

Muography: Detector optimization and Cultural heritage applications

Doctoral dissertation presented by

Maxime Lagrange

in fulfilment of the requirements for the degree of Doctor in Sciences.

Thesis support committee

Dr. Andrea Giammanco (Supervisor)	UCLouvain, Belgium
Dr. Eduardo Cortina Gil (President)	UCLouvain, Belgium
Dr. Christophe Deleare (Secretary)	UCLouvain, Belgium
Dr. Pietro Vischia (Jury member)	Universidad de Oviedo, Spain
Dr. Tommaso Dorigo (Jury member)	Luleå University, Sweden

July, 2025

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Dr. Andrea Giammanco, for his unwavering support throughout my PhD journey. His guidance has gone far beyond the professional sphere, extending also to the personal level. He placed his trust in me, which was profoundly motivating, and encouraged me to step outside my comfort zone by pushing me to present at numerous conferences. His mentorship has been a constant source of inspiration and a model I aspire to follow. I sincerely hope that one day I can offer my own students the same level of guidance and encouragement.

I am also deeply grateful to all the members of my defense committee for taking the time, despite their busy schedules, to read my work and provide valuable suggestions that undoubtedly improved this thesis.

My sincere thanks go to the entire TomOpt collaboration for welcoming me into their pioneering detector design optimization efforts. I have learned immensely from their expertise, which has greatly shaped the scientist I am today. I would particularly like to thank Giles Strong and Tommaso Dorigo for their patience and guidance during my early, inexperienced days as a PhD student, as well as Pietro Vischia, who generously invited me on several occasions to his institute in Oviedo to advance our optimization studies. Our collaboration has since grown into a genuine friendship, something I deeply value.

I am forever indebted to the CP3 Muography team, who became a second family to me. Samip, Marwa, Ishan, and Eduardo welcomed me warmly when I first joined CP3, making me feel at home from the very beginning. Vishal, Zahraa, and Sumaira, although you joined later, working with you and attending conferences together has been a real pleasure. Beyond the scientific collaborations and the publications we produced, I cherish the fact that we became more than colleagues, we became close friends.

I would also like to express my gratitude to all those with whom I had the chance to collaborate. Florian, I learned an enormous amount from you in machine learning, and I will always appreciate the work we accomplished together. From the DLR team, Sarah, Angel, Felix, and Bojra, thank you for hosting me during a three-week stay that proved extremely fruitful and enriching. Pablo, thank you for inviting me to Santander to present my work and for sharing your research activities with me. To all of you, I can say that our collaborations grew into meaningful friendships.

Finally, I would like to thank all members of CP3 and IRMP. Many of you became true friends, and you were an essential part of what made my PhD adventure so enjoyable. I cannot list everyone by name for fear of leaving someone out, but please know that you all mean a great deal to me. Special thanks go to Julien, Guillaume, and Wael for their enduring friendship.

Needless to say, I wish to thank my family for their endless love and support. You are the reason I have come this far in life, and everything I have been able to accomplish is thanks to you. To Clémentine, the love of my life, I owe more than words can express.

Abstract

Particle physics explores the Universe's building blocks, forces, and evolution. Though often seen as abstract, its technology, such as accelerators and detector, have profoundly shaped medicine, industry, and communication. The Large Hadron Collider at CERN exemplifies the scale and sophistication of this instrumentation, while many spin-offs, from synchrotron radiation to imaging tools like CT, PET, MRI, and proton therapy, have entered daily life. Muography is a more recent technique that leverages naturally occurring cosmic-ray muons instead of artificial radiation sources to create radiographic images. Since its first success in 1955, it has advanced through better detectors, simulations, and reconstruction methods. By exploiting how muons are absorbed or deflected when traversing matter, muography enables non-invasive imaging of large and dense structures inaccessible to traditional methods. This capability has unlocked applications in volcanology, archaeology, civil engineering, nuclear waste monitoring, and security screening, and is increasingly explored as a complementary tool to conventional imaging. This thesis advances muography toward greater reliability and broader use. It addresses detector optimization, systematic benchmarking of reconstruction algorithms, improved methods for inferring muon momentum, and the extension of muography to new domains. The overarching aim is to reduce the gap with established modalities such as X-ray or CT imaging, which benefit from standardized practices and commercial deployment. By contributing a more unified framework, this work aspires to help establish muography as a mature, widely adopted imaging technique for science and industry.

Associated Publications:

- [12] Sarah Barnes, Anzori Georgadze, Andrea Giammanco, Madis Kiisk, Vitaly Kudryavtsev, Maxime Lagrange, Olin Lyod Pinto, “Cosmic-Ray Tomography for Border Security”, MDPI Instruments, vol. 7, pp. 13, 2023.
- [111] Giles C Strong et al. “TomOpt: differential optimisation for task- and constraint-aware design of particle detectors in the context of muon tomography”, Machine Learning: Science and Technology 5.3, p. 035002, 2024.
- [47] Andrea Giammanco, Marwa Al Moussawi, Matthieu Boone, Tim De Kock, Judy De Roy, Sam Huysmans, Vishal Kumar, Maxime Lagrange, Michael Tytgat, “Toward using cosmic rays to image cultural heritage objects”, iScience, vol. 28, pp. 2589, 2025.
- [39] Florian Bury, Maxime Lagrange, "Scattering-Based Machine Learning Algorithms for Momentum Estimation in Muon Tomography", MDPI Particle, vol. 8, 2025.

Thesis Jury

Prof. Eduardo Cortina Gil	President	Université Catholique de Louvain
Prof. Pietro Vischia	Jury member	Universidad de Oviedo and ICTEA
Prof. Tommaso Dorigo	Jury member	Luleå University of Technology
Prof. Christophe Deleare	Secretary	Université Catholique de Louvain

Contents

Introduction	1
1 Primer on Muography	5
1.1 Cosmic Rays	5
1.1.1 Primary Cosmic Rays	6
1.1.2 Primaries Interaction with Atmosphere	8
1.1.3 Muon production	11
1.2 Muon Interaction with Matter	14
1.2.1 Energy Loss	14
1.2.2 Multiple Coulomb Scattering	17
1.3 Muon Imaging	20
1.3.1 Muon Transmission Tomography	20
1.3.2 Muon Scattering Tomography	23
1.4 Applications of muon tomography	25
1.4.1 Cargo Scanning	25
1.4.2 Industrial Applications	27
1.5 Detection Systems	30
1.5.1 Scintillators	32

1.5.2	Gaseous detectors	33
1.6	Challenges	33
1.6.1	Muon Momentum estimation	33
1.6.2	MST reconstruction algorithms	34
2	Monte Carlo simulation for Muon Scattering Tomography	37
2.1	Cosmic muon Simulation	38
2.1.1	Differential muon flux	39
2.1.2	Cosmic Muon Flux Simulation with the CRY software	41
2.2	Passage of Particles Through Matter	42
2.2.1	GEANT4 software	42
2.2.2	Multiple Coulomb Scattering	43
2.2.3	Energy Loss	44
2.3	Detector Response	45
2.3.1	Muon detection in GEANT4	46
2.3.2	Spatial Resolution	47
2.3.3	Detection efficiency	49
3	Inverse problem in Muon Scattering Tomography	51
3.1	Muon Scattering Tomography Reconstruction Algorithms . . .	52
3.1.1	Scattering density	53
3.1.2	Point Of Closest Approach algorithm	55
3.1.3	Binned Clustered Algorithm	61

3.1.4	Angle Statistics Reconstruction	63
3.2	Characterization of MST imaging	69
3.2.1	Metrics	69
3.2.2	Lumiosity	75
3.2.3	Gaussian smoothing	80
3.2.4	Noise in MST reconstruction	83
3.3	Deep-learning approach	94
3.3.1	Data generation	95
4	Muon momentum estimation	99
4.1	Muon Momentum in Muon Tomography	100
4.1.1	Muon Momentum in Muon Transmission Tomography	100
4.1.2	Muon Momentum in Muon Scattering Tomography . .	102
4.1.3	Muon momentum measurement	103
4.2	Muon momentum inference via machine learning methods . .	105
4.2.1	Deviations as features	106
4.2.2	Deviations as a Sequence	109
4.2.3	Deviations as a geometry	110
4.2.4	Model Comparison	112
4.3	Real-Life Experiment Considerations	113
4.3.1	Hit Detection Efficiency	113
4.3.2	Hit Spatial Resolution	115

5	Cultural heritage study	117
5.1	Simulation study	118
5.2	Small scale objects: Muon Scattering Tomography	121
5.3	Medium scale objects: Scattering versus Transmission Tomography	124
5.3.1	Transmission tomography	124
5.3.2	Scattering tomography	126
5.4	Large scale objects: Muon Transmission Tomography	127
6	TomOpt	129
6.1	Software description	130
6.1.1	Cosmic muon generation	131
6.1.2	Passive volumes and multiple scattering	132
6.1.3	Detector modeling and hit recording	134
6.1.4	Trajectory fitting and Point-of-Closest-Approach finding	137
6.1.5	Volume inference	138
6.1.6	Loss functions and optimisation	140
6.1.7	Update cycle	142
6.1.8	Current limitations	144
6.2	Benchmark: Fill-level estimation at metal refinery	146
6.2.1	Description	146
6.2.2	Detector and volume	147
6.2.3	Inference method and loss function	148

6.2.4	Optimisation process and results	153
6.2.5	Comparison to human baselines	158
Conclusion and prospects		161
A Appendix		165
A.1	Coordinate system	165
A.2	Simulation scenario: Elementary cubes	166
A.3	Muograph software	168
A.4	Deviations' basis	169
A.5	Network hyperparameters	170
A.6	TomOpt	172

Introduction

Particle physics, alongside astrophysics and cosmology, is a foundational field in our quest to understand the Universe, its fundamental components, the forces governing them, and its evolution over the past 14 billion years. Although often perceived as abstract or removed from everyday life, particle physics has profoundly shaped modern technology, particularly in how we communicate, diagnose illness, and manufacture goods. At the heart of this scientific endeavor is the development of highly sophisticated instrumentation. To observe and study elementary particles and their interactions, researchers rely on enormous and intricate detector systems, some as large as buildings. A prime example is the Large Hadron Collider (LHC) at CERN, which explores conditions similar to those just after the Big Bang. It consists of a 27-kilometer ring of superconducting magnets, operating at temperatures just a few degrees above absolute zero and producing magnetic fields billions of times stronger than Earth's.

While accelerators and detectors are central to fundamental research, their technological spin-offs have had far-reaching impacts across society. Today, only a small fraction of accelerators are used for basic research, and the vast majority serve medical, industrial, and environmental purposes. Synchrotron radiation, for example, is harnessed in numerous applications, from medical imaging and cancer therapy to materials testing and semiconductor development.

Similarly, the detector technologies pioneered for particle physics experiments have enabled major advances in other scientific domains. Medical imaging has benefited enormously, giving rise to tools such as computed tomography (CT), positron emission tomography (PET), magnetic resonance imaging (MRI), and proton therapy, all of which draw on principles and components developed in high-energy physics.

Beyond these applications lies a more recent and rapidly evolving technique: muography. Unlike conventional imaging methods that use artificial particle sources such as X-rays or protons, muography leverages naturally occurring cosmic-ray muons to generate radiographic images. Since its first successful demonstration in 1955, the field has grown considerably, aided by advances in detector technology, simulation, and reconstruction methods. Muography techniques, which rely on the absorption or deflection of muons as they pass through matter, enable the non-invasive imaging of large and dense structures that are otherwise inaccessible. This has opened up new avenues of interdisciplinary research in areas as diverse as volcanology, archaeology, civil engineering, nuclear waste monitoring, and security screening.

This thesis explores several key aspects of muography with the aim of advancing it toward greater robustness and broader applicability. It addresses the optimization of detector configurations, the characterization and benchmarking of reconstruction techniques, the development of improved methods for inferring muon momentum, and the exploration of new application domains. The overarching motivation is to help close the gap between muography and more established imaging modalities, such as X-ray or CT imaging, which benefit from standardized methodologies and widespread commercial deployment. By contributing to a more unified and rigorous framework, this work aspires to move muography closer to being a reliable and mature imaging technique, capable of broader scientific and industrial adoption.

Chapter 1 introduces the foundational physics and methodologies of muography. It provides an overview of its diverse applications across disciplines and surveys popular detector technologies. The chapter concludes with a discussion of the current limitations and challenges in the field, which set the stage for the investigations and contributions presented in the remainder of the thesis.

Chapter 2 addresses the central role of Monte Carlo simulations in muon tomography. As in other areas of particle physics, simulations are indispensable for designing detectors, developing reconstruction algorithms, and interpreting data. This chapter details the simulation tools used throughout this thesis to model detector behavior and evaluate reconstruction performance.

Chapter 3 reformulates muon scattering tomography as an inverse problem and introduces a quantitative framework to evaluate reconstruction performance. It

reviews common reconstruction algorithms and proposes a systematic characterization of the noise sources that affect 3D image quality. This analysis reveals previously undocumented limitations and motivates the application of machine learning as a means of improving image fidelity through denoising and restoration.

Chapter 4 tackles one of the main limitations in muon tomography: the lack of momentum measurements. Unlike detectors in high-energy physics experiments, which are often equipped with strong magnetic fields or high timing resolution, typical MST setups do not provide momentum information. Yet, even coarse momentum estimates can significantly enhance both transmission and scattering tomography, improving background rejection and signal purity, and enhancing the reconstruction algorithms' performance. This chapter reports on an exploration of how muon momentum can be inferred from scattering through known materials and demonstrates how machine learning models can further boost estimation accuracy.

Chapter 5 presents a new application of muon tomography to cultural heritage. While muography has primarily been used to image large archaeological structures, this chapter extends its applicability to smaller objects, on the order of a few meters, which pose different challenges in terms of muon flux, absorption, and scattering. It demonstrates that both scattering and transmission tomography can be adapted to these constraints and implemented within the logistical realities of museum environments, where conventional radiation imaging techniques are often impractical.

Chapter 6 introduces TOMOPT, a software framework developed within the MODE collaboration [37] to support end-to-end detector design optimization. The tool combines a fully differentiable simulation pipeline with gradient-based optimization techniques powered by automatic differentiation. This chapter outlines how TOMOPT enables task-aware design by directly optimizing performance metrics tied to reconstruction goals. A benchmark case study demonstrates its effectiveness, showing that the optimized configurations can match or surpass manually designed detectors for MST tasks.

Chapter 1

Primer on Muography

This chapter provides the essential background for understanding muon tomography, covering both scattering and transmission techniques. It begins with an overview of cosmic muons and their production on Earth through interactions between primary cosmic rays and atmospheric atoms in Section 1.1. Next, the physical processes governing muon propagation through matter are explored in Section 1.2. The principles of muon imaging, based on scattering and absorption, are then introduced in Section 1.3, followed by a review of detector technologies in Section 1.5. Applications across various sectors are discussed in Section 1.4, and the chapter concludes with an examination of the challenges facing muography in Section 1.6.

1.1 Cosmic Rays

The discovery of cosmic rays goes back to the early 20th century, when scientists were trying to get a better understanding of the conductivity of air. The question of why a charged conductor placed in an isolated container slowly loses its charge was first investigated by Charles Augustin de Coulomb in 1785 and remained unsolved for more than a hundred years. Since air is made of neutrally charged gases, how can it show conductive properties? Later, the discovery of natural radioactivity and its ionization radiations α -, β -rays eventually answered the inquiry. The many studies of ionization and absorption processes permitted the characterization of the various sources of radioactivity: radioactive substances in the Earth, radioactive emanations in the air, and radiation pollution around and within the electroscope. The study of the influence of natural radioactivity on the conductivity of the air led scientists to measure the ionization rate at increasing altitudes, expecting it to decrease as

the distance from the natural radiation at the surface of the Earth increases. In 1912, after improving the detection performance of the electrometers used to measure radiation, Victor F. Hess performed several measurement campaigns at various altitudes and established that the ionization rate increases with altitude. The same year, he summarizes his study as follows [56]: "The results of these observations seem to be explained by the assumption that a radiation of high penetration power hits our atmosphere from the top, which causes also in their lower layers a fraction of the observed ionization in the closed detectors." Hess's discovery opened the door to many other discoveries in particle and nuclear physics, particularly the discovery of cosmic muons by Anderson between 1935 and 1937. The designation *cosmic muons* can be misleading, as muons themselves do not originate from the cosmos but are produced on Earth by the interaction of cosmic radiations – primary cosmic rays – with the atmosphere. Section 1.1.1 provides a characterization of those primary cosmic rays, and Section 1.1.2 expands on their interaction with the atmosphere. Finally, Section 1.1.3 characterizes the resulting muon flux within the atmosphere and at sea level.

1.1.1 Primary Cosmic Rays

Since the early discovery of cosmic rays by Hess in 1912, the understanding and characterization of cosmic radiation have taken a key place in the landscape of experimental and theoretical physics and led to many studies of their intensity in different locations around the world. The evidence that this radiation had penetration power even higher than the known gamma radiation, which was known to have a higher penetration power than alpha and beta radiations (the other two known sources of radiation at the time), suggested that they too were extremely high-energy gamma radiations. More generally, it was natural to ask the question of the nature of these radiations as beta rays i.e., electrons, alpha rays i.e. ${}^4\text{He}$ nuclei emissions, were known at the moment through radioactivity studies. The first experiment that came to mind was to check whether or not cosmic radiation would feel the effect of the Earth's magnetic field, by measuring the cosmic ray flux at various geomagnetic latitudes. Although preliminary results had been previously published by Compton and his collaborators, his 1933 survey summarizes the main results obtained from eight different expeditions and shows that variations in the cosmic radiation

flux amount up to 14% at sea level and even 33% at an altitude of 4360m [33]. This last result suggests that the latitude effect on cosmic rays is due to phenomena that occur at great altitudes and that a significant fraction of primary particles that reach the top of the atmosphere are charged particles. Later in the second half of the 1930s, the scientific community moved towards a more and more convincing hypothesis that primary radiations are made of protons rather than electrons or positrons, although the underlying mechanisms of particle production were only discovered years later. The characterization of the primary cosmic ray is now close to perfect, as the development of detection techniques has made tremendous progress since the 1930s, and while it still occupies an eminent role today, problems have shifted to other fronts or more specific details. Latest outer-space measurements from AMS [24], DAMPE [26], CALET [25], and PAMELA [30] experiments offer a clear picture of the direction, energy, and composition of the primary cosmic rays before reaching our atmosphere.

Beyond the usual distinction between "primary" cosmic rays – cosmic particles that reach the top of the Earth's atmosphere – and "secondary" cosmic rays – particles produced by the interaction of the primary cosmic rays with the atmosphere's molecules –, primary cosmic rays can be further subdivided into two "primary" and "secondary" groups. The primary group includes stable particles and nuclei with a sufficiently large mean lifetime emitted "directly" by the astrophysical sources, while the secondary group includes particles produced by interactions of the primaries with the interstellar medium. Protons, light nuclei, and electrons synthesized in stars can be considered the true primary component, while heavier nuclei that do not derive from stellar nucleosynthesis are considered secondaries. Unifying measurements from a large number of experiments produces an inclusive energy spectrum of primary particles, as shown in Figure 1.1. The flux's intensity is $10^4 \text{ m}^{-2}\text{s}^{-1}\text{sr}^{-1}$ but quickly falls with kinetic energy E , following a $E^{-\gamma}$ dependence. The striking features of the energy differential flux are the knee where at a few PeV γ changes from 2.7 to 3.0, the second knee at 100 PeV where γ changes from 3.0 to 3.3 and the ankle at a few EeV where γ changes from 3.3 to 2.5. Protons constitute approximately 90% of the charged primary radiation, the rest being light nuclei and nuclei of intermediate mass, although the exact ratio between the various components depends on the range of energy considered, as well as electrons and positrons. Gamma rays can be separated into those coming from astrophysical sources and a diffuse flux from galactic and extragalactic origins.




Figure 1.1: Energy spectrum and composition of primary cosmic rays. Shown are measurements of the intensity of neutral and charged cosmic rays, multiplied by kinetic energy squared. The data for charged cosmic rays have been extracted from the Cosmic Ray Database [69]. Figure taken from [76].

The distribution of charged cosmic rays' arrival direction is mostly isotropic, but anisotropies below and around the knee due to the distribution of sources and properties of the Galactic magnetic fields have been observed.

1.1.2 Primaries Interaction with Atmosphere

The interaction of the primary cosmic rays with atoms in the atmosphere produces a high multiplicity of secondary particles through successive and entangled strong and electromagnetic processes. This intricate cascade of particles' decay and creation is called an Extensive Air Shower and is represented in Figure 1.2, where the 3 components- hadronic, muonic, and electromagnetic- are




Figure 1.2: Interaction of primary cosmic rays with molecules of the atmosphere leading to the creation of an extensive air shower. The main interactions and decay processes giving rise to muonic, hadronic, and electromagnetic components are highlighted with dashed arrows.

highlighted. The initial interaction of the primaries with atmospheric molecules and nuclei produces secondary hadrons that share most of the primaries' kinetic energy. As those high-energy hadrons -mesons and baryons- propagate through the atmosphere, they further decay, and the charged component gives rise to an electromagnetic shower by radiating bremsstrahlung photons or by electron-positron pair production. This constitutes the electromagnetic component of the shower. Finally, the secondary mesons and hadrons decay into muons and neutrinos, making up the muonic component.

As the shower develops longitudinally, its components progressively lose energy as it gets redistributed to more and more secondary particles, while another part of the energy gets lost through the propagation medium. The properties of the secondary cosmic rays, namely energy and direction distribution, as well as their particle composition, are governed by a complex system of equations involving the particle's decay probability and interaction cross-section, and the properties of the atmospheric medium. The analytical particle transport equations can be found in [108]. Atmospheric depth summarizes the vertical profile of the atmospheric density and can be approximated by the following equation under an isothermal approximation:

$$X = \int_h^\infty \rho(h) dh \simeq X_0 e^{-h/h_0}, \quad (1.1)$$

where X_0 is the atmospheric depth at sea level (1030 g cm^{-2}) and h_0 scale height of the atmosphere. For practical use, all coefficients should be evaluated more precisely to interface the three atmospheric layers. The shower size, typically expressed in terms of the number of charged particles, primarily depends on the energy and inclination of the primary particle, as well as the altitude at which the first interaction occurs. Its evolution follows a characteristic pattern: starting from low values near the initial interaction point, it increases to a maximum at a specific atmospheric depth before gradually decreasing toward the ground (approximately 1000 g cm^{-2} in atmospheric depth), as shown in Figure 1.3. The longitudinal profile of the shower is governed by the underlying physical processes driving its development, including both hadronic and electromagnetic interactions. These processes are influenced by the relevant cross sections, the multiplicity of particles produced in each interaction, and the relative proportions of different particle species. Additionally, kinematical constraints dictate the production mechanisms of individual particles. Profiles of this nature can be effectively described using a semi-empirical function, as proposed by Gaisser and Hillas [45]:

$$N(X) = N_{\max} \left(\frac{X}{X_{\max}} \right)^{X_{\max}/\lambda} e^{(X_{\max}-X)/\lambda}, \quad (1.2)$$

where X is the atmospheric depth in g cm^{-2} , $N_{\max}$ the number of particles corresponding to the maximum $X_{\max}$, and λ a characteristic depth (usually assumed to be 70 g cm^{-2}).




Figure 1.3: Average number of $e^+ + e^-$, $\mu^+ + \mu^-$ and hadrons as a function of the atmospheric depth in g cm^{-2} along the shower axis. The profiles are obtained from AIRES simulation [103] of proton, iron, and photon-induced extensive air showers of energy $E = 1 \text{ EeV}$ and zenith angle $\theta = 45^\circ$ with respect to the vertical. Figure taken from [76].

1.1.3 Muon production

Muons are elementary particles of the lepton family and are naively referred to as the *heavy* cousin of the electron. While they share the same quantum numbers: an elementary charge $-e$ and a spin of $-1/2$, the muon's mass is $105.6 \text{ MeV}/c^2$, making it about 200 times heavier than the electron. Unlike electrons that are stable particles, muons have a lifetime of $\tau \sim 2.2 \mu\text{s}$. For relativistic muons of kinetic energy 4 GeV , the observed lifetime get dilated by the Lorentz factor $\gamma = 1/\sqrt{1 - \beta^2} \approx 20$, with $\beta = v/c$, leading to a decay length of around 24 km . Cosmic muons originate from the decay process of charged pions and kaons produced during the development of the extensive air shower:

$$\begin{aligned}\pi^\pm &\rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \\ K^\pm &\rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu),\end{aligned}$$

with respective branching ratio of 100% and 64%. The key features of cosmic rays' muonic component are the dependence of intensity on the altitude, the energy and angular distribution, and charge states- the ratio between positively and negatively charged muons. If we exclude low-energy particles and focus on an energy threshold of the order of GeV, muons are the most abundant charged particles at sea level. Most muons are produced at high altitudes in the atmosphere, around 10-15 km, and progressively lose their kinetic energy through ionization as they travel. Assuming a specific energy loss of $2 \text{ MeV g}^{-1}\text{cm}^2$ and an atmospheric thickness of 1000 g cm^{-2} , such energy loss amounts to 2 GeV. Taking into account various processes (production, energy loss, decay), the muon flux can be parametrized as a function of its kinetic energy and direction, e.g., using the Gaisser model [44]:

$$\begin{aligned}\frac{dN}{dE d\Omega dt dS} &= 0.14 E^{-2.7} \left(\frac{1}{1 + \frac{1.1E \cos \theta}{115 \text{ GeV}}} - \frac{0.054}{1 + \frac{1.1E \cos \theta}{850 \text{ GeV}}} \right) \\ &\times \text{cm}^{-2}\text{s}^{-1}\text{sr}^{-1}\text{GeV}^{-1}.\end{aligned}\quad (1.3)$$

Various other parameterizations have been proposed in recent years, utilising data accumulated by many experiments. As shown in Figure 1.4, the energy distribution of cosmic muons from all directions at sea level peaks around 4 GeV.

Measurements of the muon flux as a function of the zenith angle made it possible to establish that the intensity can be represented by a relationship of the type:

$$I(\theta) = I(\theta = 0) \cos^n(\theta), \quad (1.4)$$




Figure 1.4: Distribution of atmospheric muons' energy at sea level obtained through a horizontal surface with the CRY cosmic-ray particle shower library [51].

with $I(\theta = 0)$ the intensity of vertical muons (i.e., with projected zenith angle $\theta = 0$) and an exponent n dependent on the momentum range considered, with an average value of 1.85 [94]. It should be noted that at first order, the muon flux is symmetrical along the azimuthal angle φ . Many muography experiments based on the absorption of cosmic muons through a volume of interest utilize the free sky muon flux to infer the object's density through the measurement of the transmission ratio. In such a context, having an accurate model of the cosmic muon flux is of paramount importance. Monte Carlo tools, such as CRY [51], EcoMug [82], propose models of cosmic ray primaries and secondaries based on experimental and simulated data and are a precious tool for such experiments.

Despite the many efforts to model the cosmic muon flux, it is subject to fluctuations that remain difficult to properly assess. The development of an extensive air shower involves numerous processes that are influenced by macroscopic atmospheric properties, particularly temperature and pressure. Atmospheric pressure experiences both periodic and aperiodic variations: periodic changes include daily fluctuations (day/night cycle) and seasonal cycles, while aperiodic changes are primarily due to local weather conditions. A comprehensive study of the temporal variations in cosmic ray flux can be found in [38]. On a daily timescale, atmospheric pressure is affected by tidal forces and the thermal expansion of air due to temperature increases. These variations typically have amplitudes on the order of a few mbar. Over a year, additional small

variations in pressure arise from seasonal changes and the heating or cooling of the ground. These variations are usually of low amplitude and highly dependent on geographic location. In contrast, aperiodic variations resulting from local weather phenomena can exhibit much larger amplitudes—up to 10–20 mbar—and can sometimes obscure the underlying periodic variations. For the muonic component of cosmic rays, the corresponding variation in flux is approximately 0.1–0.2 % mbar. The flux of cosmic rays is also influenced by the Earth’s magnetic field, independently of atmospheric conditions. The Earth’s magnetic field, which extends up to several thousand km from the Earth’s surface, acts as a deflecting shield for many primary charged particles, significantly altering their trajectories. Its intensity ranges from 60 to 30 μT , and it can be approximated by a dipole inclined relative to the Earth’s rotational axis. This inclination results in variations of the magnetic field across different geographic coordinates. Although the field is weaker at large distances, its wide spatial extent still results in substantial deflection of charged particles. Variations in magnetospheric conditions can lead to sudden increases in the cosmic muon flux by a few percent, as observed in the GRAPES-3 experiment [70].

1.2 Muon Interaction with Matter

As muon tomography aims at inferring material properties using muons, the knowledge of the physics governing their interaction with matter comes as a fundamental requirement. Muons, being charged leptons, are subjected to both weak and electromagnetic interactions, with the latter driving muons’ behavior when propagating through matter. On a macroscopic level, we are mostly interested in two electromagnetic processes: Multiple Coulomb Scattering and Energy Loss. This section describes those two physical processes and highlights the relation between the muons’ and the medium’s properties.

1.2.1 Energy Loss

As a charged particle propagates through matter, it progressively loses its kinetic energy via electromagnetic interactions with atoms of the medium. As the incident particle undergoes millions of interactions on a microscopic level,




Figure 1.5: Mass stopping power for positive muons in copper as a function of momentum $p = \beta\gamma/Mc$. Vertical bands indicate the boundaries between regimes. Given the average cosmic muons' momentum of 4 GeV/c, the relevant regime for muography is the minimum ionisation. Figure is taken from [76].

the total amount of energy loss at a macroscopic level has a statistical nature. The energy loss probability distribution follows a Landau-Vavilov distribution, and the mean energy loss rate is defined as the *mass stopping power* and described by the Bethe equation [76]:

$$-E \left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \quad (1.5)$$

where x is the opacity unit in g cm^{-2} , K a proportionality coefficient, z the charged number of the incident particle, Z the atomic number of the material, A the atomic mass number of the material, m_e the electron mass, I the mean excitation energy, $W_{\max}$ the maximum energy transfer to an electron in a single collision, and $\delta(\beta\gamma)$ the density effect correction. Equation 1.5 is valid in the momentum region $0.1 < \beta\gamma < 1000$ with an accuracy of a few percent. Figure 1.5 shows how the mass-stopping power fluctuates as a function of the incident particle's momentum.

