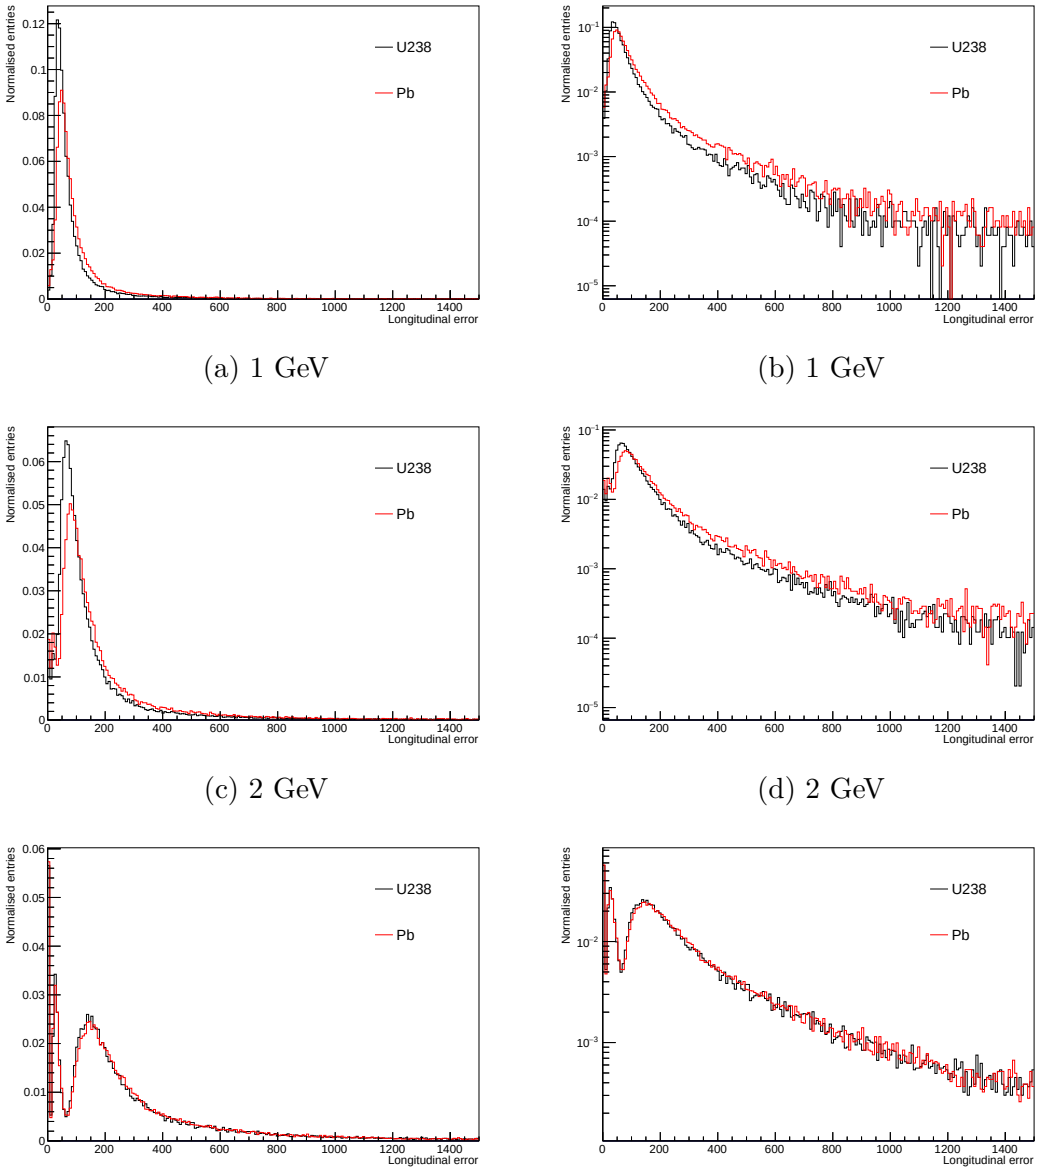

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.8: Longitudinal error for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

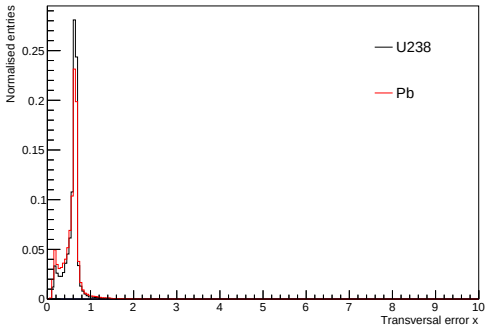

(a) 1 GeV

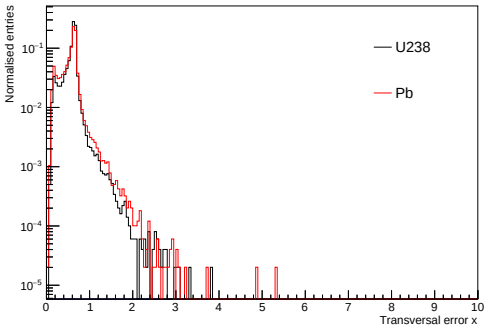

(b) 1 GeV

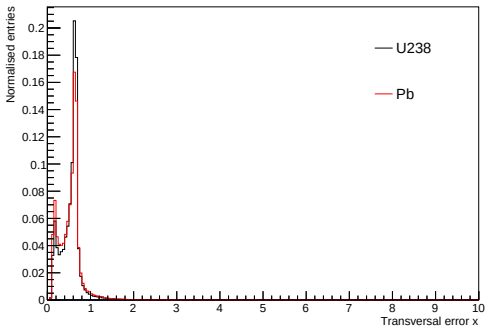

(c) 2 GeV

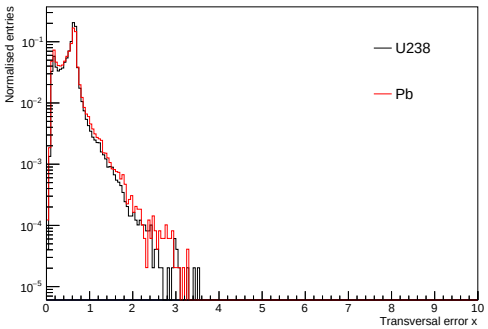

(d) 2 GeV

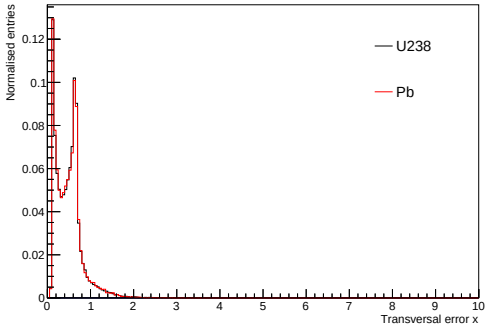

(e) 10 GeV

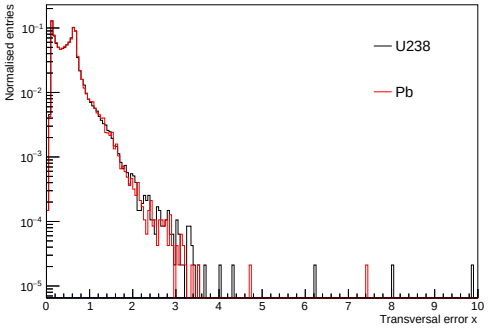

(f) 10 GeV

Figure A.9: Transversal error in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

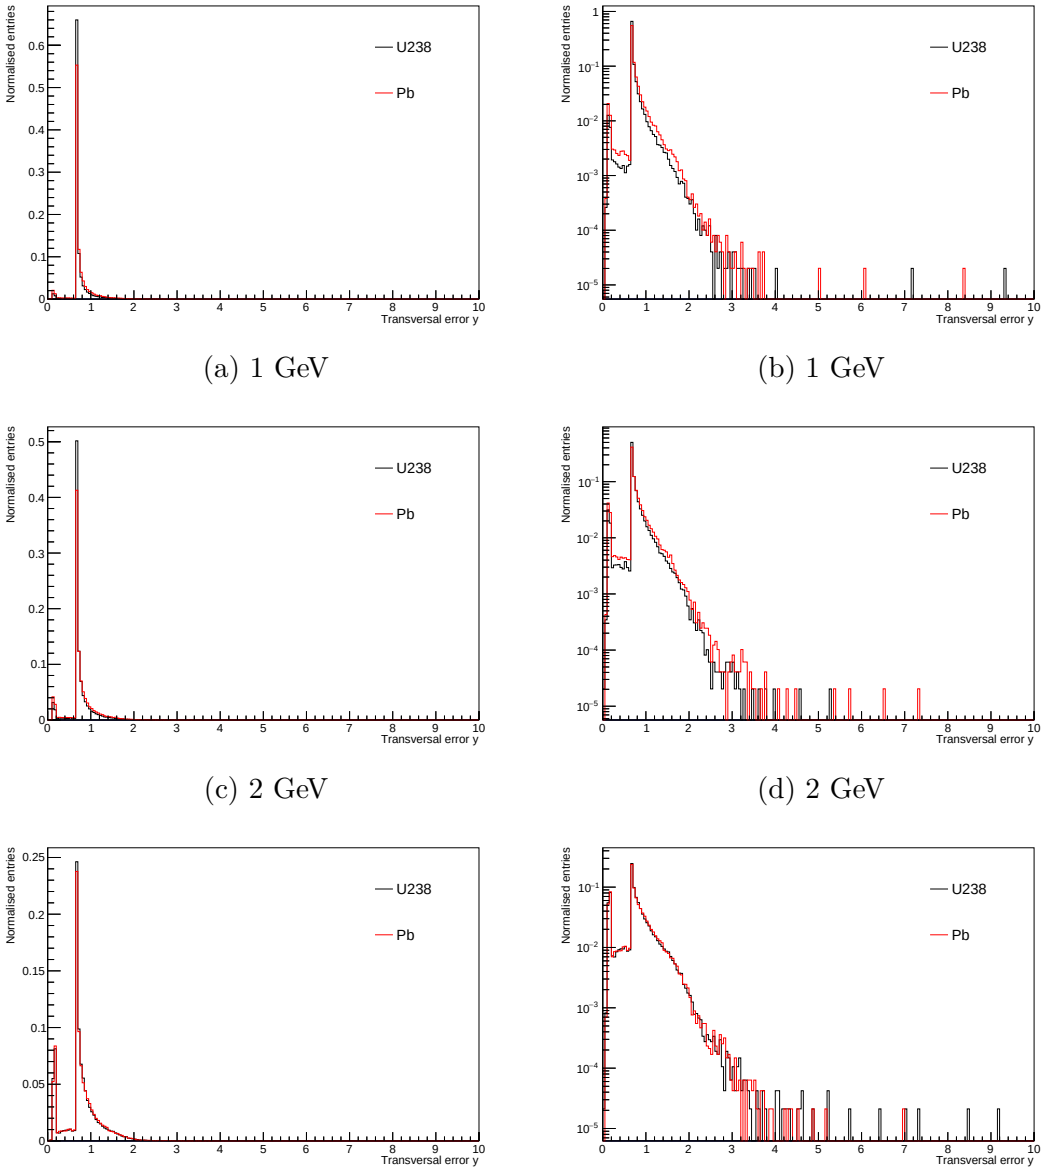

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.10: Transversal error in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

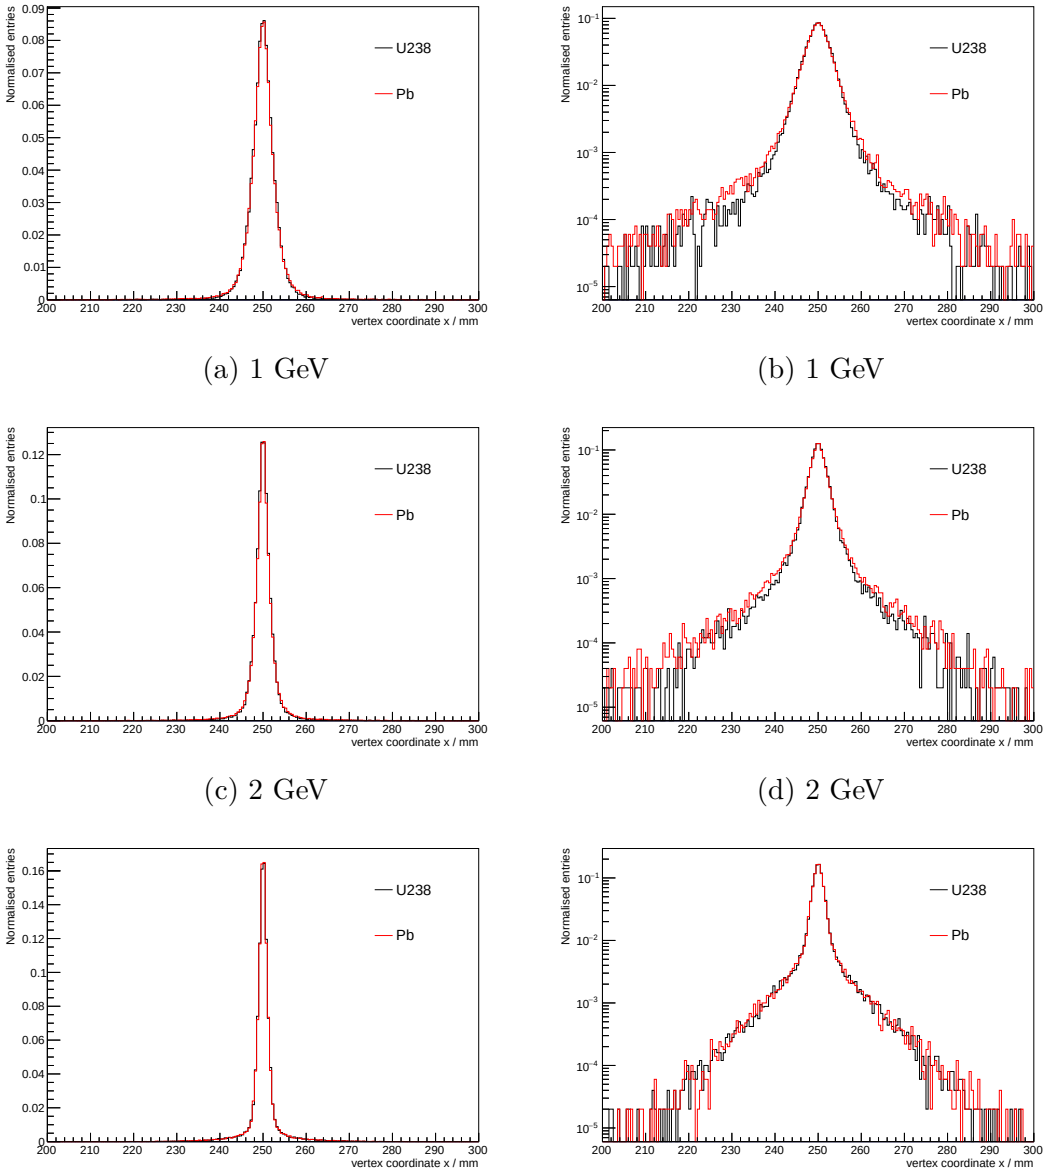

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.11: Coordinate x of the reconstructed vertex for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

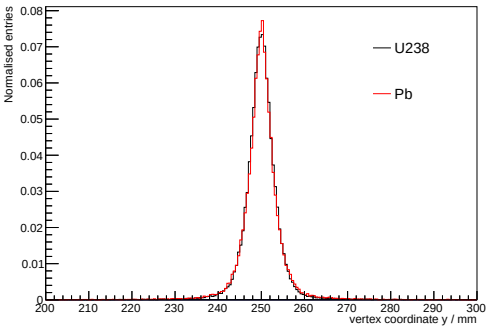

(a) 1 GeV

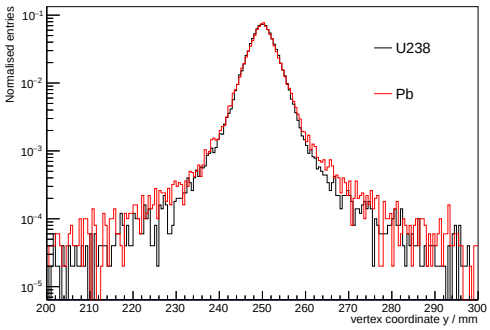

(b) 1 GeV

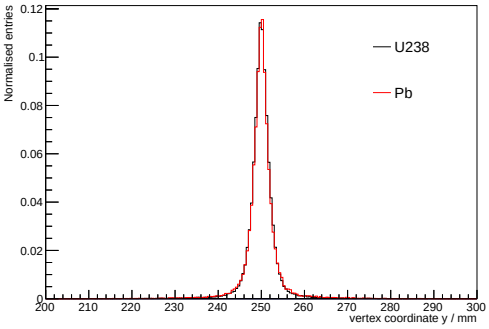

(c) 2 GeV

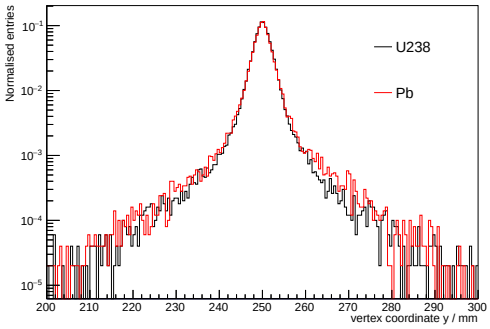

(d) 2 GeV

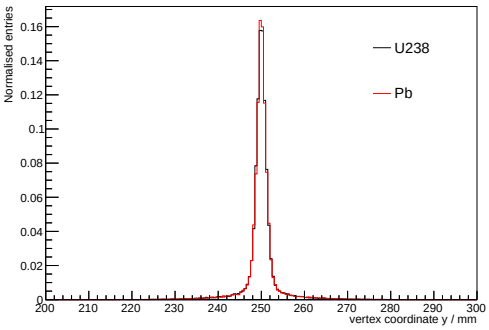

(e) 10 GeV

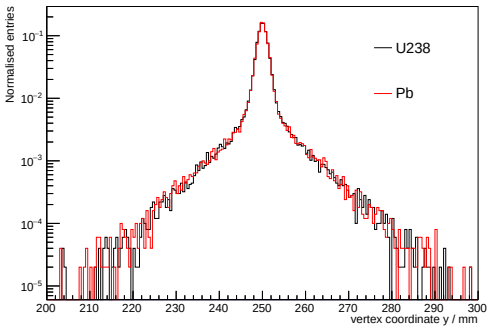

(f) 10 GeV

Figure A.12: Coordinate y of the reconstructed vertex for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