In a regime close to the minimum ionization where radiative effects can be neglected, typically below a few hundred GeV for muons, the mean stopping power can be described by:

$$-\frac{dE}{dx} = a(E) + b(E) \cdot E, \quad (1.6)$$

where $a(E)$ is the electronic stopping power due to atoms' excitation and ionization, and $b(E)$ is the energy-scaled contribution from radiative processes—bremsstrahlung, pair production, and photo-nuclear interactions. Tables of $a(E)$ and $b(E)$ are provided by the Particle Data Group for pure elements [83]. While the estimation of the mass stopping power is rather straightforward for pure materials – materials made of a single atomic element – it gets more complicated for compound materials. In a first approximation, it can be done relying on Bragg's additivity law, which states that the mass stopping power of a mixture $E\langle\frac{dE}{dx}\rangle$ can be computed as a linear combination of the mass stopping power of its constituents:

$$E\left\langle\frac{dE}{dx}\right\rangle = \sum_j w_j \frac{dE}{dx}|_j, \quad (1.7)$$

where $dE/dx|_j$ is the mean rate of energy loss (in MeV g cm^{-2}) in the j -th element and w_j the weigh fraction of the j -th element. Equation 1.5 can be inserted in Equation 1.7 to find expressions for $E\langle Z/A \rangle$, $E\langle I \rangle$ and $E\langle \delta \rangle$; e.g. $E\langle Z/A \rangle = \sum_j n_j Z_j / n_j A_j$. However, computing $E\langle I \rangle$ this way is an underestimate, as compound electrons are more tightly bound than in the free elements. It is then recommended to use the tables given in Reference [34]. It should be noted that Equation 1.5 gives the *mass stopping power* in $\text{MeV g}^{-1} \text{cm}^2$. The amount of energy loss per unit of length is the *linear stopping power* $\rho dE/dx$ in MeV cm^{-1} , and is the product of the mass stopping and the density of the medium ρ in g cm^{-3} . As such, muons are not only sensitive to the type of material they traverse – A and Z terms in Equation 1.5 – but also the material's density.




Figure 1.6: Quantities used to describe the multiple Coulomb Scattering: the projected scattering angle θ_{plane} , the projected displacement y_{plane} , and the geometrical path length x .

1.2.2 Multiple Coulomb Scattering

When a charged particle passes near an atom, it is deflected from its original trajectory due to the Coulomb interaction arising from the strong electric field of the atomic nucleus. In the classical approximation, and neglecting strong inelastic interactions, the cross-section for single scattering is given by Rutherford's Law: $\frac{d\sigma}{d\Omega} \propto 1/\sin^4(\theta/2)$. As the particle traverses a medium, it experiences not just a single deflection but countless individual scatterings off the nuclei of atoms. The cumulative effect of these interactions results in a net multiple scattering angle and an associated displacement. Figure 1.6 illustrates the key quantities used to characterize multiple Coulomb scattering: the projected scattering angle θ_{plane} , the transverse displacement y_{plane} , and the geometric path length x .

The projected scattering θ_{plane} is defined as the projection of the space scattering angle. Assuming an incoming muon with direction $\vec{d}_{\text{in}}$ exiting the material with a direction $\vec{d}_{\text{out}}$, the space scattering angle θ_{space} and its projections in the xy and xz plane are defined as:

$$\begin{aligned}
\theta_{\text{plane}|xy} &= \cos^{-1} \left(\frac{\vec{d}_{\text{in},xy} \cdot \vec{d}_{\text{out},xy}}{||\vec{d}_{\text{in},xy}|| \cdot ||\vec{d}_{\text{out},xy}||} \right), \quad \theta_{\text{plane}|xz} = \cos^{-1} \left(\frac{\vec{d}_{\text{in},xz} \cdot \vec{d}_{\text{out},xz}}{||\vec{d}_{\text{in},xz}|| \cdot ||\vec{d}_{\text{out},xz}||} \right) \\
\theta_{\text{space}} &= \cos^{-1} \left(\frac{\vec{d}_{\text{in}} \cdot \vec{d}_{\text{out}}}{||\vec{d}_{\text{in}}|| \cdot ||\vec{d}_{\text{out}}||} \right) = \tan^{-1} \left(\sqrt{\tan^2 \theta_{\text{plane}|xy} + \tan^2 \theta_{\text{plane}|xz}} \right) \\
&\approx_{\theta < 20^\circ \text{deg}} \sqrt{\theta_{\text{plane}|xy}^2 + \theta_{\text{plane}|xz}^2}
\end{aligned}$$

Many successive statistical processes obey the Central Limit Theorem: the sum of many random small displacements follows a Gaussian distribution. The statistical outcome is a multiple scattering angle whose distribution is a 98% Gaussian core with elongated tails corresponding to "hard" scatter events given by:

$$f(\theta|\mu = 0, \sigma = \theta_0) = \frac{1}{\sqrt{2\pi}\theta_0} \exp \left(-\frac{\theta_{\text{plane}}^2}{2\theta_0^2} \right). \quad (1.8)$$

The standard deviation of the scattering angle distribution θ_0 is given by the Lynch & Dahl [105] equation:

$$\theta_0 \approx \frac{13.6 \text{MeV}}{\beta pc} \sqrt{\frac{x}{X_0}} \left(1 + 0.038 \ln \frac{x}{X_0 \beta^2} \right) \quad (1.9)$$

where p is the particle's momentum in MeV/c, x the length of the muon path in the medium, and X_0 the radiation length of the material in m computed as the ratio of X_0 in g cm^{-2} and the material's density ρ in g cm^{-3} . While other parametrization of scattering angles exists, the one proposed by Lynch & Dahl offers a satisfactory precision in the energy regime of muography, where the atmospheric muon spectrum peaks at 4 GeV. Additionally, density and radiation length appear separately in the equation, making this expression suitable for inverse problems based on density estimation. Figure 1.7 shows that the




Figure 1.7: Radiation length of pure elements and compounds (left y -axis in blue) and standard deviation of the scattering angle distribution $\theta_{0,\text{plane}}$ (right y -axis) of 5 GeV muons through 20 cm of material as a function of the material's density. Scattering angle values are obtained from Equation 1.9.

scattering of 5 GeV muons through 20 cm of material leads to measurable deflections.

The radiation length is a characteristic of a material defined as the mean distance over which a high-energy electron loses all but $1/e$ of its energy by bremsstrahlung. It reflects the microscopic characteristic of the material and can be defined as follows [116] for a single atomic element:

$$X_0 = 716.408 \text{ g cm}^{-2} A / [Z^2(L_{\text{rad}} - f(Z)) + ZL_{\text{rad}}] \quad (1.10)$$

With L_{rad} , L'_{rad} and f the Coulomb corrections given in Reference [116] for pure elements.

Matter around us is rarely made of pure elements but rather of compounds whose chemical composition is difficult to estimate with absolute precision. As such, it is recommended to use tabulated data provided by the Particle Data

Group [76]. Nevertheless, the radiation length of mixtures and compounds can be estimated relying on Bragg’s additivity law:

$$1/X_0 = \sum w_j/X_j, \quad (1.11)$$

with w_j the weigh fraction of the j -th element.

1.3 Muon Imaging

Muography exploits the abundant natural flux of muons to produce radiographic images of structures and objects by collecting individual muon position and direction data collected by the detection apparatus. This section introduces the two main categories of muography methods, namely muon transmission (also called absorption) and scattering tomography, which respectively rely on the energy loss and multiple Coulomb scattering mechanisms.

1.3.1 Muon Transmission Tomography

The first practical application of muography dates back to the 1950s, when E.P. George measured the mass of the overburden above the Guthega-Munyang tunnel in New South Wales. At the time, modern particle detection technology had yet to emerge, so he relied on Geiger counters—devices that could not measure directional information but only give an estimate of the charged particles’ rate coming from all directions. By moving the device along the tunnel, he was able to measure the dependence of the muon flux on the position. This experiment laid the foundation for muon absorption tomography: using muon flux measurements to infer an object’s internal structure. The principle behind this technique is rather straightforward—the more material a muon encounters along its path, the more kinetic energy it loses according to Equation 1.6, increasing the likelihood of its absorption. Consequently, an excess in muon flux indicates a region of lower density or the presence of a void, while a decrease suggests denser material. Since this pioneering effort, advances in particle detection technologies—driven by high-energy physics re-

search worldwide—have enabled more sophisticated experimental setups that incorporate positional information from the detectors.

Muon transmission tomography method aims at measuring the transmission, defined as the ratio between the muon flux upstream and downstream of the target object. An example of a detection system measuring the downstream flux is presented in Figure 1.8. The upstream-or "free sky"-flux is usually measured by rotating the detection system away from the target along the azimuthal angle, exploiting the φ -symmetry of the cosmic muon flux. In order to reduce systematic effects due to the apparatus, both measurements are conducted with the same detection system. Typical detection apparatus measures individual muons' direction (θ, φ) , the zenith and azimuthal angle, assuming a spherical coordinate system.




Figure 1.8: Example of muon transmission experiment, where the detector records the muons' direction after they traverse the volume of interest. The measured muon fluxes are used as input data for the muography reconstruction.

The effective muon transmission ratio $T(\theta, \varphi)$ is then computed as the ratio between the flux downstream the target $N_{\text{object}}(\theta, \varphi)$ and the free sky flux $N_{\text{free sky}}(\theta, \varphi)$:

$$T(\theta, \varphi) = \frac{N_{\text{object}}(\theta, \varphi)}{N_{\text{free sky}}(\theta, \varphi)} \quad (1.12)$$

To extract information about the density of the object, one compares the muon transmission ratio with expected fluxes under the uniform material assumption. Such expected flux is obtained by running a Monte-Carlo simulation, utilizing realistic models of the cosmic muon flux as well as muon propagation models. If a prior on the target object's composition is available, it can conveniently be used to measure density from opacity maps, defined as the integral of the densities along the muon path in a straight-line approximation: $O = \int \rho dl$, where O is the opacity, ρ the density, and dl is a differential distance. The transmission ratio measured from one point of view intrinsically produces 2D images, but measurements from multiple viewpoints can be correlated to obtain 3D density distributions of the scanned target [81]. Advanced algorithms based on the simultaneous algebraic reconstruction technique (SART) enable the generation of a 3D density map, taking into consideration the correlation between the muon zenith angle and the muon's kinetic energy [63]. A major advantage of muon transmission tomography is its ability to image extremely large structures, far beyond the size limits of conventional imaging techniques. As long as the muon flux emerging from the object can be measured within a reasonable time frame, a transmission map can be constructed. Thanks to the broad momentum spectrum of cosmic muons, ranging from hundreds of MeV up to several TeV, this technique is able to probe massive geological formations, such as volcanoes. To reach the detector after passing through such large volumes of rock, muons must possess sufficient energy to overcome the material's opacity, defined as $X : [\text{g.cm}^{-2}] = \int_{\mathcal{L}} \rho(x) dx$, where $\rho(x) : [\text{g.cm}^{-3}]$ is the local density along the muon's path $\mathcal{L}$. This requirement sets a minimum kinetic energy $E_{\min}(X)$ for the muons that can successfully traverse the medium. Taking the example of a 1 TeV muon propagating through standard rock, the distance d after it loses all its energy is given by the ratio of the continuous slowing down approximation range (CSDA range) to the material density ρ . Using numerical values provided by the Particle Data Group [83], we obtain a distance $d = \text{CSDA}_{\text{rock}} / \rho_{\text{rock}} = 2.453 \times 10^5 / 2.650 \approx 1 \text{ km}$. This 1 TeV energy falls in the far tail of the cosmic muon spectrum, meaning only a small fraction of incoming muons have sufficient energy to make it through 1 km of rock. Nevertheless, by collecting data over extended periods—often several months—statistically significant measurements of muon transmission can still be obtained, enabling imaging of vast structures.

1.3.2 Muon Scattering Tomography

The Muon Scattering Tomography (MST) concept was first introduced by Konstantin N. Borozdin *et al.* in a two-page-long publication in Nature [17]. While it has been a rather shy start for MST (the paper was published in the "brief communication" section of the Nature journal, in between "Deepwater variability in the Holocene epoch" and "Hunting and fox numbers in the United Kingdom"), it marked the dawn of a new area in muon imaging. The authors demonstrated that by measuring muons' scattering through the target, one can produce 3D images that reflect the object's density distribution. Because muons can travel through longer distances and denser materials than photons¹, MST shows a way better penetration power than X-ray imaging, which opens the door to a wide range of new applications involving large volumes made of dense high- Z materials. MST revolves around the measurement of individual muon deflection through the VOI and exploits the relation between muons' scattering angle distribution and the material's property to generate 3D images of the object's underlying structure. As shown in Equation 1.9, the standard deviation of the muon scattering angle distribution is tightly connected to the material radiation length X_0 , which itself depends on both density and atomic number Z . While inverting Equation 1.9 makes it possible to obtain an estimate of the material's X_0 , it would either require a measurement of the muons' momentum or the accumulation of enough data to cover the wide range of the muon's momentum spectrum. Unfortunately, neither condition is satisfied in most MST experiments: tracking detectors do not allow measurement of the muon's momentum, unlike in many high-energy physics experiments which utilize a magnetic field, and even a week of data-taking does not allow coverage of the whole cosmic muon energy spectrum. As a consequence, the various reconstruction algorithms used to produce the 3D images have to rely on approximations, which will be discussed in Section 3.1. To measure the scattering angle $\Delta\theta$, the muons' direction must be evaluated before and after interacting with the volume of interest: it is done by sandwiching the VOI between 2 sets of detectors, as shown in Figure 1.9. For each muon, its position through each detection plane i is recorded (x_i, y_i) . After applying a linear track fit, the incoming and outgoing direction $(\theta_{in,out}, \varphi_{in,out})$ are retrieved. Ultimately, the

¹ A typical 9 MeV X-ray beam is attenuated by a factor of 10 after passing through just 4 cm of lead, whereas an average cosmic muon with an energy of 4 GeV can travel through 2.6 m of lead before being stopped




Figure 1.9: Example of detection setup for MST experiment. Two sets of detectors record the muons' direction before and after the volume of interest, enabling the measurement of the scattering angle $\Delta\theta$.

inputs to the image reconstruction algorithms are $(xy_{in,out}, \theta_{in,out}, \varphi_{in,out}, \Delta\theta)$.

Since the true path of muons traversing the VOI cannot be known, reconstruction algorithms make assumptions and model muon paths through the VOI and associate muon deflection or passage within a certain region of space in the VOI. Path estimation algorithms usually require the initial and final positions and directions of the particles. Predicting the path is crucial for accurate image reconstruction. Those algorithms will be discussed in Section 3.1. Similarly to muon transmission tomography, the inference of the VOI's density and material composition constitutes an inverse problem and will be treated in Section 3. Since MST requires detecting both incoming and outgoing muons' direction, the target object size is limited (1-10m), especially compared to muon transmission tomography: detectors are expensive, and instrumenting large areas does not constitute a viable option. On the other hand, the typical acquisition time of muon scattering tomography is typically smaller—of the order of hours up to days—than for transmission. Applications of both techniques will be presented in the following sections.

1.4 Applications of muon tomography

In practical terms, muon tomography can be applied to almost any object, provided muons can pass through it and a detector can be positioned nearby. The key question, therefore, is not whether muography is feasible, but when conventional imaging techniques become inadequate or impractical. Is muography more cost-effective? Can it deliver results within a reasonable time frame? In many cases, muography becomes the method of choice precisely where other technologies fall short. This section highlights a range of applications involving both scattering and transmission muography. While not exhaustive, it provides a representative overview of the field. For a more in-depth review, see [58, 16]. The examples discussed here underscore the primary advantages of muography: its remarkable ability to look through dense materials and its non-invasive, passive nature.

1.4.1 Cargo Scanning

Border authorities employ a range of technologies to intercept illegal goods such as drugs, weapons, ammunition, and even special nuclear materials. Current screening methods for containers and cargo include ion mobility spectrometry (IMS), gas chromatography IMS (GC-IMS), and X-ray imaging. IMS and GC-IMS are typically used at checkpoints to inspect smaller items, often in desktop configurations, and are not well suited for scanning large cargo or freight vehicles. X-ray systems, especially those based on linear accelerators [8], are capable of penetrating container walls using dual-energy pulses (commonly 4 and 6 MeV or 6 and 9 MeV) delivered at frequencies up to 400 Hz. However, their use involves high-energy ionising radiation, which is subject to strict regulatory controls, making their deployment challenging. Millimeter wave technology, widely used in airport body scanners, is limited to detecting objects concealed in clothing or thin materials and lacks the ability to penetrate metal containers. While active inspection techniques can offer deeper penetration, they also introduce challenges related to radiological safety, including potential material activation. Moreover, both passive and active methods can be rendered ineffective if the contraband is heavily shielded. These limitations underline the need for alternative technologies that are safer, more effective, and capable of identifying material composition without rely-




Figure 1.10: Reconstructed 3D image of a cargo container holding a 1 m³ iron cube, generated using simulated data and visualized with the QIM3D software [41]. After just 30 minutes of data acquisition, MST produces a clear and detailed image of the internal structure. Unlike X-rays or γ -rays, muons can penetrate dense materials such as iron without significant absorption. The resulting resolution is sufficient to distinguish key features within the cargo.

ing on hazardous radiation.

MST presents a promising alternative to conventional inspection techniques. By analysing how cosmic muons get deflected as they traverse materials, MST can reveal and identify concealed objects, even those shielded by dense materials that typically defeat standard imaging methods, as illustrated in Figure 1.10. As a passive, radiation-free technique, it poses no health risks and does not require the constraining safety protocols associated with ionising radiation techniques. Pioneering efforts in applying MST to cargo security were led by Los Alamos National Laboratory and its spin-off, Decision Sciences [87, 13]. Using a detector system based on drift tubes, they demonstrated the ability to detect high-Z materials within time frames compatible with operational constraints at border checkpoints, typically around 15 minutes. Since these early developments, MST has advanced significantly in both detector technology and image reconstruction methods, making its implementation in

real-world border control increasingly viable. Ongoing projects aim to optimise system design and operation, including the integration of muon scanners into inspection facilities, selection and configuration of detectors [111], and development of fast reconstruction algorithms, image processing pipelines, and intuitive user interfaces. Chapters 3 and 6 will explore these aspects in greater detail. Recent research also expands MST’s capabilities by incorporating muon absorption analysis and the detection of secondary particles to enhance material discrimination. Among the active initiatives, the SilentBorder project [107] — part of the EU’s Horizon 2020 programme and the funding source for this PhD research — is a notable example. Additionally, private-sector interest is growing; companies like Muon Systems [75] have launched cargo inspection initiatives with support from national governments, such as Spain’s. A more in-depth review on Cosmic-Ray Tomography for Border Security can be found in Reference [12].

1.4.2 Industrial Applications

While being a recent technique, muography has raised interest in several industrial sectors. Its ability to deeply penetrate matter to image objects in a non-invasive way offers new possibilities that no other technology can claim.

Monitoring Blast Furnaces In industries that transform raw materials, such as minerals, petroleum, and rubber, production often involves chemical reactions under extreme conditions of temperature and pressure. The equipment used to transport raw materials, sustain reactions, supply energy, and extract products operates within these harsh environments, making both process monitoring and equipment maintenance particularly challenging. Muon imaging offers a compelling, non-destructive solution in this context. Thanks to its high penetration capability and passive nature, it allows for internal inspection without interrupting operations. Unlike traditional methods, it requires neither artificial radiation sources nor physical contact with the equipment, making it ideal for in-situ diagnostics in demanding industrial settings. The diversity of processes across the industrial sector means that imaging needs vary significantly from one application to another. One illustrative case is the inspection of ladle furnaces, where muon tomography is used to assess the distribution of material density within the blast furnace burden during operation. This tech-

nique enables continuous monitoring of wall wear, facilitating more precise planning of maintenance shutdowns and leading to substantial economic and energy savings. Given the considerable size of such furnaces—often around 15 meters in diameter—and the high density of the materials involved, muon tomography stands out as the only viable non-invasive method. These challenges have prompted the launch of European-funded initiatives to evaluate and develop muon imaging technologies for industrial use [32, 117]. Interestingly, given the scale of blast furnaces, which can reach 10 meters in diameter and 30 meters in height, both muon scattering and transmission tomography are regarded as promising imaging approaches for such large and dense structures.

Monitoring Insulated Pipes In the heavy industry and oil sectors, pipelines play a critical role in transporting gases and liquids throughout the production process. For example, a medium-sized petroleum refinery can operate with up to 200 km of piping. Over time, these pipes are subject to corrosion, which gradually reduces wall thickness. Once this degradation exceeds a critical threshold, the structural integrity of the pipeline—and potentially the entire facility—is at serious risk. To mitigate this, routine inspections are scheduled during planned technical shutdowns to assess pipe conditions and evaluate potential risks. A variety of non-destructive testing (NDT) techniques are commercially available for this purpose, including gamma radiography and ultrasound tomography. The required resolution varies depending on the pipe’s geometry and can range from a few millimeters to several centimeters. However, many pipelines are thermally insulated with layers of rock wool to prevent heat loss, which can hinder the effectiveness of certain NDT methods. In this context, muon tomography offers a promising alternative.

Nuclear Fuel Monitoring Muon tomography is particularly well-suited for the characterization of nuclear waste. Nuclear waste containers are typically designed with heavy shielding to contain radioisotopes and attenuate radiation, making them difficult to assess using conventional non-destructive techniques. In contrast, cosmic ray muons possess sufficient penetration power to probe the interior of virtually any waste container, and the extended mea-




Figure 1.11: Overview of instrumentation of the dry storage cask for the momentum-informed muon tomography system using a Cherenkov muon spectrometry (left), and MST reconstruction output shown as the average scattering angle across the cask (right), where the two missing fuel assemblies can be identified. Figure reproduced from [11].

surement times they require are generally acceptable given the long-term nature of waste monitoring. The EU-funded CHANCE project [21] (Characterization of Concealed Nuclear Waste for its Safe Disposal in Europe) explored the potential of MST for spent fuel monitoring. It successfully demonstrated that high-Z materials such as lead and uranium can be distinguished after just a few days of data collection [43, 125]. Additional research [11] has focused on imaging large dry storage casks—typically around 4 meters in height—showing that missing fuel assembly can be identified using MST, as shown in Figure 1.11. They concluded that the usual acquisition time of roughly 90 days can be significantly reduced by incorporating muon momentum information. Despite these advances, further investigation is needed to determine the sensitivity of muon tomography to smaller-scale material failures, such as cracks, imperfections, and swelling in structural components.

Civil Engineering Concrete is the most widely used construction material, and ensuring its structural integrity is critical for the maintenance and safety of buildings and infrastructure. Over time, aging and environmental stress can lead to internal defects such as cracks and voids, which may compromise the strength of a structure and, in extreme cases, result in catastrophic failures,

such as the collapse of the Morandi Bridge in 2018. Reliable, non-invasive inspection techniques are therefore essential. Traditional non-destructive testing methods—including X-rays, ultrasound, and radar—are commonly used to inspect concrete. However, each of these has limitations: X-rays have low penetration depth and require strict safety protocols, ultrasound lacks true 3D imaging capability, and radar struggles with complex internal geometries. Moreover, these methods often involve logistical constraints that make them impractical for large or remote structures. Again, thanks to the penetrating power of cosmic ray muons and the possibility of deploying portable detection systems, muon imaging offers a promising alternative by enabling in-situ inspection of large concrete structures without the need for artificial radiation sources or disassembly. Both muon scattering and transmission tomography have been explored for this purpose. Each method has its strengths: MST generally offers higher precision and faster data acquisition, while absorption techniques are logistically simpler, requiring only one detector placed on a single side of the structure. Recent investigations by Lynkeos Technologies and the German Bundesanstalt für Materialforschung und -prüfung in Berlin have further evaluated the use of muon tomography for reinforced concrete inspection [77]. Their studies involved embedding standard reinforcement bars within concrete blocks and comparing MST to conventional methods, as shown in Figure 1.13. While muon imaging does not match the resolution and acquisition time of X-rays, its ability to be deployed easily in the field, without regulatory hurdles, makes it a highly attractive option for large-scale structural health monitoring.

1.5 Detection Systems

For both muon scattering and muon transmission tomography, the role of the detection system is to measure the muon direction and position with reasonable precision. Unlike most particle physics detection systems that deal with high event rates-e.g. the proton bunch spacing of the Large Hadron Collider is 25 ns- the cosmic muon flux has a rather low event rate $\sim 1\text{cm}^{-2}\text{s}^{-1}$ which does not impose strong constraint on the time resolution of the apparatus, as one only needs to discriminate between two successive muons.

For MST, the muon position and trajectory must be reconstructed with a high accuracy (ideally $< 1\text{ mm}$ and about a few mrad) before (upstream) and af-




Figure 1.12: Sketch of the 20 cm thick concrete block with reinforced metal bars and a 50 mm thick Styrofoam block to be imaged (a), the image obtained with ultrasound (b), and the MST image (c), corresponding to an acquisition time of the order of weeks. Muography successfully detected and identified the built-in features of the concrete block: the Styrofoam block appears as a lower density region (white), while the reinforced metal bars appear as a higher density region (black). Figure reproduced from [77].

ter (downstream) the volume of interest. It requires at least two sets of two position-sensitive detector panels, as shown in Figure 1.9. If the symmetry of the problem allows it, detectors are placed horizontally to better exploit the verticality of the cosmic muon flux, thus granting shorter acquisition time. Muons being the most abundant charged particles at sea level, muon tomography detectors can function without particle identification modules and reject background particles such as electrons based on the quality of the track linear fit.

For muon transmission tomography, the muon trajectory is only measured once, downstream of the object. For experiments imaging large structures like pyramids or volcanoes, a rather small detector (with active detection surface of the order of 1 m^2) is placed at a large distance from the target, such that its acceptance covers the region of interest, as illustrated in Figure 1.8. For a given spatial resolution of a single detection plane, the angular resolution scales linearly with the distance between the two external panels. The optimal detector configuration becomes a trade-off between detector acceptance and angular resolution. Apart from measurement precision considerations, other practical aspects such as power consumption and sensitivity to external conditions are critical, especially in the context of in situ deployment. This section will review various detector technologies used in muon scattering and transmission tomography.




Figure 1.13: Detector panel of the MURAVES [46] experiment, made of 32 scintillator bars and wavelength-shifting fibers, that record the muon position along the direction perpendicular to the scintillator’s arrangement (a). Model of double-fiber layers (b). Figures reproduced from [35] and [12] respectively.

1.5.1 Scintillators

Scintillator-based detectors have played a fundamental role in nuclear and particle physics and constitute an ideal choice in muon tomography: they can be molded into various shapes to adapt to the foreseen detector geometry, and can be designed with a reasonable price/performance ratio. Their detection mechanism is based on scintillation, whereby materials emit light of a specific wavelength when traversed by particles or radiation. The emitted light is then collected with wavelength-shifting fibers and amplified by a photomultiplier. Scintillator-based detectors particularly suit muon transmission tomography, where the constraint on spatial resolution is more relaxed than for MST. Additionally, their robustness and low sensitivity to temperature and pressure variations make them adapted for operating in harsh conditions. The measurement of muons’ X and Y positions is done by superimposing two layers with orthogonal fiber placement. Figure 1.13a shows how scintillator bars can be assembled to form position position-sensitive detection surface.

For MST applications, a higher spatial resolution is required, which can be difficult to achieve for a scintillator bar arrangement given their individual thickness. Instead of bars, scintillating fibers of thin diameter (~ 2 mm) can be arranged in mats as shown in Figure 1.13b. Lynkeos technologies [77] and the collaboration between GSCan and the University of Tartu [4] adopted this technology and were able to achieve a spatial and angular resolution of the order of $100\text{ }\mu\text{m}$ and 1 mrad .

1.5.2 Gaseous detectors

Various gaseous detector technologies have already been used in many muography experiments due to their high resolution capabilities. Their detection principle relies on the collection of ionizing charges produced in the gas by the charged particle, using high-voltage electric fields. This type of detector has been used by various MST groups involved in homeland security projects [74], or nuclear casks imaging [85]. In muon transmission tomography, examples include MicroMegs [18] in the ScanPyramids project [73], Multi-Wire Proportional Chambers in the study of the Sakurajima volcano [80], and Resistive plate chambers for the imaging of the Puy-de Dôme [20].

1.6 Challenges

While many applications of muon tomography across various interdisciplinary domains have flourished over the past decades, the field still suffers from a few limitations that will be addressed in this section.

1.6.1 Muon Momentum estimation

Many muon transmission tomography experiments are affected by background contamination, primarily from low-energy muons that scatter and enter the detector acceptance—sometimes even from the rear—mimicking muons that have passed through the target. This contamination leads to an overestimation of the actual muon flux through the object, which biases the density measurement. In the context of imaging geological structures, only the most energetic muons (typically above 100 GeV) are capable of traversing hundreds of meters of rock without being stopped, making them the true carriers of useful information. However, these high-energy muons represent only a small fraction of the total cosmic muon flux and are vastly outnumbered by lower-energy muons, which form the dominant background. Some background muons, especially those entering the detector from behind, can be rejected using timing information—provided the detector has sufficient time resolution, as demonstrated in [60]. The challenge is greater for low-energy muons that enter the detector from the front, having been scattered off the surrounding environment.

The MURAVES experiment [46] tackled this issue by inserting a 60 cm thick lead wall between the third and final detection planes. Since low-energy muons scatter more than high-energy ones, they can be rejected based on the residuals of their track linear fit. While this method reduces low-energy contamination, it remains sub-optimal as the added scattering in the lead block degrades the overall track quality. In this context, incorporating muon momentum measurements into telescope systems, even with modest resolution, would allow more effective filtering of background muons and significantly enhance signal quality.

In the case of MST, momentum information becomes even more critical. The scattering angle distribution of a muon is directly dependent on its momentum. For example, a low-energy muon passing through a low-density material can produce the same scattering angle as a high-energy muon traversing a high-density material. Since MST image reconstruction and material classification algorithms rely on measured scattering angles, their accuracy is fundamentally limited by the absence of momentum data.

To overcome these limitations, Chapter 4 presents a detailed study of muon momentum inference techniques based on the analysis of scattering angles through known high-density material.

1.6.2 MST reconstruction algorithms

Since the first introduction of MST in scientific literature, a plethora of algorithms have been developed in the quest for obtaining the best density regression and material identification capabilities. Several research groups have tried to improve existing algorithms and adapt them to their specific use case, some focusing on material identification or density regression. However, this rapid increase in research groups, new ideas, and applications has not been accompanied by the development of standardized data formats and common tools for analysis. Each application being unique, whether it is in terms of required acquisition time, size and composition of the object of interest, or detector technology and configuration, efforts led to significantly different algorithms that ultimately compute different metrics reflecting different aspects of the target material properties. To this day, while almost every *muographer* has their own favorite reconstruction algorithm, no exhaustive benchmark comparing

the various reconstruction techniques and backed up with quantitative performance estimators exists.