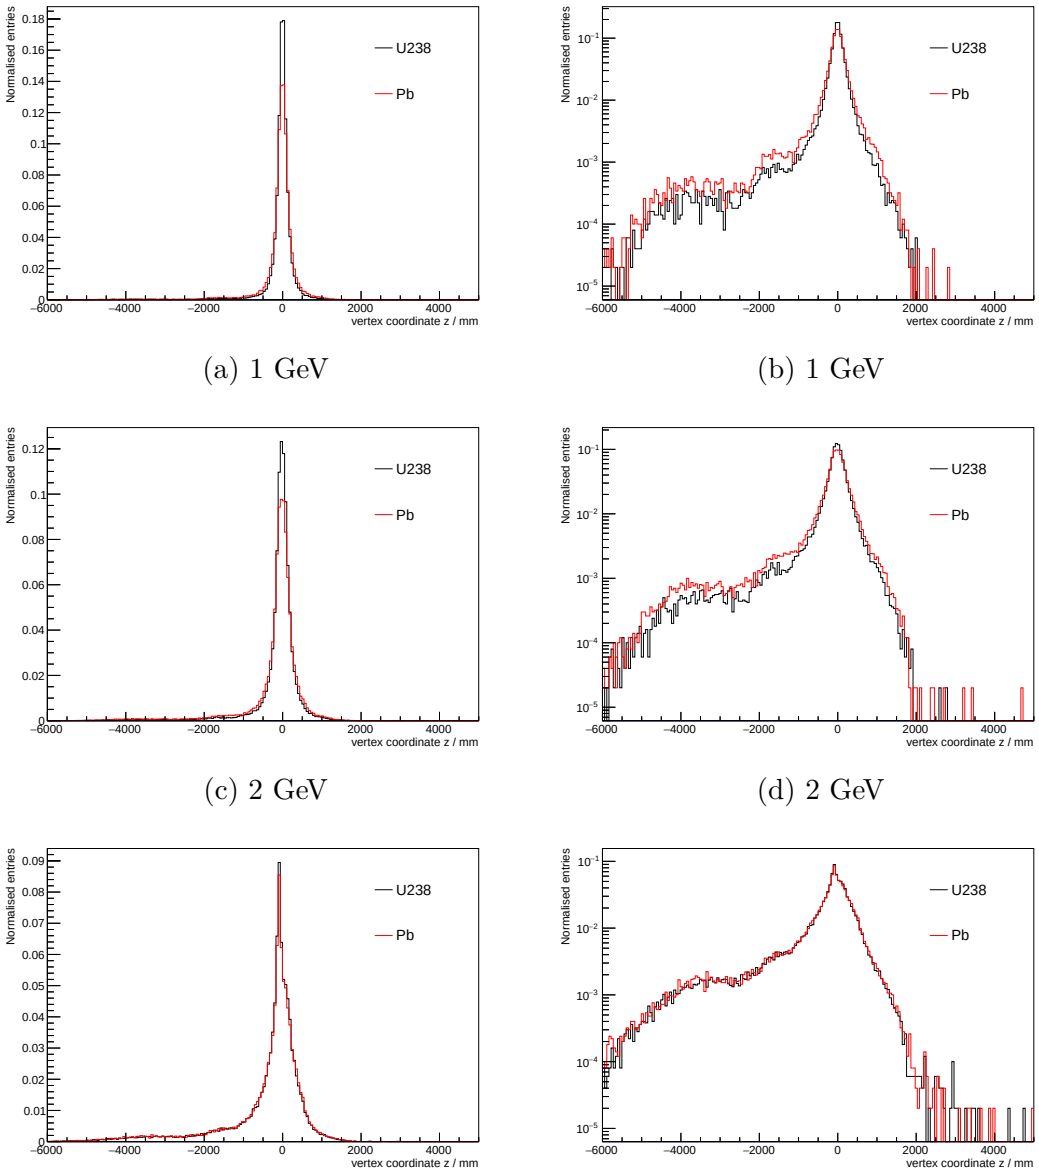

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.13: Coordinate z of the reconstructed vertex for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

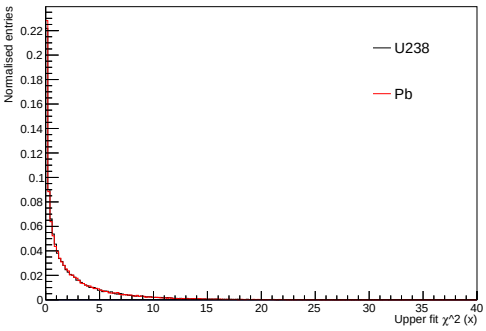

(a) 1 GeV

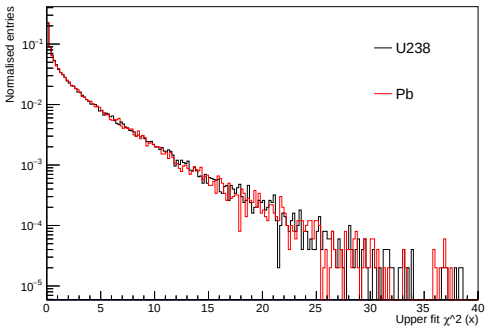

(b) 1 GeV

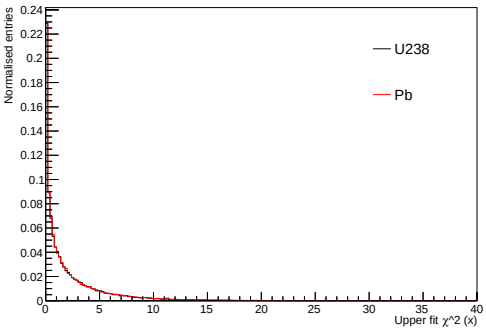

(c) 2 GeV

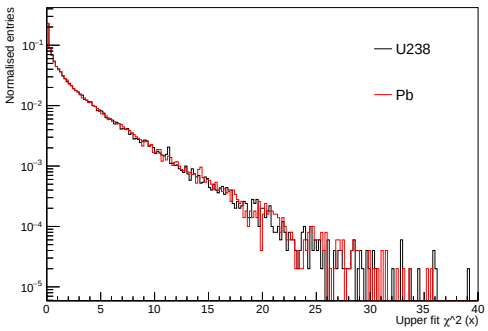

(d) 2 GeV

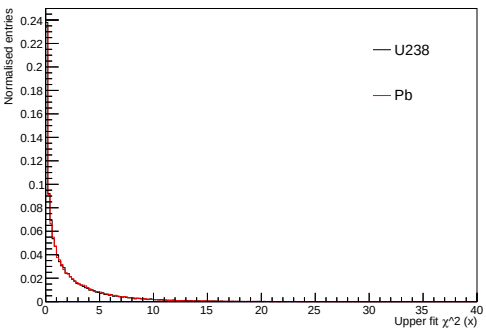

(e) 10 GeV

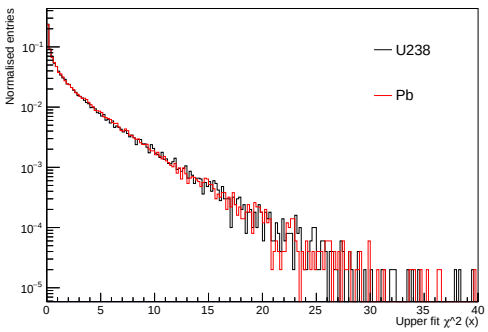

(f) 10 GeV

Figure A.14: χ^2 from fit of upper track in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

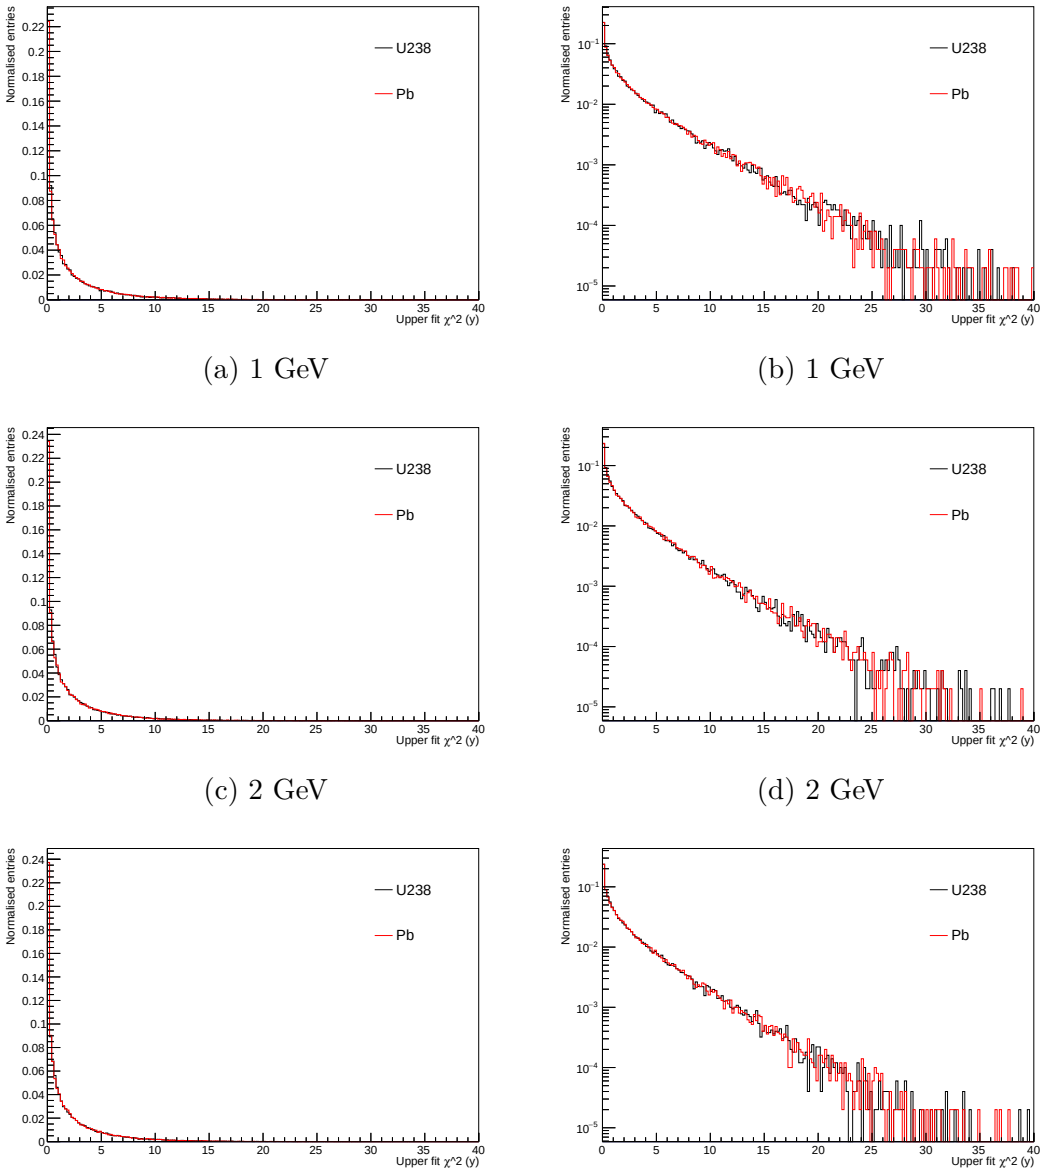

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.15: χ^2 from fit of upper track in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

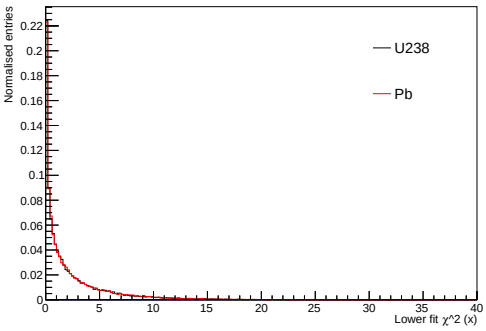

(a) 1 GeV

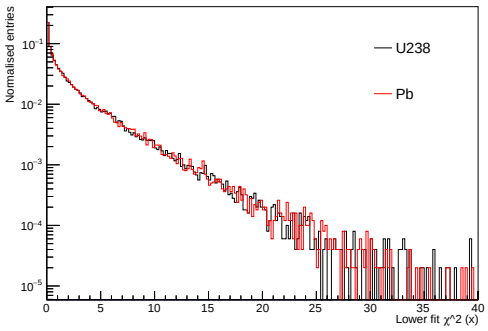

(b) 1 GeV

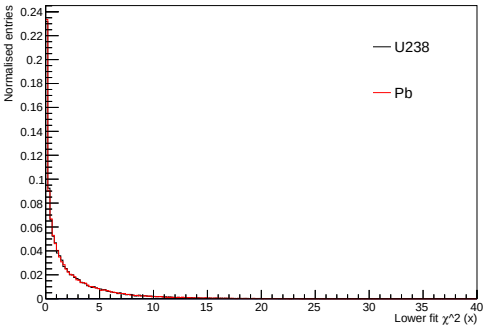

(c) 2 GeV

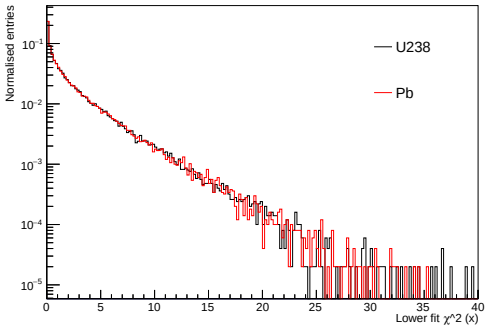

(d) 2 GeV

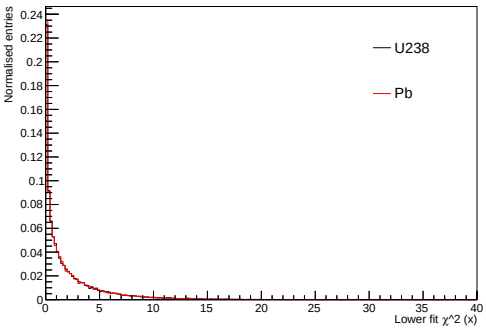

(e) 10 GeV

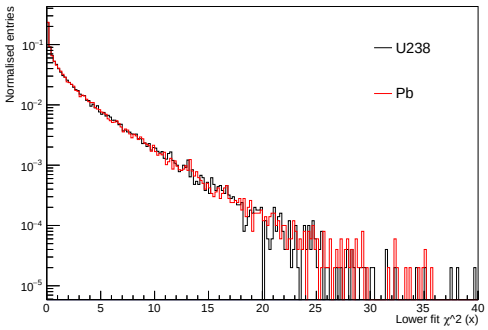

(f) 10 GeV

Figure A.16: χ^2 from fit of lower track in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

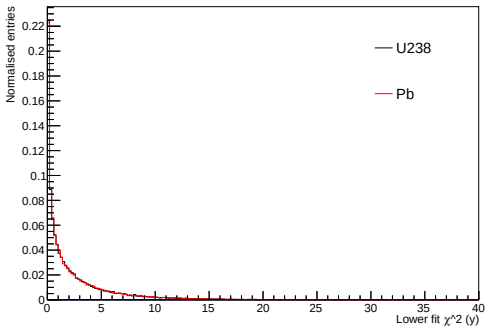

(a) 1 GeV

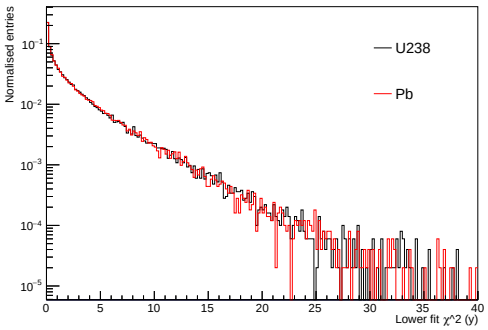

(b) 1 GeV

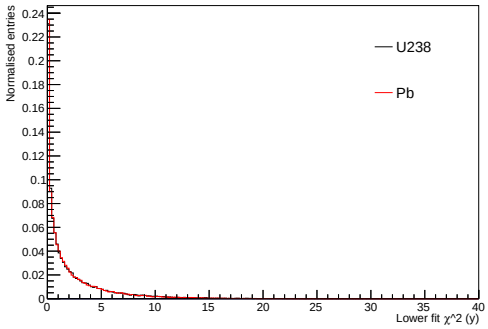

(c) 2 GeV

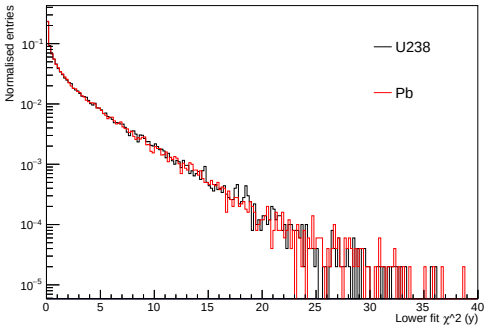

(d) 2 GeV

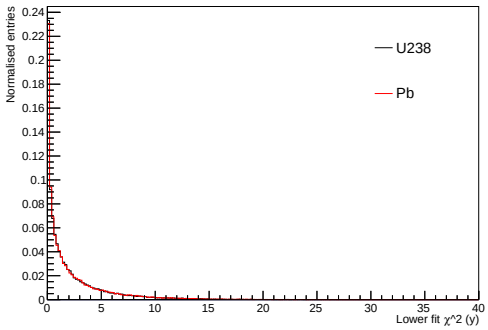

(e) 10 GeV

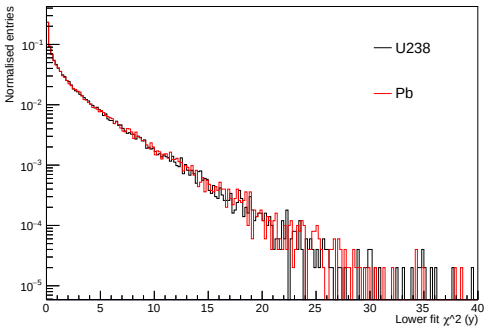

(f) 10 GeV

Figure A.17: χ^2 from fit of lower track in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

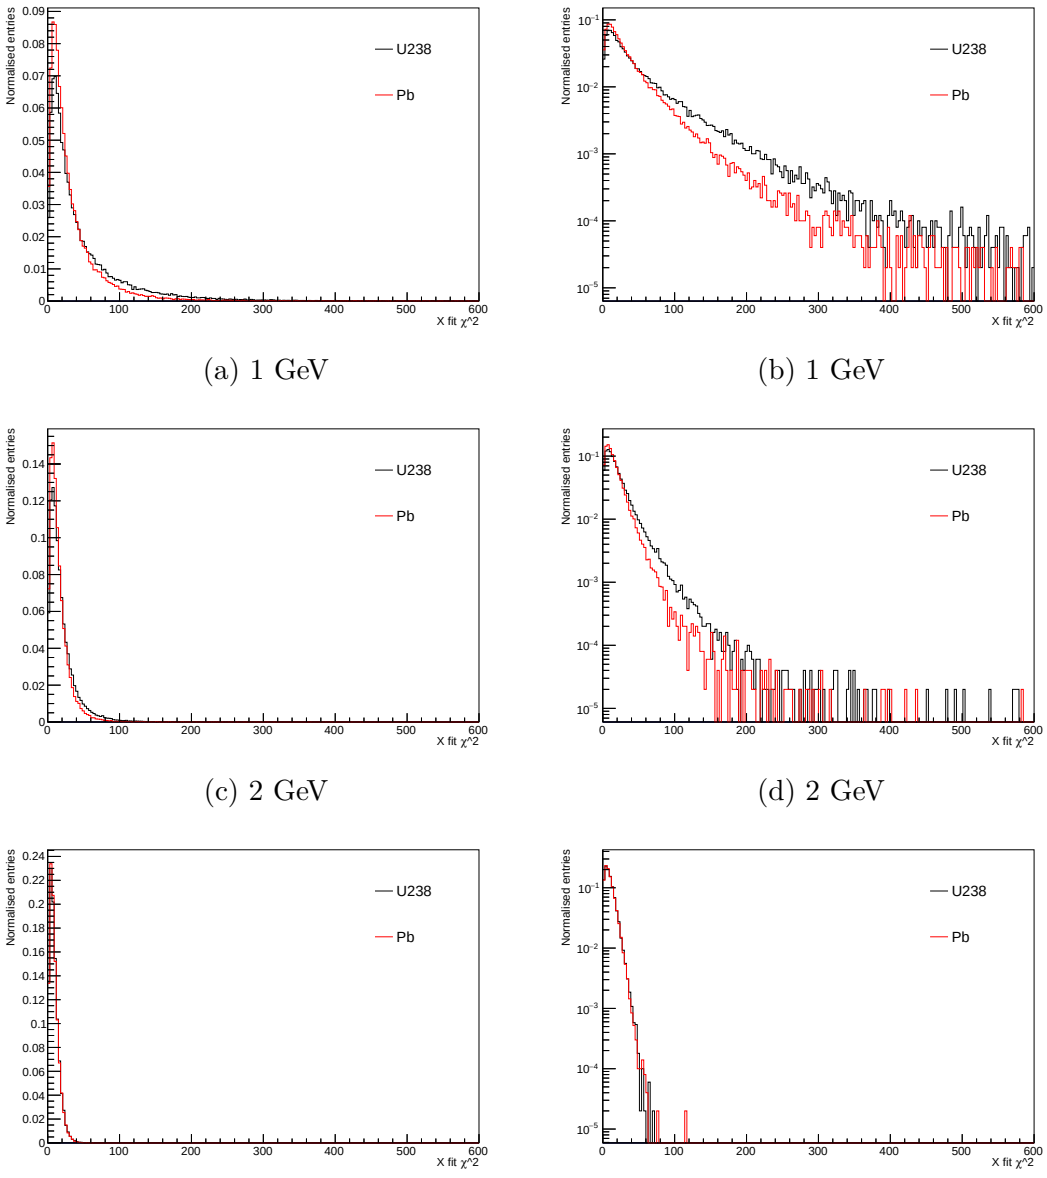

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.18: Six points fit χ^2 in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

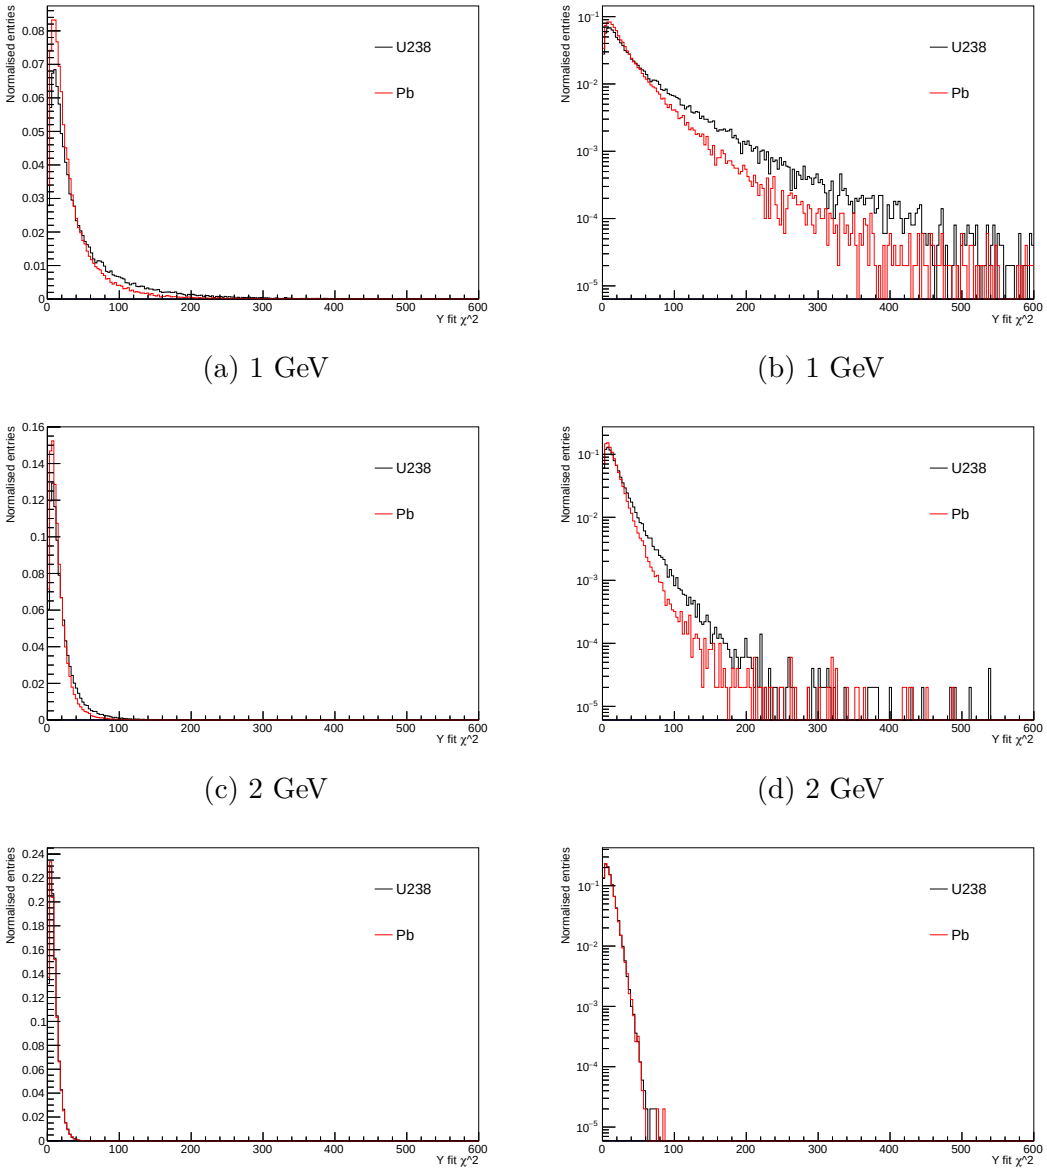

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.19: Six points fit χ^2 in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

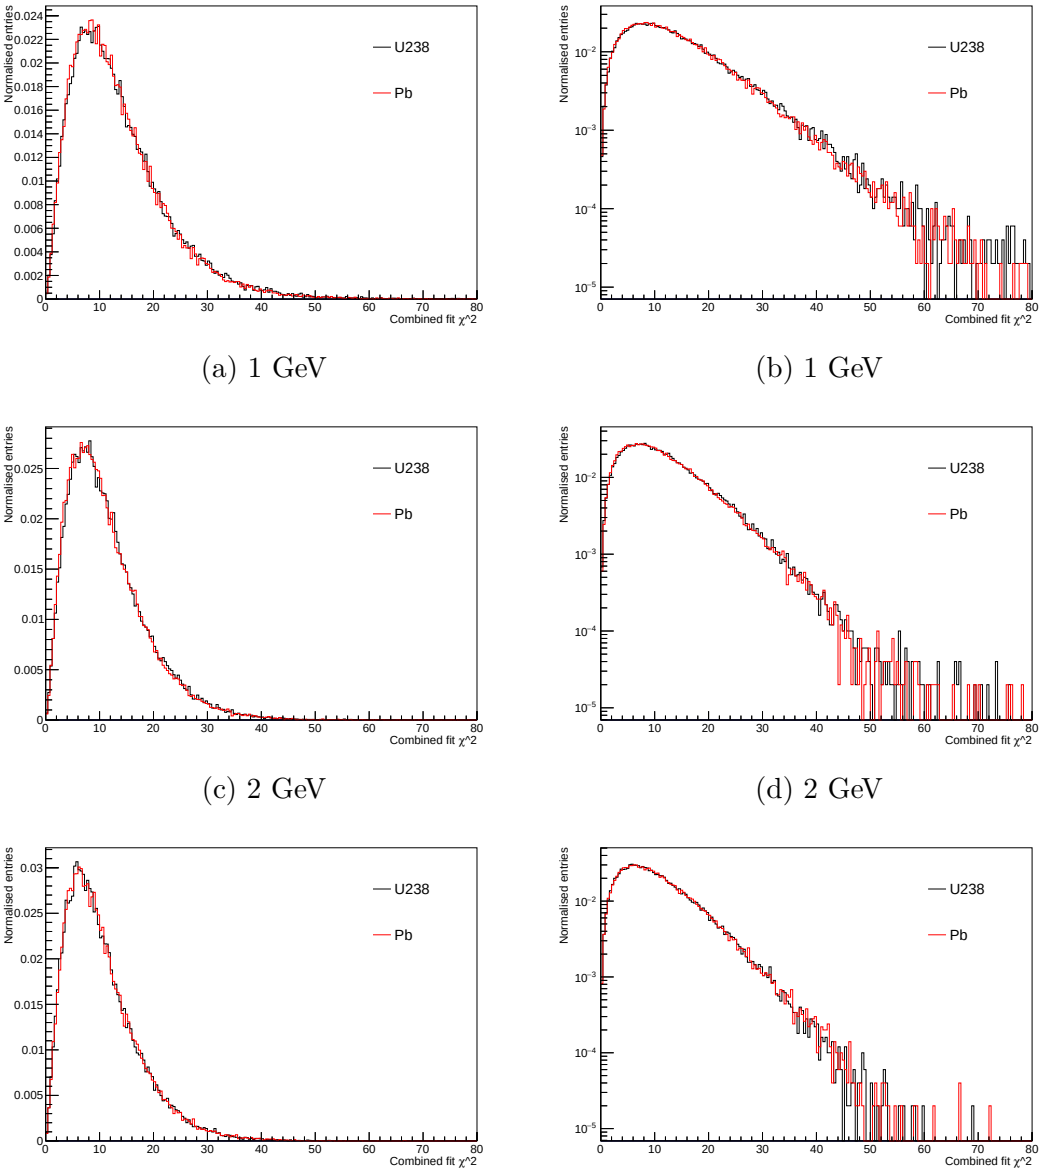

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.20: Combined fit χ^2 for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

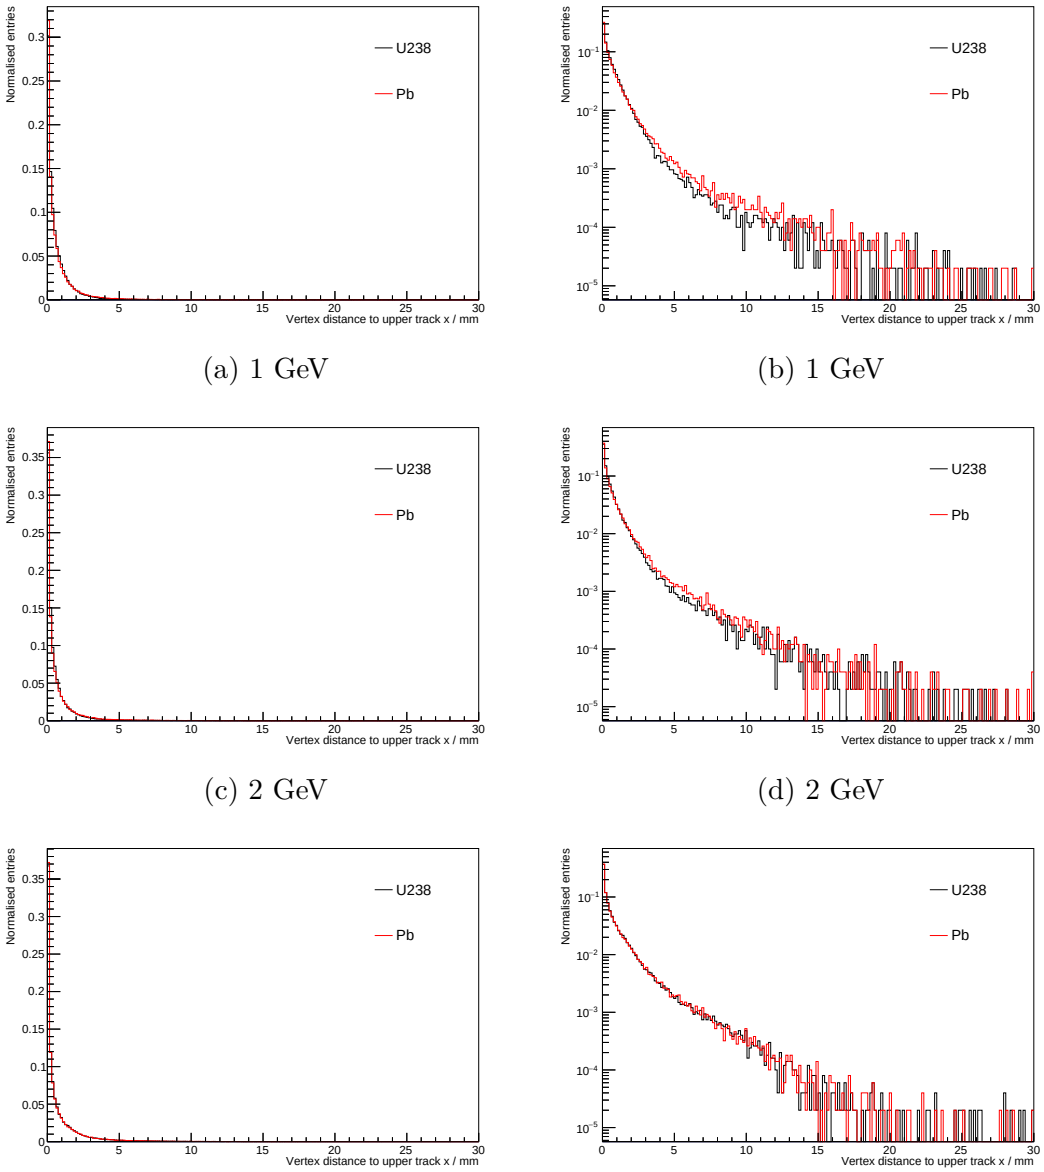

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.21: Distance between the reconstructed vertex and upper track in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

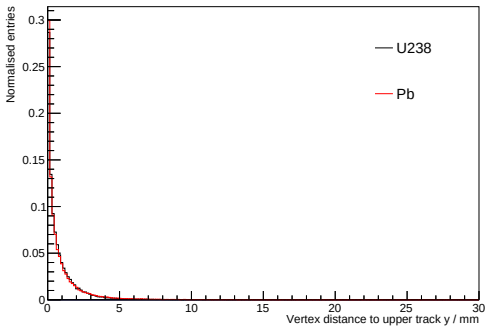

(a) 1 GeV

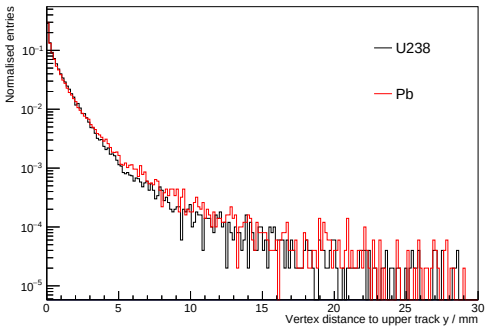

(b) 1 GeV

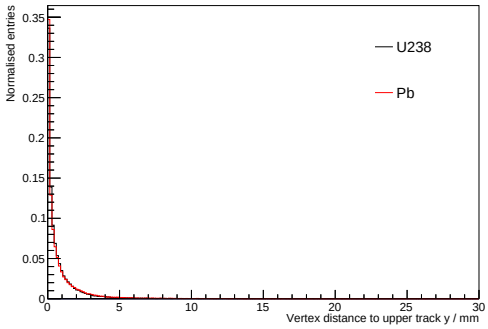

(c) 2 GeV

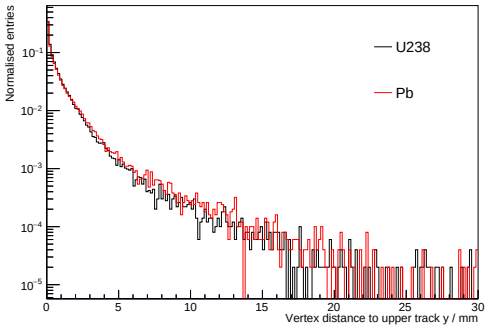

(d) 2 GeV

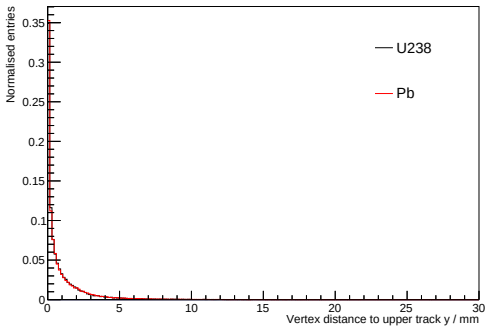

(e) 10 GeV

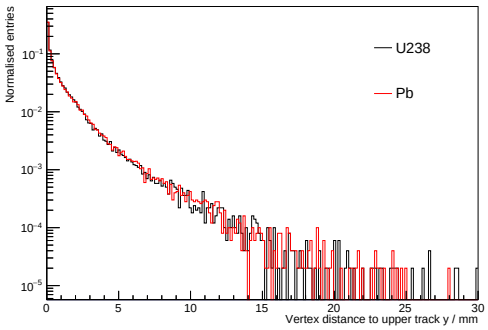

(f) 10 GeV

Figure A.22: Distance between the reconstructed vertex and upper track in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

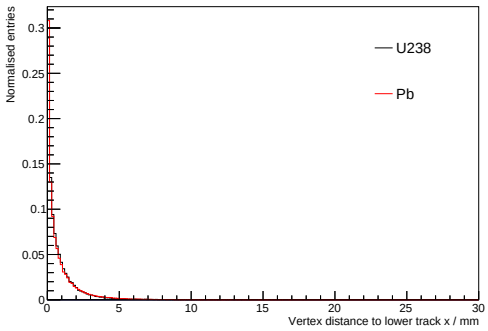

(a) 1 GeV

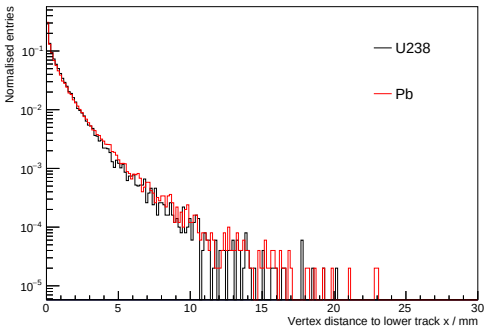

(b) 1 GeV

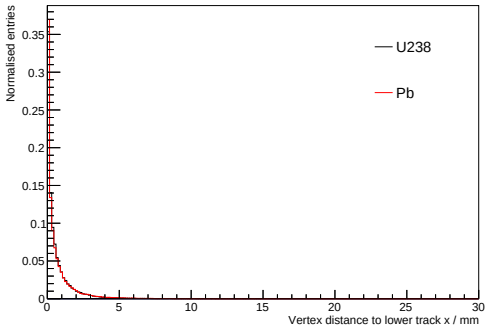

(c) 2 GeV

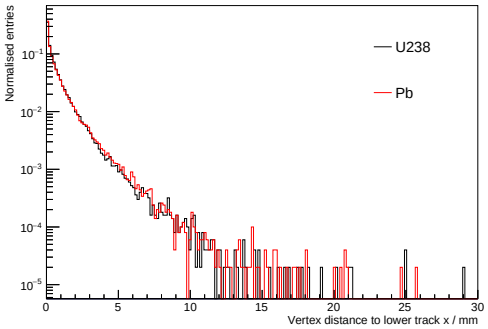

(d) 2 GeV

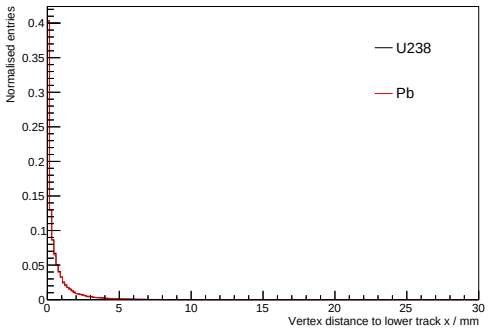

(e) 10 GeV

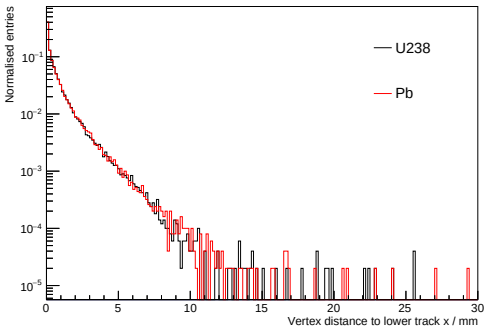

(f) 10 GeV

Figure A.23: Distance between the reconstructed vertex and lower track in x for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