To this day, there is no data analysis software or library shared among the community that compiles the most commonly used inference and image reconstruction algorithms. This is quite different from the high-energy physics community, where software solutions are shared among the community (e.g, ROOT, pyROOT). Even though individual tomography algorithms are relatively straightforward compared to full-scale particle-physics workflows, implementing and validating them from scratch is time-consuming and resource-intensive. As a result, most groups expend considerable effort re-coding similar routines rather than building on a shared foundation.

To address this issue, Chapter 3 will provide an extensive review and characterization of MST reconstruction algorithms, and introduce metrics that allow for meaningful performance comparison between them. These methods have been implemented in a Python-based data analysis framework called `Muograph` (further described in Appendix A.3) to facilitate community-wide adoption and benchmarking.

Monte Carlo simulation for Muon Scattering Tomography

Simulations have become an essential tool in physics, particularly for exploring complex systems and phenomena that are difficult to study analytically. Among these, Monte Carlo methods are especially valuable for addressing problems with inherent probabilistic behavior and for modeling systems with many coupled degrees of freedom. The core principle of this approach is the use of repeated random sampling combined with statistical analysis to estimate physical outcomes. The choice of simulation toolkit must be tailored to the specific research objectives, aiming at an appropriate balance between computational accuracy and performance, as a more realistic modeling of physical processes typically improves simulation fidelity but also increases computational cost.

In particle physics, Monte Carlo simulations are fundamental to detector design, data analysis, and the modeling of particle–matter interactions. Muon tomography is no exception; it heavily relies on such simulations to evaluate detector and image reconstruction algorithms’ performance prior to experimental implementation and to provide a reference framework for interpreting results.

A typical simulation workflow for muon tomography involves three main stages. First, muons must be generated with realistic properties—assigning each a kinetic energy and direction that reflect the cosmic muon flux, while preserving the correlation between momentum and direction. Second, the simulation must accurately model the propagation of muons through matter, accounting for how their trajectories and energies are altered by interactions with different materials. Maintaining sufficient precision at this stage is critical to avoid introducing bias into the reconstruction algorithms. Finally, the simulation must




Figure 2.1: Typical muon scattering tomography simulation steps: In red, the generation of cosmic muons with initial direction (θ, ϕ) and kinetic energy E , in yellow, the recording of the muon position $x_{\text{rec}}, y_{\text{rec}}$ through detector active area, and in grey, the propagation the muon through matter resulting in a new energy and direction E' and (θ', ϕ') .

model the detection process. The structure and interconnections of these three components are illustrated in Figure 2.1 and described in detail in the following sections. Finally, an example of a simulation scenario will be presented.

2.1 Cosmic muon Simulation

The first step of any muon tomography simulation is the generation of the cosmic muons. By generation, one means creating a collection of muons with an initial position x, y, z , a direction (θ, ϕ) , and kinetic energy E . It is done by relying on Monte Carlo simulations that model the development of Extensive Air Showers (EAS) in the atmosphere.

As discussed in Section 1.1.3, the development of EAS is governed by a complex system of equations involving the particle's decay probability and interaction cross-section and the properties of the atmospheric medium. As such, Monte Carlo simulations are essential to model such a convoluted system with many coupled degrees of freedom. The development of Monte Carlo sim-

ulation tools was motivated by the many studies and experiments involving cosmic rays such as IceCube [28], KM3NET [29], or the Pierre Auger Observatory [31], to name a few, which need an accurate modeling of the development of extensive air showers. Such effort gave birth to Monte Carlo toolkits like CORSIKA [54] (COsmic Ray SIMulation for KAscade), which simulate the evolution of EAS in the atmosphere initiated by photons, protons, nuclei, or any other particle. It tracks each particle as it propagates through the atmosphere, following its interactions, decays, or absorption. While tracking charged particles other than muons can be useful in muon transmission tomography to evaluate the background they generate in the detector, it becomes irrelevant in muon scattering tomography applications where background contamination from other charged particles is not significant. In such a context, the most popular approach is to use parametric generators like CRY [51] or ECMUG [82], where cosmic muons are generated according to a parametrization of their differential flux, based on either experimental data or results from cosmic-ray air shower generators simulating the full cascade of secondary particles initiated by the primary cosmic rays. This section first introduces the mathematical tools relevant to parameterizing the cosmic muon flux and presents Monte Carlo tools for cosmic muon generation.

2.1.1 Differential muon flux

The differential flux J is a relevant mathematical tool to describe the cosmic muon flux at sea level and can be parametrized as:

$$J = J(t, E, \theta, \varphi) = \frac{dN}{dE \cdot dt \cdot d\Omega \cdot dS_n}, \quad (2.1)$$

Where dN is the number of muons, with kinetic energy between E and $E + dE$, passing through a surface element dS_n orthogonal to the direction of motion within time interval dt and the solid angle $d\Omega = \sin \theta \, d\theta \, d\varphi$ with θ and φ the zenith and azimuthal angle defined in Appendix A. Many parameterizations are available in the literature, some of them correcting the shortcomings of the Gaisser model [44] from Equation 1.3, namely neglecting muon decay ($E_{kin} < 100/\cos \theta$ GeV) and the curvature of the Earth ($\theta < 70^\circ$):

$$J = 0.14 \left[\frac{E}{\text{GeV}} \left(1 + \frac{3.64 \text{ GeV}}{E(\cos \theta^*)^{1.29}} \right) \right]^{-2.7} \times \left[\frac{1}{1 + \frac{1.1E \cos \theta^*}{115 \text{ GeV}}} + \frac{0.054}{1 + \frac{1.1E \cos \theta^*}{850 \text{ GeV}}} \right], \quad (2.2)$$

with,

$$\cos \theta^* = \sqrt{\frac{(\cos \theta)^2 + P_1^2 + P_2(\cos \theta)^{P_3} + P_4(\cos \theta)^{P_5}}{1 + P_1^2 + P_2 + P_4}} \quad (2.3)$$

where $P_1 = 0.102573$, $P_2 = -0.068287$, $P_3 = 0.958633$, $P_4 = 0.0407253$ and $P_5 = 0.817285$.

To have an efficient simulation code, one must make sure that a large proportion of the generated particles will get within the detection system's acceptance. To achieve this, the surface from which the particles are generated must fit the geometry of the experiment. Given that most muon scattering tomography detectors consist of 2 sets of horizontal detection planes, it is convenient to generate muons from a flat, horizontal generation surface placed directly above the upstream detectors, at a fixed position z_{gen} and with an span that matches the active area of the first upstream detector. The differential flux through a horizontal surface J_z becomes:

$$J_z = \frac{dN}{dE \cdot dt \cdot d\Omega \cdot dS_n} \cdot \cos \theta = J \cdot \cos \theta \quad (2.4)$$

Such a strategy is called "flat sky generation", where the origin position of each muon is uniformly distributed over the generation surface, the azimuthal angle is uniformly distributed between $[0, 2\pi]$, and the energy E and zenith angle θ are sampled from the differential flux in Equation 2.4. For other muon tomography applications relying on the detection of horizontal muons, it can be relevant to use cylindrical or hemispherical generation surfaces, as proposed by the EcoMug generator [82]. Simulation results produced in this thesis are generated using flat sky generation thanks to the CRY software [51], as described in the following Section.




Figure 2.2: Comparison of the muon differential flux at sea level from CRY [51], as simulated by the MCNPX [124] software, and experimental data from [91]. CRY simulated data shows a good agreement with the experiment in the 200 MeV - 3 TeV range. Figure taken from [51].

2.1.2 Cosmic Muon Flux Simulation with the CRY software

The CRY generator [51], simulates the air showers initiated by primary cosmic rays, but unlike CORSIKA, which performs simulations in real time, it instead relies on precomputed tables generated through MCNPX [124] simulation of protons with energies ranging from 1 GeV to 100 TeV injected at the top of the atmosphere at an altitude of ~ 31 km. Although CRY produces all secondary particles—not just muons—it offers a significant speed advantage over CORSIKA. Users can generate particle shower distributions at three selectable altitudes (sea level, 2100 m, and 11,300 m) and adjust parameters such as geomagnetic cutoff and solar cycle effects. Naturally, such speed advantage comes with trade-offs, including the exclusive use of protons as primary particles, a limited energy range, and a simplified atmospheric model. Despite such shortcomings, CRY reproduces quite well the muon differential flux measured at sea level, as shown in Figure 2.2.




Figure 2.3: Distribution of atmospheric muons energy (left) and zenith angle (right) at sea level obtained through a horizontal surface with the CRY cosmic-ray particle shower library [51].

The CRY generator will be used throughout this thesis. After defining the vertical position at which muons are generated z_{gen} , and the span of the generation surface, CRY returns a collection of muons with energy E , direction (θ, φ) and position (x, y, z_{gen}) . CRY assumes a uniform distribution of the azimuthal angle φ and of the initial position (x, y) while it preserves the correlation between muons' kinetic energy E and zenith angle θ . The resulting angular and energy distributions are shown in Figure 2.3.

2.2 Passage of Particles Through Matter

After the generation of the cosmic muons $\mu(E, \theta, \varphi, x, y, z_{\text{gen}})$, they have to propagate through the volume of interest. The various electromagnetic interactions with the medium will ultimately result in a change of direction $(\Delta\theta, \Delta\varphi)$ and kinetic energy ΔE . The challenge is to reproduce the muon behavior as it propagates through the material while maintaining a realistic computation time. Given the relatively low size ($\sim 1 - 10 \text{ m}^3$), the GEANT4 software was deemed to be a good candidate as it proposes an accurate particle transport within a reasonable computation time.

2.2.1 GEANT4 software

The GEANT4 software [1] exploits the efficiency, flexibility, and oriented-object aspect of the C++ programming language to provide an efficient simu-

lator of the passage of particles through matter. Its areas of application include high energy, nuclear, and accelerator physics, as well as studies in medical and space science. As such, it incorporates models of the many processes particles are subjected to: elastic and inelastic scattering, photon interaction, X-ray production, hadronic processes, etc... In the context of muon tomography, we are mostly interested in the modeling of multiple Coulomb scattering and energy loss. When it comes to the modeling of detectors, more processes are relevant, but are specific to the technology used. While GEANT4 is a great tool to model scintillation-based detectors, which involves processes like electron-positron pair creation, ionization, and photon transport, it becomes more challenging for most gaseous detector technologies as it lacks the ability to properly model electric and magnetic fields, thus requiring linking it with dedicated software like Garfield [2]. In this thesis, GEANT4 was only used to simulate the propagation of muons through matter, which involves two main components: multiple Coulomb scattering and energy loss.

2.2.2 Multiple Coulomb Scattering

Multiple Coulomb Scattering Algorithms can be classified as either *detailed* or *condensed*. In detailed simulation algorithms, every interaction experienced by a particle is explicitly modeled. This approach is effectively exact, yielding results equivalent to solving the transport equation. However, such simulations are only practical when the number of collisions is relatively small—a condition typically met in special cases such as thin foils or low-density gases. In denser media like solids or liquids, where particles undergo a very large number of interactions, detailed simulations become computationally inefficient. To address this, high-energy physics simulations often employ condensed algorithms. Instead of modeling each interaction individually, only the cumulative effects of multiple collisions over a single step or track segment are simulated. The quantities typically computed include the net energy loss, total displacement, and change in direction of the charged particle. This philosophy is particularly relevant for muon scattering tomography applications, where the path length of particles within the material can reach up to 1 meter. Each simulation code implements its own algorithm to estimate angular deflection, path length correction, and spatial displacement following a step. These algorithms are approximate by nature and represent a significant source of uncertainty in transport simulations.

In GEANT4, the modeling of the multiple Coulomb scattering is more complex than the description given in Section 1.2.2. The main difference is the distinction between the *geometrical path length* z and the *true path length* t . On its way to point B from point A and in the absence of a magnetic field, the geometrical distance is simply the straight line that links A to B. The path length of an actual particle, however, is usually longer than the geometrical path length due to multiple scattering. The true path length t is used to sample the scattering angle from a modified Highland-Lynch-Dahl formula similar to Equation 1.9, which is accurate to 11% or better for $10^{-3} < t/X_0 < 100$. A thorough review of the implementation can be found in the GEANT4 physics manual [27].

2.2.3 Energy Loss

In GEANT4, energy loss is modeled using either a continuous or discrete approach. This implementation is designed to obtain a balance between simulation accuracy and computational efficiency, as modeling every individual interaction between the incident particle and the atoms of the medium becomes impractical for large thicknesses of material. To manage this, GEANT4 introduces a kinetic energy threshold for secondary particles, denoted as T_{cut} . Below the energy threshold, the energy loss is continuous, and above it, the energy loss is simulated by the explicit production of secondary particles. Given the differential cross-section per atom (atomic number Z) for the ejection of a secondary particle with kinetic energy T by an incident particle of total energy E :

$$\frac{d\sigma(Z, E, T)}{dT}, \quad (2.5)$$

the mean rate of energy loss is given by:

$$\frac{dE_{\text{soft}}(E, T_{\text{cut}})}{dx} = n_{\text{at}} \cdot \int_0^{T_{\text{cut}}} \frac{d\sigma(Z, E, T)}{dT} T dT, \quad (2.6)$$

with n_{at} the number of atoms per volume in the material. If there are several processes providing energy loss for a given particle, the total continuous part of the energy loss is the sum:

$$\frac{dE_{\text{soft}}^{\text{tot}}(E, T_{\text{cut}})}{dx} = \sum_i \frac{dE_{\text{soft},i}(E, T_{\text{cut}})}{dx} \quad (2.7)$$

These values are pre-calculated and stored in dE/dx tables and used at run time to obtain the continuous mean energy loss. In contrast, the production of secondaries with kinetic energies above T_{cut} is sampled by each concrete energy loss process. Energy loss being a stochastic process, the actual value of energy loss, whether it is continuous or discrete, is estimated using fluctuation models according to the estimated number of interactions of the particle in a step. A thorough review of the implementation can be found in the GEANT4 physics manual [27].

2.3 Detector Response

In muon scattering tomography, a muon’s position is recorded both before and after it passes through the volume of interest. As with all metrological processes, these measurements are subject to uncertainty, and modeling this uncertainty is essential for a realistic simulation of the detector response.

As discussed in Section 1.5, a variety of detector technologies exist, each exhibiting specific response and signal formation when traversed by muons. While GEANT4 is a widely used tool for detailed simulations of such effects, enabling the modeling of detector components and signal generation, it remains a labor-intensive approach, requiring careful specification of detector geometries and materials. While scintillator-based detectors have been successfully implemented in GEANT4 [53], gaseous detectors present greater challenges due to their reliance on ionization charge collection in high-voltage electric fields—an effect not natively supported by the GEANT4 framework.

For the development of reconstruction algorithms and imaging techniques, simulating a specific detector technology is not strictly necessary. To remain

broadly applicable, our focus is instead placed on the influence of general detector performance parameters—namely, spatial resolution σ_{xy} and detection efficiency ϵ . These factors are modeled independently of GEANT4 to streamline the simulation process while capturing their impact on reconstruction accuracy. Section 2.3.1 describes how the true muons’ positions are retrieved by GEANT4 and how they are used to compute the reconstructed position. Finally, Sections 2.3.2 and 2.3.3 justify the modeling of spatial resolution and efficiency effects, respectively.

2.3.1 Muon detection in GEANT4

As shown in Figure 2.1, muon positions must be recorded before and after interacting with the volume of interest. As GEANT4 propagates muons, it can retrieve muon positions by implementing sensitive volumes that record particle positions at each propagation step. To remain agnostic regarding the detector technology, we implement these sensitive volumes as non-physical and infinitely thin horizontal detection panels at a given vertical position z and with a given area. Such implementation is, of course, not perfectly realistic as it neglects the scattering of muons through the detector material, which would tend to degrade positional measurements. These panels record the true muon’s position as a continuous variable distributed across the panel’s area. An example of a typical detection setup in GEANT4 is shown in Figure 2.4. The reconstructed hits xy_{rec} are then derived from the true hits xy_{true} , by incorporating spatial resolution and detection efficiency effects as follows:

- **Detection efficiency:** A detection probability is assigned to each muon hit, based on the individual panel efficiency ϵ_{panel} . The muon-wise detection probability is computed as the product of the individual hit efficiency $P_\mu = \prod_{i=1}^n \epsilon_i$, assuming a full-coincidence trigger logic. To reject or not an event, one generates a random variable from a uniform distribution $X \sim \mathcal{U}(0, 1)$ and compares it with the detection probability. Only events with $X < P_\mu$ are kept.
- **Spatial resolution:** The reconstructed position is computed from the true muon position and the detector spatial resolution $xy_{rec} = f(xy_{true}, \sigma_{xy})$.




Figure 2.4: Example of muon scattering tomography setup in GEANT4, with 3 detection planes above and below the volume of interest. The latter consists of 4 cubes made of Aluminum (blue), copper (orange), iron (yellow), and lead (red). Muons are generated from a generation surface placed directly above the top detection plane, and their position are recorded by each of them.

The modeling of spatial resolution and efficiency is described in the following sections.

2.3.2 Spatial Resolution

Detectors record the particle's position as it passes through its active detection surface. Whatever the technology at use, the position measurement comes with an error $\delta_{x,y}$ defined as the difference between the true particle position xy_{true} and the reconstructed position xy_{rec} for a detection surface placed at a given z position:

$$\delta_{x,y} = xy_{rec} - xy_{true} . \quad (2.8)$$

Our goal is to obtain the reconstructed particle's position xy_{rec} from the true particle's position obtained in GEANT4, using a realistic modeling of $\delta_{x,y}$. To do so, we rely on simple assumptions to remain as general as possible:

1. The vertical position z of a flat, horizontal active detection surface is assumed to be known with absolute precision. The error on the position measurement $\delta_{x,y}$ is evaluated at z . We neglect any misalignment effects.
2. The error on the position measurement is considered independent of the particle's position $\delta_{x,y} \neq \delta_{x,y,z}(x,y)$. This constitutes an optimistic assumption, as real detectors may exhibit spatially varying resolution across their active area, particularly near the edges..
3. The error on the x and y position measurement are assumed to be identical: $\delta_x = \delta_y = \delta_{xy}$.
4. The positional error δ_{xy} is also assumed to be independent of the muon's zenith angle θ . In practice, the path length d of a muon within the detector active area, whether it is gaseous or scintillator-based, depends on the muon's zenith angle θ , with $d = \Delta a / \cos \theta$, where a is the active detection area thickness along the vertical direction. This angular dependence is neglected here.
5. The error on the position measurement δ_{xy} is an independent and identically distributed variable: the error associated with each muon has the same probability distribution and is disconnected from one another. In a realistic scenario, such an assumption is incorrect; the position measurement of a muon can be affected by the passage of a previous muon. Events where multiple muons cross the detector within a time interval smaller than the detector timing resolution are also possible.
6. The error on position measurement distribution obeys the Central Limit Theorem: it constitutes a sum of many random variables and therefore approaches a Normal distribution. It is then modeled as a continuous Normal distribution $\delta_{xy} \sim \mathcal{N}(\mu, \sigma)$ with mean μ and standard deviation σ .
7. Finally, positional error δ_{xy} is assumed to be unbiased: $\delta_{xy} \sim \mathcal{N}(\mu = 0, \sigma)$.

Ultimately, the muon's reconstructed position through a horizontal detection plane at a given vertical position z is computed as:

$$xy_{rec} = xy_{true} + \mathcal{N}(\mu = 0, \sigma = \sigma_{xy}), \quad z_{rec} = z,$$

with σ_{xy} the spatial resolution of the detection plane.

2.3.3 Detection efficiency

The other main parameter characterizing a detector's performance is the efficiency ϵ , defined as the probability of recording a signal if a particle passes through. Of course, this value must be as close as possible to 1 to ensure reasonable acquisition time. The panel's efficiency ϵ is modeled as a constant following a similar logic as for the spatial resolution:

1. The detection efficiency does not depend on the particle's position on the detection plane $\epsilon \neq \epsilon(x, y)$. For scintillators, the efficiency of light collection can vary from one end of the scintillator bar to another. For gaseous detectors, efficiency can be affected by the inhomogeneity of the gas mixture between the center and the side of the detectors.
2. The detection efficiency is assumed to be constant throughout the data taking. In reality, detection efficiency can fluctuate as it can depend on the temperature, atmospheric pressure, and quality of the gaz mixture, which is known to degrade over time.

Inverse problem in Muon Scattering Tomography

Two problems are considered inverses of each other when solving one requires knowledge of the other. More practically, an inverse problem consists of determining the causes from observed effects. This contrasts with a direct problem, in which the effects are predicted based on known causes. This distinction captures the core challenge of muon scattering tomography (MST): given the observed deflections of muons, what materials caused this scattering, and how are these materials distributed within the volume of interest?

A common difficulty in solving inverse problems lies in the need for a thorough understanding of the corresponding direct problem. Fortunately, in the case of muography, the physical processes governing muon propagation through matter have been extensively studied and are well understood.

Inverse problems are generally divided into two categories: restoration and reconstruction [9]. In restoration problems, the domain of the measurements and the domain of the unknowns are the same, as in image or signal denoising. In contrast, reconstruction problems involve different domains, such as inferring a 3D structure from 2D projections. MST combines both approaches. First, the reconstruction phase generates a 3D image from measured muon features—namely, the incoming and outgoing positions and directions. This can be followed by a restoration phase, where the resulting image is enhanced to reduce noise and artifacts inherent in MST reconstruction methods. Section 3.1 explores the main reconstruction algorithms found in the literature, describing how they transform measured muon data into 3D scattering density estimates and how these relate to material properties. To quantify the reconstructed image quality, we reformulate the MST reconstruction task as the minimization of an error functional between the predicted scattering density and the ground

truth radiation length. Section 3.2 analyzes the resulting scattering density maps, highlighting the presence of noise and artifacts, and connecting these to reconstruction techniques, detector characteristics, and acquisition time. Such a description will lay out the motivations for the use of deep learning-based approaches for image restoration, aiming to close the gap between reconstructed and true material distributions. This will be treated in Section 3.3.

3.1 Muon Scattering Tomography Reconstruction Algorithms

In muon scattering tomography, we are interested in predicting the density and radiation length of the volume of interest. The measured quantities used as input or the reconstruction algorithms are the muon position (xy_{in}, xy_{out}) , direction $(\theta, \varphi)_{in}, (\theta, \varphi)_{out}$ before and after the volume of interest, as well as the scattering angle $\Delta\theta$, to produce scattering density maps: the volume of interest is divided into 3-dimensional pixels, *voxels*, and a scattering density value is associated with each of them. Finally, the material's properties, namely radiation length and density, are inferred from the scattering density map. Since the true path of muons traversing the volume of interest (VOI) cannot be known, reconstruction algorithms make assumptions and model muon paths depending on the model through the VOI and associate muon deflection or passage within a certain region of space in the VOI.

After introducing the concept of scattering density, this section presents three reconstruction algorithms with various levels of complexity, and for each of them, a thorough description of how scattering densities are computed from the measured muon position and directions is provided. Distributions of scattering density predictions will be generated from a unique simulated data set to allow for a meaningful comparison between each of them. First, the point of closest approach (POCA), which was the first muon scattering tomography algorithm developed. Then, the Binned Clustered Algorithm, which is an extension of POCA that integrates a clustering metric and can conveniently utilize momentum information if available. Finally, the Angle Statistic Reconstruction, which employs a fundamentally different approach, where muon events are considered as tracks instead of point-like objects. The three algorithms are benchmarked using simulated data generated from the procedure

described in A.2, which will allow us to draw quantitative conclusions on their respective performance.

3.1.1 Scattering density

The concept of scattering density λ , or scattering strength, was first introduced in [102] and is defined as the *mean squared scattering of nominal muons per unit of depth*. Its expression is obtained by simplifying the equation of the standard deviation of the muons' scattering angle distribution σ_θ in Equation 1.9:

$$\sigma_\theta \cong \frac{13.6 \text{ MeV}}{p} \sqrt{\frac{x}{X_0}}. \quad (3.1)$$

By establishing a nominal momentum p_0 , the mean squared scattering of nominal muons per unit of depth of a material λ_{mat} is defined as:

$$\lambda_{\text{mat}} \equiv \left(\frac{13.6 \text{ MeV}}{p_0} \right)^2 \frac{1}{X_0} \cong \sigma_{\theta, \text{mat}}^2 \quad (3.2)$$

with $\sigma_{\theta, \text{mat}}$ the standard deviation of the scattering angle distribution. Throughout this manuscript, we will use a slightly different convention of the scattering density:

$$\begin{aligned} \sqrt{\lambda_{\text{mat}}} &= \frac{13.6 \text{ MeV}}{p_0} \sqrt{\frac{1}{X_0}} \propto X_0^{-1/2} \\ &\cong \sigma_{\theta, \text{mat}}. \end{aligned} \quad (3.3)$$

This formulation of the scattering density is designed for inference based on measured scattering angles. For monochromatic muons, the scattering density of any material is rather easy to infer from the standard deviation of the measured scattering angles, as shown in Figure 3.1 (left). But muon scattering tomography relies on the measurement of cosmic muons, whose energy ranges from a few 100 MeV up to a few TeV. Figure 3.1 (right) shows that the separation between various materials becomes less pronounced, making any material




Figure 3.1: Scattering angle distribution of muons through 10 cm of Aluminum and Iron, with 4 GeV/c momentum (left), and with realistic momentum distribution (right).

identification task more challenging.

In the context of such scattering angle-based inference, it is relevant to highlight the dependence between the set of measured scattering angles and the muon and materials properties:

$$\sigma_{\theta, \text{mat}} = f(\theta, p, X_0), \quad (3.4)$$

where X_0 is the material radiation length, $p = [p_0, \dots, p_n]$ the momentum and $\theta = [\theta_0, \dots, \theta_N]$ the zenith angle of the N muons. The zenith angle indirectly enters the equation as the muon path x within a given thickness scales with $\cos(\theta)$. This formulation assumes that scattering occurs within a homogeneous material. In practice, especially in MST applications, muons traverse the entire VOI, encountering varying thicknesses of multiple materials. Since the exact muon trajectory within the VOI is not known, reconstruction algorithms must rely on assumptions to assign scattering densities to specific regions. As a result, the measured scattering in one part of the VOI can be influenced by materials located elsewhere, a phenomenon discussed in detail in Section 3.2.4. The following sections will describe a few scattering density reconstruction algorithms, from the muon path modeling to the metric computation, and demonstrate their capabilities with simulated data generated scenario

I (further described in Section A.2), corresponding to 8 hours of acquisition time.

3.1.2 Point Of Closest Approach algorithm

Point of Closest Approach (POCA) is a commonly used muon path model that assumes that muon scattering occurs at a single point in space, which is located at the closest point between the incoming and outgoing muon trajectories. This approach was first introduced in [102] and gave birth to many reconstruction algorithms, with various levels of complexity. Because the POCA acronym can be a bit misleading—some use it to refer to the closest geometrical point between two tracks, while others use it to refer to any reconstruction muon scattering algorithm relying on the geometrical assumption mentioned above—we will explicitly refer to the POCA algorithm or POCA. This section introduces the POCA metric and presents two ways of using it to produce scattering density predictions of the VOI.

Metric computation As its name indicates, this approach assumes that muon scattering occurs at a single point in space, which is located at the closest point between the incoming and outgoing muon trajectories. In 2D space, this point would be the intersection of two straight lines. However, in 3D space, two nonparallel lines might not intersect, and the closest point is located in the middle of the (only) segment perpendicular to both lines. Its position can be computed from the parametric equations of the incoming and outgoing tracks, solving the following equation system:

$$\begin{aligned}
 A &= p_{in} + t_1 \cdot d_{in} \\
 B &= p_{out} + t_2 \cdot d_{out} \\
 \overrightarrow{AB} \cdot d_{in} &= 0 \\
 \overrightarrow{AB} \cdot d_{out} &= 0,
 \end{aligned}
 \tag{3.5}$$




Figure 3.2: Sketch depicting how the POCA is reconstructed from the incoming and outgoing muon entry point p_{in} , p_{out} and direction $\vec{d}_{in}$, $\vec{d}_{out}$ (a). Example of POCA reconstruction shown with the actual muon path as generated by GEANT4, for 1 GeV muons traversing 26 cm of lead (b). The left and right plots show successful and failed POCA reconstruction, respectively, where the POCA is placed inside and outside the material.

where $p + t \cdot d$ is the parametric equation of tracks, and p_{in} , d_{in} , and p_{out} , d_{out} are the measured points and directions of incoming and outgoing tracks, respectively. An illustration of Equation 3.5 is provided in Figure 3.2a. In this manuscript, we will refer to these geometrical points as POCAs.

But how precise can a reconstruction method based on such a simplistic single scattering hypothesis be, as we know that muons undergo not one but multiple scatterings? POCAs tend to be located close to the actual muon path within the material when the latter contains large deflections, as shown in Figure 3.2b, left. In some cases, the POCAs are not only reconstructed far from the muon path, but also outside of the volume of interest, as seen in Figure 3.2b, right. The cause of such mis-reconstruction is the small difference in angle between the incoming and outgoing muon direction, as the closest point between two parallel tracks is not mathematically defined. This constitutes a major shortcoming, since mis-reconstructed POCA cannot be exploited during scattering density inference, resulting in the loss of the muon event. POCA can be computed for each muon recorded by the apparatus throughout data taking, to obtain a distribution of 3D points $POCAs = [P_0, ..., P_i, ..., P_n]$, with $P_i = (x_i, y_i, z_i)$. To illustrate how the spatial distribution of POCAs is




Figure 3.3: Distribution of POCAs obtained from simulated muon scattering tomography data of scenario I, corresponding to 8h of acquisition time. Details about data generation can be found in A.2.

connected to the material properties, we compute the POCAs from the simulated muon scattering tomography data of 4 cubes of various materials, namely Lead, Iron, Copper, and Aluminum, as described in Appendix A.2. Looking at the spatial distribution of the POCAs presented in Figure 3.3, it becomes clear that the density of POCAs is correlated with the presence or absence of materials. The 4 cubes that constitute the VOI become visible as a region with large POCA density, while empty regions exhibit smaller density. The obtained 3D image remains quite noisy, not only because of the spatial resolution of the detector, but also because of the inherent uncertainty on the POCA location caused by the single scattering assumption.

At this stage, we remain quite far from making quantitative predictions about the VOI's material composition. The following paragraphs present two different yet complementary ways of producing scattering density predictions from the POCAs.