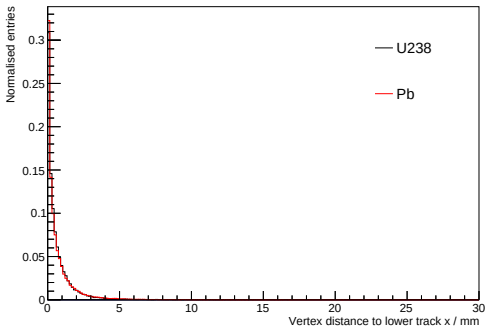

(a) 1 GeV

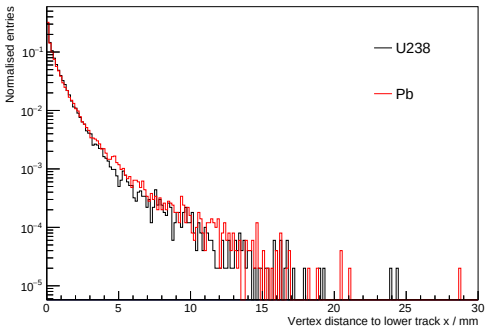

(b) 1 GeV

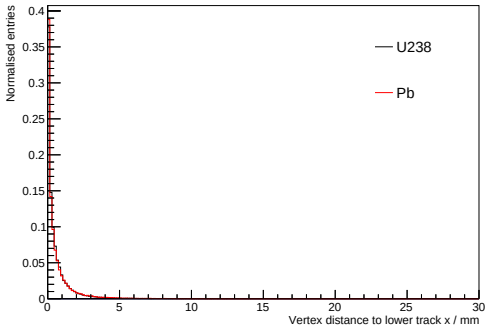

(c) 2 GeV

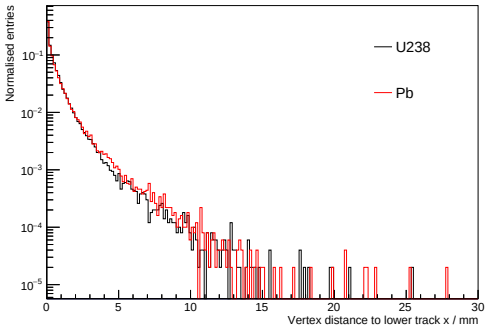

(d) 2 GeV

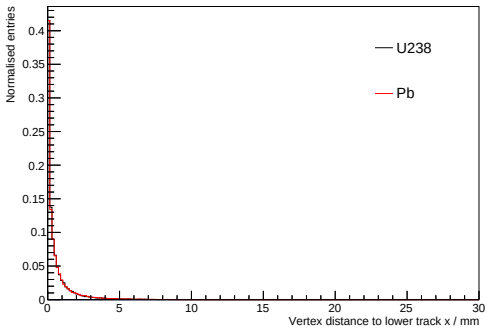

(e) 10 GeV

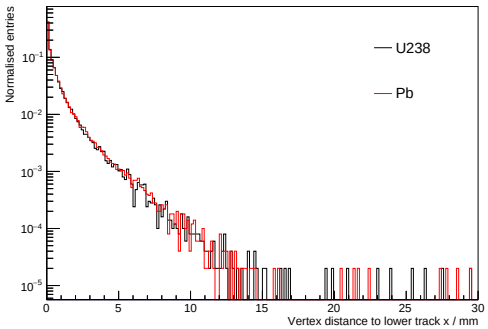

(f) 10 GeV

Figure A.24: Distance between the reconstructed vertex and lower track in y for uranium and lead plates, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

A.2 Simulation of a block

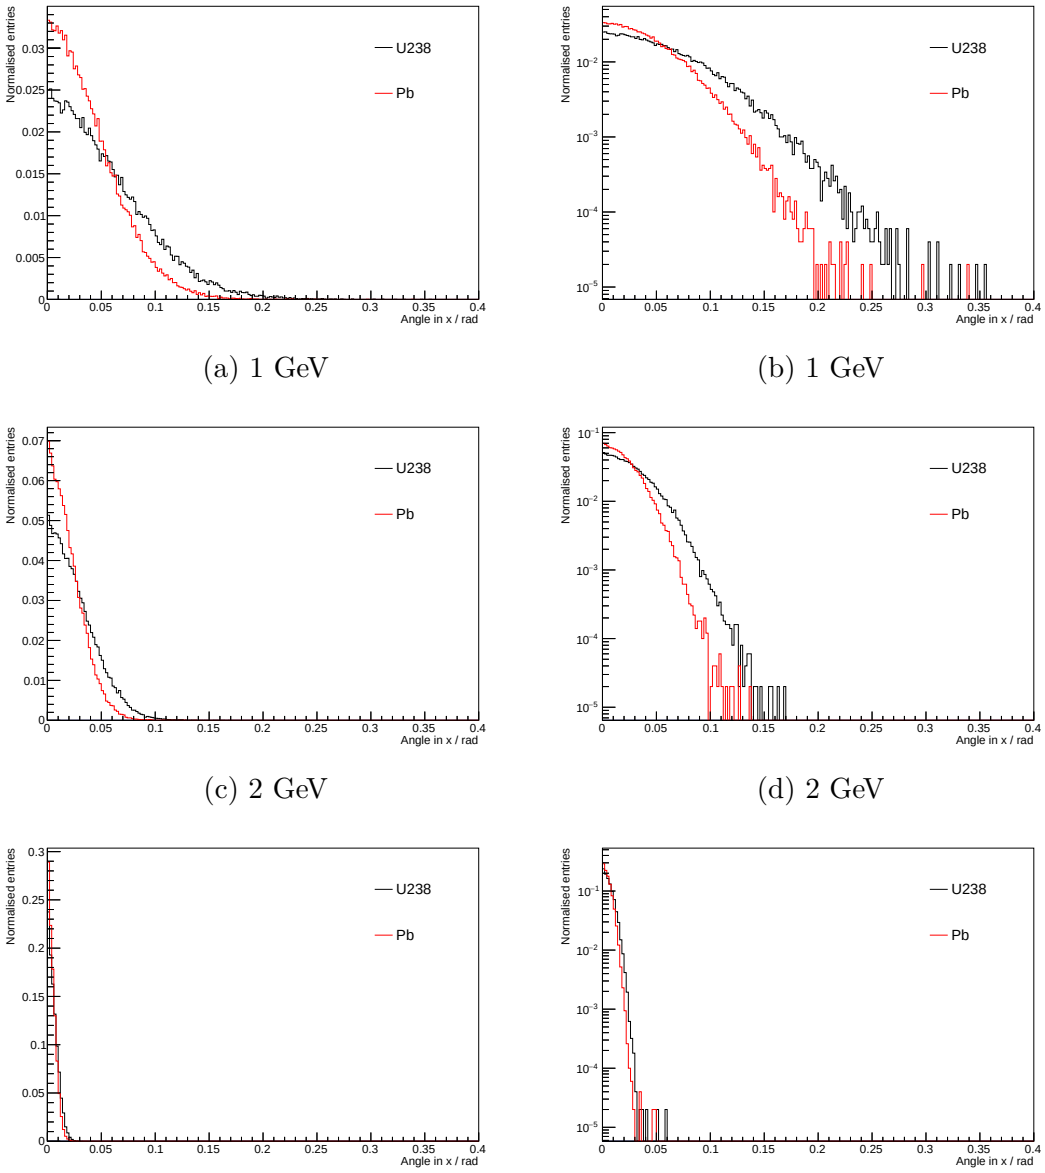

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.25: Projected angle in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, in linear (left) and logarithmic (right) scales, for different muon energies. The histograms are normalised.

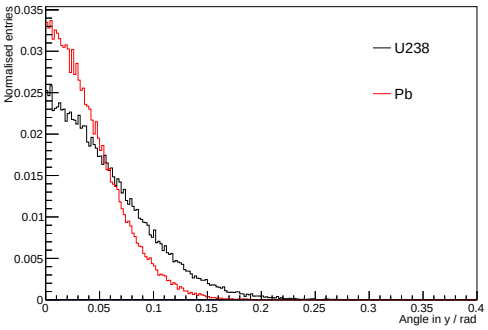

(a) 1 GeV

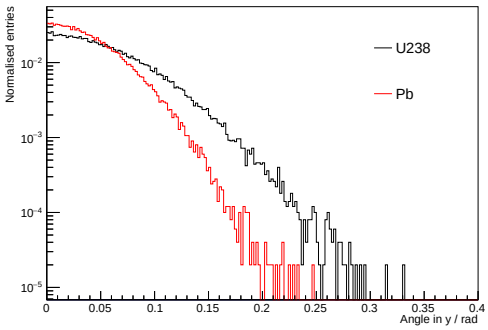

(b) 1 GeV

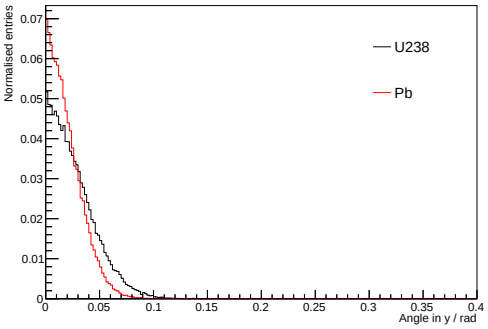

(c) 2 GeV

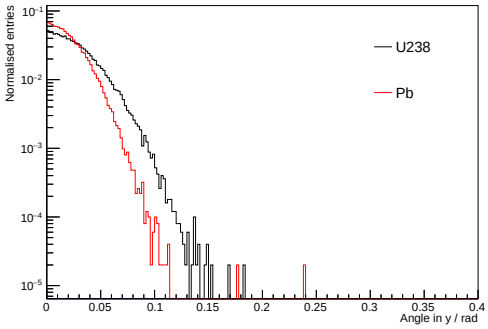

(d) 2 GeV

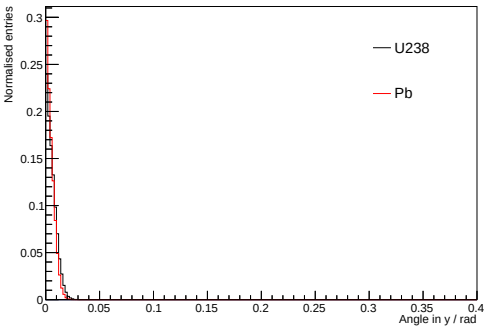

(e) 10 GeV

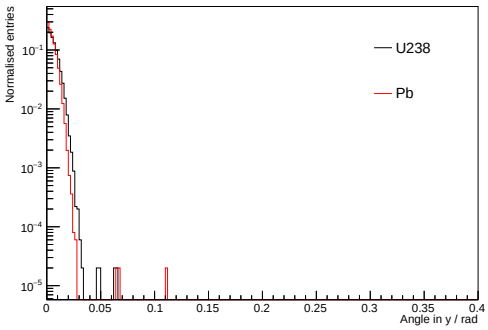

(f) 10 GeV

Figure A.26: Projected angle in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, in linear (left) and logarithmic (right) scales, for different muon energies. The histograms are normalised.

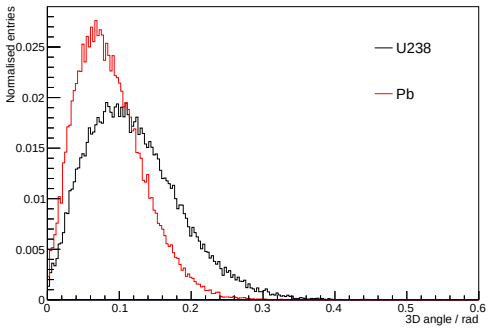

(a) 1 GeV

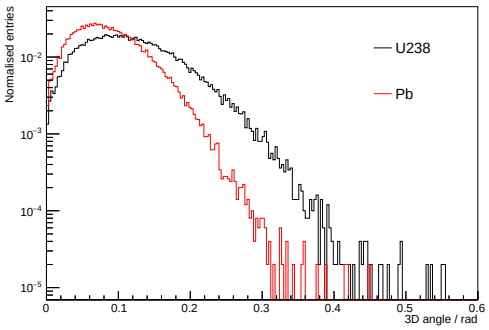

(b) 1 GeV

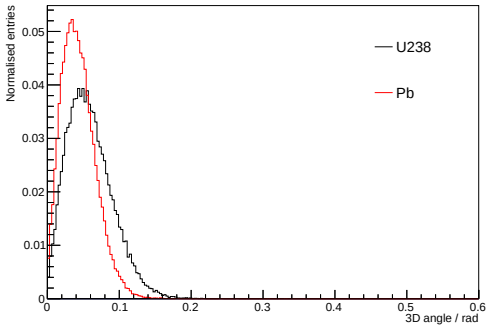

(c) 2 GeV

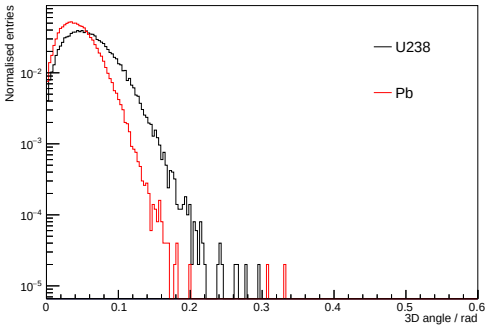

(d) 2 GeV

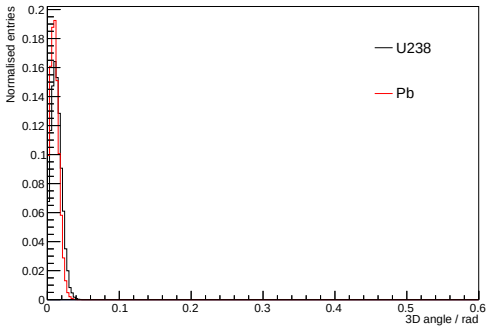

(e) 10 GeV

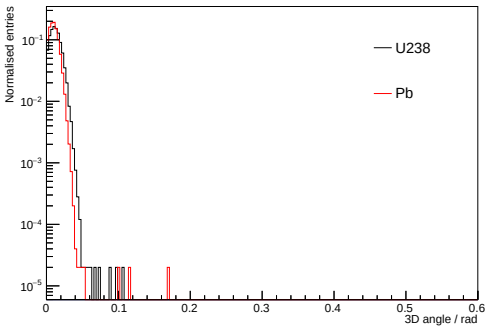

(f) 10 GeV

Figure A.27: 3D angle for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, in linear (left) and logarithmic (right) scales, for different muon energies. The histograms are normalised.

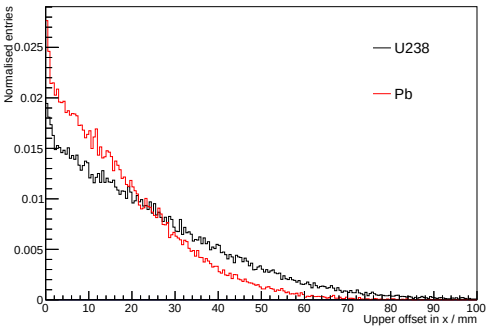

(a) 1 GeV

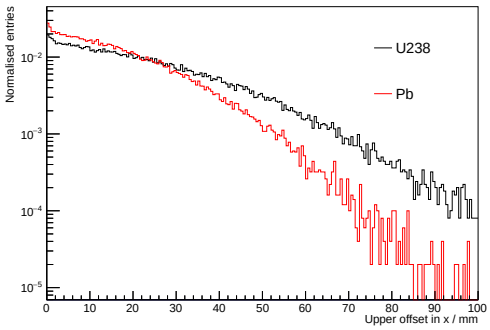

(b) 1 GeV

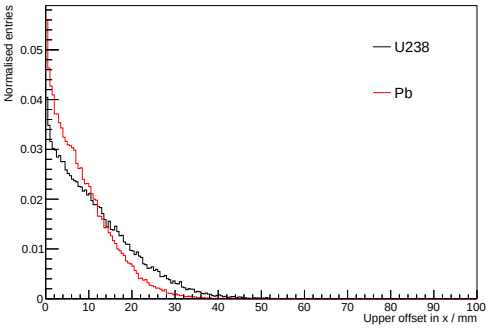

(c) 2 GeV

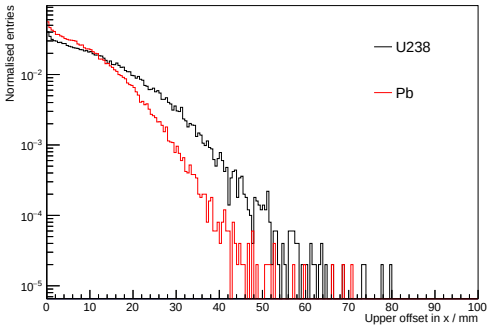

(d) 2 GeV

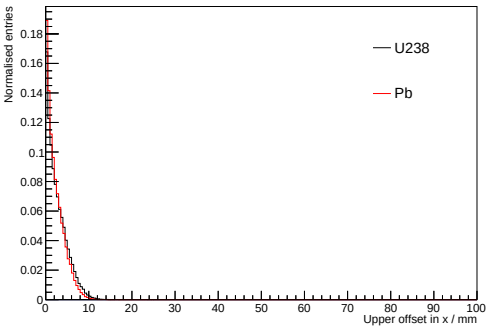

(e) 10 GeV

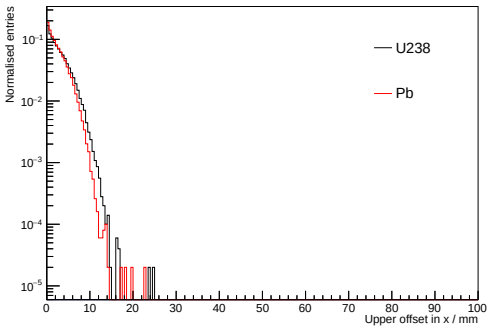

(f) 10 GeV

Figure A.28: Upper offset in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

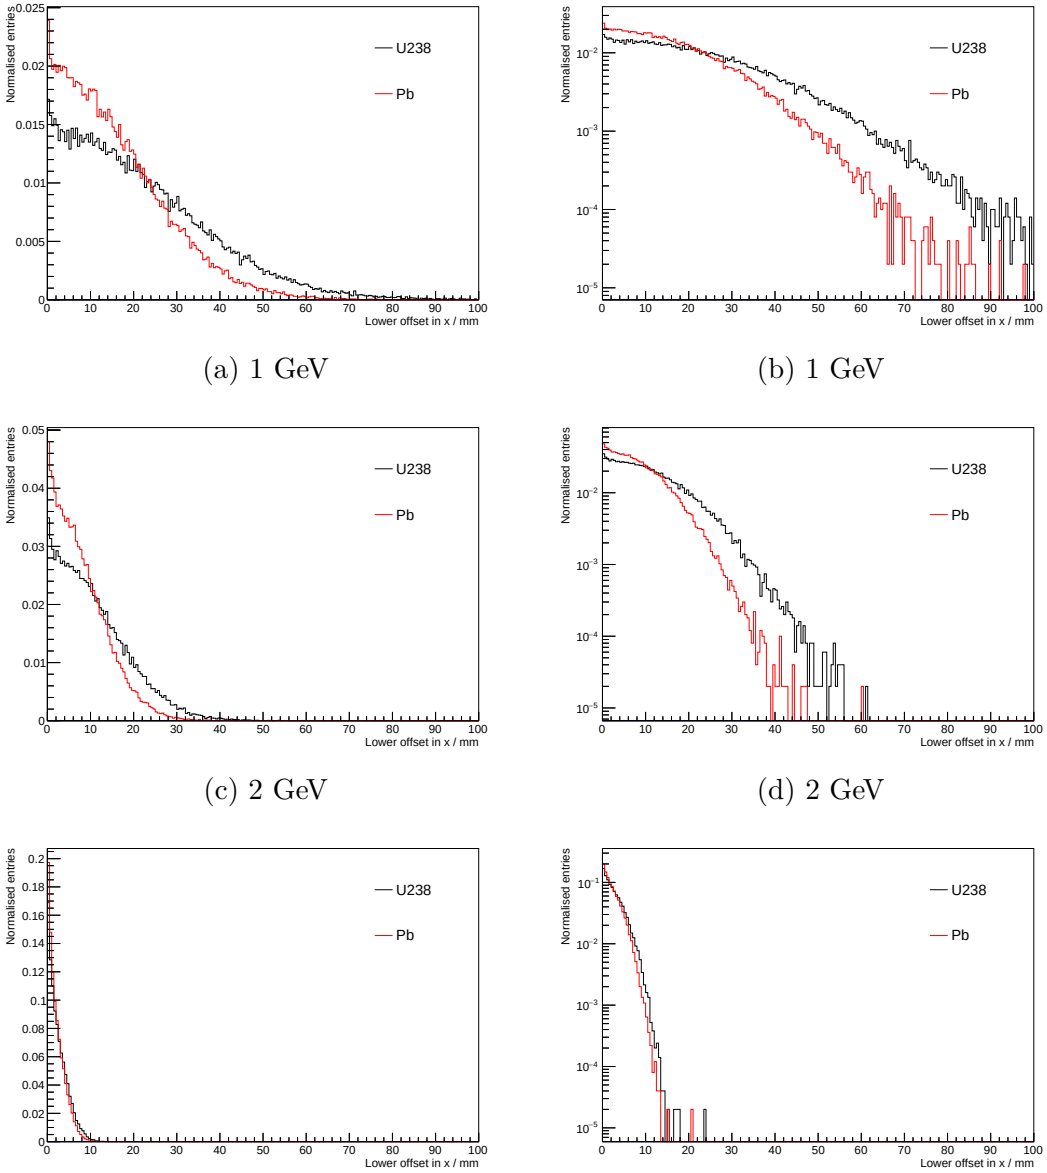

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.29: Upper offset in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