Scattering density predictions Now that the POCA's have been computed, they can be used to compute scattering density predictions of the VOI. To remain consistent throughout this manuscript and to ensure meaningful comparisons between reconstruction algorithms, the VOI is subdivided into 1 cm^3 *voxels* (3-dimensional pixels), and a scattering density prediction is attributed to each of them. The first method proposed to infer voxel scattering density was introduced in [102], and combines POCA's with scattering angle measurement. The voxel scattering density is estimated as:

$$\hat{\lambda}_{POCA} = \frac{1}{N} \sum_{i=1}^N \left(\frac{p_i^2}{p_0^2} \cdot \frac{\Delta\theta_i^2}{2L_i} \right), \quad (3.6)$$

where N is the number of POCA's within the voxel, p_i/p_0 the i -th muon estimated momentum over the average momentum of the cosmic muon flux ($\sim 3 \text{ GeV}$), L_i the muon path length within the voxel, and $\Delta\theta_i$ the i -th muon scattering angle. If no estimate of the muon momentum is available, the p_i^2/p_0^2 term is set to 1. Additionally, for small voxel sizes $\sim 1 \text{ cm}^3$, the muon path length within the voxel can be set to a constant, as the location of each POCA is not precise enough to provide a meaningful estimate of the actual muon path within the voxel. The scattering density then becomes the root mean square of the set of measured scattering angles:

$$\hat{\lambda}_{POCA} = \frac{1}{N} \sum_{i=1}^N \left(\Delta\theta_i^2 \right) = \text{RMS}(\Delta\theta). \quad (3.7)$$

Figure 3.5a shows how the scattering density prediction values computed from Equation 3.7 are distributed for the VOI from Scenario I (described in Appendix A.2). While the $\hat{\lambda}$ values corresponding to each material exhibit different values, their respective distributions overlap and cannot easily be separated. To better illustrate the relation between the scattering density estimator $\hat{\lambda}_{POCA}$ and the voxels' radiation length, we plot the correlation between the average $\hat{\lambda}_{POCA}$ of the 1000 voxels for each material and the associated inverse squared radiation length in Figure 3.4a. The resulting 3D image presented in




Figure 3.4: Average normalized scattering density predictions $\hat{\lambda}_{\text{POCA}}$ obtained from Equation 3.6 (a), and $\hat{\lambda}_{\text{Npoca}}$ from Equation 3.8 (b), as a function of inverse squared radiation length. Results and residuals of a linear fit are shown on both plots.

Figure 3.5 remains noisy, due to the statistical fluctuation in the $\text{RMS}(\Delta\theta)$. Such an effect can be circumvented by increasing the muon statistic and applying a post-processing routine to the predicted scattering density map. This procedure will be described in Section 3.2.3. On the other hand, this algorithm shows a good background discrimination as voxels that do not contain materials have their scattering density close to zero, with marginal fluctuations.

Another estimator of the scattering density is simply the number of POCA per voxel N :

$$\hat{\lambda}_{Npoca} = N \quad (3.8)$$

While being straightforward and not utilizing the measurement of the muons' scattering angle, this method has proved to be quite reliable. Figure 3.5a shows how the number of POCAs per voxel is distributed for the VOI from scenario I. The distributions appear to be better separated than the one computed from Equation 3.7, and do not contain outlier events affecting the reconstructed image's contrast. Consequently, the relation between the scattering density estimator $\hat{\lambda}_{Npoca}$ and the voxels' radiation length shows a better correlation than $\hat{\lambda}_{\text{POCA}}$, as shown in Figure 3.4b. On the other hand, the reconstructed image presented in Figure 3.6 exhibits more background in voxels that do not contain material.




Figure 3.5: Normalized scattering density predictions of 1cm^3 voxels obtained from Equation 3.7, for a volume of interest containing four $10 \times 10 \times 10\text{cm}^3$ cubes of Lead, Copper, Iron, and Aluminum and an acquisition time of 8 hours. Scattering density predictions of voxels containing material are shown as a histogram in (a) while (b) shows all voxel predictions averaged along the vertical axis. More details about the data generation can be found in Appendix A.2.




Figure 3.6: Scattering density predictions as the number of POCA's per 1cm^3 voxel, for a volume of interest containing four $10 \times 10 \times 10\text{cm}^3$ cubes of Lead, Copper, Iron, and Aluminum and an acquisition time of 8 hours. Scattering density predictions of voxels containing material are shown as a histogram in (a), while (b) shows all voxel predictions averaged along the vertical axis. More details about the data generation can be found in Appendix A.2.

3.1.3 Binned Clustered Algorithm

The Binned Clustered Algorithm (BCA) was first introduced in [114] in the context of spent nuclear fuel imaging, aiming at providing scattering density predictions that can effectively separate high- Z from low- Z materials by exploiting the spatial clustering of POCA. This section introduces the clustering metric and how it is exploited to produce scattering density predictions.

Metric computation The main idea behind the BCA algorithms is the computation of a voxel-wise estimator of the scattering density, based on the weighted cluster density of the POCA within that voxel. This metric m_{ij} is computed for each pair of POCA (i, j) within a given voxel:

$$\tilde{m}_{ij} = \frac{(\Delta\theta_i p_i) \cdot (\Delta\theta_j p_j)}{m_{ij}}, \quad (3.9)$$

with (p_i, p_j) the momentum and $(\Delta\theta_i, \Delta\theta_j)$ the scattering angle of muons i and j , and m_{ij} the geometrical distance between the POCA location of muon i and j . BCA can conveniently be run in *momentum-informed* mode: if an estimate of the muon momentum is available, it can be used in the computation of $\tilde{m}_{ij}$. Otherwise, it is run in *momentum-agnostic* mode, where the p_i and p_j values are set to 1. The $\tilde{m}_{ij}$ values are computed voxel-wise, for each pair of muon tracks $(i, j) \forall i \neq j$ whose POCA lands within the voxel's volume and are stored to form a distribution L :

$$L = [\tilde{m}_{01}, \tilde{m}_{02}, \dots, \tilde{m}_{ij}, \dots, \tilde{m}_{n-1}], \quad (3.10)$$

with n the total number of POCA within the voxel. A voxel containing n POCA thus contains $n(n-1)/2$ $\tilde{m}_{ij}$ values. Dense material with low radiation length produces many POCA points per voxel, and their associated scattering is large, leading to large $\tilde{m}_{ij}$ values. Low-density and large radiation length materials, on the other hand, tend not to scatter muons, leading to a small number of POCA per voxel and smaller $\Delta\theta$, thus smaller $\tilde{m}_{ij}$ values. Figure 3.7a and 3.7b show how the $\tilde{m}_{ij}$ values are distributed for voxels of various materials, with and without momentum information. It is interesting to note that utilizing the momentum information in Equation 3.11 allows for




Figure 3.7: Natural logarithm of the weighted metric values in momentum-agnostic (a) and momentum informed-mode (b) for $10 \times 10 \times 10 \text{ cm}^3$ cubes of Aluminum, Copper, Iron, and Lead. Utilizing the true muon momentum in the computation reduces the overlap between the different $\ln(\tilde{m}_{ij})$ distributions. Details about data generation can be found in Appendix A.2.

substantial reduction of the overlap between the $\tilde{m}_{ij}$ distributions associated with different materials.

Scattering density predictions The final scattering density prediction of each voxel $\hat{\lambda}$ is computed as the median of their respective $\tilde{m}_{ij}$ distribution L :

$$\hat{\lambda}_{\text{BCA}} = \text{med}(L) \quad (3.11)$$

Once all voxel scattering density predictions are obtained, they are normalized such that they range between 0 and 1. The flow of the algorithm is detailed in pseudo-code in Algorithm 1. The $\hat{\lambda}_{\text{BCA}}$ scattering density map shown in Figure 3.8a contains a major limitation: the scattering density of voxels without POCAs cannot be inferred, as no $\tilde{m}_{ij}$ values can be computed. Similarly, the scattering density prediction of voxels with only a few POCAs (~ 3) is not statistically significant. Consequently, the scattering density of low-density materials can become difficult to estimate as they tend not to produce enough POCAs. To overcome this limitation, we simply assign to those empty voxels the lowest scattering density value encountered in the VOI. Figure 3.8a shows that for low-density materials like aluminum, around half of the 1000 aluminum voxels contain less than 4 POCAs and get a scattering density of

0, against around 40 for lead voxels. Section 3.2.3 will address this issue by proposing post-processing procedures. As demonstrated in Figure 3.10, the $\hat{\lambda}_{\text{BCA}}$ estimator shows a strong correlation with the inverse squared radiation length, even more so than $\hat{\lambda}_{Npoca}$ and $\hat{\lambda}_{\text{POCA}}$.

Algorithm 1 Workflow of the Binned Clustered Algorithm in pseudo code.

Before the scan ($n, \Delta\theta_{\text{cut}}$)

Sort tracks by descending scattering angle.

Compute POCA location of each track.

Create an empty array to store voxels scattering density predictions S .

Voxel scan

for each voxel k **do**:

Assign an empty list L_k to the current voxel.

for each pair of tracks with their POCA location in k **do**:

Compute $\tilde{m}_{ij}$ from Equation 3.11.

Append $\ln(\tilde{m}_{ij})$ to L_k .

Set scattering density of current voxel to the median of L_k .

At the end of the scan (N_{min})

Retrieve the lowest scattering density value $S_{\text{min}} \leftarrow \min(S)$.

Set scattering density of voxels with less than N_{min} POCA to S_{min} .

Re-normalize scattering density within the $[0, 1]$ interval.

3.1.4 Angle Statistics Reconstruction

The Angle Statistics reconstruction (ASR) algorithm was first introduced in [109], in an attempt to overcome the two major failings of algorithms based on POCA: the poor reflection of the path of the muon and the waste of muon events having their POCA located outside the VOI. To overcome these two issues, the authors adopted an approach similar to backprojection algorithms [93], which are commonly encountered in traditional tomography and often used as a starting point for more advanced algorithms.




Figure 3.8: Scattering density predictions obtained from the BCA algorithm 1, for a volume of interest containing four $10 \times 10 \times 10 \text{ cm}^3$ cubes of Lead, Copper, Iron, and Aluminum and an acquisition time of 8 hours. Scattering density predictions of voxels containing material are shown as a histogram in (a), while (b) shows all voxel predictions averaged along the vertical axis. The minimum number of POCA's per voxel is set to 4. More details about the data generation can be found in Appendix A.2.




Figure 3.9

Figure 3.10: Average normalized scattering density predictions $\hat{\lambda}_{BCA}$ obtained from Equation 3.11 as a function of inverse squared radiation length.

Metric computation Similarly to BCA, ASR relies on the voxelization of the volume of interest to apply a score to each individual voxel. Since ASR does not utilize POCA, it has to use a different approach to determine which voxels to apply a score to, given the incoming and outgoing muon tracks. The idea is to assign a score to each voxel that is close enough to both the incoming and outgoing muon tracks. While the initial implementation proposed in [109] uses a characteristic distance to decide if a voxel is close to the muon tracks, we went for a slightly simpler practice: only voxels that get traversed by both the incoming and outgoing tracks receive a score for the current muon event. We refer to those voxels as triggered voxels. Figure 3.11 shows how the number of triggered voxels decreases as the muon’s scattering angle increases. Once the triggered voxels have been identified, a metric m_i is assigned to them based on the muon’s scattering angle $\Delta\theta_i$ and momentum p_i , if an estimate is available:

$$m_i = \Delta\theta_i \cdot p_i . \quad (3.12)$$

If no momentum estimate is available, p_i is set to 1. Similarly to BCA, each voxel receives an empty list L during initialization, which is later filled iteratively with the metric m_i , as muons get detected:

$$L = [m_0, \dots, m_n] \quad (3.13)$$

Figure 3.12a and 3.12b show how the metric values m_i are distributed for different materials, with and without the momentum information. While the distributions associated with different materials overlap, the 75% quantile of the m_i distributions are separated. The scattering density of each voxel is then computed as a function of the metric values’ list L .

Scattering density predictions At the end of the scan, the voxel scattering density of the voxel k is computed as a quantile of its m_i distribution L_k :

$$\hat{\lambda}_{\text{ASR}} = \text{quantile}(L_k) \quad (3.14)$$




Figure 3.11: Illustration of the ASR voxel triggering for a low scattering angle (left) and a large scattering angle event. As the scattering angle increases, the number of triggered voxels gets smaller.




Figure 3.12: Normalized natural logarithm of the metric values in momentum-agnostic (a) and momentum-informed mode (b) for $10 \times 10 \times 10 \text{ cm}^3$ cubes of Aluminum, Copper, Iron, and Lead. Utilizing the true muon momentum in the computation reduces the overlap between the different $\ln(m_i)$ distributions. Details about data generation can be found in Appendix A.2.

Algorithm 2 Workflow of the Angle Statistics Algorithm in pseudo code.

Before the scan

Each voxel receives an empty list L .

Create an empty array to store voxels scattering density predictions S .

Muon scan

for each muon i **do**:

Find the voxels $V = [v_0, \dots, v_n]$ traversed by both incoming and outgoing tracks.

for each voxel v_j in V **do**:

 Compute m_i from Equation 3.12.

 Append m_i to L_j .

At the end of the scan

for For each voxel j **do**:

 Scattering density prediction of j is the 75% quantile of S_j .

The optimal quantile value was found to be 75% [109]. The scattering density map of $\hat{\lambda}_{\text{ASR}}$ presented in Figure 3.13 shows an exceptional smoothness compared to the POCA-based algorithm. While the number of POCA per voxel does not exceed 20, the average number of m_i values per voxel amounts to 140 for ASR, which ultimately leads to a more statistically significant scattering density prediction. The counterpart of this effect is the vertical smearing of the predictions, which will be discussed in Section 3.2.4. Similarly to BCA, $\hat{\lambda}_{\text{ASR}}$ shows a good correlation with the inverse squared radiation length, as shown in Figure 3.14.

So far, we have made mostly qualitative assertions about the scattering density reconstruction performance of the algorithms. In Section 3.2, we introduce estimators borrowed from inverse problem theory that can properly quantify the quality of the inference. It is worth mentioning that there is no consensus on which quantitative value better describes the material properties. Each article uses a different figure of merit to assess the superiority or not of the proposed algorithm, uses different simulated data obtained from different acquisition times (which translates into different numbers of muons), different simulation tools, and the code is not properly explained. We will avoid those




Figure 3.13: Normalized scattering density predictions obtained from the ASR algorithm 2, for a volume of interest containing four $10 \times 10 \times 10 \text{ cm}^3$ cubes of Lead, Copper, Iron, and Aluminum and an acquisition time of 8 hours. Scattering density predictions of voxels containing material are shown as a histogram in (a), while (b) shows all voxel predictions averaged along the vertical axis. More details about the data generation can be found in Appendix A.2.




Figure 3.14: Average normalized scattering density predictions $\hat{\lambda}_{\text{ASR}}$ obtained from Equation 3.14 as a function of inverse squared radiation length.

bad practices, make the effort to unify reconstruction techniques, and get a better understanding of their limitation and respective performance.

3.2 Characterization of MST imaging

Up to this point, we have introduced several algorithms that infer scattering density maps of the volume of interest (VOI) based on measured muon positions and directions. Although these algorithms differ in their underlying assumptions and the specific features they use during inference, they are all subject to similar sources of noise that affect the quality of their predictions.

In this section, we aim to systematically characterize the imaging performance of muon scattering tomography (MST). Specifically, we seek to address the following questions: What are the primary sources of noise? How do detector configuration and geometry influence reconstruction quality? Which reconstruction algorithm yields the most accurate results?

We begin by introducing evaluation metrics commonly used in machine learning and image processing to assess the fidelity of scattering density predictions. We then examine several key sources of noise: the impact of the detector’s acceptance, formalized through the concept of luminosity, the uncertainties introduced by assumptions made about muon paths, and the precision of the muon tracks measurement. To mitigate the fluctuations of the scattering density predictions due to limited statistics, we propose the use of Gaussian smoothing kernels. This detailed examination of these intertwined sources of noise and error highlights the inherent complexity of the MST inverse problem. This motivates the exploration of deep learning approaches for enhancing the reconstructed images, a topic that will be addressed in the final section of this chapter. Algorithms and metrics discussed in this chapter have been implemented in an open-source Python-based software called `Muograph`, which is further described in Appendix A.3.

3.2.1 Metrics

Quantitative metrics are essential for evaluating how closely reconstruction results match the ground truth: they provide an objective basis for assessing

performance. However, in the context of muon scattering tomography (MST), there is currently no standard framework for image quality assessment. One commonly used metric is the Area Under the Receiver Operating Characteristic (ROC) Curve (AUC), applied to binary classification tasks. For example, studies such as [114, 109, 117] use AUC to distinguish between high- and low-density materials. While AUC may offer some insight, it is a limited indicator of overall image quality. It treats each voxel as a binary classification problem and fails to capture the continuous nature of scattering density predictions across the volume. Moreover, it does not account for image noise or incorporate the spatial coherence of the reconstruction. The challenge of finding appropriate metrics is compounded by the diversity of objectives in MST applications. Depending on the use case, the key interest may be detecting minimal differences in atomic number (Z), resolving the smallest spatial features, or reducing the acquisition time necessary to test a specific hypothesis. In all cases, meaningful and robust performance metrics are vital for optimizing algorithms and comparing results.

This Section will address this pending issue in the muon scattering tomography community by proposing several quantitative figures of merit inspired by what is done in image processing and medical imaging.

Mean Squared Error The goal of a quantitative performance metric is to compare two signals by providing a quantitative score that describes the degree of similarity/fidelity or, conversely, the level of error/distortion between them. Usually, it is assumed that one of the signals is a pristine original (voxels radiation length ground truth), while the other is distorted or contaminated by errors, in our case, the scattering density predictions. For more than 60 years, the mean squared error (MSE) has served this purpose in the fields of signal processing [122] and optimisation. In the context of MST, assuming a 3D array of n voxels, $\hat{Y}$ the vector of predicted scattering density values associated with each voxel, and Y the vector of the voxels' radiation length values, the mean squared error is computed as:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2 \quad (3.15)$$

The MSE has many attractive features:




Figure 3.15: Inverse squared radiation length $X_0^{-1/2}$ of scenario I.

- It is simple. It is parameter-free and inexpensive to compute, with a complexity of only one multiplication and two additions per sample. It is also memory-less—the squared error can be evaluated at each sample, independent of other samples.
- It satisfies the following convenient conditions, and allow for consistent, direct interpretations of similarity: non-negativity $\text{MSE}(\hat{Y}, Y) > 0$, identity $\text{MSE}(\hat{Y}, Y) = 0$ if and only if $\hat{Y} = Y$, symmetry $\text{MSE}(\hat{Y}, Y) = \text{MSE}(Y, \hat{Y})$.
- It can be evaluated on a subset of the voxel grid, to separately assess the quality of signal (where the target material is located) and background (where there is no material) reconstruction. It is particularly interesting for imaging tasks that focus on a specific region of the VOI, e.g., in nuclear waste characterization.




Figure 3.16: Normalized scattering density predictions as the number of POAs per voxel, for a volume of interest containing four $10 \times 10 \times 10 \text{ cm}^3$ cubes of Lead, Copper, Iron, and Aluminum and an acquisition time of 8 hours (left). Normalized scattering density predictions as the number of POAs per voxel obtained from the same data set, but filtering out muon events with scattering angle $< 2^\circ$ deg. While the right-most picture appears of better quality due to its lower background level, the scattering density predictions of material voxels are closer to the radiation length ground truth: $\text{MSE}_{\text{signal}} = 0.114 < \text{MSE}_{\text{signal}} = 0.165$.

The only requirement for computing the MSE is the availability of the ground truth vector Y , representing the voxels' radiation length values—i.e., the target of the inference process. While it is rather straightforward to extract from the simulation scenario from Appendix A.2, as the placement of the cubes was carefully chosen to align with the voxel grid, it can be difficult to obtain for more complex shapes. Given the definition of scattering density in Equation 3.3, the MSE will be computed by comparing the normalized scattering density predictions $\hat{\lambda}$ from different reconstruction algorithms with the normalized inverse square root of the radiation length, $X_0^{-1/2}$, illustrated in Figure 3.15. To better characterize algorithm performance, three variants of MSE will be computed: the total MSE, calculated over all voxels; $\text{MSE}_{\text{signal}}$, computed only over voxels containing material; $\text{MSE}_{\text{background}}$, computed over voxels containing no material. As shown in Figure 3.16, evaluating these three metrics separately is crucial. A lower total MSE may result from reduced background contamination rather than from more accurate inference of radiation lengths. Depending on the specific application—e.g., whether the focus is on identifying fine structural details or minimizing false positives in empty regions—it may be more relevant to optimize one metric over the others.

Structural Similarity Index Image similarity assessment is a fundamental issue in many signal processing applications. Traditionally, mean squared error has been the standard not only for the evaluation of image quality, but also for the design and optimization of signal processing algorithms and systems. It constitutes an appealing estimator because it is simple to calculate, has clear physical meanings, and is mathematically convenient in the context of optimization, but is not very well matched to perceived visual quality. The structural similarity index [123] was developed to serve that purpose, by exploiting the dependency between neighboring pixels that carries important information about the structure of the objects in the visual scene. If we consider one of the signals Y as a reference, then the similarity measure can serve as a quantitative measurement of the quality of the second signal $\hat{Y}$. The SSIM separates the task of similarity measurement between an image $\hat{Y}$ and its reference Y into three comparisons: luminance l , contrast c , and structure s :

$$\text{SSIM}(Y, \hat{Y}) = [l(Y, \hat{Y})]^\alpha \cdot [c(Y, \hat{Y})]^\beta \cdot [s(Y, \hat{Y})]^\gamma, \quad (3.16)$$

where $\alpha > 0$, $\beta > 0$, and $\gamma > 0$ are parameters used to adjust the relative importance of the three components. The resultant SSIM index is a decimal value between -1 and 1 , where 1 indicates perfect similarity, 0 indicates no similarity, and -1 indicates perfect anti-correlation. First, the luminance comparison of Y and $\hat{Y}$ is computed as:

$$l(Y, \hat{Y}) = \frac{2\mu_Y\mu_{\hat{Y}} + C_1}{\mu_Y^2 + \mu_{\hat{Y}}^2 + C_1}, \quad (3.17)$$

where μ_Y is the average voxel value, and C_1 a constant included to avoid instability when $\mu_Y^2 + \mu_{\hat{Y}}^2$ is very close to zero. The contrast comparison function takes a similar form:

$$c(Y, \hat{Y}) = \frac{2\sigma_Y\sigma_{\hat{Y}} + C_2}{\sigma_Y^2 + \sigma_{\hat{Y}}^2 + C_2}, \quad (3.18)$$

where $C_2 = (K_2 L^2)$ with $K_2 \ll 1$, and σ_Y an unbiased estimate of the voxels standard deviation:

$$\sigma_Y = \left(\frac{1}{N-1} \sum_{i=1}^N (Y_i - \mu_Y)^2 \right)^{1/2}. \quad (3.19)$$

Finally, the structure comparison function is defined as follows:

$$s(Y, \hat{Y}) = \frac{\sigma_{Y\hat{Y}} + C_3}{\sigma_Y \sigma_{\hat{Y}} + C_3}, \quad (3.20)$$

with $C_3 = C_2/2$, and $\sigma_{Y\hat{Y}}$ defined as:

$$\sigma_{Y\hat{Y}} = \frac{1}{N-1} \sum_{i=1}^N (Y_i - \mu_Y)(\hat{Y}_i - \mu_{\hat{Y}}). \quad (3.21)$$

For image quality assessment, it is better advised to evaluate the SSIM index locally rather than globally. In practice, it is done by computing μ_Y , σ_Y and $\sigma_{Y\hat{Y}}$ within a local $N \times N \times N$ window, which is moved voxel-by-voxel over the entire image. The overall image quality is then computed as the mean of the SSIM values of each local window j :

$$\text{MSSIM} = \frac{1}{M} \sum_{j=1}^M \text{SSIM}_j \quad (3.22)$$

In the following sections, the MSSIM will be evaluated for a window size of $N = 11$ voxels, taking as reference the inverse squared radiation length ground truth from Figure 3.15. Depending on the application, it is also possible to compute a weighted average of the different samples in the SSIM index map to focus on specific regions of the 3D image.

3.2.2 Lumuosity

We introduce a new physical quantity to describe the convolution of the muon flux and detector acceptance. Looking at the symmetries of a problem is always important, and often gives meaningful insight about the underlying physics. First, let's consider the symmetries in the cosmic muon flux, independently of any measurement system. As discussed in Section 1.1, cosmic muon flux depends only on the muon's energy and zenith angle. Thus, we observe a symmetry along the azimuthal angle: The measured muon flux along a given direction (θ, φ) does not depend on φ . Of course, this reasoning is valid only in outdoor conditions, where no building or geological structure absorbs muons. Now let's consider a typical muon scattering tomography system. For a muon to be detected, it has to leave a hit in every detection plane. The detection system then has a given acceptance. The effective solid angle depends on both the size and the distance between the two external detectors, and also it continuously varies on the detector surface. In most muon applications, including muon radiography and monitoring, at least two-fold coincidence muon detectors are installed, and the target objects are placed between two detectors. According to the cosine-squared model, the expected muon counts only depend on the active surface area of the detectors. However, they also depend on the distance between the two detectors because their measurable solid angle becomes larger as the distance decreases. To consider the geometry and configuration of detectors in cosmic muon flux estimation, we developed a new approach named the effective solid angle.

Let's assume a muon scattering tomography detection system, which consists of two sets of two detector planes that are vertically aligned, where the two external panels are separated by a distance z . Figure 3.17a illustrates the solid angle at a given position O on the top detector panel. A muon entering the top detector in O can only be detected if it reaches the bottom plane: its direction (θ_x, θ_y) obeys $\theta_x \in [\theta_{x,min}, \theta_{x,max}]$, $\theta_y \in [\theta_{y,min}, \theta_{y,max}]$, with $\theta_{yx,min,max} = \tan(xy_{min,max}/z)$. One can generalize to any incoming position O_i on the top detector. For the sake of illustration, we represent this "cumulative" solid angle on a 2D projection in Figure 3.17b by connecting every entry point O on the top panel to every exit point M on the lower panel. While this representation remains naive, as it is a discrete approach that does not take into account the muon flux dependence on the zenith angle, it hints that some




Figure 3.17: Illustration of the detector solid angle given the muon position on the top panel O (a). Illustration of the generalization of the detector solid angle for various muon positions O_i on the top detector panel (b).

regions of the VOI receive more muons than others. Typically, we expect the regions at the left and right edges of the VOI to get traversed by fewer muons than the center.

To quantify this effect, we introduce the concept of *lumuosity* $\mathcal{L}_\mu$ defined as the density of muon tracks per unit of volume, given a certain detector configuration and acquisition time Δt . To illustrate how such a quantity is distributed across the VOI, we run a dummy simulation using the detector configuration described in Appendix A.2. By voxelizing the VOI, we can simply count how many muons pass through each voxel for a given acquisition time. Figure 3.18 shows how the lumuosity values are distributed across the VOI, for the muon scattering tomography detector presented in Appendix A.2.




Figure 3.18: Measured lumuosity of the simulated detection setup from Appendix A.2, averaged over the horizontal direction (a) and vertical direction (b). The average lumuosity within a 10 cm^3 volume in the center of the detector acceptance (A) reaches up to $33 \text{ cm}^{-3} \text{ h}^{-1}$, and falls to $15 \text{ cm}^{-3} \text{ h}^{-1}$ in the corner (B).

Effect on reconstructed images We have established that the lumuosity is not uniformly distributed across the volume of interest. Is that effect important? Should it be taken into account during reconstruction? How does it affect the reconstructed images? To answer those questions, we will run a simple experiment. Let's try to image the same objects but placed at different locations within the VOI. In scenario I, we place a $10 \times 10 \times 10 \text{ cm}^3$ cube of copper at the center of the VOI ($x, y, z = 0, 0, -440 \text{ mm}$, region A in Figure 3.18), and in scenario 2 the same cube is placed in a corner of the VOI ($x, y, z = 250, 250, -400 \text{ mm}$, region B in figure 3.18). The average lumuosity in the cube region is $\mathcal{L}_{\text{center}} \sim 33 \text{ cm}^{-3} \text{ h}^{-1}$ for scenario 1 and $\mathcal{L}_{\text{corner}} \sim 15 \text{ cm}^{-3} \text{ h}^{-1}$ for scenario 2: their reconstructed image are expected to be different. The associated reconstructed images, computed as the number of POCA per voxel and reported in Figure 3.19, perfectly demonstrate how a change in lumuosity affects the scattering density predictions. Not only is the average scattering density prediction within the cube region way smaller in B ($\bar{\lambda}_{\text{POCA}}(B) = 17$) than in A ($\bar{\lambda}_{\text{POCA}}(A) = 36$), but their spatial distribution also differs, as shown in Figure 3.20. The difference between $\bar{\lambda}_{\text{POCA}}(A)$ and $\bar{\lambda}_{\text{POCA}}(B)$ is a direct consequence of the difference lumuosity values: as region A *sees* more muons than region B, it exhibits a large number of POCAs.




Figure 3.19: 2D projections of scattering density predictions as the number of POCA per voxel for a $10 \times 10 \times 10 \text{ cm}^3$ cube of Copper placed at $x, y, z = 0, 0, -440 \text{ mm}$ (left) and at $x, y, z = 250, 250, -400 \text{ mm}$ (right).




Figure 3.20: Average scattering density as the number of POCA per voxel $\bar{\lambda}_{\text{POCA}}$ along the horizontal x (left) and vertical z direction (right) of voxels within region A and B. The $\bar{\lambda}_{\text{POCA}}$ values in region B are systematically lower, due to the lower luminosity. While the shape of the $\bar{\lambda}_{\text{POCA}}$ distributions along the vertical is similar, they are different in the horizontal direction.

But what about the difference in the spatial distribution of $\bar{\lambda}_{\text{POCA}}$ observed in Figure 3.20 (left), where $\bar{\lambda}_{\text{POCA}}(A)$ is symmetrical with respect to the center of the cube, while $\bar{\lambda}_{\text{POCA}}(B)$ is not? This difference is the combination of two effects: the difference in luminosity across the horizontal direction and the correlation between the voxel horizontal position and the muon zenith angle. The zenith angle distribution of muons passing through regions A and B presented in Figure 3.21 clearly illustrates this last point, and shows that the range of zenith angle values gets smaller as we move away from the center of the de-

tector acceptance.




Figure 3.21: Muon zenith angle (left) and projected zenith angle in the XZ plane (right) corresponding to region A (center of the detector acceptance) and B (corner of the detector acceptance). As one moves away from the center of the detector acceptance, the range of zenith angle values gets smaller.

The consequence of the lumosity effect is determinant, as identical objects with the same material composition can exhibit different scattering densities, which is critical in the context of a material identification task that cares about the relative scattering density of objects within the VOI. From an inverse problem standpoint, it means that the scattering density λ associated with a voxel v located at a given position (x, y, z) not only depends on its material composition and the directions and momenta of muons passing through it (as established in Equation 3.23), but also on its position within the detector acceptance. As such Equation 3.23 must be modified:

$$\lambda_v = f(x, y, z, X_0, p, \theta), \quad (3.23)$$

with X_0 is the material radiation length, $p = [p_0, \dots, p_n]$ the momentum, $\theta = [\theta_0, \dots, \theta_N]$ the zenith angle of the N muons, and (x, y, z) denotes the position of the voxel within the detector acceptance. To our knowledge, the concept of lumosity has never been assessed in the literature, and no traditional reconstruction algorithm can leverage it during scattering density inference.

3.2.3 Gaussian smoothing

Section 3.2.2 showed that for the considered detector geometry, the lumosity reaches up to $33 \text{ cm}^{-3}\text{h}^{-1}$. For an acquisition time of 8 hours, it amounts to an integrated lumosity of 264 cm^{-3} . The resulting voxel-wise scattering density predictions are thus derived from a relatively low number of muon events, leading to non-negligible statistical uncertainty. As shown in Figure 3.6, the number of POCA events per cm^3 voxel does not exceed 20, making the estimation of the scattering density $\hat{\lambda}_{POCA}$ —based on the root mean square of the scattering angle—highly imprecise. Moreover, rare hard-scattering events can significantly bias the scattering density estimates, producing abnormally high values in a few voxels. This effect is clearly illustrated in Figure 3.5, where a small number of voxels exhibit extreme scattering densities, subsequently pulling down the remaining predictions. Naturally, this type of statistical noise depends on both the voxel size and the total number of detected muons. To obtain reliable evaluations of performance metrics such as MSE and MSSIM, these outlier-driven predictions must be regularized.

A straightforward method for mitigating this noise is to apply a smoothing procedure, in which each voxel's scattering density prediction is recalculated as a function of its neighbors. The Gaussian averaging operator is widely considered optimal for this purpose. In image processing, it corresponds to the convolution of an image with a Gaussian kernel. For a 3D volume, this operation is defined as:

$$G(x, y, z, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\left(\frac{x^2+y^2+z^2}{2\sigma}\right)}, \quad (3.24)$$

where x and y represent distances along the horizontal axes, z along the vertical axis, and σ denotes the standard deviation of the Gaussian distribution. Each voxel's new value becomes a weighted average of its local neighborhood, with the voxel itself receiving the highest weight and neighboring voxels contributing less as their distance increases. Figure 3.22 illustrates the impact of Gaussian smoothing on reconstructed images. As the value of σ increases, the predicted scattering densities become more homogeneous, reducing noise but




Figure 3.22: Example of 3D Gaussian smoothing applied on MST images, with increasing standard deviation values (in cm).

also blurring fine details and reducing image sharpness.

A question arises: what is the appropriate degree of smoothing? To explore this, we focus on the task of material identification, where the accuracy of the scattering density prediction is critical. For each of the algorithms described in Section 3.1, and using the simulated scenario detailed in Appendix A.2, we apply Gaussian smoothing to the normalized scattering density predictions and evaluate both the $\text{MSE}_{\text{signal}}$ and $\text{MSE}_{\text{background}}$. As shown in Figure 3.23, two trends emerge for all algorithms: as the standard deviation increases, the error on the scattering density prediction of material voxels decreases while the one of empty voxels increases. Gaussian smoothing operates as a weighted local averaging filter, which helps reduce noise by allowing each voxel’s value to benefit from neighboring values. As a result, the scattering density estimates in material regions become more robust, as the effective sample size increases through the inclusion of neighboring information. However, this benefit comes at the cost of blurring the material boundaries. At these interfaces, the inclusion of background voxel values in the averaging process leads to an underestimation of material voxel scattering densities, and vice versa. This loss of spatial precision ultimately reduces image sharpness, which is problematic in applications requiring accurate shape or feature identification. Notably, for a standard deviation $\sigma = 1\text{cm}$, matching the voxel size, all algorithms exhibit similar MSE values. While the minimum $\text{MSE}_{\text{signal}}$ is observed around $\sigma = 5\text{cm}$, this level of smoothing introduces excessive blurring: the reconstructed shapes lose their angular features and appear spherical rather than cubic. This illustrates the trade-off between denoising and structural fidelity. To better capture the perceptual quality of the images, we also compute the Mean




Figure 3.23: Mean squared error of the 1 cm voxels scattering density predictions for increasing Gaussian smoothing standard deviation values. The algorithms' names POCA, POCA RMS, BCA, and ASR refer to Equations 3.8, 3.7, 3.11, 3.14, respectively.

Structural Similarity Index (MSSIM), shown in Figure 3.24a. Two distinct behaviors emerge: for the POCA algorithms, the similarity index first drops for $\sigma \in [0, 1]$, and reaches a local maximum at $\sigma \approx 2$. For ASR and BCA, the similarity index improves until it reaches a maximum also at $\sigma \approx 2$. The MSSIM highlights that excessive smoothing degrades image structure, reaffirming that moderate smoothing offers the best compromise between noise reduction and feature preservation.

The insights provided by the MSE metrics demonstrate that applying Gaussian smoothing significantly reduces the error in the scattering density estimates for material voxels. A standard deviation of 1 cm (corresponding to one voxel) emerges as a practical compromise, yielding substantial improvements in MSE_{signal} across all algorithms, while preserving object geometry with minimal distortion. However, the optimal degree of smoothing is not universal as it depends on several factors, including the choice of reconstruction algorithm, the specific objective of the measurement (e.g., detection vs. identification), the detector's luminosity, the voxelization granularity, and the acquisition time. As a result, there is no one-size-fits-all value for the Gaussian filter's standard deviation. Instead, the optimal smoothing parameter must be tuned and validated for each specific use case. For MST tasks where a prior on the object structure is known, different Gaussian filters' standard deviation




Figure 3.24: MSSIM of the 1 cm voxels scattering density predictions for increasing Gaussian smoothing standard deviation values. MSSIM value of 1 indicates perfect similarity, and 0 indicates no similarity (a). Illustration of the image distortion for large values of Gaussian smoothing standard deviation (b): no smoothing (left) and 10 cm standard deviation smoothing (right) .

values can be used in different regions of the VOI to enhance the scattering density predictions while preserving the features of the image.

3.2.4 Noise in MST reconstruction

This section examines the limitations and artifacts present in MST-reconstructed images, which arise from multiple sources: measurement errors due to limited detector spatial resolution, insufficient muon statistics, and approximations in the modeling of muon trajectories through matter. Our goal is to qualitatively characterize the behavior of each reconstruction algorithm under these constraints, supported by the quantitative performance metrics introduced in the previous section. It is important to note that we do not seek to identify a universally optimal algorithm, as performance is inherently application-dependent and the simulated data used here do not capture the full range of possible MST imaging scenarios.

Smearing A notable feature of MST-reconstructed images is the presence of smearing, where the predicted scattering density in one voxel is influenced by neighboring voxels. This effect persists even with perfect muon track reconstruction. To illustrate it, we use the simulation setup described in Appendix A.2, placing a $10 \times 10 \times 10 \text{ cm}^3$ copper cube at the center of




Figure 3.25: Scattering density predictions from the POCA algorithm (left, Equation 3.7) and the ASR algorithm (right, Equation 3.14), averaged along the y -axis. The simulated cube location is outlined in white. Both algorithms exhibit smearing effects, but with distinct spatial distributions resulting from their differing path reconstruction assumptions.

the detector acceptance, positioned at $(x, y, z) = (0, 0, -440)$ mm. The detector panels are assumed to have ideal spatial resolution, ensuring flawless track reconstruction to better isolate intrinsic limitations of the reconstruction algorithms. Figure 3.25 presents the scattering density predictions from the ASR and POCA algorithms. Both display vertical smearing, though it is less pronounced for POCA. To quantify this effect, we plot the vertical scattering profiles in Figure 3.26. Notably, in both cases, the scattering density within voxels containing material is not uniformly distributed, despite the cube’s homogeneous radiation length. POCA tends to overestimate scattering in the lower half of the cube while underestimating it in the upper half. In contrast, ASR yields a more symmetrical distribution centered around the cube’s center $Z = -440$ mm. Additionally, voxels adjacent to the cube but containing no material also show non-zero scattering density values: POCA affects primarily the voxels below the cube, while ASR affects those both above and below.

Smearing also appears in the horizontal plane, albeit to a lesser extent. As shown in Figure 3.27 and 3.28, voxels near the cube that should be empty still receive non-zero $\hat{\lambda}$ values. Within the cube, the scattering density distribution remains non-uniform: POCA predicts higher values at the cube’s edges than at the center, whereas ASR shows the opposite trend. The difference in smearing behavior between the two algorithms stems from their underlying models of




Figure 3.26: Scattering density predictions from the POCA algorithm (left, Equation 3.7) and the ASR algorithm (right, Equation 3.14), averaged along the x and y -axis. The smearing is way more pronounced for ASR than for POCA.

the muon path. POCA-based algorithms assume that scattering occurs at a single point along the muon's trajectory. As a result, the accuracy of the inferred scattering location decreases with the length of material traversed, making the POCA estimate less reliable for deeper regions. Since muons predominantly travel from top to bottom, the upper portions of objects tend to register fewer POCA points than the lower regions, edges, or corners. In contrast, the ASR algorithm distributes scattering contributions across multiple voxels along the muon's path. Each voxel receives a weight proportional to the measured scattering angle. Due to the predominantly vertical muon flux, this distribution results in a smoother, vertically elongated scattering profile. This inherent difference in path modeling explains the distinct smearing patterns observed in the reconstructed images.

These observations highlight a broader issue: a voxel's predicted scattering density depends not only on its own radiation length and muon properties but also on the radiation length of surrounding voxels. Since the detected muon flux is predominantly vertical (with an average zenith angle of 20° , as shown in Figure A.3), this smearing effect is especially pronounced along the z -axis. As a result, Equation 3.23, which describes the factors influencing scattering density, should be extended to incorporate the influence of neighboring voxels' radiation lengths X_0^* :




Figure 3.27: Scattering density predictions obtained using the POCA algorithm (left, Equation 3.7) and the ASR algorithm (right, Equation 3.14), averaged along the z -axis. The simulated cube position is indicated in white. In both cases, the horizontal smearing is significantly less pronounced than the vertical smearing.




Figure 3.28: Scattering density predictions obtained using the POCA algorithm (left, Equation 3.7) and the ASR algorithm (right, Equation 3.14), averaged along the y and z -axis. The simulated cube position is indicated in white. In both cases, the horizontal smearing is significantly less pronounced than the vertical smearing.

$$\lambda_v = f(x, y, z, X_0^*, X_0, p, \theta) , \quad (3.25)$$

with X_0 is the material radiation length, $p = [p_0, \dots, p_n]$ the momentum, $\theta = [\theta_0, \dots, \theta_N]$ the zenith angle of the N muons, and (x, y, z) denotes the position of the voxel within the detector acceptance. This added complexity underscores the difficulty of the inverse problem in radiation length inference. None of the reconstruction algorithms presented in Section 3.1 currently account for this inter-voxel dependence, particularly in the absence of prior knowledge about the VOI's composition.

Detector resolution effects All reconstruction algorithms presented thus far rely on the same input: the muon's incoming and outgoing positions and directions. Consequently, the accuracy of the reconstructed scattering density inherently depends on the precision with which these quantities are measured. Measurement uncertainty in spatial or angular coordinates propagates through the reconstruction process and can significantly affect the final output. This raises a key question: what constitutes an acceptable spatial and angular resolution for a given detector configuration? In a fixed detector geometry, the angular resolution is governed by both the spatial resolution of the detection planes and the spacing between them. Determining a reasonable angular resolution, and understanding how reconstruction performance scales with it, are therefore critical for detector design and evaluation. To address these questions, we conduct a series of simulations based on Scenario 1 (described in Appendix A.2), systematically varying the spatial resolution of the detection planes. As established in the previous section, a Gaussian smoothing with a standard deviation of 1 cm is applied to all reconstructed scattering density predictions before normalization. The quality of each reconstruction is then assessed using three metrics: the MSE on voxels containing material ($\text{MSE}_{\text{signal}}$), the MSE on empty voxels ($\text{MSE}_{\text{background}}$), as well as the mean structural similarity index MSSIM, using the inverse squared radiation length as reference. Results are shown in Figure 3.29 and Figure 3.30, where the algorithms' names POCA, POCA RMS, BCA, and ASR refer to Equations 3.8, 3.7, 3.11, 3.14, respectively.




Figure 3.29: Scattering density MSE evaluated over material voxels (a) and empty voxels (b), taking the inverse squared radiation length as reference.




Figure 3.30: MSSIM evaluated on all scattering density predictions for various detector spatial resolution values and taking the inverse squared radiation length as reference.

Figure 3.29b highlights that the scattering density of empty voxels increases for all algorithms as spatial resolution degrades. This is a direct consequence of the reduced angular resolution: muon tracks that should appear straight are reconstructed as exhibiting scattering, thus assigning non-zero scattering angles to otherwise empty voxels. This trend is corroborated by the decline in MSSIM, which indicates a decreasing structural similarity to the ground truth. The behavior of voxels containing material, shown in Figure 3.29a, is more nuanced. For spatial resolutions in the range of 0–2 mm, the prediction error actually decreases for certain algorithms, particularly BCA and POCA. This is due to the fact that limited resolution "blurs" the POCA locations, smoothing out sharp fluctuations caused by the single-scattering assumption and resulting in more homogeneous predictions. While counterintuitive, this effect acts as a form of implicit regularization. However, beyond 2 mm, the reconstruction error increases steadily, ultimately surpassing the error in the ideal case. In contrast, ASR and POCA RMS, which rely more directly on scattering angle measurements rather than POCA clustering, show more stable performance over the range of tested resolutions. Nevertheless, for spatial resolutions exceeding 6 mm, the aluminum cube in the bottom-left corner of the volume becomes increasingly difficult to detect. At 10 mm, most material voxels are lost in the background noise for all algorithms except POCA RMS, which exhibits a relatively high tolerance to degradation in spatial and angular resolution, as illustrated in Figure 3.31.

As expected, the performance of MST reconstruction algorithms is strongly influenced by the spatial and angular resolution of the detector. To account for this, the scattering density dependency expressed in Equation 3.25 should be extended to explicitly incorporate the detector's spatial σ_{xy} and angular resolution σ_θ :

$$\lambda_v = f(x, y, z, X_0^*, X_0, p, \theta, \sigma_{xy}, \sigma_\theta), \quad (3.26)$$

Acquisition time The number of detected muons is a critical factor in Muon Scattering Tomography (MST), as low muon statistics result in high-variance and less reliable scattering density predictions. Each reconstruction




Figure 3.31: Scattering density predictions as 2D projection in the XY plane, for different detector spatial resolution.

algorithm responds differently to limited acquisition time, with varying levels of tolerance to sparse data. Importantly, the number of reconstructed muons is not determined solely by acquisition time, but also by the detector configuration—larger acceptances enable the detection of more muons per unit time. To evaluate the sensitivity of the algorithms to reduced statistics, we ran each of them on datasets corresponding to varying acquisition times and computed the MSE_{signal} , as shown in Figure 3.32. It is worth noting that simulations may underestimate the challenge of low-statistics scenarios, as we omit background contamination from other particles, errors due to wrong hits matching, non-uniform efficiency across the detection surface, and the distortion of the muon flux due to absorption by surrounding buildings.

As acquisition time decreases, the MSE_{signal} increases for all algorithms, primarily due to greater variance in the voxel-wise predictions. This effect is particularly pronounced for algorithms that rely on the spatial clustering of POCA, such as BCA (Equation 3.11), as demonstrated in Figure 3.33. When




Figure 3.32: Scattering density MSE evaluated over material voxels taking the inverse squared radiation length as reference, for different acquisition time values.




Figure 3.33: Normalized scattering density predictions obtained from the BCA algorithm for 8 hours (a) and 2 hours (b) acquisition times. The standard deviation of the scattering density predictions increases as the acquisition time decreases.




Figure 3.34: Scattering density predictions as 2D projection in the XY plane, for different acquisition times.

the number of POCAs per voxel drops below three, clustering-based metrics become unreliable, rendering BCA ineffective in low-statistics regimes as illustrated in Figure 3.34. In contrast, POCA RMS and ASR demonstrate superior robustness in low-statistics conditions. These methods maintain recognizable cube shapes even with reduced acquisition time. Regarding background noise, ASR, POCA, and POCA RMS exhibit a similar trend: increasing the number of muons leads to a reduction in noise levels across empty voxels.

As expected, the performance of MST reconstruction is highly sensitive to muon count. To reflect this dependence, Equation 3.26 should be extended to explicitly include the number of muons N contributing to each voxel's scattering density estimate.

$$\lambda_v = f(x, y, z, X_0^*, X_0, p, \theta, \sigma_{xy}, \sigma_\theta, N), \quad (3.27)$$

Throughout this section, we have shown that scattering density predictions are influenced by multiple factors, such as detector resolution, muon statistics, and reconstruction assumptions, that are difficult to correct analytically within current algorithmic frameworks. To address these limitations, the next section introduces a denoising strategy based on deep learning, aiming to improve image quality by learning to mitigate these complex, non-linear effects directly from data.

3.3 Deep-learning approach

The previous section demonstrated that, although the physics of muon propagation through matter is relatively straightforward, the associated inverse problem of scattering density reconstruction remains highly complex. Current reconstruction algorithms face multiple sources of noise, many of which cannot be addressed through analytical corrections alone. One notable challenge arises from the anisotropic nature of the muon flux: muons predominantly arrive from vertical directions, and the resulting lack of horizontal muons introduces asymmetries that cannot be resolved analytically—this directional bias is inherent to the measurement setup. To address these challenges, statistical reconstruction methods have been proposed. These approaches have shown promising performance in mitigating some reconstruction artifacts [101, 120], and can even incorporate prior information about the object’s composition. Nevertheless, they remain limited in their ability to compensate for all forms of noise and image degradation.

In such complex scenarios, where traditional methods reach their limits, machine learning offers a compelling alternative by uncovering subtle patterns and structures in the data. A variety of strategies can be employed, ranging from supervised learning, where models are trained to minimize the error between predicted outputs $\hat{Y}$ and ground truth Y , to unsupervised methods that extract compact representations of the input data in a latent space.

Several machine learning approaches have already been explored in the MST literature. In the context of material classification, random forests and boosted decision trees have been applied to the outputs of scattering algorithms to identify high-density materials [104, 125, 90]. A more general method was presented in [95], which uses t-distributed stochastic neighbor embedding (t-SNE), an unsupervised non-linear dimensionality reduction technique, to map MST data into a 2D latent space, thereby facilitating material classification tasks. Another promising avenue is the application of image super-resolution techniques, which are well established in both computed tomography [7] and general image processing [66]. These methods aim to enhance low-quality or noisy images by training models to recover high-resolution representations. Notably, the authors of [78] employed a conditional Generative Adversarial Network (cGAN) to upsample 2D MST images, achieving quality metrics comparable to those obtained after 21 days of acquisition from data collected in just one day.

While machine learning is sometimes perceived as a catch-all solution, it is important to maintain realistic expectations. Directly inferring a 3D radiation length map from raw muon measurement data (typically large sets of high-dimensional vectors) is highly ambitious and, at present, impractical. However, deep learning has proven quite effective for 3D image denoising across several domains. Recent architectures, particularly 3D U-Nets, have shown strong performance in removing Gaussian, Poisson, and speckle noise in medical imaging, point clouds, and volumetric reconstructions [127, 59], often surpassing classical approaches.

In consultation with the authors of [95], we concluded that such deep learning models could be well suited for the restoration task in MST, that is, learning to map a noisy reconstructed scattering density map to its underlying 3D radiation length ground truth. To this end, we initiated the development of a novel framework based on advanced 3D deep learning architectures, combined with synthetic data generation, to process full-volume scattering density maps at the voxel level. Convolutional Neural Networks (CNNs), particularly 3D CNNs, are well equipped for this task, as they excel at modeling noise distributions while preserving essential structural features. This work was initiated during a three-week research visit to the DLR Institute in October 2025. The first critical step, detailed in Section 3.3.1, involved generating a large, representative dataset—a time-consuming but essential task. The training of the 3D Skip-Connection Attention (SCA3D) model [92] is currently ongoing.

3.3.1 Data generation

Data generation is the first prerequisite for any machine learning workflow. In the context of muon scattering tomography (MST), generating a training dataset from real detector measurements is highly impractical. Acquiring a single MST image typically requires several hours to days, particularly for large objects. Consequently, assembling a sufficiently large and diverse dataset from real-world data is infeasible. To overcome this limitation, we rely on synthetic data, similarly to the methodology introduced in Appendix A.2. The synthetic dataset is constructed in several stages. First, detector-level data is generated using the GEANT4 simulation toolkit. This includes the true muon hits on the simulated detector planes. Next, reconstructed hits are computed by applying user-defined spatial and angular resolutions and detector efficiency, following

the approach detailed in Section 2.3.2. Finally, scattering density maps are produced using the reconstructed muon hits and track information. This step-wise process preserves the causal link between detector configuration and the resulting reconstructed image, ensuring that changes in resolution are consistently propagated through the pipeline.

To facilitate the modeling of complex geometries within GEANT4, we employ the B2G4 framework [96]. This Blender-to-Geant4 (B2G4) pipeline integrates Blender’s powerful 3D modeling capabilities [55] with GEANT4’s Monte Carlo simulation environment. Unlike conventional GEANT4 workflows, which require manual definition of object geometry and material properties, B2G4 streamlines this process. Complex, heterogeneous structures can be modeled in Blender and automatically imported into GEANT4, significantly reducing development time and enabling the mass generation of diverse scenes. Shapes, materials, and orientations are randomly sampled to create a wide range of scenarios. The corresponding voxelized ground truth for each scene is generated via a ray-tracing algorithm.

The final stage of the pipeline involves applying the scattering density reconstruction algorithms using the Muograph software, a Python-based tool developed during this thesis [98]. This modular and extensible software suite handles the entire reconstruction process and ensures compatibility with the data generated through B2G4. A schematic overview of the complete data generation and reconstruction pipeline is provided in Figure 3.35.

Both the ground truth and the predicted scattering density maps are represented as 3D images of size $64 \times 64 \times 64$ voxels, with each voxel corresponding to a physical size of 1 cm. These reconstructions reflect data collected over an acquisition time of approximately 8 hours. At this stage of the project, the dataset includes only a limited selection of shapes and materials, with a focus primarily on high-density objects embedded in an air background. We plan to introduce more diverse scenarios, including a broader range of materials and more complex geometrical configurations. In particular, future datasets will include scenes in which low-density materials are enclosed within higher-density ones, to better reflect the broad range of scenarios encountered in MST.




Figure 3.35: Sketch of the data generation pipeline.

Chapter 4

Muon momentum estimation

The direct physical processes underlying muon transmission and muon scattering tomography—primarily energy loss and multiple Coulomb scattering—are intrinsically dependent on the momentum of the incident particle. As a result, having prior knowledge of the muon’s momentum can substantially enhance the accuracy of the associated inverse problems: inferring material density in transmission tomography and scattering density in MST. However, unlike high-energy physics experiments that employ precise timing detectors or strong magnetic fields to estimate particle momentum, muon tomography setups typically lack such instrumentation. This absence introduces a fundamental limitation to the reconstruction accuracy achievable in practical applications.

To address this challenge, Section 4.1 explores how integrating a prior estimate of muon momentum can improve both scattering and transmission tomography, particularly in enhancing signal purity and the quality of the reconstructed inference. Building on this motivation, Section 4.2 investigates a range of machine learning approaches to regress muon momentum from detector-level data. These range from classical techniques leveraging engineered features with fully connected networks to more sophisticated architectures such as recurrent neural networks (RNNs), graph neural networks (GNNs), and transformers. The study also accounts for practical constraints, including realistic hit reconstruction efficiency and resolution, and emphasizes the importance of producing momentum estimates that can directly support and improve tomography reconstruction algorithms. The work presented in this section was carried out in close collaboration with Florian Bury and originated from a discussion we had during the Second MODE Workshop on Differentiable Programming for Experiment Design in 2022. In particular, I took charge of

detector design, simulation, data generation, and feature selection, while the training and choice of architectures are to be credited to Florian.

4.1 Muon Momentum in Muon Tomography

Both muon scattering and transmission tomography can benefit significantly from knowing the muon’s momentum, even if the measurement is imperfect. This section first outlines how momentum information improves signal quality and material discrimination. It then introduces a method for inferring muon momentum based on measurements of scattering through a known material.

4.1.1 Muon Momentum in Muon Transmission Tomography

In geoscience, muon transmission tomography is used in several large-scale experiments to probe the internal structure of kilometer-scale volcanoes, such as the Puy-de-Dôme in France [3], Mount Vesuvius in Italy [46], and La Soufrière in Guadeloupe [48]. Given the large thicknesses of rock that muons have to traverse before reaching the detection system, only high-momentum muons (roughly above 100 GeV) can traverse the object without getting absorbed and thus constitute most of the signal. On the other hand, low-momentum muons (5 GeV and below), which represent around 90% of the total muon flux, can scatter and enter the detector acceptance without interacting with the region of interest, thus mimicking high-momentum muons and biasing the muon flux absorption measurement, as shown in Figure 4.1. Since the flux of high-zenith-angle muons (near-horizontally) decreases by 2 to 4 orders of magnitude after a few kilometers of rock, it becomes comparable with the flux of the non-penetrating muons and makes the rejection of such events determinant. Current strategies to mitigate this problem include shielding detectors with lead blocks [46, 48], applying selection cuts based on track quality [79], or using coarse time-of-flight systems to discard slow-moving muons [86]. Nonetheless, directly measuring muon momentum within the detection system could offer a significantly more effective way to identify and reject non-penetrating muons, further enhancing the reliability of transmission-based re-

constructions.




Figure 4.1: An illustration of the different kinds of muons entering a muon detector and their acceptance: high-momentum muons traversing the target volcano (top) and low-momentum muons entering the detector by coincidental scattering, where the muon’s direction is indicated by the arrow (middle and bottom).

Measuring muon momentum can also improve transmission tomography in scenarios involving small-scale objects. Chapter 5 introduces a new application of muon transmission imaging for cultural heritage preservation, where the objects of interest are typically only a few meters in size. In this regime, muons traverse relatively thin layers of material and thus experience limited energy loss. For example, a 1 GeV muon has a continuous slowing-down approximation (CSDA) range of over 2 meters in standard rock (density 2.65 g cm^{-3})¹, which is significantly thicker than most cultural heritage objects. Since about 75% of the cosmic muon flux consists of muons with energies above 1 GeV, only a small fraction of muons are actually absorbed within such objects. This limits the contrast available from transmission-based attenuation measurements and makes accurate density reconstruction more difficult, especially when acquisition times are constrained to a few days. In this scenario, the inclusion of a muon momentum measurement module in the detection system can also be beneficial as it allows for computation transmission ratio over a subset of the cosmic muons (1 GeV and below), which are typically more

¹CSDA range of 1 GeV muon in standard rock is $5.534 \times 10^2 \text{ g cm}^{-2}$, corresponding to $2.088 \times 10^2 \text{ cm}$. CSDA values from [83].




Figure 4.2: Scattering angle distribution of muons through 10 cm of Aluminum and Iron, with realistic momentum distribution (left), and scattering angle distribution multiplied by the muon momentum (right).

likely to get absorbed within such thickness.

4.1.2 Muon Momentum in Muon Scattering Tomography

Muon momentum measurement is also advantageous for muon scattering tomography (MST), though for different reasons than in transmission tomography. Unlike transmission imaging, MST is less impacted by background contamination from low-energy muons, such as those shown in Figure 4.1, since muons are detected both before and after passing through the VOI. However, momentum information plays a crucial role in material discrimination. In MST, scattering angles are used to infer the material properties within the VOI. As shown in Figure 4.2, the raw scattering angle distributions for cosmic muons passing through iron and aluminum significantly overlap, making them difficult to distinguish. When these angles are scaled by the muon momentum, the resulting distributions become more clearly separated, enabling better material identification. Although precise momentum measurements are difficult to obtain, even a coarse estimate can meaningfully enhance image quality. For instance, Reference [11] demonstrates that incorporating momentum information into reconstruction algorithms can considerably reduce the data acquisition time required for complex targets, such as nuclear fuel casks.

4.1.3 Muon momentum measurement

Momentum measurement of charged particles is a fundamental aspect of many particle physics experiments. In high-energy physics, such as at the LHC, momentum is typically inferred from particle trajectories in strong magnetic fields. However, this approach is not feasible for muon tomography. Generating a high magnetic field requires large, complex, and power-intensive magnets, introducing substantial logistical and financial challenges. Moreover, magnetic fields can interfere with image reconstruction by making it more difficult to distinguish between deflections caused by multiple Coulomb scattering and those induced by the magnetic field itself. To overcome this limitation, two main techniques have been explored in the literature. One approach relies on detecting Cherenkov radiation emitted by muons in specialized detectors [10, 22], and the other infers momentum based on multiple Coulomb scattering through passive layers of dense, high- Z materials such as lead or iron [117, 110, 79, 5]. Figure 4.3(a) illustrates a simulated momentum measurement module composed of three dense scattering layers (e.g., lead) of thickness Δs , interleaved with eight detector planes separated by a gap Δg . These detectors are implemented as infinitely thin detectors that record the exact position of the muon and are assumed to contribute negligible scattering, which is an acceptable approximation, as real detectors are generally thinner and less dense (e.g., gas mixtures or scintillators) than the passive materials. The analysis is restricted to muon tracks that register a hit in all eight detector layers. From these eight hits, four linear track segments can be reconstructed, allowing the calculation of three scattering angles $\Delta\theta_i$. Contrary to the typical muon scattering tomography task, where one wants to infer the density from the measured scattering angle, momentum can be estimated by inverting Equation 1.9:

$$p = \frac{13.6 \text{ MeV}}{\theta_0} \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \left(\frac{x}{X_0} \right) \right], \quad (4.1)$$

where muons are assumed to be relativistic ($\beta \simeq 1$). In this context, the scattering material X_0 is pre-defined, and the distance traversed by the muon x is estimated from the track measurement. Measuring the scattering angle of the

muons through N layers of material gives a rough estimate of the scattering angles' standard deviation θ_0 :

$$\theta_0 = \frac{1}{\sqrt{2}} \theta_{RMS}^{scatt} = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^N (d\theta_i^{scatt})^2}. \quad (4.2)$$

Together with Equation 4.1, the scattering angles' standard deviation θ_0 provides an estimation of the muon momentum, as illustrated in Figure 4.3 (b). The mismeasurement tail at low momentum originates from the fact that the muon energy loss through 4 cm of lead has the same order of magnitude as its momentum, for example, a 200 MeV muon loses about 55 MeV through 4 cm of lead. This could be circumvented by the use of Kalman filters, as demonstrated in Ref. [6]. While such an approach has satisfactory results, it remains suboptimal as only a fraction of the information accumulated by the detector is used. Only scattering angle information can be plugged in Equation 1.9, representing three measurable parameters, while the detectors measure the muon position at eight locations. In Section 4.2, different machine learning algorithms are trained to improve the momentum prediction accuracy. The dataset used during the training follows the setup of Figure 4.3(a). The CRY library is used to generate the muon flux, and the muon propagation and scattering are handled by GEANT4.