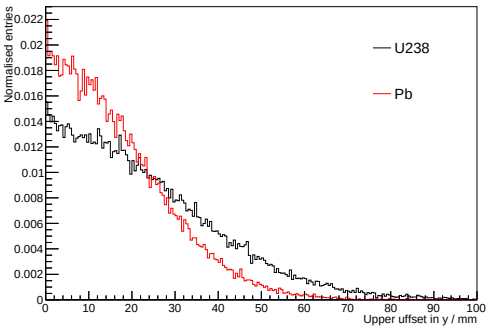

(a) 1 GeV

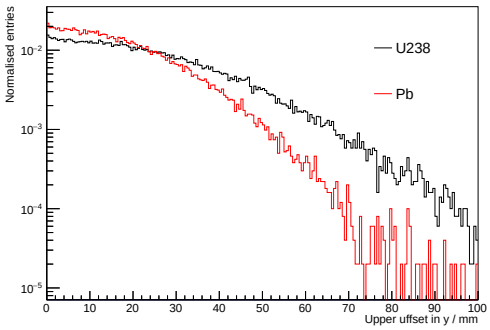

(b) 1 GeV

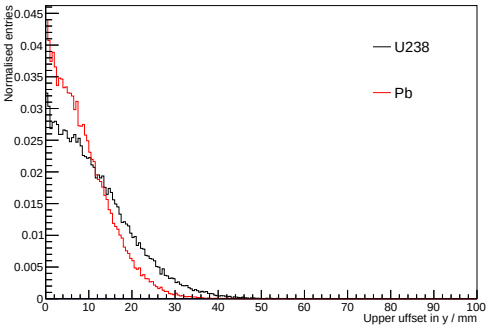

(c) 2 GeV

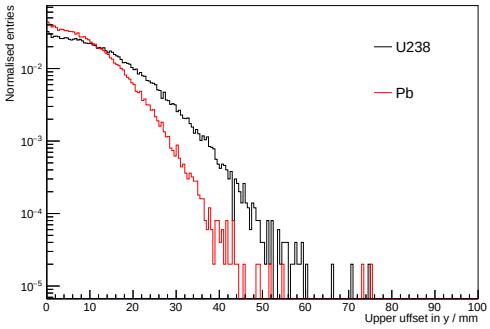

(d) 2 GeV

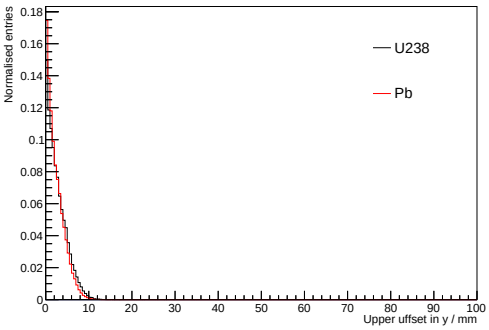

(e) 10 GeV

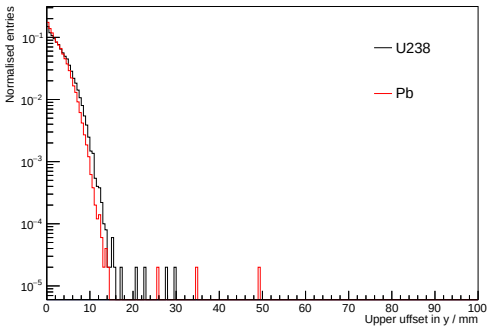

(f) 10 GeV

Figure A.30: Lower offset in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

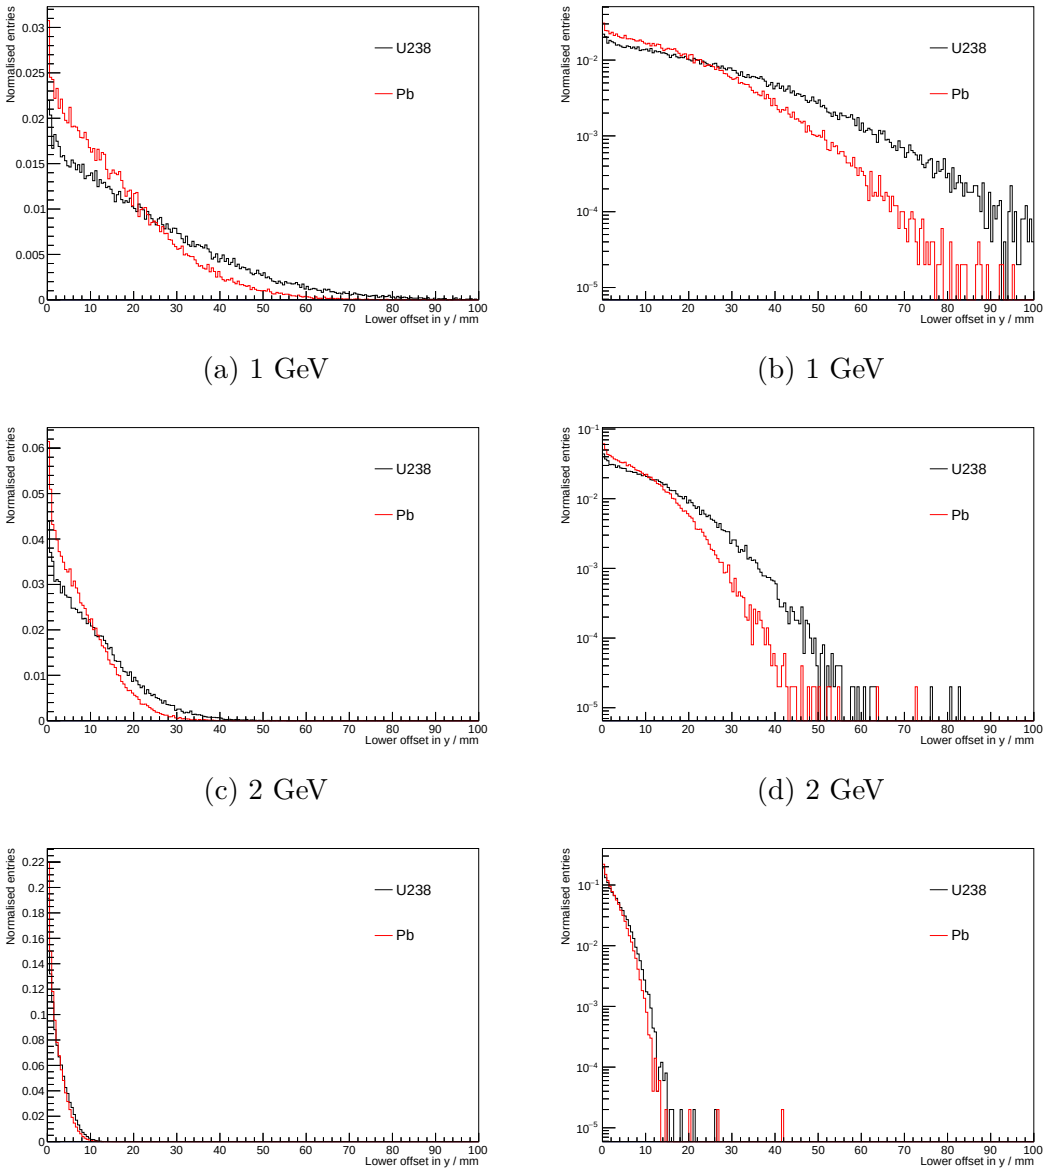

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.31: Lower offset in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

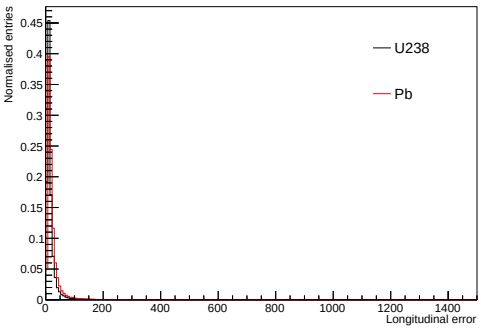

(a) 1 GeV

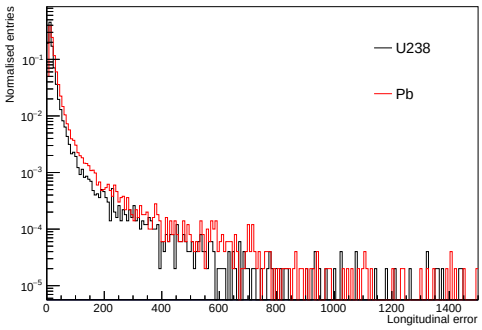

(b) 1 GeV

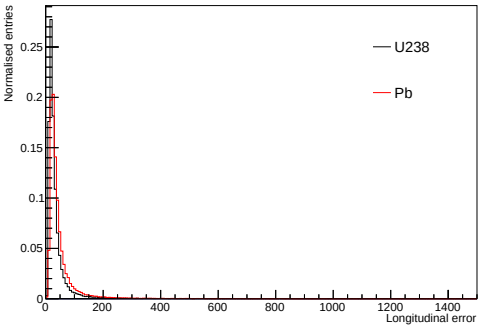

(c) 2 GeV

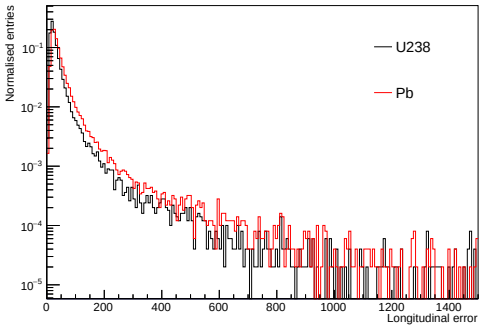

(d) 2 GeV

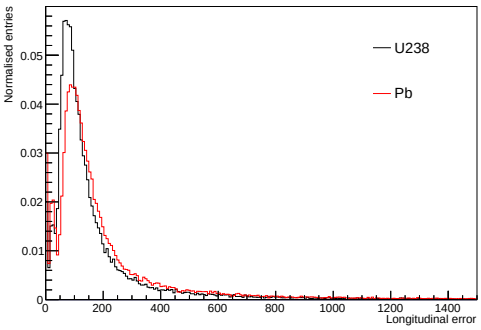

(e) 10 GeV

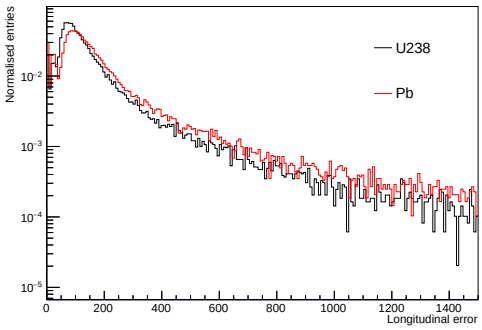

(f) 10 GeV

Figure A.32: Longitudinal error for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

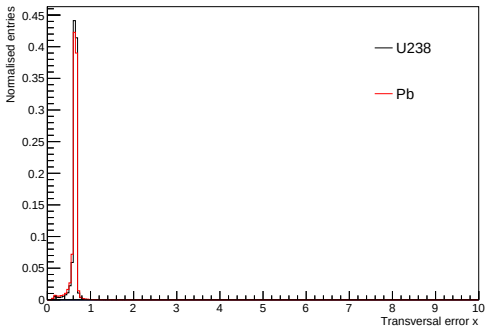

(a) 1 GeV

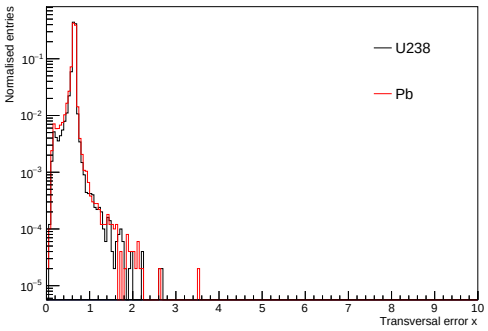

(b) 1 GeV

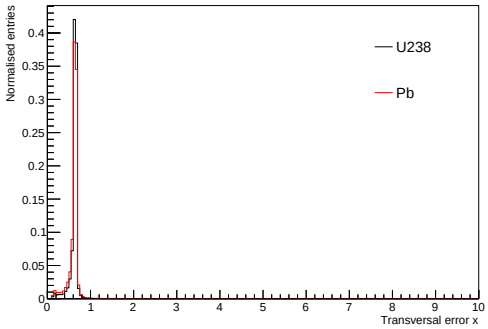

(c) 2 GeV

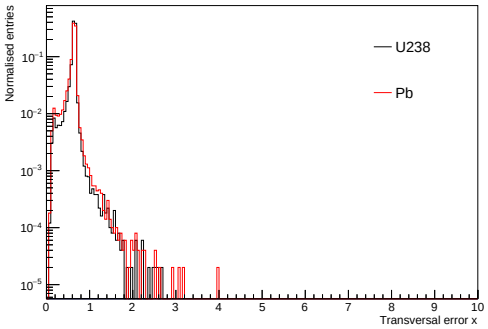

(d) 2 GeV

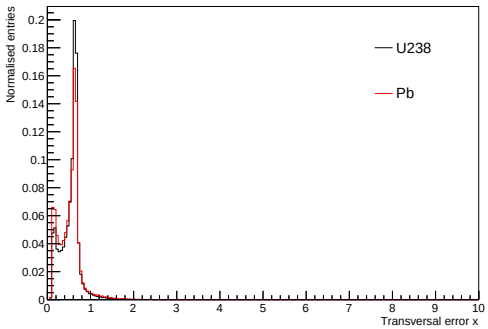

(e) 10 GeV

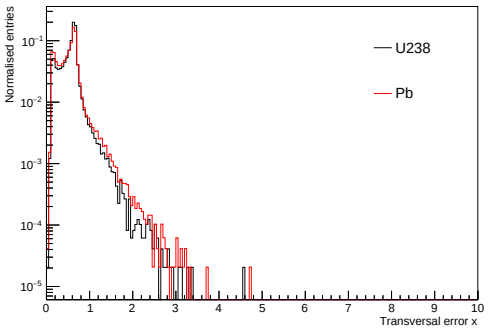

(f) 10 GeV

Figure A.33: Transversal error in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

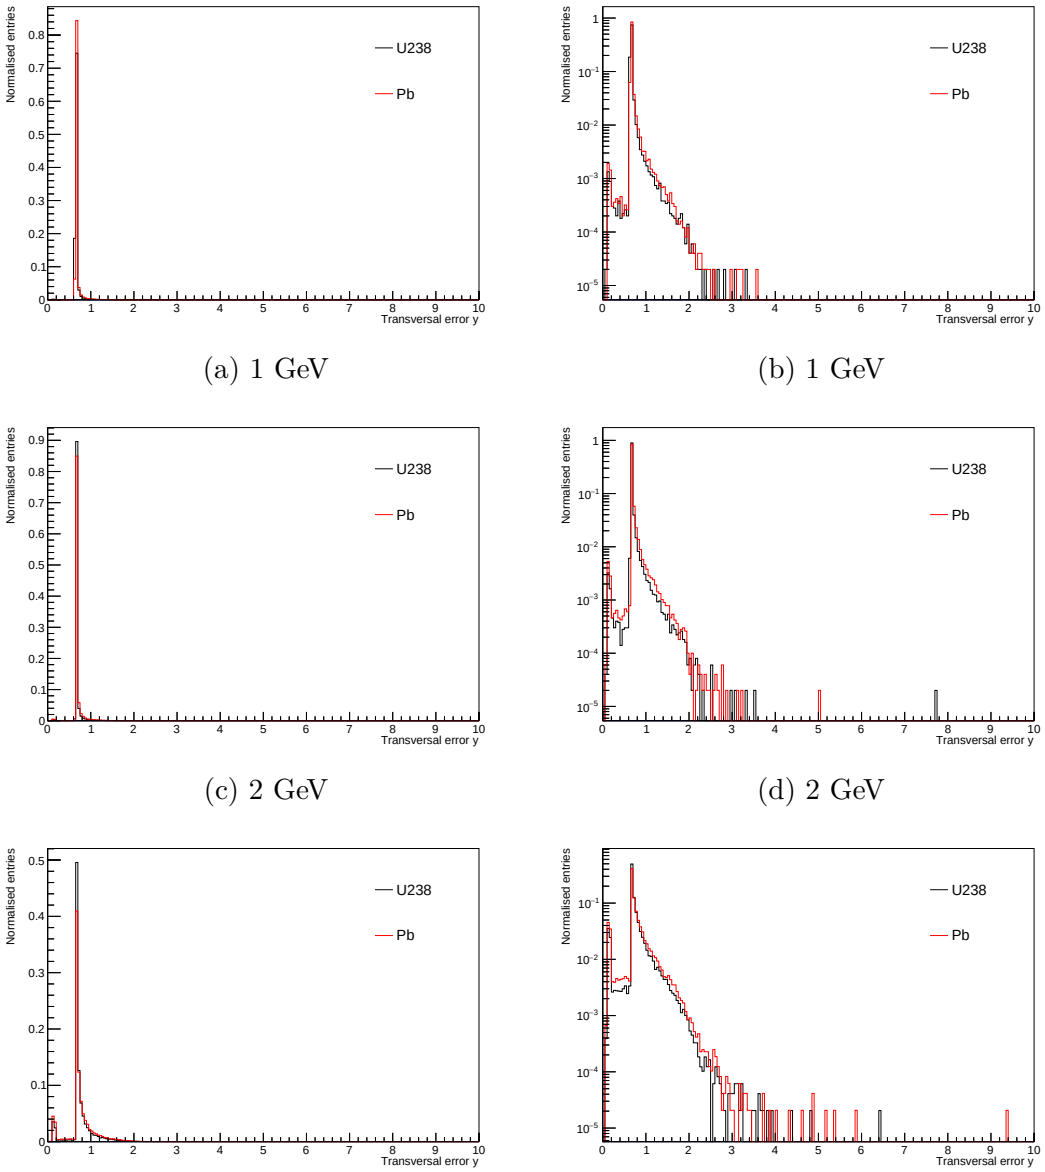

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.34: Transversal error in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