Figure 4.3: (a) The experimental setup used in this paper: three layers of $\Delta s = 4$ cm thick lead panels serve as scattering material, separated by $\Delta g = 25$ cm of air, and the hits are detected by eight two-dimensional planes—agnostic of the detector technology. (b) The result of the muon momentum inference using Equation 4.1. The Mean Absolute Percentage Error (MAPE) is shown at the top of the Figure. Reproduced from [39].

4.2 Muon momentum inference via machine learning methods

Although physically motivated, using the sole scattering angles in Equation 1.9 represents a limited use of the information available from the detector. At the lowest level, all the information that can be gathered is contained within the hits registered in the different active layers. The ambition of this section is to explore different ML algorithms and compare their performance.

Early machine learning (ML) applications in particle physics relied on boosted decision trees (BDTs) and shallow neural networks (NNs)[97, 36]. These models performed well with limited training data but required physicists to hand-craft physically motivated input variables, such as invariant mass or angular distances, to enhance classification performance. This is analogous to using scattering angles derived from Equation 4.1. With the advent of deep learning and the availability of high-fidelity simulations, physicists began supplementing these expert features with low-level detector data, improving performance significantly, especially in tasks like jet tagging at the LHC [71]. This opened new possibilities for measurements previously considered impractical.

A further leap came when model architectures began to reflect the physical symmetries of the problem. Convolutional neural networks (CNNs)[61] exploited translational symmetry by modeling jets as 2D images. Recurrent neural networks (RNNs)[57, 14] and recursive NNs [23] leveraged hierarchical structures of jet constituents. However, CNNs suffer from resolution loss due to pixelation, and RNNs face challenges with long-range dependencies, despite advances like LSTMs and GRUs. To overcome these limitations, newer models have embraced more suitable symmetries. Notably, ParticleNet[88], a graph neural network (GNN), handles the sparsity of jet constituents more naturally, while ParT[89], a transformer-based model [118], uses self-attention to capture dependencies without fixed ordering. Both models naturally incorporate permutation symmetry, aligning well with the unordered nature of particle sets in jets.

This section mirrors this historical progression: we begin with DNNs and gradually explore more advanced architectures like RNNs, GNNs, and transformers, each offering improvements in handling low-level data and model symmetry. All models are implemented in PyTorch[84], using PyG[42] for GNNs and trained with Lightning[40]. Hyperparameters are detailed in Appendix A.5 (Tables A.1–A.4). Training used the RAdam optimizer [65], a batch size of 1024, MSE loss, and a ReduceLROnPlateau scheduler, with learning rates of 10^{-3} or 10^{-4} . Roughly one million muons were generated using the setup in Figure 4.3(a); 70% were used for training, 10% for validation, and 20% for testing and visualization to ensure unbiased evaluation.

4.2.1 Deviations as features

As a first step beyond the analytical approach of Equation 4.1, we train a deep neural network (DNN) to regress the muon momentum using selected input features. Since the momentum distribution, shaped by the flux in Figure 1.4 and modulated by detector acceptance, spans over several orders of magnitude, the model is trained to predict $\log_{10} P$, a choice further discussed later. For direct comparison with Figure 4.3(b), the initial inputs are limited to the three scattering angles shown in Figure 4.3(a). These inputs are transformed via $\log_{10}$, then standardized by mean subtraction and scaling to unit variance. As shown in Figure 4.4(a), the DNN already outperforms the analytical estimate from Equation 4.1, demonstrating its ability to extract more information

from the same inputs. Further performance gains are achieved by enriching the input set with additional features, such as deflection angles and track residuals (Figure 4.5). Using the same DNN architecture and preprocessing steps, this expanded input set significantly improves regression accuracy, as illustrated in Figure 4.4(b).




Figure 4.4: Muon momentum regression using a DNN with a different set of features: only scattering angles (a); scattering angles, deflection angles, and fit residuals (b), as shown in Figure 4.5. Reproduced from [39].




Figure 4.5: An example of variables that can be extracted from muons' hits in the detector layers. Green crosses represent measured hits, dashed line segments join two consecutive hits, the dotted line is the track linear fit, and the solid rectangles are the scattering material blocks. **(Left)** Deflection angles are the angles between two consecutive segments. **(Middle)** Scattering angles are the angles between two segments on each side of the scattering material. **(Right)** Residuals are the distances between the hits and their orthogonal projections on the fitted linear track. Reproduced from [39].

Applying a $\log_{10}$ to both inputs and targets came from the empirical observation that spans several orders of magnitude, and that the angles scale non-linearly from it. High-momentum muons suffer very few deflections, and it becomes hard for the DNN to distinguish them above 100 GeV. The $\log_{10}$ was then chosen to disentangle them, effectively “zooming in” on that high-momentum region. This, however, should not come as a surprise since Equation 1.9 implies that $\log_{10} \theta \propto -\log_{10} P$, as shown in Figure 4.7.




Figure 4.6: Representation of the muon’s deviations in 3D (a) and in XZ and YZ projections (b). The hits measured by the detector planes are used to fit a linear track and then projected onto that track. Reproduced from [39].




Figure 4.7: The distributions of the different DNN features as defined in Figure 4.5 as a function of the momentum: scattering angle (a), deflection angle (b) and the fit residuals (c). The result of a linear fit on the logarithmically scaled variables is shown in red. Reproduced from [39].

The results in Figure 4.4(b) indicate that raw detector hits may provide richer information for momentum inference. However, the key information is encoded in how the muon scatters relative to an ideal straight-line trajectory, as described by Equation 1.9. High-energy muons produce nearly linear tracks, whereas low-energy muons exhibit greater deviation, often forming curved or helical paths around the fitted track, as illustrated in 3D in Figure 4.6(a) and projected in Figure 4.6(b). This behavior reflects the fact that muons scatter within a cone defined by a zenith angle from Equation 1.9, and a uniform azimuthal distribution, giving rise to a rotational pattern around the track axis. To capture this, rather than using only the scalar distance from each hit to the fitted track, we use the full 2D vector in the plane orthogonal to the track, which we refer to as the deviation.

4.2.2 Deviations as a Sequence

To feed the deviations into a recurrent neural network (RNN), an ordering must be defined. A natural choice is to follow the muon’s trajectory from top to bottom. To exploit the azimuthal symmetry in Equation 1.9, deviations are expressed in cylindrical coordinates (r, θ) relative to the track direction, which defines the z -axis. This coordinate choice improves regression performance in sequential models. To address the logarithmic dependence of deviation magnitude on momentum (see Figure 4.5), we apply the log-modulus transformation: $f(x) = \text{sign}(x) \log_{10}(|x| + 1)$. Unlike a pure logarithm, this function preserves the sign and avoids distortion near zero, effectively enhancing resolution for high-energy muons without losing shape information, as illustrated in Equation A.4. Finally, tracks are rotated so the initial deviation has $\theta = 0$, standardizing the orientation. The ordered sequence of (r, θ, z) is passed through the RNN, and the final hidden state is fed into a DNN in a many-to-one setup. The resulting regression performance is shown in Figure 4.8(a).




Figure 4.8: Muon momentum regression using advanced ML algorithms. Reproduced from [39].

Unlike RNNs, transformers are naturally permutation-invariant, so positional information must be injected to preserve the sequential order of deviations along the muon track. In language processing, this is typically done via positional encodings—originally using sine and cosine functions, later extended to learnable or relative encodings. To fully leverage the 3D geometry, the hit positions projected onto the linear track are processed by a DNN to produce a learnable position encoding. Likewise, the deviations (r, θ) are passed through another DNN to produce feature embeddings. Both position and feature inputs are batch-normalized, summed, and passed to a transformer encoder. Following the CaiT architecture [115], a summary token interacts with the encoder output through multiple cross-attention layers before being passed to a final DNN. The resulting regression performance is shown in Figure 4.8(b).

Viewing deviations as a sequence also opens the door to identifying low-momentum muons in transmission tomography, as they show increasing scattering toward the exit of dense materials. This trend is partially visible in Figure A.4(d) for low-momentum muons. However, the effect is negligible for high-energy muons, limiting its practical utility. Since this relates more to background rejection than to momentum inference, it falls outside the scope of this study.

4.2.3 Deviations as a geometry

To use CNNs with geometric data, the deviations from Figure A.4 must be discretized into pixels. However, due to the wide dynamic range of deviations

(from 0.1 mm at 100 GeV to several cm at 1 GeV, as shown in Figure 4.7), achieving sufficient resolution would require thousands of pixels per dimension, exceeding memory limits for 3D CNNs. A workaround using 2D projections (x - y and x - z), inspired by multi-view CNNs [113], reduced memory by allocating one pixel per layer along z . Even so, two resolutions were needed per view to cover the full momentum range. While this approach yielded reasonable regression results, it barely outperformed Equation 4.2, likely due to projection, binning losses, and image sparsity. Still, it confirmed the geometric signal could be harnessed. To better handle sparsity, a graph neural network (GNN) was used instead. Hits were rotated and translated so the track aligns vertically at $(x = 0, y = 0)$, and the log-modulus was applied to the radial deviations to preserve shape while compressing scale. Hits became graph nodes, with edges linking each node to its nearest neighbors (typically two, except at the edges). Figure A.5 shows sample graphs with and without the log-modulus transformation.

Convolutions in a GNN can be reformulated in the message passing scheme, as implemented in `PyG`, where edge features of $\mathbf{x}_i^{(k)}$ at node i and layer k can be obtained from the previous layer $k - 1$ as

$$\mathbf{x}_i^{(k)} = \gamma^{(k)} \left(\mathbf{x}_i^{(k-1)}, \bigoplus_{j \in \mathcal{N}(i)} \varphi^{(k)}(\mathbf{x}_i^{(k-1)}, \mathbf{x}_j^{(k-1)}, \mathbf{e}_{ij}) \right), \quad (4.3)$$

where $\mathcal{N}(i)$ is the set of neighbours of node i , $\mathbf{e}_{ij}$ is the edge attribute between nodes i and j , $\bigoplus$ is a differentiable aggregation operation, and γ and φ are optional, complex, and differentiable functions, typically DNNs. In particular, for this application, the model used is `edgeConv` [121], specifically designed for point cloud convolutions, where Equation 4.3 is particularized to

$$\mathbf{x}_i^{(k)} = \bigoplus_{j \in \mathcal{N}(i)} \varphi^{(k)}(\mathbf{x}_i^{(k-1)} || \mathbf{x}_j^{(k-1)} - \mathbf{x}_i^{(k-1)}). \quad (4.4)$$

In this particular message passing mode, there is no edge attribute and the DNN φ acts on the concatenated tensor of the node tensor and its difference with other connected nodes. In our graphs (Figure A.5), node features correspond to hit positions, and mean pooling is used for aggregation. As in [121], outputs from each layer are concatenated using "jumping knowledge" [126], and residual connections are added to aid training [64]. For the final graph-to-vector step, various permutation-invariant pooling strategies were tested (mean, min,

max, std, or combinations). However, this permutation symmetry may limit performance, as seen in Figure 4.8c. More expressive, order-aware aggregation methods [19, 52] could offer further improvements.

4.2.4 Model Comparison

The real measure of the regression performance, besides the prediction bias, resides in the prediction resolution. This is illustrated as a function of the momentum in Figure 4.9 for the different models developed in this section. The sizes and inference times of the different models are in Table 4.1. The best model, the transformer (TNN), has a relative resolution below 10% and an inference time not far from the DNN thanks to GPU parallelization.




Figure 4.9: The relative resolution of the true momentum as a function of the predicted momentum. The true momentum resolution ΔP is estimated from the standard deviation of the true momentum distribution for muons whose predicted momentum falls within a range around the x -axis marker within 10%. The uncertainty band illustrates the limited number of muons used to calculate the standard deviation. Reproduced from [39].

Table 4.1: A comparison of the size of the different ML models in terms of the number of parameters and memory, and inference time per muon (in a batch size of 1024 on a GPU).

Model	# params	Size [MB]	Time [μ s]
DNN	19 K	0.07	12
RNN	165 K	0.63	50
TNN	420 K	1.61	17
GNN	100 K	0.40	25

4.3 Real-Life Experiment Considerations

Besides good performance, there are a few other considerations for an ML model to be truly useful in a real-life environment. Section 4.2 only focused on an ideal detector, assuming all hits were detected and their position measured with infinite precision. In addition, the predicted momentum does not come with any prediction resolution, which would improve muon tomography inference. This section will address how these experimental factors can be measured or improved by re-training the models of Section 4.2 with the same simulated data, but this time incorporating a more realistic detector response.

4.3.1 Hit Detection Efficiency

Missing hits must be considered when inferring momentum. One solution is to veto tracks that don't leave hits in all detector layers, but this reduces usable data. For instance, with eight layers and 95% efficiency, a typical value in muon tomography, about 34% of tracks would be discarded. Allowing tracks with missing hits, especially in the last layers, could increase acceptance and improve statistics. To evaluate this, modified datasets were created from the original training set in Section 4.2, where one or more hits are randomly removed and marked as missing. A veto is applied to tracks with fewer than six hits. Models are trained on tracks with zero, one, or two missing hits. The impact on performance is shown in Figure 4.10.




Figure 4.10: The relative resolution of different models as a function of the predicted momentum. Each ML model has been trained on a combination of datasets with 8 (all hits detected), 7, and 6 hits. In datasets with missed hits, the track fit, computation of angles in Figure 4.5, and deviations such as in Figure A.4 and A.5, are performed with detected hits only. When two hits are missed, in around 50% of the cases, the formula of Equation 4.1 cannot be applied as no scattering angle can be measured, and these muons are excluded. Reproduced from [39].

Any hit missed on either side of a scattering material block will make it impossible to measure the scattering angle, but Equations 4.1 and 4.2 can still be used with fewer values, albeit with limited precision. However, in some unfortunate cases, no scattering angles can be measured and no momentum can be inferred, which is another limitation of that method.

The DNN model is more resilient to missed hits because more deflection angles and residuals can be measured. The main drawback of DNNs is that the number of input features is fixed. When a hit is missed, one alternative is to use a zero value for the features associated with the missed hit; this is called *zero-padding*. An increase in prediction resolution, as observed in Figure 4.10, is expected from the loss of information, but the mixed training setting seems to be able to mitigate that.

The transformer, as well as the GNN, can handle variable-length inputs and does not necessitate zero-padding. There is, however, a larger degradation in the resolution of the transformer compared to the DNN, which indicates that the model struggles to handle the change in the deviations' representation.

Comparing the case with no missed hit with Figure 4.9 additionally shows that the mixed training causes confusion. This seems to be linked to the cylindrical representation, and especially the polar angle, on Figure A.4 (d) that presents large variations when the track parameters obtained from a fit with fewer hits are modified. The GNN, on the other hand, even if less performant overall, suffers less from that effect as it uses a Cartesian representation (Figure A.5 b). The DNN is also less affected because it relies more on the deviation angles.

4.3.2 Hit Spatial Resolution

To simulate the hit resolution, a Gaussian smearing is applied to the hit position in both the x and y directions. The width is chosen as 1 mm. This is a typical, albeit slightly pessimistic, value for resistive plate and drift wire chamber detectors. The features and deviations are extracted from these smeared detector hits, and the different models are then trained on this modified dataset. The regression results are shown in Figure 4.11. The loss of performance matches the expectation from Figure 4.7c, as even for a muon with infinite momentum, the smeared hits will emulate deviations of a muon with a momentum of a couple dozen GeVs. While ML models perform better than Figure 4.11b because they leverage more features, they hit a plateau relatively early in the training, and it is unlikely that the information loss from the training can be recovered. In the transformer case in Figure 4.11c, if the training is performed for longer, the overall performance decreases, but it is able to predict higher-momentum muons.




Figure 4.11: The result of the regression for a detector with 1 mm spatial resolution using Equation 4.1 (a), the DNN of Section 4.2.1 (b), and the transformer of Section 4.2.2 (c). Reproduced from [39].

To go beyond the stage of the proof of concept, there are several avenues that should be addressed. First, the hyperparameters of the advanced algorithms have been manually selected and could be further optimized by a wider automatic search. In particular, the graph neural network presented here is only one model out of a wide range of models proposed in the literature. It might also be worth exploring whether both the sequence and geometry aspects of the deviations could be tackled by a combined model. Second, the use of the deviations from a linear track may have been successful, but it still represents a processing of the raw hit position data that the ML algorithms might not need. There might be a way to relax that step to improve both performance and inference time. Third, the use of the quantile regression can be used to obtain an estimate of the prediction resolution, and by extension, a point-wise description of the posterior distribution. This would probably be better tackled by a generative model more suited to probability density estimation.

Chapter 5

Cultural heritage study

This chapter introduces a novel application of muon scattering tomography in the context of cultural heritage objects. While the application of MST in large archaeological structures has already proven to be successful, we tackle the problem of smaller structures, where the absorption of muons through the object alone is not sufficient to give conclusive results.

In recent decades, imaging techniques based on elementary particles have been extensively employed in the analysis of cultural heritage. The most widely used methods rely on X-rays [72], enabling the visualization of sculptures and artifacts to reveal their internal composition, including different material volumes, joints, past restorations, inner structures, and reinforcements. X-ray imaging is also used for non-invasive examination of mummies, reliquaries, and tombs, as well as for dating wooden artifacts through dendrochronology. However, X-ray techniques have significant limitations, particularly when applied to large objects or those composed of dense materials such as compact stone or metal, where standard X-ray penetration is insufficient. While higher-energy X-rays in the MeV range and neutron imaging can partially address this issue, their applicability remains constrained by the need to transport valuable objects to specialized imaging facilities—an option often impractical due to factors such as size, weight, fragility, poor preservation state, or the risk of damage during transport. These challenges affect both movable and immovable heritage, including monuments, monumental statues, architectural decorations, and individual sculpted objects. Nonetheless, imaging remains essential for a deeper understanding of their composition, construction techniques, past interventions (such as gluing, reinforcements, and additions), current condition, and interaction with the environment.

Table 5.1: The three scenarios for scattering and transmission imaging explored in this study.

Scenario	Volume	Size [cm ³]	Technique
1	Statue	$80 \times 30 \times 30$	Scattering
2	Statue	$160 \times 60 \times 60$	Scattering and transmission
3	Limestone block	$70 \times 70 \times 40$	Transmission

This chapter explores muon tomography as a promising alternative for cultural heritage studies, extending its application into the relatively unexplored domain of moderately sized objects, where conventional radiation-based methods fall short. We propose that both scattering and absorption muography offer complementary advantages in terms of detector cost, imaging capability, and the scale of the target objects. To illustrate the potential and limitations of muography in this context, we present a preliminary imaging study based on simulated muon tomography data. Section 5.1 presents the three imaging scenarios and describes the choices of volume of interest and their materials. A description of the simulation framework and the simulated detection setup is also provided. Section 5.2 presents the performance of muon tomography for the three scenarios, respectively. The simulation study presented in this chapter was carried out in collaboration with Marwa Al Moussawi, who should be credited for the implementation of the `GEANT4` simulation. A more comprehensive version of this study is available in [47].

5.1 Simulation study

To investigate the potential of muography for material identification, we consider two volumes of interest: a marble statue containing a hidden cylinder made of bronze and a block of limestone with an anomalous region of high water content. These models are used in three imaging scenarios, as shown in Table 5.1. Their implementation in the `GEANT4` software and the simulated detection system are further described in this section.

Statues In order to have a complex and realistic shape to use as representative of a typical statue, the geometrical model of the African statue shown in Figure 5.1 was implemented in GEANT4 as a mesh file obtained from a high-resolution CT scan (voxel size 121 μm) performed at the HECTOR scanner of the Ghent University Center for X-ray Tomography [68]. Although the actual statue is uniformly composed of wood, marble was selected as the primary reference material for this study due to its higher density. This choice allows for the discernment of internal features even in a relatively small statue while also enabling realistic scaling, as large marble statues are common in cultural heritage. Wood, on the other hand, is very light and is thus challenging for muography, and we used it to explore the limits of the method. As GEANT4 conveniently allows users to implement custom composite materials, we implemented a typical marble composition with relative mass composition: 40% CaO , 25% SiO_2 , 30% MgCO_3 and 5% Al_2O_3 , and with a density of 2.71 g cm^{-3} [119]. To investigate the potential of muography for material identification, which may be relevant, e.g., to identify the inner tube system of a monumental fountain, we simulate the insertion of hidden cylinders of bronze within its internal structure. Two statue models were generated to assess the performance of muon scattering and transmission tomography under different conditions: one maintaining the original statue size and another scaled by a factor of two. Correspondingly, the radius of the inserted bronze cylinders was set to 5 cm in the original model and 10 cm in the scaled model. The scattering and transmission simulation setups are sketched in Figure 5.2, and possess an angular resolution of 2.6 mrad.

Limestone block The ability to assess the moisture content of sculptures and structures is crucial for both preventive conservation and the development of appropriate treatment strategies. Water plays a key role in various degradation mechanisms, including freeze-thaw cycles, salt crystallization, biological growth, hygric/hydric dilation, corrosion, and the reduction of mechanical strength. In this context, muography presents a valuable tool as an in situ, non-destructive remote-sensing technique that provides direct sensitivity to density—and, consequently, to moisture content. In cases where the volume is not altered and only mass is changing (as in the case of stones), adding moisture




Figure 5.1: Picture of the wooden statue at the AfricaMuseum of Tervuren that we used as a model. From project ToCoWo (<https://tocowo.ugent.be/>). Reproduced from [47].




Figure 5.2: GEANT4 simulation setup; green and blue panels represent the scattering and absorption setups, respectively. The horizontal surface from which the muons are generated is represented by the red line. Reproduced from [47].

to a stone object of density ρ_{dry} results in a new apparent density ρ' according to:

$$\rho' = \rho_{\text{dry}} \times (1 + \text{MC}), \quad \text{MC} = \frac{\rho_{\text{water}}}{\rho_{\text{dry}}} \times \text{porosity} \times \text{saturation}, \quad (5.1)$$

with MC the mass fraction of water with respect to the dry material mass, i.e. $M_{\text{water}} = \text{MC} \times M_{\text{dry}}$. The porosity is defined as the ratio of the pore volume to the stone volume [vol.%], and the saturation is the ratio of the volume of water to the pore volume [%]. The target volume to image is a $70 \times 70 \times 40 \text{ cm}^3$ block of calcium carbonate limestone with 40% porosity and 1% saturation, which translates to a 0.5% moisture content. The exact composition that we simulated consists of 99.5% CaCO_3 and 0.5% H_2O , with the percentages given in fraction of mass, and corresponds to a density of 1.63 g cm^{-3} . Inside the limestone block is placed a 15 cm radius sphere of the same material but with a higher water content, corresponding to an 80 % saturation and thus a 36% moisture content. The exact composition that we simulated consists of 64% CaCO_3 and 36% H_2O , with the percentages given in fractions of mass. The resulting density of the high moisture content sphere is 1.95 g cm^{-3} . The simulation setup is sketched in Figure 5.3.

5.2 Small scale objects: Muon Scattering Tomography

In this initial scenario, the simulated statue measures $80 \times 30 \times 30 \text{ cm}^3$, which is a relatively small size by muography standards and makes the scattering method the most suitable. To ensure adequate image resolution and consistency with the MST characterization presented in Section 3.2, a voxel size of 10 unit mm is chosen. Unlike the cube dataset used in Section 3.2, where the objects were small relative to the detection panels, the statue occupies nearly the full volume of the detector's acceptance. As a result, the extremities of the statue (head and feet) fall in regions of lower lumuosity compared to the central body, which lies within the detector's highest lumuosity region. This




Figure 5.3: GEANT4 simulation setup for the third scenario. The horizontal surface from which the muons are generated is represented by the red line. The parallelepiped and spherical volumes contain calcium carbonate limestone with two different water content values and respective densities 1.63 g.cm^{-3} and 1.95 g.cm^{-3} . Reproduced from [47].

non-uniform muon flux introduces spatial bias in the scattering density predictions obtained from algorithms that care about the spatial clustering of the POCAs. To mitigate this, we avoid algorithms that rely on POCA clustering (e.g., POCA and BCA), which are particularly sensitive to such variations in muon statistics. Instead, we employ the POCA RMS algorithm (as defined in Equation 3.6), which emphasizes scattering angle magnitude rather than spatial clustering. Additionally, a Gaussian smoothing with a standard deviation of $\sigma = 2 \text{ cm}$ is applied to enhance image quality and reduce prediction variance. The predicted scattering densities are shown in Figure 5.4.

Eight hours of data acquisition are sufficient to discriminate between marble and bronze. The cylinder's position and dimensions can be retrieved. The true position of the 10 cm long and 5 cm radius cylinder is $x, y, z = (0, 0, -71 \text{ cm})$.




Figure 5.4: Normalized scattering density predictions obtained from the POCA algorithm (Equation 3.6) for the small-scale statue. Results are shown as slices of 3 voxel layers along the vertical direction z .

5.3 Medium scale objects: Scattering versus Transmission Tomography

A relevant question for planning a muography investigation of an object is whether scattering or absorption is the most appropriate method. Based on which is chosen, the arrangement of detectors is rather different as illustrated in Fig 5.2. The simulated statue's size is $160 \times 60 \times 60 \text{ cm}^3$, which is a challenging regime for both absorption and scattering muography. Because the scattering technique requires a set of detection planes above and below the object, the overall price scales faster with the detection area than with the absorption technique. The statue's thickness is 60 cm and is rather thin by absorption muography standards, as muons lose only a small fraction of their kinetic energy, making their absorption within the object unlikely.

5.3.1 Transmission tomography

The thickness of the object being rather small compared to absorption muography standards, it is rather unlikely for a muon to get absorbed within the object. To compensate for this effect, only low-energy muons with kinetic energy $E_{\text{kin}} < 1 \text{ GeV}$ were selected. Although the current detector setup does not permit direct measurement of the muons' kinetic energy, several techniques have been proposed in the literature. Some rely on detecting Cherenkov radiation using dedicated detectors [10, 22], while others exploit the multiple Coulomb scattering of muons in passive layers of high-Z, high-density materials [117, 90, 79, 5], as discussed in Section 4. In this study, we assume that momentum selection is achieved using an absorption-based filtering technique, as described in [100], which is both cost-effective and practical for low-energy muon selection. The reconstruction of the density map within the volume of interest (VOI) follows the routine detailed below:

- **Pre-processing:** Filter out muon event with $E_{\text{kin}} < 1 \text{ GeV}$.
- **Reconstruction algorithm:** The voxel-wise density predictions ρ are computed as the ratio between the recorded muon flux through the object and the free sky flux $\rho = N_{\text{object}}/N_{\text{freesky}}$. Both N_{object} and N_{freesky} are estimated using a custom back-projection algorithm inspired by [15]




Figure 5.5: Density predictions as the muon transmission ratio for the absorption technique and using the back projection algorithm, for Scenario II. Predictions are shown by slices of 3 voxel layers along z (vertical axis). Reproduced from [47].

and [109], by simply counting the number of muons passing through each voxel.

- **Post-processing:** The 3D transmission map is also smoothed using a 3D Gaussian smoothing kernel with standard deviation $\sigma = 1$ voxel, and are presented in Fig 5.5.

The presence of the bronze cylinder can be identified in Fig 5.5. The true position of the 20 cm long and 10 cm radius cylinder is $x, y, z = (0, 0, -31 \text{ cm})$. While the image resolution in the horizontal (x - y) plane is satisfactory, the reconstruction is noticeably smeared along the vertical (z) axis. This vertical blurring prevents the precise localization of the cylinder's position along that dimension. Such an effect can be circumvented by combining reconstructed images with different rotations, at the cost of extending the data-taking period.




Figure 5.6: Normalized density predictions for the scattering technique using the Binned Clustered Algorithm in momentum-agnostic mode, for Scenario II. Predictions are shown by slices of 3 voxel layers along z (vertical axis).

5.3.2 Scattering tomography

In this scenario involving a greater material thickness, the POCA RMS algorithm from Equation 3.6 performs poorly, failing to provide reliable material identification. The increased path length through dense material reduces the precision of POCA location estimates compared to Scenario I. To mitigate this limitation, a cut on the scattering angle is applied: only events with scattering angles exceeding three times the detector’s angular resolution are retained. This filtering step helps remove low-angle scattering events that are less informative in the context of thick materials. BCA proves particularly effective in this regime, where longer muon paths naturally lead to larger average scattering angles. As a result, BCA becomes more advantageous than simple POCA, especially when the scattering angles are sufficiently large to support meaningful clustering. The resulting scattering density predictions, shown in Figure 5.6, confirm that eight hours of data acquisition are sufficient to detect the presence of the bronze cylinder and to localize it with reasonable accuracy. The true position of the bronze bar is $x, y, z = (0, 0, -810)$ mm.

The scattering method outperforms absorption, given that it is sensitive to marble and that it shows a better image resolution not only in the horizontal plane but also along the vertical plane.

5.4 Large scale objects: Muon Transmission Tomography

This section aims to demonstrate that the use of muon transmission tomography can be meaningful in the context of water content assessment. The density predictions shown in Figure 5.7 are obtained using the same back-projection and Gaussian smoothing algorithms as in Section 5.3.1, with a kernel size of 2 voxels. Our results indicate that eight hours of data acquisition are sufficient to detect the presence of a higher-density region within the limestone block. The image resolution in the x and y directions is adequate for estimating the radius of the high-moisture-content sphere. However, the resolution along the vertical z-axis is relatively poor, preventing precise determination of the sphere's z-position. Nonetheless, in a practical measurement scenario, combining data from multiple viewpoints would enable triangulation of the anomalous region's location.




Figure 5.7: Transmission ratio for the transmission technique using the back-projection algorithm. The high water-content 15 cm radius sphere with density 1.95 g cm^{-3} and located at $x, y, z = -150, 0, -260 \text{ cm}$, can be easily identified and exhibits a smaller transmission ratio. Reproduced from [47].