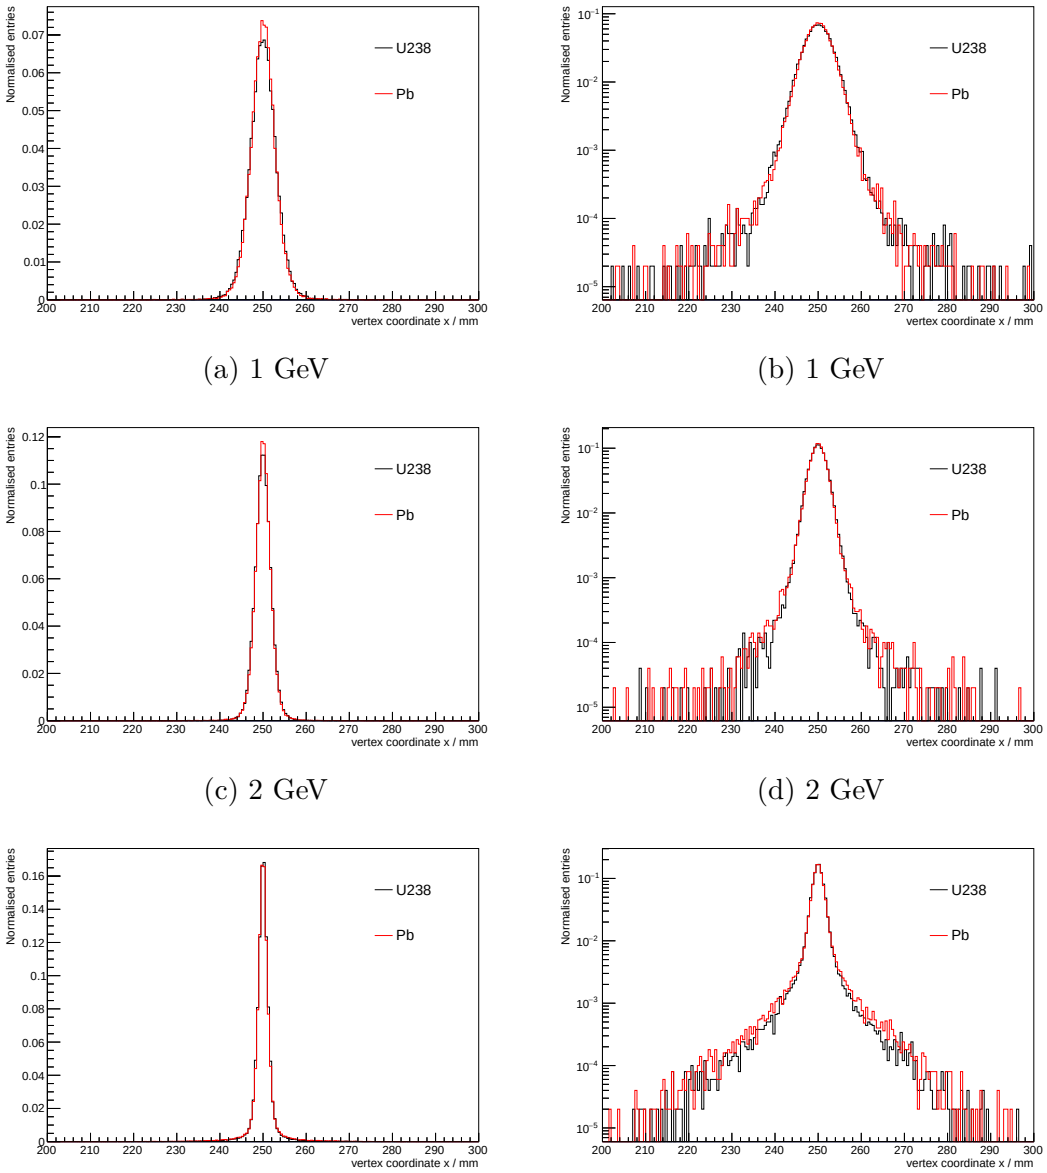

(a) 1 GeV

(b) 1 GeV

(c) 2 GeV

(d) 2 GeV

(e) 10 GeV

(f) 10 GeV

Figure A.35: Coordinate x of the reconstructed vertex for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

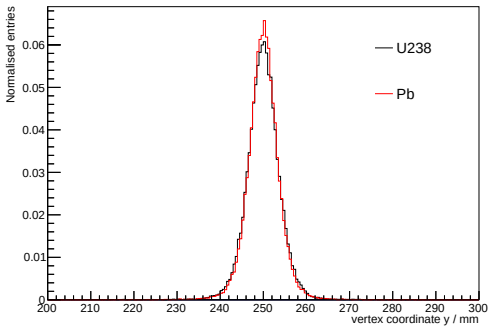

(a) 1 GeV

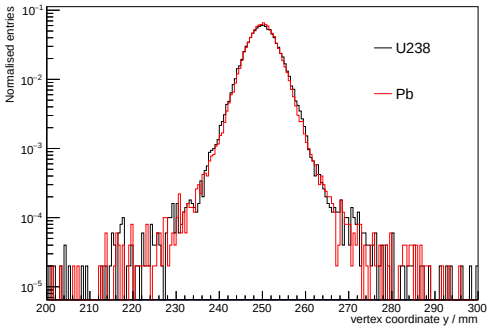

(b) 1 GeV

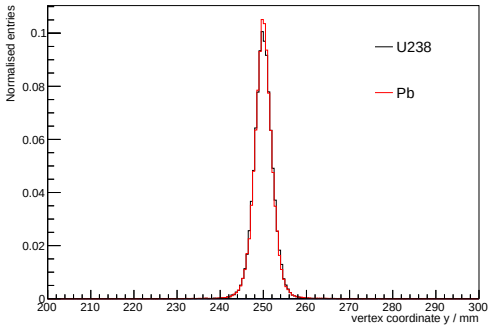

(c) 2 GeV

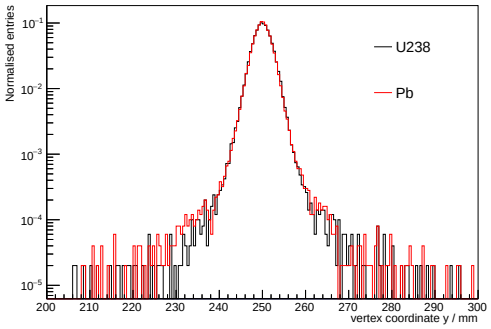

(d) 2 GeV

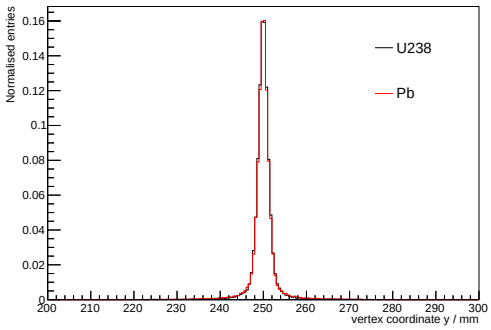

(e) 10 GeV

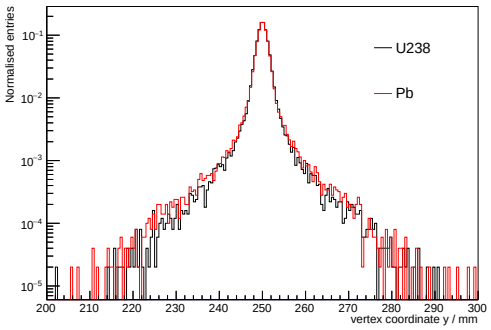

(f) 10 GeV

Figure A.36: Coordinate y of the reconstructed vertex for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

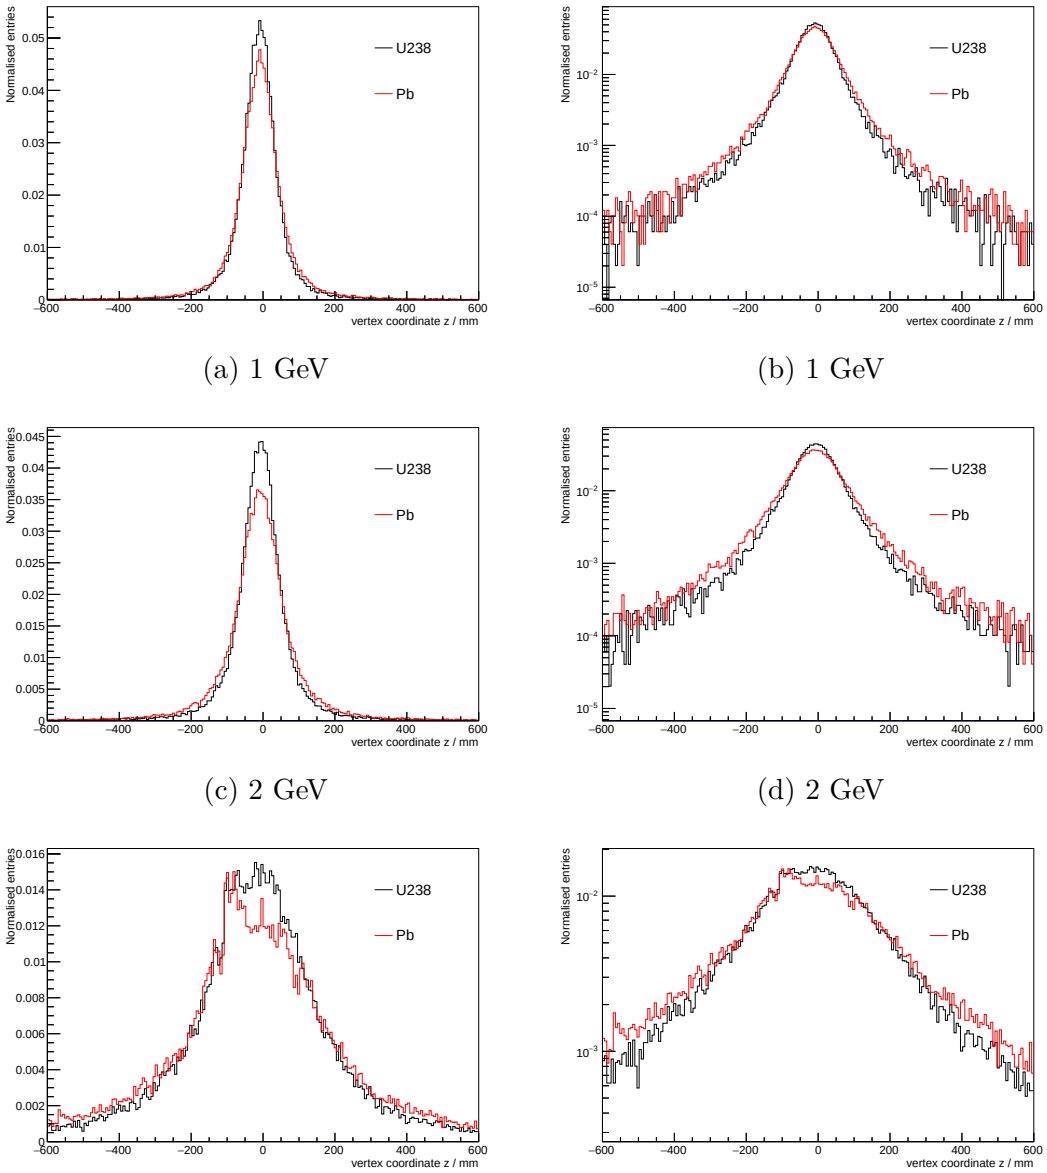

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.37: Coordinate z of the reconstructed vertex for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

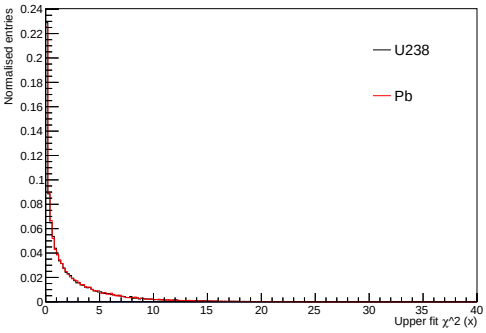

(a) 1 GeV

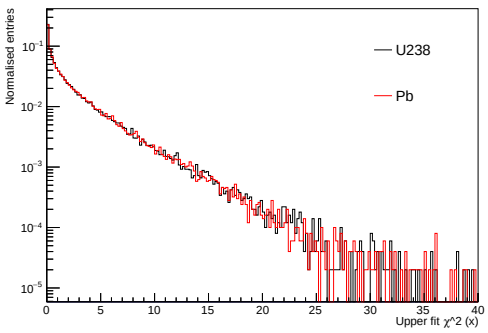

(b) 1 GeV

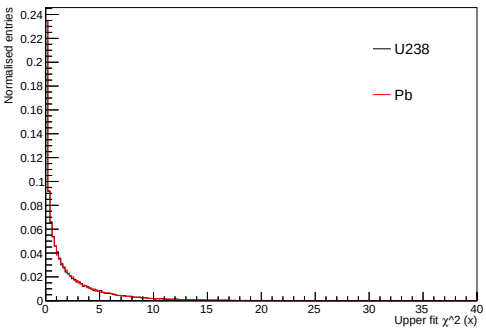

(c) 2 GeV

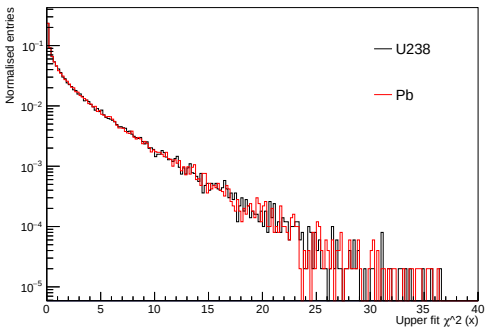

(d) 2 GeV

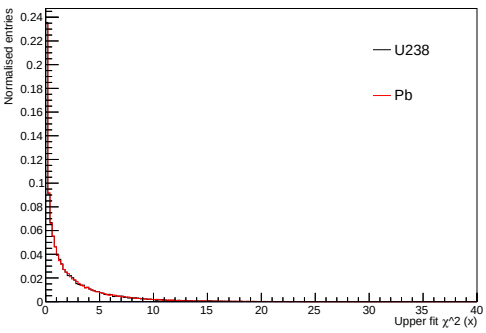

(e) 10 GeV

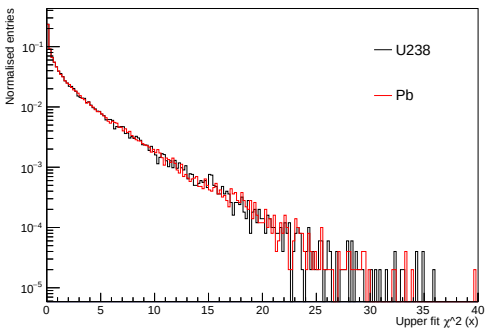

(f) 10 GeV

Figure A.38: χ^2 from fit of upper track in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

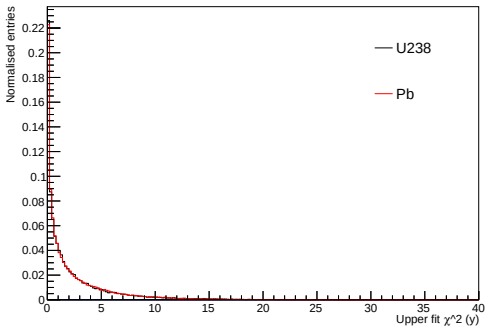

(a) 1 GeV

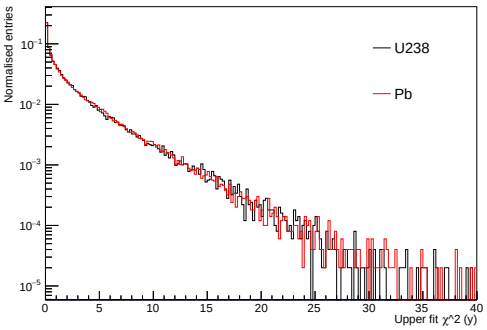

(b) 1 GeV

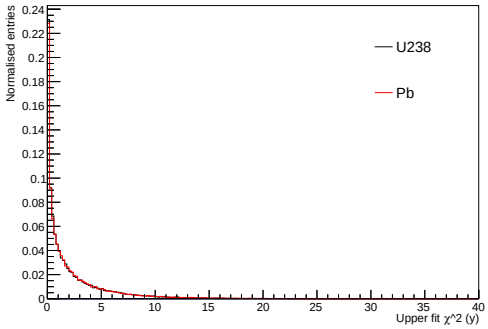

(c) 2 GeV

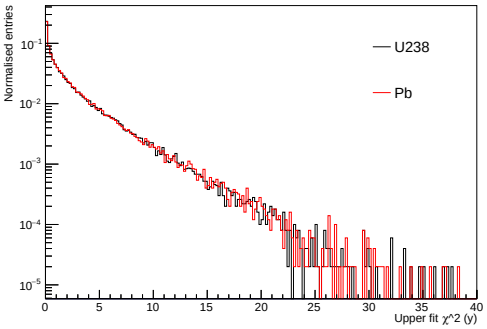

(d) 2 GeV

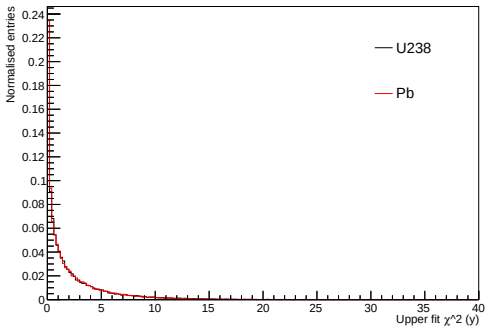

(e) 10 GeV

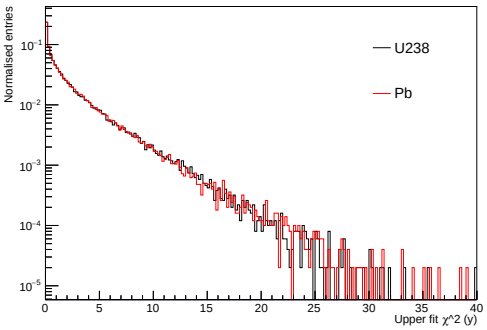

(f) 10 GeV

Figure A.39: χ^2 from fit of upper track in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

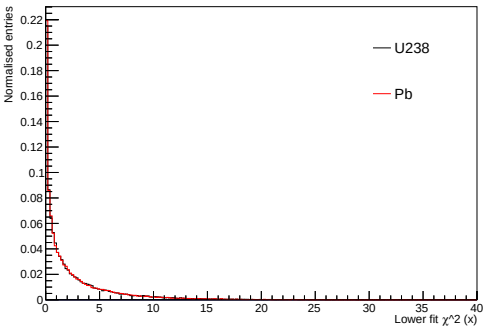

(a) 1 GeV

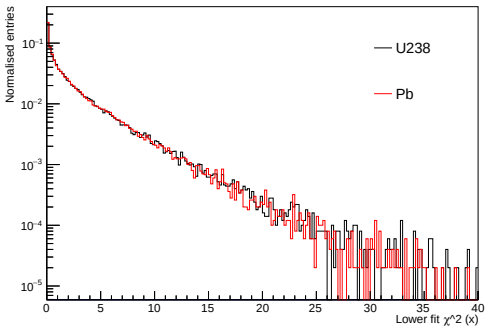

(b) 1 GeV

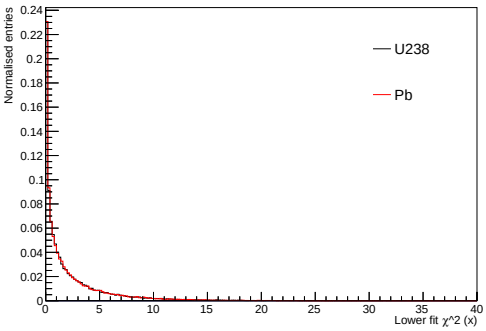

(c) 2 GeV

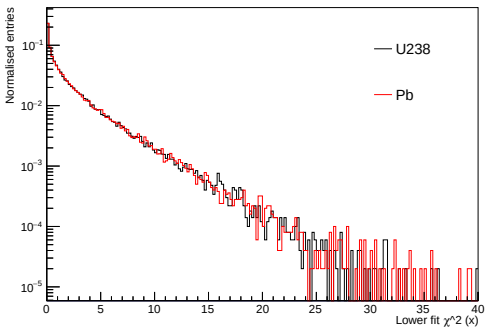

(d) 2 GeV

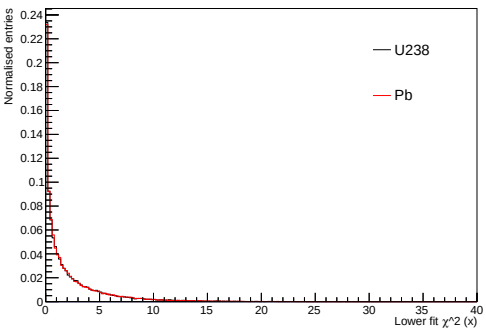

(e) 10 GeV

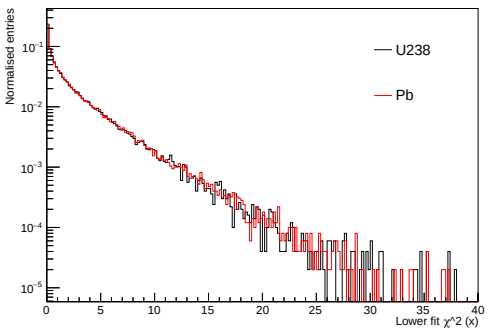

(f) 10 GeV

Figure A.40: χ^2 from fit of lower track in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

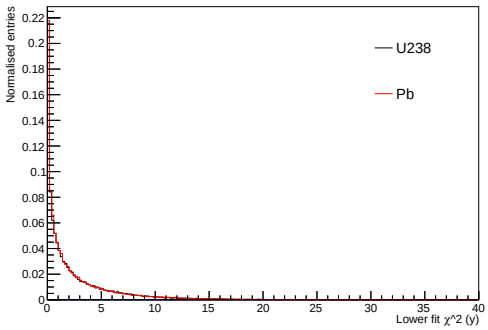

(a) 1 GeV

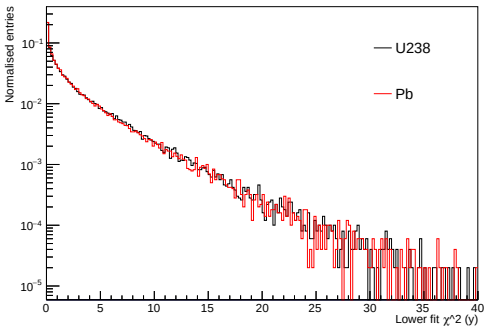

(b) 1 GeV

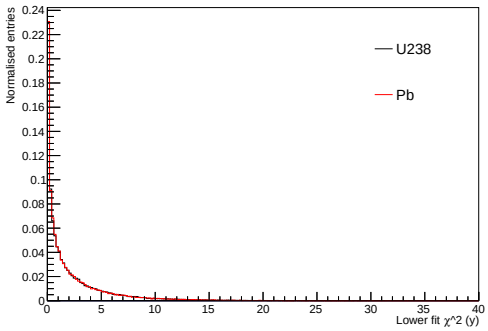

(c) 2 GeV

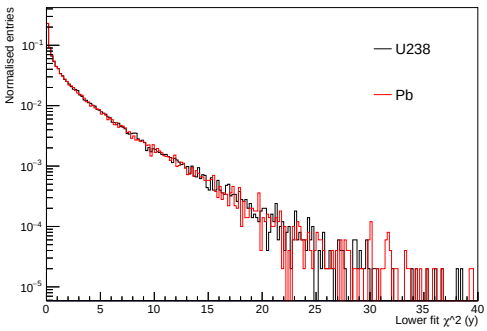

(d) 2 GeV

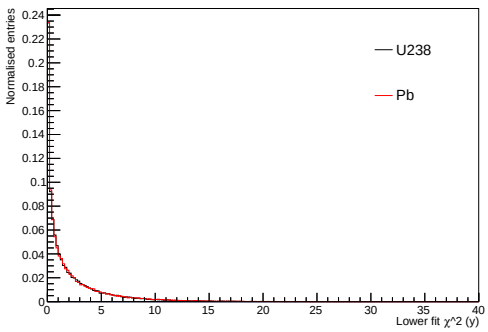

(e) 10 GeV

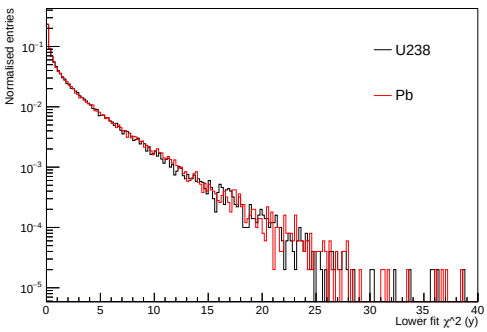

(f) 10 GeV

Figure A.41: χ^2 from fit of lower track in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

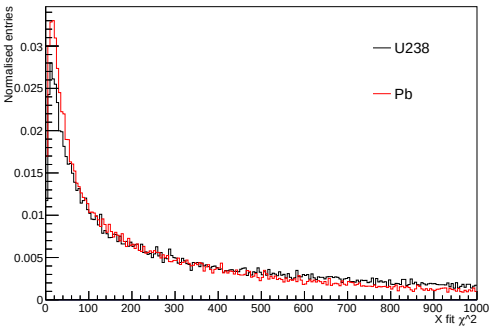

(a) 1 GeV

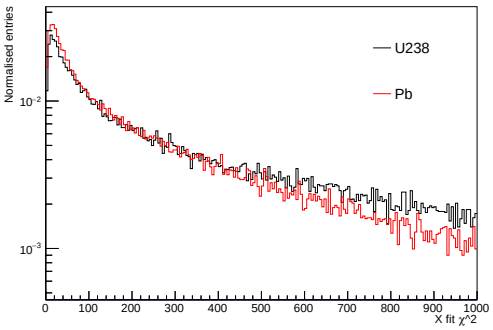

(b) 1 GeV

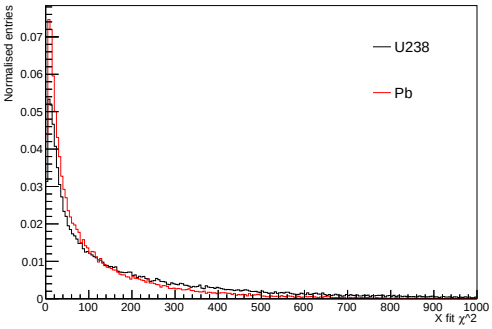

(c) 2 GeV

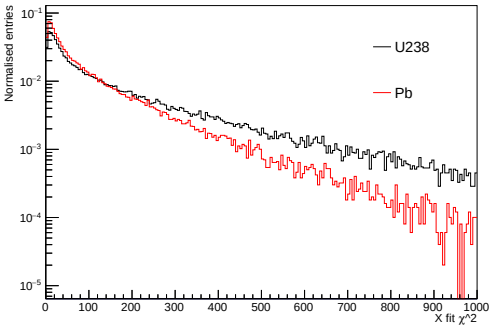

(d) 2 GeV

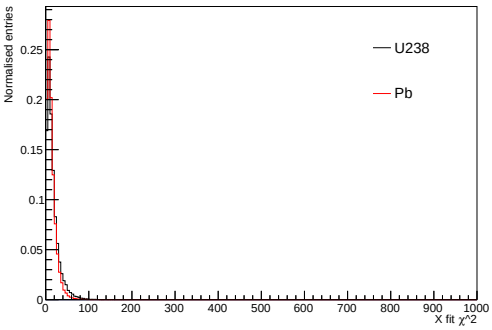

(e) 10 GeV

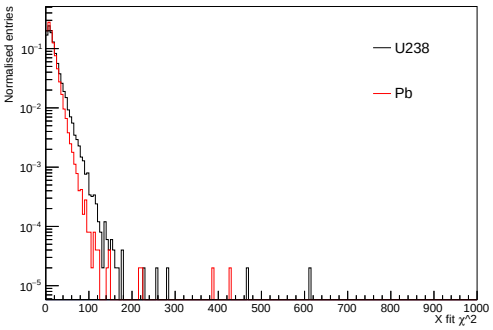

(f) 10 GeV

Figure A.42: Six points fit χ^2 in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

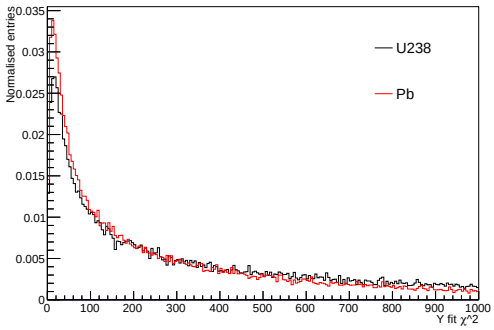

(a) 1 GeV

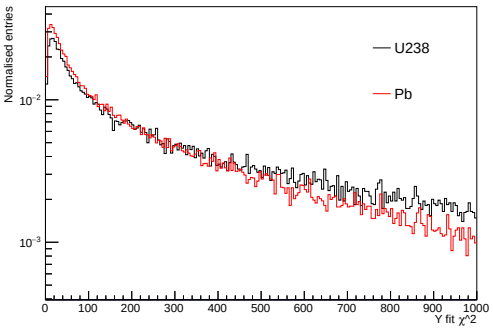

(b) 1 GeV

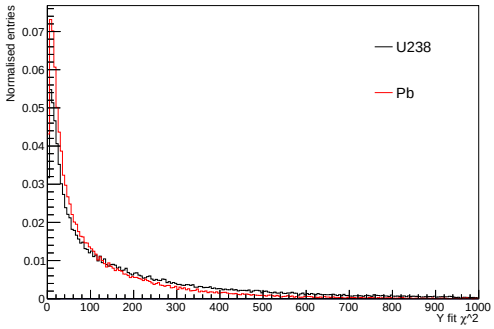

(c) 2 GeV

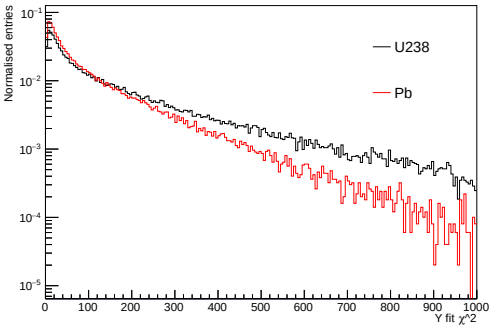

(d) 2 GeV

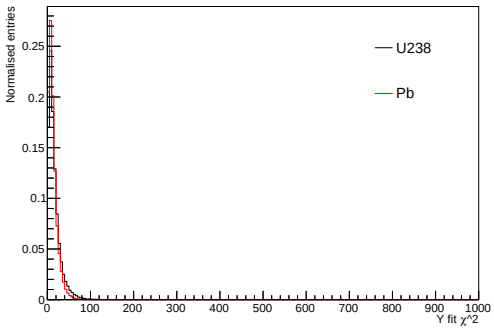

(e) 10 GeV

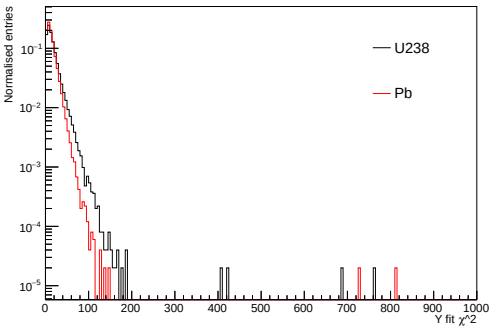

(f) 10 GeV

Figure A.43: Six points fit χ^2 in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

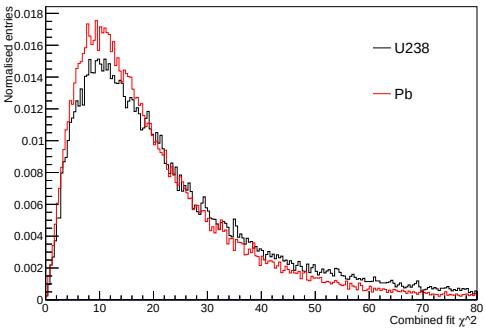

(a) 1 GeV

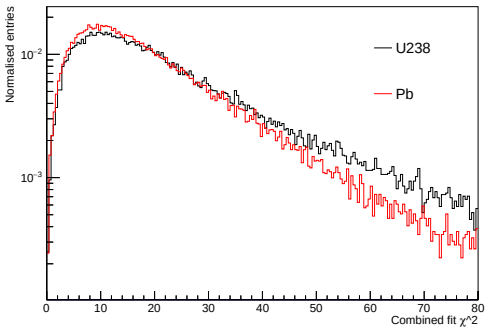

(b) 1 GeV

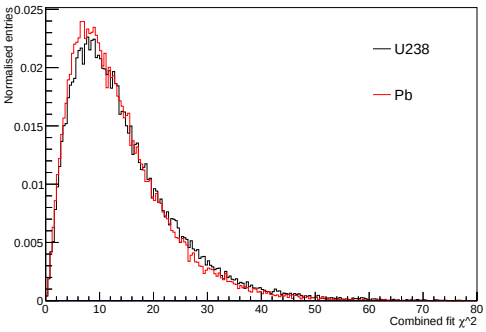

(c) 2 GeV

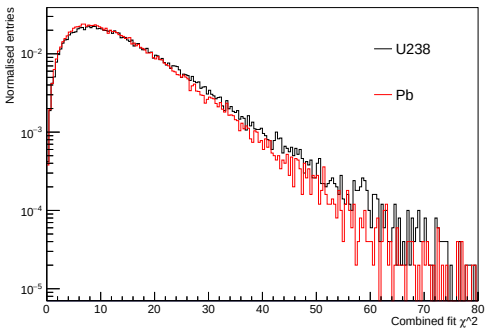

(d) 2 GeV

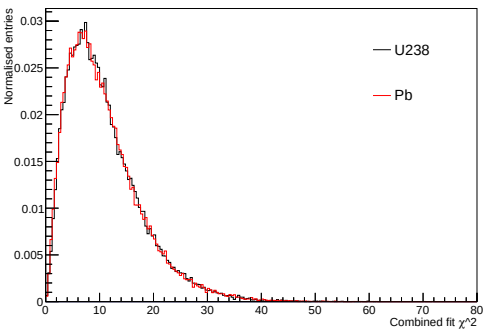

(e) 10 GeV

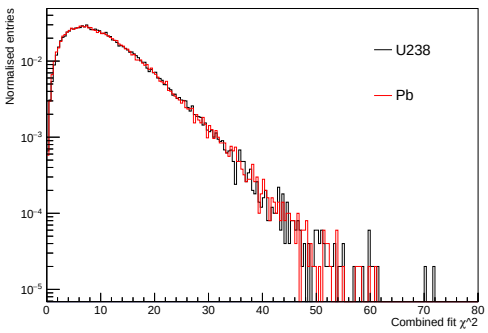

(f) 10 GeV

Figure A.44: Combined fit χ^2 for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

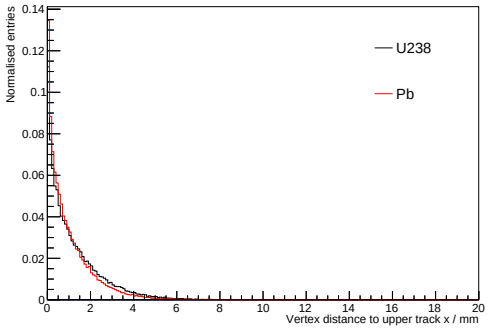

(a) 1 GeV

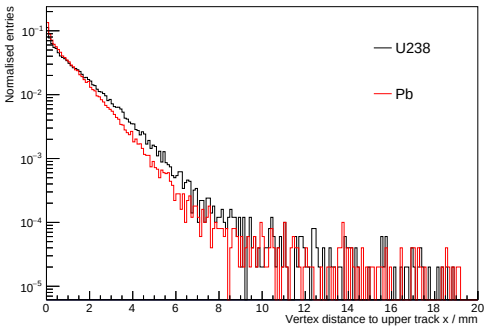

(b) 1 GeV

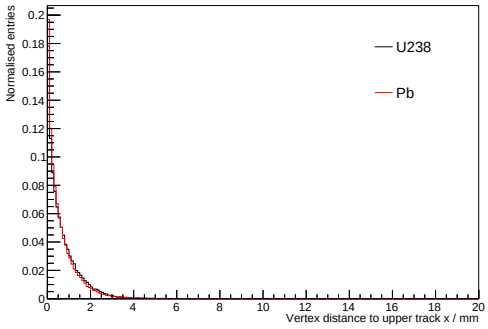

(c) 2 GeV

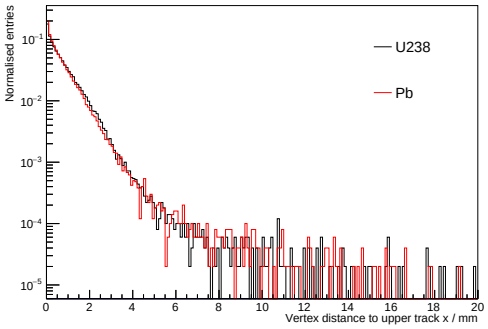

(d) 2 GeV

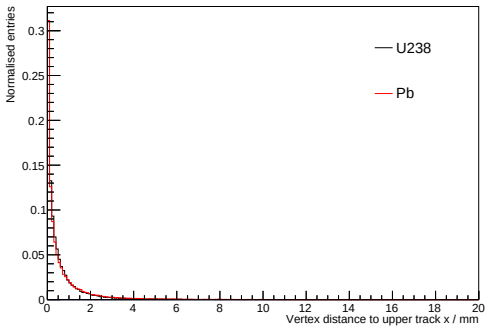

(e) 10 GeV

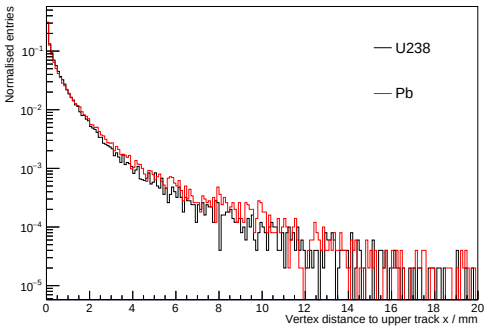

(f) 10 GeV

Figure A.45: Distance between the reconstructed vertex and upper track in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