TomOpt

Optimizing the design of instruments that rely on the interaction of radiation with matter is an inherently complex challenge. The dimensionality of the parameter space, which includes geometry, detection technologies, materials, data acquisition strategies, and information extraction techniques, is vast and highly interdependent, making full manual optimization a challenging task. Yet, this complexity also offers an opportunity: substantial performance gains are potentially achievable if one can systematically explore this space and optimize configurations based on an objective function directly aligned with the instrument’s final goals. Traditionally, the stochastic nature of radiation-matter interactions has rendered classical modeling approaches insufficient for such tasks. However, the advent of fully differentiable pipelines and deep learning now enables the joint optimization of all design parameters. Once the dimensionality of the design space exceeds just a few parameters, brute-force grid searches become infeasible. In such cases, automated optimization techniques powered by automatic differentiation can be employed to navigate toward optimal configurations. Automatic differentiation allows gradients of complex objective functions to be computed efficiently through the chain rule, even as these functions are embedded within arbitrarily intricate simulation pipelines. Muon tomography detectors are especially well suited to this type of end-to-end optimization. They typically involve relatively simple physical processes, reconstructing just one particle per event, and allow for effective parametric modeling of the full data-generation process. This makes them an ideal testbed for developing and validating automated, objective-driven detector optimization strategies. TOMOPT represents a concrete step in this direction within the MODE collaboration [37], aiming to explore the practicality and scalability of detector optimization via automatic differentiation. Muon scattering tomography, in particular, offers a tractable and informative use case, balancing

optimization potential with manageable complexity. It allows us not only to explore meaningful improvements in detector performance but also to deepen our understanding of the requirements, feasibility, and limitations of end-to-end, task-aware detector design.

Section 6.1 introduces how TOMOPT performs detector optimization by defining task-specific objective functions and using gradient descent to improve performance. Section 6.2 then presents a benchmark example of this approach applied to a representative muon scattering tomography task. The work presented in this chapter is the result of the TOMOPT collaboration. My primary contributions include the development of the muon propagation and generation models, maintaining of the code, as well as the design and implementation of the inference methodology used in the furnace ladle benchmark. Following the departure of the principal investigator, Dr. Giles Strong, in 2023, I also assumed responsibility for coordinating group meetings and ensuring continuity of the project. In addition, I contributed to extending TOMOPT to applications in cargo scanning—an ongoing research direction initiated during my involvement.

6.1 Software description

The simulation and optimisation routines used for this study are structured into a software package referred to as TOMOPT. The code for TOMOPT is available open-source on Github [112]. In this section, we will attempt to provide the reader with a sufficient understanding of the program.

TOMOPT is a highly modular Python-based package that provides the full suite of tools and resources required for the investigation of the general problem of optimization of a scattering tomography detector. Its modular design allows each stage of the process to be performed by inheriting classes written to be most suitable for the particular task being studied. The key feature of TOMOPT is the fully differentiable nature of the detector modeling and inference pipeline, implemented using the automatic differentiation framework PYTORCH [84]. In this Section, we describe the capabilities and workflow of TOMOPT. Figures A.6, A.7 and A.8 in Appendix A.6 gives an overview of the whole optimization chain. The algorithmic implementation of the various simulation and optimization steps can be found in [111].

6.1.1 Cosmic muon generation

In the context of detector optimisation, it is essential to have a realistic and flexible cosmic muon source that can adapt to the various muon scattering tomography applications. TOMOPT samples muons from a horizontal generation surface S_{gen} with customisable energy and angular distributions, as it was presented in Section 2.1. The user is free to provide their own parameterisation of the differential cosmic muon flux, expressed as:

$$\Phi_z(E, \theta) = \frac{dN}{dE \cdot d\theta \cdot d\varphi \cdot dS_z} \text{ m}^{-2} \cdot \text{sr}^{-1} \cdot \text{GeV}^{-1}, \quad (6.1)$$

where dN is number of muon crossing a horizontal surface element dS_z , with energy between E and $E + dE$, zenith angle between θ and $\theta + d\theta$, and azimuthal angle between φ and $\varphi + d\varphi$. Most of cosmic muon flux parametrizations are given as:

$$\Phi(E, \theta) = \frac{dN}{dE \cdot dS_n \cdot d\Omega} \text{ m}^{-2} \cdot \text{sr}^{-1} \cdot \text{GeV}^{-1}, \quad (6.2)$$

where dN is number of muon crossing a surface element dS_n perpendicular to the muon direction, with energy between E and $E + dE$, and solid angle between Ω and $\Omega + d\Omega$. Because $d\Omega = \sin\theta \, d\theta \, d\varphi$ and $dS_z = dS_n \cos\theta$, Φ_z can be obtained from Φ as:

$$\Phi_z = \Phi \cdot \sin\theta \cdot \cos\theta. \quad (6.3)$$

The present version of TOMOPT embeds two default differential muon flux parameterisations: Guan et al. [49] and Shukla et al. [106]. After providing the desired generation surface span dx, dy , energy range E_{range} , zenith angle range θ_{range} , and number of muons N_{muon} , their position is uniformly distributed on S_{gen} .




Figure 6.1: Representation of the generation surface in TOMOPT (left), muon zenith angle (middle), and energy (right) distribution sampled from the Guan et al. [49] parameterisation. Reproduced from [111].

For computing efficiency purposes, it might be relevant to choose θ_{range} upper bound $< 70^\circ$ since horizontal muons are unlikely to reach both the passive volume and the detector layers. Once muons are generated, they will be propagated along negative z direction. The simulation of their propagation through passive and active materials is described in 6.1.2 and 6.1.3 respectively.

6.1.2 Passive volumes and multiple scattering

Once muons are generated, TOMOPT models their passage through the volume of interest (*passive volume*) and their interaction with matter. The passive volumes of interest are modelled in 3D using layers of voxels, and one can assign to each voxel a material (specified by its radiation length X_0). TOMOPT has its own database of X_0 values for various pure elements computed assuming nominal material density at 20°C . Alternatively, users can easily provide their custom materials' radiation lengths and densities. They may also man-

ually specify the layout of materials inside the volumes of interest, or instead provide a customised generator class to provide randomly generated volumes that are representative of the current task.

For computational efficiency via vectorisation, batches of many muons traverse the volume simultaneously. Each muon is defined using its momentum, (x, y, z) position, and (θ, φ) trajectory, where $0 \leq \theta < \pi/2$ ¹ is the deviation from the z axis, and $0 \leq \varphi < 2\pi$ is the azimuthal angle in the xy plane. Their propagation through the volume is handled by a simple model recommended by the Particle Data Group [76]. It assumes a small-angle Gaussian scattering and does not take into account non-Gaussian tails from rare large-angle scatterings. The model consists of picking two random scattering angles independently in orthogonal directions, as well as correlated transverse displacements. It begins by computing the width of the angular distribution θ_0 . For high-energy muons with momentum p travelling a distance d , the PDG model assumes:

$$\theta_0 = \frac{0.0136 \text{ GeV}}{p} \sqrt{n_{X_0}} \times (1 + 0.038 \ln n_{X_0}) , \quad (6.4)$$

where $n_{X_0} = d/X_0$, d is the distance traversed by the muon in the material in meters, and X_0 is the material radiation length in meters. The angular scattering is sampled from a zero-centred Gaussian distribution of this width with random numbers z_1 and z_2 taken from a normal distribution of mean 0 and variance 1:

$$\delta\theta = \theta_0 \times z_2. \quad (6.5)$$

The lateral displacement due to scattering is computed as:

$$\delta x = \theta_0 \left(z_1 \frac{d}{\sqrt{12}} + z_2 \frac{d}{2} \right). \quad (6.6)$$

¹Muons travelling upwards are removed, since all muons in the batch must share the same z position.

We show in subsection 6.1.8 that, although embedding such a model in a step-by-step algorithm affects the precision of the Coulomb scattering simulation compared to, for instance, highly accurate but slow models such as those provided in GEANT 4 [1], the optimisation-focused nature of the procedure we present in this paper requires prioritising computational speed over precision.

The angular scattering is sampled according to Equation 6.5 and subsequently compute the displacement due to the scattering according to Equation 6.6: this procedure is performed twice per scattering call, to compute scatterings in perpendicular directions (x_μ and y_μ) in the muon reference frame. The muon-frame scatterings must then be converted into the volume frame using 3D rotation matrices. As mentioned earlier in this section, muons are propagated by steps of distance dr along their flight path. The default value of dr is set to 1 cm, but can be changed to accommodate the voxel size, the volume of interest size, and the remaining distance to the bottom of the voxel. The optimal choice of dr is discussed in Section 6.1.8.

6.1.3 Detector modeling and hit recording

The detector models implemented in TOMOPT are designed with a flexible parameterisation scheme, allowing adaptation to a wide range of use cases without relying on fixed, a priori assumptions regarding optimal configurations. The fundamental building block of the detector is referred to as a *panel*. In the xy -plane, each panel is characterised by a spatial resolution $\sigma_{\max}$, a detection efficiency $\epsilon_{\max}$, and a fixed cost per unit area. This design allows users to simulate detector panels that are either custom-built or commercially sourced.

Panels may be freely positioned within the three-dimensional space (xyz), constrained only to predefined spatial regions termed *detector layers*. During the optimisation process, the objective is to learn five parameters per panel: its position in xyz -space and its spatial extent in the xy -plane (i.e., xy -span). Currently, the number of panels within each detector layer is held constant throughout optimisation, due to the challenges associated with gradient-based optimisation of discrete variables. Muon trajectories are recorded in the form of spatial coordinates, or *hits*, measured in xyz by the panels. The impact of spatial resolution on these recorded positions is dependent on the specific

detector technology employed. In the current implementation of TOMOPT, a simplified model is used wherein the position of a hit is drawn from a Gaussian distribution $\mathcal{G}(\mu = xy_{\text{true}}, \sigma = \sigma_{xy,\text{eff}})$, centered on the true xy -position of the muon, xy_{true} . The effective resolution $\sigma_{xy,\text{eff}}$ is determined by the panel's maximum resolution σ_{max} and the muon's position relative to the panel's centre.

In practice, a muon either registers a hit on a panel or does not, depending on the intrinsic efficiency of the panel. For the sake of differentiability within optimisation, TOMOPT adopts a probabilistic approximation: all hits are assumed to be recorded, but are weighted by an effective hit efficiency ϵ_{hit} . The overall detection efficiency of a muon, denoted ϵ_{μ} , is defined as the probability that the muon produces at least two hits before and after traversing the passive volume². Both the individual hit efficiencies and the aggregated muon detection efficiency ϵ_{μ} are integral to various stages of inference and optimisation, and will be discussed in detail in subsequent sections.

Optimisation mode Optimising the parameters of detector panels requires that muon trajectories be differentiable with respect to those parameters. While trajectory reconstruction is naturally differentiable with respect to the z -coordinate of the recorded hits, differentiation with respect to the xy -position and xy -span of the panels poses greater challenges. These parameters influence whether or not a muon is detected by a panel, rather than how it is recorded. Since the resolution and efficiency are otherwise held constant, the optimisation process becomes insensitive to muons that could have been detected but were excluded due to panel placement or size.

To circumvent this issue, TOMOPT employs a relaxed, non-physical approximation during optimisation. In this mode, every panel records a hit for every muon at the panel's z -position, regardless of whether the muon trajectory intersects the physical bounds of the panel. The spatial resolution $\sigma_{xy,\text{eff}}$ and detection efficiency ϵ_{eff} of each such hit are then modulated based on the muon's distance from the panel centre. These quantities decrease smoothly as a function of radial distance, scaled by the panel's spatial extent, thereby introducing a continuous dependence of the hit characteristics on the panel parameters.

²A minimum of two hits on either side of the passive volume is required to reconstruct the muon's trajectory. Additional hits contribute to reducing the uncertainty in the reconstructed parameters.

The effective resolution and efficiency are computed as:

$$\sigma_{xy,\text{eff}} = \sigma_{\text{max}} / w(xy_{\text{true}}), \quad \epsilon_{\text{eff}} = \epsilon_{\text{max}} \cdot w(xy_{\text{true}}), \quad (6.7)$$

where $w(xy_{\text{true}})$ represents the panel response function evaluated at the true xy -position of the muon. This function is designed to scale between 0 and 1, attaining its maximum when the muon passes through the center of the panel. In the current implementation, w is modeled as a pair of independent double-sigmoid functions in the x and y directions. The individual sigmoid values are used to determine the spatial resolution in x and y , while their product yields the hit efficiency. These are further scaled to ensure that hits at the panel center achieve the full nominal resolution and efficiency.

This approach yields detector behaviour that closely approximates the physical system for muons within the panel boundaries, while providing a smooth fall-off in performance for muons outside the panel. Consequently, the recorded hits and muon trajectories remain differentiable with respect to all learnable panel parameters. For validation or deployment, this relaxed model can be replaced with a physically faithful one, in which hits are recorded only if the muon passes through the panel's active area. The smoothness of the panel response function w , i.e., the rate at which hit quality degrades near the panel edges, can be tuned, and may itself be adapted during optimisation. A smoother transition increases sensitivity to muons further from the panel, thereby enhancing gradient flow with respect to the xy -position and span. In contrast, sharper transitions restrict full detector performance to within the panel bounds, improving sensitivity to the z -position, as only hits from nearby muons contribute significantly to inference. This behaviour arises because muons outside the panel are assigned low weights, while muons inside are treated nearly uniformly, leaving the z -coordinate as the dominant source of variation. The response function w is user-definable, allowing for the construction of custom panel behaviours. Examples include anisotropic scaling in x and y , or periodic functions to emulate discretised detection structures, such as wire chambers with fixed spatial frequency.




Figure 6.2: Example of detector panel modeling with sigmoid function used during optimisation (blue) and with rectangle function used during validation (red). Reproduced from [111].

6.1.4 Trajectory fitting and Point-of-Closest-Approach finding

The objective of trajectory fitting is to reconstruct the paths of muons as they enter and exit the passive volume, based on the hits recorded by the upstream and downstream detector panels. Each trajectory is obtained through analytic minimisation of a likelihood function that takes the recorded xyz hit positions as input. In this formulation, each hit is assigned a weight that reflects both its effective spatial resolution $\sigma_{xy,\text{eff}}$ and its detection efficiency ϵ_{eff} , hits with lower resolution or lower probability contribute less to the overall fit. Uncertainty in the z -coordinate is neglected in the current implementation. From an optimisation standpoint, the accuracy of the reconstructed trajectory is maximised when the associated hits are recorded with high resolution and efficiency. Consequently, the optimal panel configuration is one that maximises the informativeness of hits for the largest number of muons. However, improving the reconstruction quality for one set of trajectories may degrade it for others, and a global compromise must be achieved. Following trajectory reconstruction, scattering-related quantities such as the incoming and outgoing angles (θ , φ), and the total angular deflection are computed in a differentiable

manner. This enables the calculation of the uncertainties associated with these quantities, obtained via propagation of the partial derivatives with respect to the recorded hit positions and their xy uncertainties.

After muon tracking, TOMOPT proceeds to an initial inference phase, aiming to localise regions of space responsible for the observed deflections. This is achieved using an extended version of the Point-of-Closest-Approach algorithm, introduced in Section 3.1. Since POCA points are derived from the reconstructed incoming and outgoing tracks, they inherit uncertainties in all three spatial dimensions. These POCA points can then be used in multiple downstream applications: for example, to generate a three-dimensional visualisation for interpretability, or to compute voxel-wise hit densities, optionally weighted by the muons' angular deflections. The choice of features and their usage depend on the specific inference strategy employed. An example of such an inference pipeline is provided in Section 6.2.3.

6.1.5 Volume inference

The specific form of these predictions, as well as the inference method employed, depends on the particular muon tomography task. A wide range of inference strategies may be adopted, provided they yield outputs that remain differentiable with respect to the muon-trajectory variables, thus preserving end-to-end differentiability for optimisation. In the current implementation of TOMOPT, volume inference via the POCA method is designed to estimate the radiation length X_0 of voxels by inverting the PDF model of multiple Coulomb scattering. The basic approach for a given batch of muons with their PoCA located in a given voxel of height δz is:

$$\theta_0^{\text{rms}} = \frac{\theta_{\text{tot.}}^{\text{rms}}}{\sqrt{2}}, \quad (6.8)$$

$$\theta_0^{\text{rms}} = \frac{0.0136 \text{ GeV}}{p^{\text{rms}}} \sqrt{n_{X_0}^{\text{rms}}}, \quad (6.9)$$

$$\begin{aligned} n_{X_0}^{\text{rms}} &= \frac{\delta z}{X_0 \cos(\bar{\theta}^{\text{rms}})}, \\ \bar{\theta}^{\text{rms}} &= \frac{\theta_{\text{in}}^{\text{rms}} + \theta_{\text{out}}^{\text{rms}}}{2}, \\ X_0 &= \left(\frac{0.0136 \text{ GeV}}{p^{\text{rms}}} \right)^2 \frac{\delta z}{\cos(\bar{\theta}^{\text{rms}})} \frac{2}{\theta_{\text{tot.}}^{\text{rms}2}}. \end{aligned} \quad (6.10)$$

Where $\theta_{\text{tot.}}^{\text{rms}}$ is the root-mean-square (RMS) of the measured scattering angles. In Equation 6.9, $\theta_{\text{in}}^{\text{rms}}$ and $\theta_{\text{out}}^{\text{rms}}$ respectively refer to the incoming and outgoing muon zenith angle. Note that the natural log term present in Equation 6.4 has been ignored to simplify the inversion. The effects of efficiency and resolution may be further included in Equation 6.10 by instead aggregating over the muons with a weighted RMS, where muon-wise weights are computed as the efficiency divided by the variance of the squared value of the variable being aggregated.

In the above inference, we only considered muons whose PoCA was located inside the voxel in question, however, we have an uncertainty on the PoCA position, and some muons may have a PoCA located outside the passive volume due to mis-reconstruction. An extension is then to consider that muons can contribute to inference in multiple voxels by considering that the PoCA could have occurred in voxels surrounding its nominal location. The PoCA uncertainty POCA_{unc} in x , y and z is computed via automatic differentiation. Then the probability density function (PDF) for a PoCA to be located within a voxel is modelled using 3 independent Gaussian distributions centered over the nominal PoCA location and scaled in the three directions according to POCA_{unc} . A “scattering probability” is then computed per voxel for each muon, by integrating the PoCA PDFs over every voxel in the passive volume. An illustration of this procedure is presented on Figure 6.3. This probability then enters into the weighted RMS as a multiplicative coefficient.




Figure 6.3: Illustration of PoCA point being extended to other neighbouring voxels with a Gaussian model. For the sake of readability, only the extension along the y dimension is shown. Reproduced from [111].

This method of using PoCA points, however, is biased to underestimate X_0 , since the entirety of the scattering is assumed to happen in one scattering event. The muons of course undergo many scattering events as they traverse through the volume. Additionally, it converts muon-wise data into voxel-wise predictions, however, the muon tomography task may not require voxel-wise predictions, but instead perhaps predictions concerning the volume as a whole. Nonetheless, it can serve as a useful and generic first-stage of inference to serve as input to a second, more task-specific, stage of inference.

6.1.6 Loss functions and optimisation

As with volume inference, the choice of loss function used during optimisation depends on the specific muon tomography task. For example, classification tasks may naturally employ a cross-entropy loss. However, in the context of detector design, the loss function must also account for practical constraints such as cost. This can be expressed either as an explicit penalty in the loss or as a hard constraint. While the cost could be incorporated directly into the loss via a scalar weighting factor, it is often the case that the user has a predetermined budget and wishes to identify the highest-performing detector configu-




Figure 6.4: Example of the cost component of the loss as a function of detector cost, for a target budget of 4 A.U. A sigmoid component around the budget provides a slowly increasing gradient as the cost approaches the target budget, beyond which the loss increases linearly. The function is fully differentiable. The steepness of the sigmoid can be adjusted to better adjust the loss behavior around the budget.

ration that adheres to this constraint. TOMOPT provides two mechanisms for incorporating cost: budget penalisation and fixed-budget optimisation.

In the budget-penalisation approach, the user specifies a target budget, and an additional term is introduced into the loss function that penalises configurations exceeding this target. The penalty function is designed to remain small for costs below or near the target, but increases rapidly for configurations that exceed the budget. This allows the optimisation process to occasionally exceed the target budget if doing so yields a substantial performance improvement, as illustrated in Figure 6.4. This mode introduces hyperparameters to control both the relative weight of the cost term in the overall loss and the functional form of the penalty curve.

In contrast, the fixed-budget mode enforces a hard constraint on the total detector cost. Here, the user specifies the total allowable budget, and this cost is distributed among panels as learnable fractional allocations. Each panel still independently learns an xy -span, but the physical area is then rescaled to ensure its cost matches the assigned fraction of the total budget. This mode offers two main advantages: (i) it guarantees strict adherence to the user-defined budget, and (ii) the loss function remains purely performance-driven, obviating the need to tune trade-off parameters between performance and cost.

Once the total loss is computed, gradients of the loss with respect to each detector parameter are obtained via backpropagation [99], allowing updates to be made using standard gradient-based optimisation algorithms. TOMOPT interfaces with the optimisers provided by the PYTORCH framework, including stochastic gradient descent (SGD)[50], Adam[62], and others. The training process is facilitated through a flexible callback system, which enables dynamic control over module behaviour and training logic without requiring changes to core module definitions.

6.1.7 Update cycle

When computing a value for the loss function, it is important to ensure that it is computed under conditions that are generally representative of those that can be encountered during actual operation, in terms of e.g., number of muons and passive volumes. In effect, this means that optimisation should be performed using a range of passive volumes (whether generated on demand, or supplied by the user as a dataset), and that these should each be scanned using a number of muons that would be reasonable to expect given the position and exposure time of the proposed detector. Additionally, rather than computing the loss using just a single passive volume, it should instead be averaged over many passive volumes in order to provide a more generalized update to the detector parameters.

The update cycle in TOMOPT is performed in the following way:

1. Figure 6.5c: Scanning of a given passive volume involves passing a batch of muons through the volume and inferring their tracks. For computational reasons, the total number of muons may be split into smaller batches. Once all muons have been processed, the volume predictions can be made and the loss computed;
2. Figure 6.5b: This process is repeated for a batch of volumes, and the sum (or average) of their losses can be used to update the detector parameters. This is referred to as an *epoch*, in keeping with Machine Learning terminology³;

³This assumes that volumes are being generated. When working with a dataset of volumes, instead, this process can be repeated several times using a volume batch size that is less than the total number of volumes in the dataset.

3. Figure 6.5a: The fit runs for a specified number of epochs before finishing. It is also possible to run epochs in validation mode, where the detector is not updated: the detector panels can therefore be considered, more realistically, as discrete detector panels rather than with an associated distributed resolution and efficiency, as discussed in Section 6.1.3.




Figure 6.5: Breakdown of the fitting procedure of detectors in TOMOPT. Reproduced from [111].

6.1.8 Current limitations

Muon propagation through matter

As described in Section 6.1.2, the implementation of multiple Coulomb scattering is rather straightforward since the computation efficiency is prioritised over precision. Scattering angle and displacement distributions are assumed to be 100% Gaussian, where in reality only % of the bulk is Gaussian and the tails follow a $1/\theta^4$ law corresponding to large scattering events. In addition, the PDG model used is not designed for a step-by-step algorithm: when it is applied to such an algorithm, it results in underestimating the scattering and displacement amplitudes more and more with each step, as shown in Figure 6.6.




Figure 6.6: Comparison of standard deviations of projected scattering angles and displacement distributions between TOMOPT and GEANT 4 and for various numbers of propagation steps.

It becomes clear that the choice of step length dr affects the precision of the multiple Coulomb scattering, and thus has to be chosen carefully. In this context, it is advised to choose a step length dr which represents about 1% of the

passive volume size. In addition to multiple Coulomb scattering, muons lose energy through ionisation, thus sub-GeV muons might decay in the passive volume. Energy loss has not been implemented in TOMOPT yet, which results in an overestimation of the detected muon flux. The creation of a more advanced propagation model is currently under investigation.

Detector and cost modeling

TOMOPT's ambition is to guide the user through the choice of detector technology and not to provide accurate simulation of particle detection, signal generation, or light transport. Thus, the difference between various detector technologies will be implemented at the level of cost modelling, meaning how the cost scales with the active detection surface and number of readout channels. Such features have not been implemented yet, but are planned to be included in future updates. Regarding detector modelling, TOMOPT is still limited to horizontal panels with continuous spatial resolution and uniform granularity. This limits the phase space of possible detector geometries.

Volume inference

For now, TOMOPT only provides a basic implementation of radiation length inference based on the inversion of the PDG scattering model utilising the Point of Closest Approach. The inversion of the scattering model Equation 6.10 assumes muon momentum to be known with a 100% precision, which is quite unrealistic in the context of a muon scattering tomography experiment. If a momentum estimation is available, the user can still modify the inversion model and replace the nominal muon momentum with an estimated momentum. If not, it can be replaced by the RMS computed over the whole cosmic muon momentum spectrum.

6.2 Benchmark: Fill-level estimation at metal refinery

The following section describes a benchmark involving the estimation of the fill level at a metal refinery, and aims at demonstrating a typical inference and optimisation chain using TOMOPT. After evaluating the performance of an initial sub-optimal detector configuration, the latter will be optimised, and its performance will be compared to two human baseline detector configurations.

6.2.1 Description

Furnace ladles are structures used to transport melted metal produced in an industrial plant, as shown in Figure 6.7. A description of a generic industrial furnace ladle model used to simulate muography measurements can be found in [67]. A typical example of the part of the manufacturing process which involves the use of the ladle could be described in a general way as follows: firstly the metal is melted in a furnace, then it is transported with a furnace ladle, and finally it is poured into moulds to create metal parts with a defined shape. These parts can later be post-processed in order to get the final product.



Figure 6.7: Example of a furnace ladle used in the foundry industry. Image by Boris Bukovský, via Pixabay.

In this process, the amount of melted metal inside the furnace is a key parameter. The lack of metal in the ladle can result in moulds that are not completely filled. Therefore, they may not be used to produce the final product correctly. On the other hand, an excessive charge of metal in the ladle leaves remnants and scraps, causing problems and losses in the manufacturing process. It has to be mentioned that usually, in the top of the melted metal inside the ladle, there is a lighter layer of slag produced in the melting process itself. This layer, prevents the use of optical methods to determine the liquid metal level. The resolution needed by the industry to optimise the manufacturing process is of the order of a few centimetres, although it depends on the distinctive features of each case.

6.2.2 Detector and volume

Active layers, i.e., the detector panels, are located above and below the measurement target, in a usual scattering muography configuration. At least two panels at different z locations providing $x - y$ positions of muons are required per tracking detector (upper, and lower detectors). As specified in subsection 6.1.3, they are modelled considering several parameters, such as position resolution $\sigma_{\max}$, detection efficiency $\epsilon_{\max}$, and cost per unit area.

Within the volume of interest between the detectors, a rectangular furnace ladle filled with liquid steel has been simulated (passive volume). In the Figure 6.8 a sketch of the furnace ladle is shown. The grey parts correspond to ladle walls, made of solid steel, the molten steel is represented in yellow, the slag layer on top of it in orange, and the remaining void volume over them in blue. The definition of materials is detailed in Table 6.1.

Material	Density, d [g cm^{-3}]	X_0 [cm]	Composition
Solid steel	7.818	1.782	<i>Fe</i> (99%), <i>C</i> (1%)
Liquid steel	7.000	1.991	<i>Fe</i> (99%), <i>C</i> (1%)
Slag	2.700	8.211	<i>CaO</i> (53.5%), <i>Al₂O₃</i> (33.5%)
Air	0.001205	30390	<i>N₂</i> (78%), <i>O₂</i> (20%)

Table 6.1: Density, radiation length and composition of the simulated furnace ladle materials.




Figure 6.8: Top view (left), and front view (right) of the rectangular furnace ladle simulated. It is composed of walls (grey), liquid steel (yellow), slag (orange), and the air volume (blue) over the liquid steel and slag. Reproduced from [111].

6.2.3 Inference method and loss function

The aim of this task is to design a detector that can provide an accurate measurement for the fill-heights of furnaces. This is essentially a single-variable regression task, and a natural choice for the loss function could be the mean squared-error (MSE) between the predicted and true fill-heights. However, as will be discussed below, MSE is not necessarily the most suitable loss for optimisation of this task.

Inference

In the spirit of an end-to-end optimisation pipeline, the inference must output the fill-height prediction and receive optimisation loss gradients based on this. In developing a suitable inference method, we noticed that the mean value of the z -positions of PoCAs ($\overline{z_p}$) was monotonic with the fill-height, due to hard scatterings occurring inside the liquid-metal part of the furnace volume, rather than the air and slag regions. This meant that it could serve as an indicator of the mass density distribution centre in the vertical Z axis (see Figure 6.9). As a benefit, this inference approach does not require computing X_0 predictions for the voxels, and thus is more computationally efficient.

This approach can be adapted into an inference method suitable for detector optimisation by computing a weighted average of the POCA z -positions, with weights determined by the properties of the corresponding muon trajectories. Specifically, the weight assigned to each POCA point is proportional to the




Figure 6.9: Muography of a simulation ladle furnace filled with steel up to 0.4 m. The bottom wall of the furnace goes from 0.2 m to 0.3 m in z axis. On the left, the front view of PoCA reconstruction. On the right, PoCA distribution in vertical axis (z) and filling level inference example: true steel level (blue), $\bar{z}_p$ (red), and h_{pred} (green) for the particular case of a bias correction on z_{pred} , with $a = 2$ and $b = -0.2$ (see section 6.2.3 for details). Reproduced from [111].

trajectory efficiency ϵ_p , and inversely proportional to the squared uncertainty in its z -position, $(\Delta z_p)^2$.

However, we observed that POCAs located within the furnace walls can introduce bias and degrade the precision of the height estimation. Additionally, the uncertainties in the xy positions of POCAs must be taken into account. Following the approach used for χ_0 inference in subsection 6.1.5, we model each POCA as a two-dimensional Gaussian distribution in the transverse (xy) plane, centred on its position and scaled according to its xy uncertainty. These distributions are integrated over the transverse voxels of the furnace. To mitigate the influence of peripheral POCAs, we apply a spatial weighting function across the xy plane. This function is constructed from a pair of double-sigmoid functions—one in x , one in y —that peak at the centre of the furnace and smoothly decay to zero at the walls. Each voxel receives a weight accordingly, and these are used to scale the integrated Gaussian values per POCA. The total xy weight for each POCA, denoted $w_{xy,p}$, is defined as the sum of these weighted integrals.

The final weight assigned to each POCA point is given by:

$$w_p = \frac{\epsilon_p}{(\Delta z_p)^2} w_{xy,p}. \quad (6.11)$$

This weight has the effect of emphasising the z positions of PoCAs that are well measured in z , located near the centre of the furnace, and have a high trajectory efficiency.