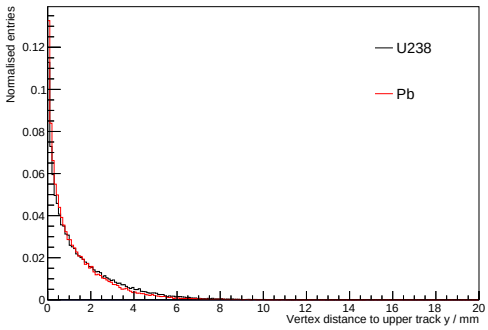

(a) 1 GeV

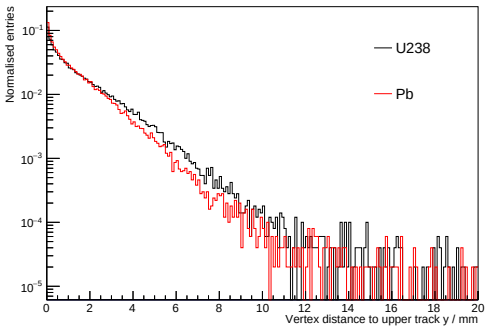

(b) 1 GeV

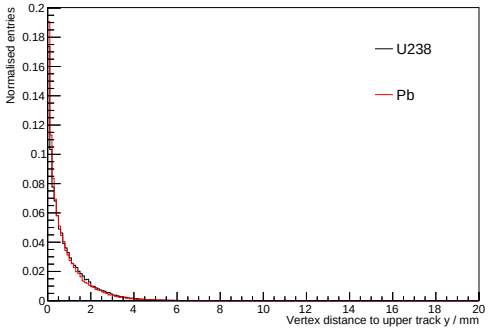

(c) 2 GeV

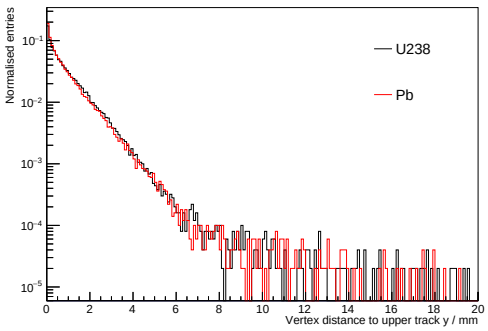

(d) 2 GeV

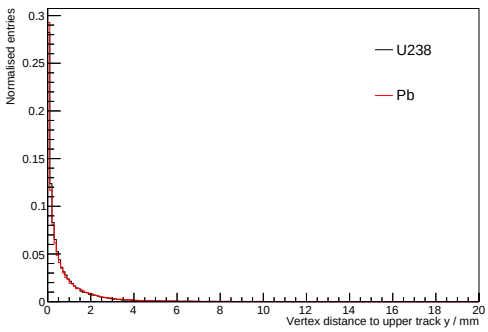

(e) 10 GeV

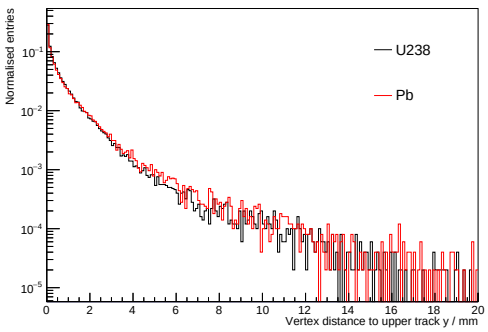

(f) 10 GeV

Figure A.46: Distance between the reconstructed vertex and upper track in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

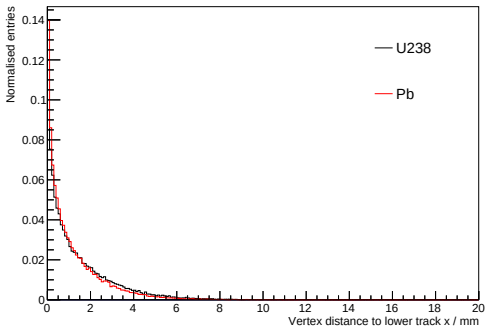

(a) 1 GeV

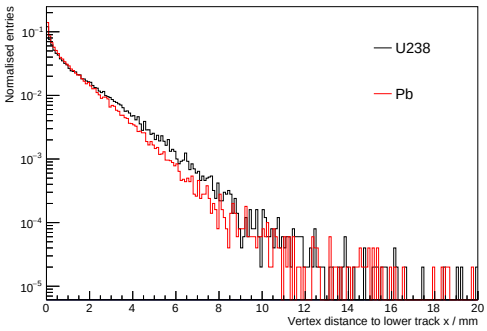

(b) 1 GeV

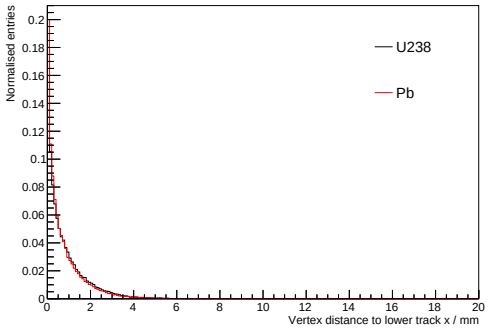

(c) 2 GeV

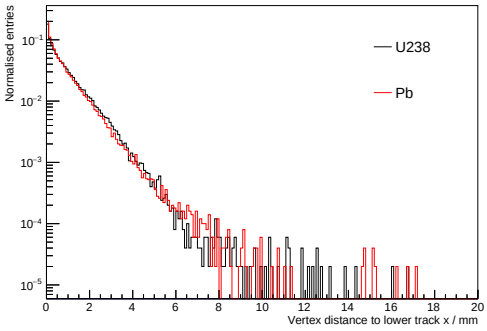

(d) 2 GeV

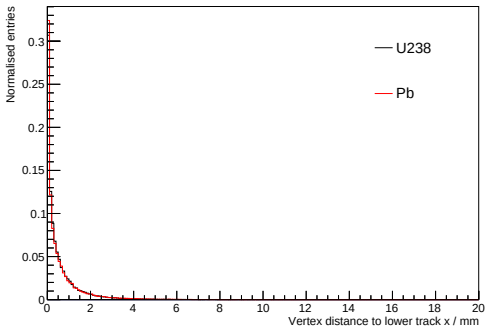

(e) 10 GeV

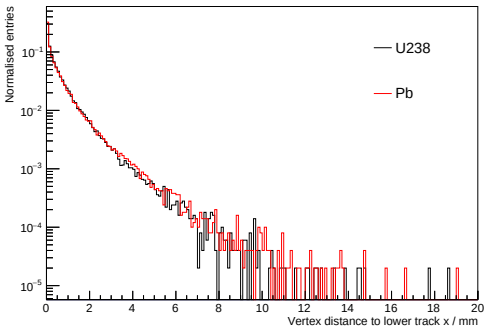

(f) 10 GeV

Figure A.47: Distance between the reconstructed vertex and lower track in x for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

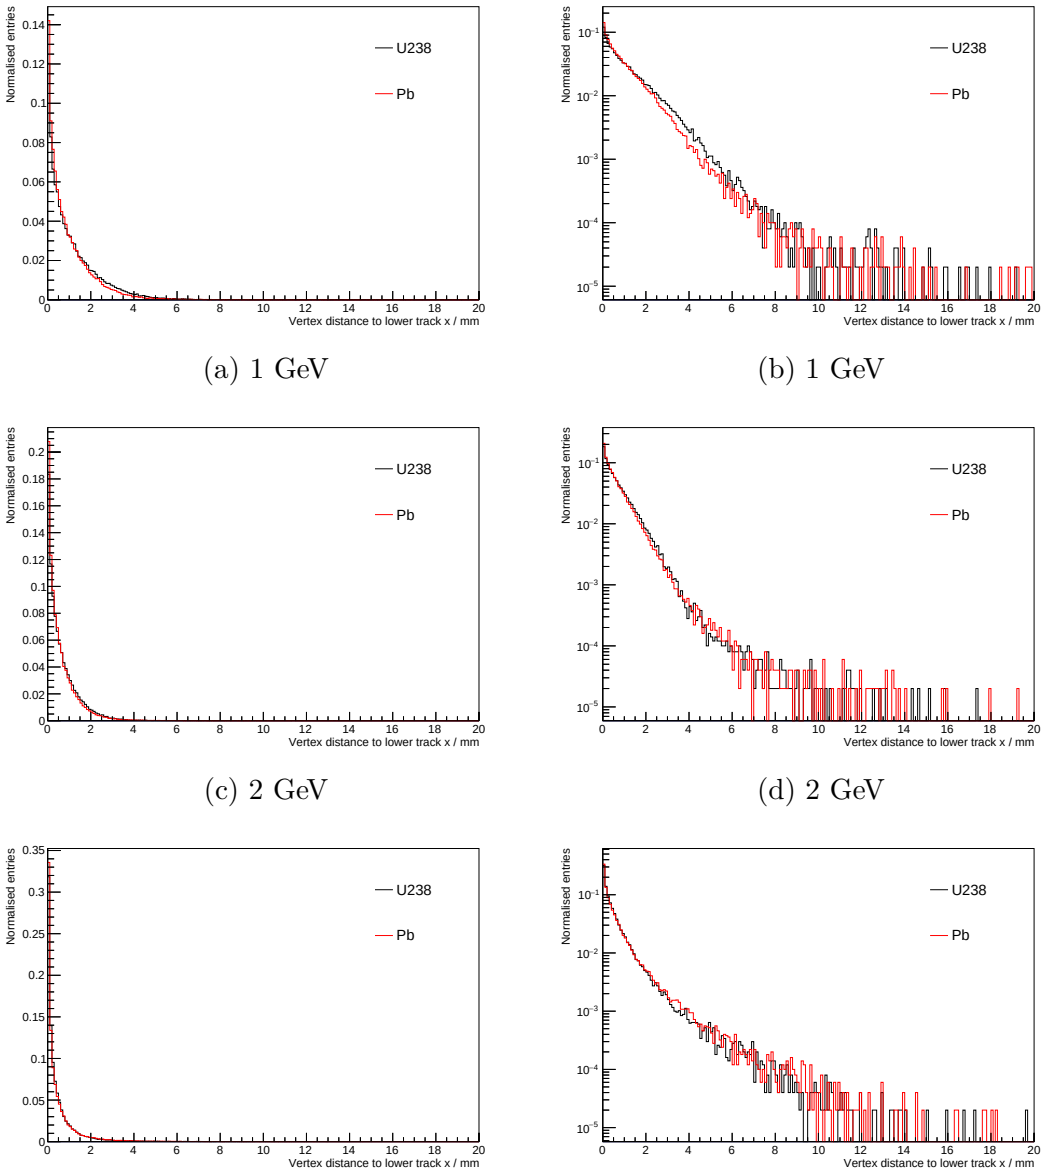

(a) 1 GeV (b) 1 GeV

(c) 2 GeV (d) 2 GeV

(e) 10 GeV (f) 10 GeV

Figure A.48: Distance between the reconstructed vertex and lower track in y for uranium and lead $10 \times 10 \times 10$ cm^3 blocks, for different muon energies, in linear (left) and log (right) scales. The histograms are normalised.

References

[1] IAEA. Strategy and methodology for radioactive waste characterization, IAEA-TECDOC-1537, March 2007.

[2] G. Jonkmans, V. N. P. Anghel, C. Jewett, and M. Thompson. Nuclear waste imaging and spent fuel verification by muon tomography. *Annals of Nuclear Energy*, 53:267–273, 2013.

[3] C. Patrignani et al. Review of Particle Physics. *Chinese Physics C*, 40(10):100001, 2016.

[4] W. R. Leo. *Techniques for nuclear and particle physics experiments: a how-to approach*. Springer Science & Business Media, 2012.

[5] V. L. Highland. Some practical remarks on multiple scattering. *Nuclear Instruments and Methods*, 129(2):497–499, 1975.

[6] G. R. Lynch and O. I. Dahl. Approximations to multiple Coulomb scattering. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 58(1):6–10, 1991.

[7] Particle Data Group. Atomic and nuclear properties of materials. http://pdg.lbl.gov/2017/AtomicNuclearProperties/.

[8] L. W. Alvarez et al. Search for Hidden Chambers in the Pyramids. *Science*, 167(3919):832–839, 1970.

[9] K. Morishima et al. Discovery of a big void in Khufu's Pyramid by observation of cosmic-ray muons. *Nature*, pages 1–18, 2017.

[10] N. Lesparre, D. Gibert, J. Marteau, J. C. Komorowski, F. Nicollin, and O. Coutant. Density muon radiography of La Soufriere of Guadeloupe volcano: comparison with geological, electrical resistivity and gravity data. *Geophysical Journal International*, 190(2):1008–1019, 2012.

[11] H. K. M. Tanaka, T. Uchida, M. Tanaka, H. Shinohara, and H. Taira. Development of a portable assembly-type cosmic-ray muon module for measuring the density structure of a column of magma. *Earth, planets and space*, 62(2):119, 2010.

[12] H. K. M. Tanaka et al. High resolution imaging in the inhomogeneous crust with cosmic-ray muon radiography: The density structure below the volcanic crater floor of Mt. Asama, Japan. *Earth and Planetary Science Letters*, 263(1-2):104–113, 2007.

[13] H. Miyadera et al. Imaging Fukushima Daiichi reactors with muons. *AIP Advances*, 3(5), 2013.

[14] TEPCO. Reactor imaging technology for fuel debris detection by cosmic ray muon. Measurement status report in Unit-1. http://www.tepco.co.jp/en/nu/fukushima-np/handouts/2015/images/handouts_150319_01-e.pdf.

[15] TEPCO. Locating fuel debris inside the unit 2 reactor using a muon measurement technology at Fukushima Daiichi nuclear power station. https://www4.tepco.co.jp/en/nu/fukushima-np/handouts/2016/images/handouts_160728_01-e.pdf.

[16] TEPCO. Locating fuel debris inside the unit 3 reactor using a muon measurement technology at Fukushima Daiichi nuclear power station. https://www4.tepco.co.jp/en/nu/fukushima-np/handouts/2017/images/handouts_170928_01-e.pdf.

[17] K. N. Borozdin, G. E. Hogan, C. Morris, W. C. Priedhorsky, A. Saunders, L. J. Schultz, and M. E. Teasdale. Surveillance: Radiographic imaging with cosmic-ray muons. *Nature*, 422(March):277, 2003.

[18] C. L. Morris et al. Tomographic Imaging with Cosmic Ray Muons. *Science & Global Security*, 16(1-2):37–53, 2008.

[19] M. Hohlmann, P. Ford, K. Gnanvo, J. Helsby, D. Pena, R. Hoch, and D. Mitra. GEANT4 simulation of a cosmic ray muon tomography system with micro-pattern gas detectors for the detection of high-Z materials. *IEEE Transactions on Nuclear Science*, 56(3):1356–1363, 2009.

[20] S. Pesente et al. First results on material identification and imaging with a large-volume muon tomography prototype. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 604(3):738–746, 2009.

[21] T. B. Blackwell and V. A. Kudryavtsev. Development of a 3D muon disappearance algorithm for muon scattering tomography. *Journal of Instrumentation*, 10(05):T05006–T05006, 2015.

[22] C. Thomay, J. J. Velthuis, P. Baesso, D. Cussans, P. A. W. Morris, C. Steer, J. Burns, S. Quillin, and M. Stapleton. A binned clustering algorithm to detect high-Z material using cosmic muons. *Journal of Instrumentation*, 8(10):P10013, 2013.

[23] C. Thomay. *Muon Scattering Tomography for Nuclear Security Applications*. PhD thesis, University of Bristol, 2015.

[24] Decision Sciences. MMPDS. https://decisionsciences.com/mmpds/.

[25] IAEA. *Classification of radioactive waste*. IAEA safety standards series No. GSG-1, 2009.

[26] H. Wimmer, J. Skrzyppek, and M. Köbl. CASTOR® and CONSTOR®: A well established system for the dry storage of spent fuel and high level waste. *VGB PowerTech*, 95(5):53–57, 2015.

[27] L. J. Schultz et al. Image reconstruction and material Z discrimination via cosmic ray muon radiography. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 519(3):687–694, 2004.

[28] L. J. Schultz et al. Statistical reconstruction for cosmic ray muon tomography. *IEEE Transactions on Image Processing*, 16(8):1985–1993, 2007.

[29] A. Clarkson et al. Characterising encapsulated nuclear waste using cosmic-ray muon tomography. *Journal of Instrumentation*, 10(03):P03020, 2015.

[30] X. D. Wang et al. The study of cosmic ray tomography using multiple scattering of muons for imaging of high-z materials. *arXiv preprint arXiv:1608.01160*, 2016.

[31] V. Anghel, G. Jonkmans, C. Jewett, and M. Thompson. Detecting high atomic number materials with cosmic ray muon tomography, 2015. US Patent 9,035,236.

[32] C. Thomay, J. J. Velthuis, T. Poffley, P. Baesso, D. Cussans, and L. Frazão. Passive 3D Imaging of Nuclear Waste Containers with Muon Scattering Tomography. *Journal of Instrumentation*, 03008(2):4–6, 2016.

[33] B. C. Rastin. An accurate measurement of the sea-level muon spectrum within the range 4 to 3000 GeV/c. *Journal of Physics G: Nuclear Physics*, 10(11):1609–1628, 1984.

[34] A. Benaglia, S. Gundacker, P. Lecoq, M. T. Lucchini, A. Para, K. Pauwels, and E. Auffray. Detection of high energy muons with sub-20 ps timing resolution using L(Y)SO crystals and SiPM readout. *Nuclear Instruments and*

Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 830:30–35, 2016.

[35] E. C. Zeballos, I. Crotty, D. Hatzifotiadou, J. L. Valverde, S. Neupane, M. C. S. Williams, and A. Zichichi. A new type of resistive plate chamber: the multigap RPC. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 374(1):132–135, 1996.

[36] E. C. Zeballos et al. A very large multigap resistive plate chamber. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 434(2-3):362–372, 1999.

[37] A. Akindinov et al. Multigap resistive plate chamber as a time-of-flight detector. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 456(1-2):16–22, 2000.

[38] The OPERA Collaboration. Momentum measurement by the Multiple Coulomb Scattering method in the OPERA lead emulsion target. *New Journal of Physics*, 013026:13, 2011.

[39] Geant4 Collaboration. Geant4. https://cern.ch/geant4.

[40] S. Agostinelli et al. GEANT4 - A simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[41] J. Allison et al. Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, 2006.

[42] C. Hagmann, D. Lange, and D. Wright. Cosmic-ray shower generator (CRY) for Monte Carlo transport codes. In *Nuclear Science Symposium Conference Record, 2007. NSS'07. IEEE*, volume 2, pages 1143–1146. IEEE, 2007.

[43] Los Alamos National Laboratory. A General Monte Carlo N-Particle (MCNP) Transport Code. https://mcnp.lanl.gov/.

[44] L. S. Waters et al. The MCNPX Monte Carlo radiation transport code. In *AIP conference Proceedings*, volume 896, pages 81–90. AIP, 2007.

[45] P. Baesso, D. Cussans, C. Thomay, J. J. Velthuis, J. Burns, C. Steer, and S. Quillin. A high resolution resistive plate chamber tracking system developed for cosmic ray muon tomography. *Journal of Instrumentation*, 8(08):P08006, 2013.

[46] P. Baesso, D. Cussans, C. Thomay, and J. Velthuis. Toward a RPC-based muon tomography system for cargo containers. *Journal of Instrumentation*, 9(10):C10041, 2014.

[47] Geant4 Collaboration. *Geant4 User's Guide for Application Developers*, 6 December 2013.

[48] R. Santonico and R. Cardarelli. Development of resistive plate counters. *Nuclear Instruments and Methods*, 187(2-3):377–380, 1981.

[49] C. Thomay, J. J. Velthuis, P. Baesso, D. Cussans, C. Steer, J. Burns, S. Quillin, and M. Stapleton. A novel Markov random field-based clustering algorithm to detect high-Z objects with cosmic rays. *IEEE Transactions on Nuclear Science*, 62(4):1837–1848, 2015.

[50] R. Brun and F. Rademakers. ROOT - an object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389(1):81–86, 1997.

[51] F. James and M. Roos. MINUIT - A system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343–367, 1975.

[52] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936.

[53] J. W. Sammon. An optimal discriminant plane. *IEEE Transactions on Computers*, 100(9):826–829, 1970.

[54] M. Welling. Fisher linear discriminant analysis. *Note from the Department of Computer Science, University of Toronto*, 2005.

[55] S. S. Wilks. *Mathematical statistics*. John Wiley & Sons Inc., New York, 1962.

[56] L. Frazão, J. Velthuis, C. Thomay, and C. Steer. Discrimination of high-Z materials in concrete-filled containers using muon scattering tomography. *Journal of Instrumentation*, 11(07):P07020–P07020, 2016.

[57] T. Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.

[58] X. Robin et al. pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC bioinformatics*, 12(1):77, 2011.

[59] E. R. DeLong, D. M. DeLong, and D. L. Clarke-Pearson. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, 44(3):837–845, 1988.