The predicted height is then taken as the weighted average of PoCA z positions:

$$z_{\text{pred.}} = \frac{\sum_p w_p z_p}{\sum_p w_p}. \quad (6.12)$$

Whilst the mean of PoCA z positions is actually more likely to correspond to the centre of the fill-material, this redefinition to the fill-height can be accomplished by the linear transformation described in the next section.

Bias correction

As discussed, the method outlined above does not yield an exact estimate of the true fill height. The mean of the POCA z -positions typically corresponds to the centre of the fill material, but with finite precision. Furthermore, the POCA approach introduces a systematic bias by attributing all muon scattering to a single point along the trajectory. Importantly, the magnitude and direction of this bias are dependent on the specific detector configuration.

Nonetheless, provided the predicted value increases monotonically with the true fill height, it is possible to correct for this bias in a configuration-dependent manner. This can be achieved by evaluating the inference model on a diverse set of volumes with known fill heights, and then adjusting the predictions, either via a continuous mapping function or using discrete bin-based corrections, so that the mean prediction aligns with the true value across the dataset. From this perspective, the key property we seek in the inference method and detector configuration is not necessarily unbiased prediction, but rather low variance across repeated measurements for the same volume, and a strong monotonic relationship with the true fill height.

Ultimately, when the detector is deployed for real measurements, it becomes essential to debias the predictions so that they correspond as closely as possible to the true fill heights. For simplicity, we adopt a linear correction model of

the form:

$$h_{\text{pred.}} = a \times z_{\text{pred.}} + b, \quad (6.13)$$

where parameters a and b are computed based on the predictions and targets from a large set of known example volumes at various fill heights. The effect of this debiasing can be seen in Figure 6.12.

In order to improve monitoring capabilities during the optimisation, we refine the values of a and b after each update to the detector and display the resulting MSE of de-biased predictions. In order to avoid having to re-scan a large number of volumes, this refinement is performed using predictions and targets collected during the current batch of volumes.

Loss function

Even after applying debiasing, predictions from the detector may still exhibit residual bias. Optimising under a standard mean squared error (MSE) loss can make the process sensitive to such residuals, potentially leading to instability during training. As previously discussed, the overarching goal of detector optimisation is to enable the inference procedure to produce predictions that are both precise (low variance) and monotonically increasing with the true fill height. A loss function that captures this goal must therefore reflect the variability in predictions at fixed fill heights. One natural choice is to minimise the within-group variance of the predictions, i.e., the spread of predictions for each true fill height. This can be expressed as:

$$\mathcal{L} = \frac{1}{N_l} \sum_l \sigma(p_l), \quad (6.14)$$

where $\sigma(p_l)$ is the standard deviation of predictions for fill-height index l . If the spread of predictions becomes smaller, the loss will decrease, and vice versa, thus the loss encourages precise measurements of the fill-heights.

However, this formulation alone is insufficient. A trivial solution that minimizes the loss would be to collapse all predictions to a single value, yielding zero intra-class variance but completely eliminating sensitivity to changes in the fill height. To prevent this degenerate outcome, we introduce a regularising term that encourages separation between the mean predictions across different

fill heights. This gives rise to the final loss function:

$$\mathcal{L} = \frac{1}{N_l} \frac{\sum_l \sigma(p_l)}{\sigma(\bar{p}_l)}, \quad (6.15)$$

where $\sigma(\bar{p}_l)$ is the standard deviation of the mean prediction for each fill-height.

This formulation, which we refer to as the “spread-over-range” loss, is minimised when the intra-class prediction variance is low, and the inter-class separation (i.e., dynamic range of the means) is high. In doing so, it encourages a model that produces stable, repeatable predictions for each fill level, while also ensuring that those predictions remain sensitive to the underlying true value. This, in turn, facilitates the de-biasing process: if predictions for each volume lie within a compact and well-separated range, one can more reliably map them back to the corresponding true fill height.




Figure 6.10: “Spread over range” loss function tested for different detection surfaces, and using the h_{pred} inference method with unweighted z_{pred} and a linear de-biasing of $a = 2$ and $b = -0.2$.

Figure 6.10 shows how the “spread-over-range” loss varies with the span of the detection panels, centred in X and Y on the furnace ladle. The loss drops sharply as panel size increases within the ladle’s interior, where the liquid steel—the muography target—is located. Beyond this region, the loss reduc-

tion slows, mainly because fewer muons intersect the target and more are scattered by the furnace walls, which distorts the inferred density along z . This effect is partly mitigated by the weighted PoCA inference (subsection 6.1.5). Despite this, the loss continues to decline as detector surface grows, reflecting the general benefit of larger detectors in muography: they collect more muons, improving inference precision. Note that the cost component of the detector surface is not included in this figure and is handled separately in the optimisation process (subsection 6.1.6).

6.2.4 Optimisation process and results

Passive-volume data

Using 10 cm voxels, we generate volumes at eight different fill-heights to simulate from ladles that are almost empty to ones that are almost full. In every case, we add a 10 cm layer of slag on top. When optimising the detector, each of these volumes will be evaluated five times (40 volumes per *volume batch*) – whilst the volumes will always be the same, stochasticity in prediction will still be present due to the muon generation. When evaluating the performance of the detectors, each volume is evaluated 12 times. 1000 muons will be used per evaluation.

Initial configuration

In order to test the optimisation capabilities of TOMOPT, we initially set the detector in a knowingly sub-optimal configuration: one in which the detector panels are off-set from the passive volume in $x - y$, thus lowering the flux muons that are well detected, and interact with the passive volume; and closely space in z , thus increasing the uncertainty in PoCA reconstruction.

The detector itself (illustrated in Figure 6.11) consists of four equally sized square panel above the passive volume, and another four panels below. These panels begin centred in $x - y$ over the corner of the passive volume, and with an (unrealistic) separation in z of 1 cm between each panel. There is a 5 cm separation between the passive volume to the nearest panels. The total cost of the detector is fixed at six arbitrary units, and panels cost one arbitrary unit

per m^2 . The available budget is initially distributed equally to all panels, thus each panel measures 86.6 cm in both x and y length. Panels are defined to have nominal efficiencies of 90 %, and $x - y$ resolutions of 0.1 mm.




Figure 6.11: Initial layout of the detector panels above and below the passive volume.

The initial performance for the detector is shown in Figure 6.12. The “raw” predictions are the result of the inference with no debiasing. We can observe a slight slope to their values with respect to the fill height, however, the large spread in predictions for the same target means that it is difficult to distinguish different fills. The de-biasing process via a linear correction is able to slightly reduce the MSE, but cannot do much more than shift the predictions. Averaging across the range of fill-heights considered, the initial detector has a mean MSE of 0.037 m^2 , and an average absolute error on fill-level prediction of 16 cm.

Our intuitive expectation is that from this position, the panels should move to cover the centre of the passive volume, and increase in z -separation. There is, however, a trade-off in increasing the z -separation: whilst it generally increases the PoCA resolution, it also decreases the hit-reconstruction for muons entering or leaving the passive volume at large angles. We therefore expect the optimisation to find a natural point at which the improved resolution is balanced by the decreased high-angle-muon flux. Additionally, we have no




Figure 6.12: Performance of the initial detector before and after the de-biasing process.

confident intuition over how the panels should be resized to best improve performance given a constant budget.

Optimisation callbacks

As mentioned in subsection 6.1.6, TOMOPT has a callback system that allows the user to augment the basic optimisation loop to include different functionality, and to help stabilise the optimisation process. The description of the callbacks used for this task can be found in [111].

Optimisation process

Five parameter families are considered: xy positions of the detector panels and the z positions of the panels are considered separately for panels above and below the volume (four parameter families); and the budget assignment per panel (one parameter family). Over the course of five epochs, we monitor the gradients received by the trainable parameters in these families and compute

nominal learning rates that correspond to position updates at a rate of 1 cm per epoch, and budget weights at a rate of 0.1 (dimensionless).

Stage one Over the course of 10 epochs, the detector is updated, using 1-cycle learning rate annealing to scale the nominal learning rates over the course of the optimisation: initially they increase to allow large updates of the detector, and towards the end they decrease to much smaller values, allowing the panels to converge and stabilise. Additionally, the sigmoid detector model is annealed from an initially smooth configuration to a sharper, more physical distribution of efficiency and resolution. On a 2018 Macbook Pro with an Intel i7, this takes about 30 minutes to run (no GPU is used).

Figure 6.13 illustrates the detector configuration before and after this stage of the optimisation process. From this we can see that the panels have indeed moved to be more central in $x - y$, and have expanded in z .




Figure 6.13: Comparison of detector configurations before and after stage one of the optimisation process. The coloured lines/squares indicate the positions and sizes of the panels, and the hatched area indicates the position and size of the passive volume. The top rows indicate panels above the passive volumes, and the bottom rows indicate panels below the passive volume.

Figure 6.14 compares the performance of the optimised detector to the initial detector, and indicates a significant improvement in both the precision and bias: now predictions are clearly distinguishable by true fill-height. The optimised detector has a mean MSE of 0.0014 m^2 , and an average absolute error on fill-level prediction of 2.8 cm.





Figure 6.14: Performance of the optimised detector after stage 1 compared to the initial detector (both shown after the de-biasing process). Heights shown include the space below the passive volume.

Stage two The previous optimisation process showed that TOMOPT can take a poorly designed detector and quickly shift it into a configuration that provides much better performance. In this second stage, we will demonstrate that TOMOPT can be used to refine detectors that are already thought to be well-performing, by altering only subsets of the total available parameters. To do so, we will use our detector from stage one.

During stage one, we saw that the panels moved to become central in $x - y$. We will take this experimental evidence, along with our domain knowledge of the problem, to assume that the ideal position for the detectors is directly above and below the passive volume. We will manually shift the panel to the centre and fix it there, allowing only the z and budget assignments to slowly change.

Once again, we compute suitable nominal learning rates for our three remaining parameter families by monitoring the gradients they receive for five epochs. Over another 10 epochs we allow the detector to further be refined, this time with smaller update steps. Figure 6.15 illustrates the detector after this refinement stage, and Figure 6.16 its performance. Mainly we can see that the panels below the passive volume have moved to cover a larger z range. While

this could come at the expense of missing some high-angle muons leaving the passive volume, the panels have slightly increased in size by decreasing the intermediate panels above the passive volume. The refined detector has a mean MSE of 0.0011 m^2 , and an average absolute error on fill-level prediction of 2.5 cm.



(a) Detector configuration after stage one optimisation process. (b) Detector configuration after stage two optimisation process.

Figure 6.15: Comparison of detector configurations after stage one (left) and stage two (right) of the optimisation process.

6.2.5 Comparison to human baselines

Having demonstrated that TOMOPT can successfully optimise detectors, we now want to see whether the resulting detector is good compared to a human-designed detector. For this we will take two baseline configurations:

- Detector panels are arranged in pairs with a 5 cm gap within each pair and a 25 cm gap between the upper and lower pairs. This layout prioritises trajectory precision by maximising the baseline for muon position measurements.
- Baseline 2: Detector panels are placed with an even 10 cm spacing. Baseline 1 potentially misses some high-angle muons due to the large distance to the second pair panels. This baseline attempt to strike a balance between precision and exposure.

These are illustrated in Figure 6.17.




Figure 6.16: Performance of the optimised detector after stage 2 compared to the initial detector (both shown after the de-biasing process). Heights shown include the space below the passive volume.

Figure 6.18 compares the performances of the stage-two optimised detector against the two baseline detectors. Whilst all three detectors offer similar levels of performance, our optimised detector is able to not only match human performance, but slightly and consistently outperforms it.

We demonstrated the capabilities of TOMOPT through a simulation of a typical industrial scenario: estimating the fill level of furnace ladles in a metal refinery. By applying task-specific inference and loss functions, we optimised a poorly designed initial detector into one that slightly outperformed two human-designed baselines in fill-height prediction precision. In this example, the optimisation was performed under a fixed-cost constraint, though TOMOPT also supports cost-minimising optimisation within a specified budget.




Figure 6.17: Configurations for the human-designed baseline detectors. The coloured lines/squares indicate the positions and sizes of the panels, and the hatched area indicates the position and size of the passive volume. The top rows indicate panels above the passive volumes, and the bottom rows indicate panels below the passive volume




Figure 6.18: Performance of the optimised detector after stage 2 compared to human-designed baseline detectors (both shown after the de-biasing process). Heights shown include the space below the passive volume.

Conclusion and prospects

This thesis has explored several key challenges and innovations in muon-based imaging, from improving reconstruction fidelity in muon scattering tomography to optimising detector design for specific tasks. Through a combination of analytical insights and machine learning approaches, we addressed fundamental limitations in current methodologies and proposed new tools, such as `Muograph` and `TOMOPT`, that offer more flexible, precise, and application-aware solutions. Together, these contributions advance the feasibility of practical muography across scientific, industrial, and cultural heritage domains.

Characterization of muon scattering tomography imaging This thesis has highlighted the challenges inherent in reconstructing scattering densities from muon scattering tomography (MST) data. We demonstrated that the inverse problem of estimating material properties is affected by multiple sources of noise, many of which are not easily mitigated by analytical corrections. In this context, we introduced the concept of *lumuosity*, a spatially dependent parameter that significantly influences the reconstructed scattering densities. Its impact should be considered both during detector design and in reconstruction algorithms. Given the limitations of existing reconstruction methods and the anisotropic nature of the predominantly vertical cosmic muon flux, we proposed machine learning as a promising tool for restoring high-quality 3D images. However, implementing this approach will require the development of a realistic, diverse training database representing the variety of materials, shapes, and imaging scenarios encountered in MST. To support reproducibility and foster community collaboration, the reconstruction algorithms and metrics developed in this work have been implemented in a Python-based software package, `Muograph`. This open-source tool aims to standardize algorithm development, facilitate benchmarking, and reduce redundant efforts, while remaining compatible with machine learning workflows.

Muon momentum estimation We showed that estimating muon momentum from deflections in known materials can be significantly enhanced through the use of machine learning. The achievable precision is constrained by the spatial resolution of the detector, which limits our ability to distinguish small-angle scatterings from instrumental uncertainty. To approach real-world applicability, further work must address additional complexities such as secondary particle production, which can contaminate muon signals and degrade tracking quality. Future efforts should also incorporate more realistic models of the detection planes, including additional scattering effects. Ultimately, these developments will need to be tailored to specific detector technologies, taking into account their particular responses and limitations. Additionally, the detector setup used is not particularly compact or cheap in terms of material, which might be a limiting factor for its practical use. In order to accommodate the logistical constraints of an in situ deployment, the detector setup has to move toward a simpler configuration, with fewer detector panels and thinner absorbers.

Muography in cultural heritage applications Muography, with its ability to non-invasively probe dense volumes, offers valuable complementary capabilities to existing cultural heritage imaging techniques. We advocate for its use in conservation science, where both absorption and scattering-based muography can provide meaningful insights. We have proposed several practical use cases from a local heritage context, including the detection of internal metallic structure in statues or density variations. Of particular interest is the potential application of muography for monitoring moisture accumulation, an important factor in the degradation of cultural materials. To our knowledge, this application remains unexplored and represents a promising avenue for future research.

Task-aware detector design optimisation We presented TOMOPT, a task-aware detector design optimization tool that demonstrates the feasibility of end-to-end, differentiable experiment modeling. This tool represents a shift in how detector systems might be designed, moving from manual, heuristic methods toward automated, performance-driven approaches. Although black-box optimizers like Bayesian optimization may be well-suited for moderate-scale problems, the high-dimensional parameter spaces of large experiments

(e.g., at the CERN LHC) necessitate scalable, differentiable optimization methods. Our work shows that such an approach is viable and offers a flexible framework for integrating measurement-aware objectives into detector design. Future developments will expand TOMOPT to support non-horizontal geometries, significantly enlarging the space of potential detector configurations. We also plan to integrate machine-learning-based physics models that account for muon energy loss, and to benchmark TOMOPT in broader application domains such as cargo inspection for border security.

Appendix A

Appendix

A.1 Coordinate system

A convenient coordinate system for muon tomography is the spherical coordinate system, where a track can be parameterized with 2 degrees of freedom: the zenith angle θ and the azimuthal angle φ . They are represented in the Figure A.1a. Given a 3D vector $\mu = (x, y, z)$ in Cartesian coordinates, the variables of interest are:

1. **Azimuthal angle φ :** The angle between the track and the horizontal x -axis, $\varphi = \text{sgn}(y) \cos^{-1} \frac{x}{\sqrt{x^2+y^2}}$.
2. **Projected zenith angle $\theta_{x,y}$:** The angle between the projected track in the XZ and YZ plane and the vertical z -axis, $\theta_{x,y} = \tan^{-1} \frac{x,y}{z}$.
3. **Zenith angle θ :** The angle between the track and the vertical z -axis, $\theta = \pi - \tan^{-1} \frac{\sqrt{x^2+y^2}}{z}$. A vertical track has $\theta = 0$, and a horizontal track has $\theta = \pi/2$. The zenith angle θ can be contained from its projections $\theta_{x,y}$ with the formulae: $\theta = \tan^{-1} \left(\sqrt{\tan^2 \theta_x + \tan^2 \theta_y} \right) \approx \sqrt{\theta_x^2 + \theta_y^2}$ for $\theta_x, \theta_y < 20^\circ$.




Figure A.1: Representation of a 3D vector μ in a spherical coordinate system (a), with the azimuthal angle ϕ , the zenith angle θ and its projections in the XZ and YZ plane θ_x and θ_y . View of the simulation scenario I, with 3 detection planes above and below the volume of interest. The latter consists of 4 cubes made of Aluminum (blue), copper (orange), iron (yellow), and lead (red). Muons are generated from a generation surface placed directly above the top detection plane, and their position are recorded by each of them.

A.2 Simulation scenario: Elementary cubes

This section describes a simple simulation scenario, where cubes of pure elements are placed within a typical muon scattering tomography detection system. The latter consists of two sets of three 1 m^2 horizontal detection planes separated by 10 cm, respectively placed before and after the volume of interest. The volume of interest is $64 \times 64 \times 64 \text{ cm}^3$ and contains four $10 \times 10 \times 10 \text{ cm}^3$ cubes made of pure elements: Lead, Copper, Iron, and Aluminum, with respective horizontal positions $(-15, 15)$, $(15, 15)$, $(15, -15)$ and $(-15, -15) \text{ cm}$. All cubes are placed on top of the lower detection panel, as shown in Figure A.1b. Muons are generated from a flat 1 m^2 surface placed directly above the upper detection panel, according to the CRY flat sky generation described in Section 2.1.2. Only full-coincidence muons (muons that leave a hit in every detection plane) are recorded. The reconstructed hits are obtained by smear-

ing the true muon position on the detection plan, as described in Section 2.3.1. Their distribution across the first three detection planes is shown in Figure A.2. Muon incoming and outgoing tracks are obtained by applying a linear fit to the reconstructed hits, and their zenith angle and scattering angle are shown in Figure A.3. The angular resolution of the apparatus amounts to 6.74 mrad, given a spatial resolution of 1 mm, and its efficiency to 85.7 % given a individual panel efficiency of 95 %. A total of 5 million muons are generated, which corresponds to ≈ 8 hours of acquisition time, leading to a total of 1,292,500 full coincidence muon events.




Figure A.2: Hits on the top, middle, and bottom panels above the volume of interest of the full-coincidence muons.




Figure A.3: Zenith angle θ (left) and scattering angle $\Delta\theta$ (right) distributions of the full-coincidence muons. Mean μ and standard deviation σ are shown in the top right corner.

A.3 Muograph software

The methods required for muon scattering tomography (MST) reconstruction have been integrated into a unified Python framework called Muograph. This software is designed to streamline MST data analysis by providing a consistent and modular approach. It takes as input the muon hits recorded on detector planes, accounts for detector effects such as efficiency and spatial resolution, and includes tools to reconstruct muon tracks from the measured hits. These reconstructed tracks can then be processed using multiple algorithms implemented within the framework, including POCA, BCA, and ASR. In addition, Muograph offers dedicated visualization tools and metrics tailored for 3D data analysis. The software is open source and publicly available at: <https://pypi.org/project/muograph/>, <https://github.com/MaximeLagrange/muograph>.

A.4 Deviations' basis




Figure A.4: The deviations of the muons' hits from a fitted linear track expressed in different bases: Cartesian (a), Cartesian with the log-modulus function applied to both x and y (b), cylindrical (c), and cylindrical with log-modulus applied to the radius r while keeping the polar angle θ constant (d).




Figure A.5: The deviations of the muons' hits from the fitted linear track, defined as a graph along the track axis (black) using different scales: linear **(a)** and with log-modulus applied on the magnitude while keeping the same angle for each point **(b)**.

A.5 Network hyperparameters

Table A.1: DNN hyperparameters.

Parameter	Value
Number of layers	6
Neurons per layer	64
Hidden activation	GELU
Output activation	None
Batchnorm	True

Table A.2: RNN hyperparameters.

Parameter	Value
Embedding	
Normalize inputs (batchnorm)	True
Neurons for each layer	32,64
Activation	ReLU
Recurrent cell	
Cell type	GRU
Dimension	64
Layers	6
Bidirectional	False
Output (DNN)	
Layers	3
Neurons per layer	64
Hidden activation	ReLU
Output activation	None

Table A.3: Transformer hyperparameters.

Parameter	Value
Embedding/position encoding	
Normalize inputs (batchnorm)	True
Neurons for each layer	32, 64
Activation	GELU
Encoder	
Dimension	64
Encoder layers/class layers	6/2
Heads	8

Table A.3: *Cont.*

Parameter	Value
Feed-forward dimension	128
Layer norm	True
Output (DNN)	
Layers	3
Neurons per layer	64
Hidden activation	GELU
Output activation	None

Table A.4: GNN hyperparameters.

Parameter	Value
Graph blocks	
Block type	EdgeConv
Feature dimension	64
Number of blocks	6
DNN feedforward dimension	128
DNN batchnorm	True
Graphnorm	True
Activation	GELU
Node aggregation	Mean
Graph pooling	
Aggregation	Mean
Jumping knowledge	Concatenation
Output (DNN)	
Layers	3
Neurons per layer	64
Hidden activation	GELU
Output activation	None

A.6 TomOpt




Figure A.6: Scan loop for a batch of passive volumes (left), scan loop for muons over a single passive volume (right) and breakdown of the fitting procedure of detectors in TOMOPT (below).




Figure A.7: Simplified representation of the various computation steps from the *fit epoch* and *scan volume* methods, from Figure A.6.




Figure A.8: Simplified representation of the various computation steps from the *fit epoch* and *scan volume* methods, from Figure A.6.

Bibliography

- [1] S. Agostinelli et al. “GEANT4 — a simulation toolkit”. In: *Nucl. Inst. Meth. A* 506 (2003), p. 250. DOI: 10.1016/S0168-9002(03)01368-8.
- [2] I. Alsamak et al. *Garfield++: a toolkit for the detailed simulation of particle detectors based on ionisation measurement in gases or semiconductors*. URL: <https://garfieldpp.web.cern.ch/garfieldpp/about/>. (accessed: 01.04.2025).
- [3] F. Ambrosino et al. “Joint measurement of the atmospheric muon flux through the Puy de Dôme volcano with plastic scintillators and Resistive Plate Chambers detectors”. In: *Journal of Geophysical Research: Solid Earth* 120.11 (2015), pp. 7290–7307. DOI: 10.1002/2015JB011969.
- [4] G. Anbarjafari et al. *Atmospheric ray tomography for low-Z materials: implementing new methods on a proof-of-concept tomograph*. 2021. arXiv: 2102.12542 [physics.ins-det].
- [5] V. Anghel et al. “A plastic scintillator-based muon tomography system with an integrated muon spectrometer”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 798 (2015), pp. 12–23. ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.06.054.
- [6] A. Ankowski et al. “Measurement of through-going particle momentum by means of multiple scattering with the ICARUS T600 TPC”. In: *Eur. Phys. J. C* 48 (2006), pp. 667–676. DOI: 10.1140/epjc/s10052-006-0051-3.
- [7] K. Armanious et al. “MedGAN: Medical image translation using GANs”. In: *Computerized Medical Imaging and Graphics* 79 (2020), p. 101684. ISSN: 0895-6111. DOI: 10.1016/j.compmedimag.2019.101684.

- [8] A. Arodzero et al. “High speed, low dose, intelligent X-ray cargo inspection”. In: *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. 2015, pp. 1–7. DOI: 10.1109/NSSMIC.2015.7581836.
- [9] S. Arridge et al. “Solving inverse problems using data-driven models”. In: *Acta Numerica* 28 (2019), pp. 1–174. DOI: 10.1017/S0962492919000059.
- [10] J. Bae and S. Chatzidakis. “Momentum-Dependent Cosmic Ray Muon Computed Tomography Using a Fieldable Muon Spectrometer”. In: *Energies* 15.7 (2022). ISSN: 1996-1073. DOI: 10.3390/en15072666.
- [11] J. Bae, R. Montgomery, and S. Chatzidakis. “Momentum informed muon scattering tomography for monitoring spent nuclear fuels in dry storage cask”. In: *Scientific Reports* 14 (Mar. 2024). DOI: 10.1038/s41598-024-57105-y.
- [12] S. Barnes et al. “Cosmic-Ray Tomography for Border Security”. In: *Instruments* 7.1 (2023). ISSN: 2410-390X. DOI: 10.3390/instruments7010013.
- [13] G. Blanpied et al. “Material discrimination using scattering and stopping of cosmic ray muons and electrons: Differentiating heavier from lighter metals as well as low-atomic weight materials”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 784 (2015). Symposium on Radiation Measurements and Applications 2014 (SORMA XV), pp. 352–358. ISSN: 0168-9002. DOI: 10.1016/j.nima.2014.11.027.
- [14] E. Bols et al. “Jet Flavour Classification Using DeepJet”. In: *JINST* 15.12 (2020), P12012. DOI: 10.1088/1748-0221/15/12/P12012.
- [15] L. Bonechi et al. “A projective reconstruction method of underground or hidden structures using atmospheric muon absorption data”. In: *JINST* 10.02 (2015), P02003. DOI: 10.1088/1748-0221/10/02/P02003.
- [16] L. Bonechi, R. D’Alessandro, and A. Giammanco. “Atmospheric muons as an imaging tool”. In: *Rev. Phys.* 5 (2020), p. 100038. DOI: 10.1016/j.revip.2020.100038.

- [17] K. N. Borozdin et al. “Radiographic imaging with cosmic-ray muons”. In: *Nature* 422 (2003), p. 277. DOI: 10.1038/422277a.
- [18] S. Bouteille et al. “A Micromegas-based telescope for muon tomography: The WatTo experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 834 (2016), pp. 223–228. ISSN: 0168-9002. DOI: 10.1016/j.nima.2016.08.002.
- [19] D. Buterez et al. “Graph Neural Networks with Adaptive Readouts”. In: *arXiv e-prints* (Nov. 2022). DOI: 10.48550/arXiv.2211.04952.
- [20] C. Cârloganu et al. “Towards a muon radiography of the Puy de Dôme”. In: *Geoscientific Instrumentation, Methods and Data Systems* 2.1 (2013), pp. 55–60. DOI: 10.5194/gi-2-55-2013.
- [21] *Characterization of conditioned nuclear waste for its safe disposal in Europe*. URL: <https://cordis.europa.eu/project/id/755371> (visited on 06/02/2025).
- [22] J. Chen et al. “Towards a muon scattering tomography system for both low-Z and high-Z materials”. In: *JINST* 18.08 (2023), P08008. DOI: 10.1088/1748-0221/18/08/P08008.
- [23] T. Cheng. “Recursive Neural Networks in Quark/Gluon Tagging”. In: *Comput. Softw. Big Sci.* 2.1 (2018), p. 3. DOI: 10.1007/s41781-018-0007-y.
- [24] AMS Collaboration. “The Alpha Magnetic Spectrometer (AMS) on the international space station: Part II — Results from the first seven years”. In: *Physics Reports* 894 (2021). The Alpha Magnetic Spectrometer (AMS) on the International Space Station: Part II - Results from the First Seven Years, pp. 1–116. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2020.09.003.
- [25] CALET Collaboration. “Observation of Spectral Structures in the Flux of Cosmic-Ray Protons from 50 GeV to 60 TeV with the Calorimetric Electron Telescope on the International Space Station”. In: *Phys. Rev. Lett.* 129 (10 Sept. 2022), p. 101102. DOI: 10.1103/PhysRevLett.129.101102.

- [26] DAMPE Collaboration. “Measurement of the cosmic ray proton spectrum from 40 GeV to 100 TeV with the DAMPE satellite”. In: *Science Advances* 5.9 (2019), eaax3793. DOI: 10.1126/sciadv.aax3793.
- [27] GEANT4 Collaboration. *Physics Reference Manual*. 2024. URL: <https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsReferenceManual/fo/PhysicsReferenceManual.pdf>.
- [28] Ice Cube Collaboration. “First year performance of the IceCube neutrino telescope”. In: *Astroparticle Physics* 26.3 (2006), pp. 155–173. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2006.06.007.
- [29] KM3NeT Collaboration. “Letter of intent for KM3NeT 2.0”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.8 (June 2016), p. 084001. DOI: 10.1088/0954-3899/43/8/084001.
- [30] PAMELA Collaboration. “PAMELA Measurements of Cosmic-Ray Proton and Helium Spectra”. In: *Science* 332.6025 (2011), pp. 69–72. DOI: 10.1126/science.1199172.
- [31] Pierre Auger Collaboration. “Properties and performance of the prototype instrument for the Pierre Auger Observatory”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 523.1 (2004), pp. 50–95. ISSN: 0168-9002. DOI: 10.1016/j.nima.2003.12.012.
- [32] European Commission et al. *Study of the capability of muon tomography to map the material composition inside a blast furnace (MuBlast) – Final report*. Publications Office, 2018. DOI: doi/10.2777/24858.
- [33] Arthur H. Compton. “A Geographic Study of Cosmic Rays”. In: *Phys. Rev.* 43 (6 Mar. 1933), pp. 387–403. DOI: 10.1103/PhysRev.43.387.
- [34] N.V. Mokhov D.E. Groom and S.I. Striganov. “Muon stopping-power and range tables: 10MeV–100 TeV”. In: *Atomic Data and Nuclear Data Tables* 78 (2001), pp. 183–356.

- [35] M. D’Errico et al. “Muon radiography applied to volcanoes imaging: the MURAVES experiment at Mt. Vesuvius”. In: *Journal of Instrumentation* 15.03 (Mar. 2020), p. C03014. DOI: 10.1088/1748-0221/15/03/C03014.
- [36] B. Denby. “Neural networks and cellular automata in experimental high energy physics”. In: *Computer Physics Communications* 49.3 (1988), pp. 429–448. ISSN: 0010-4655. DOI: 10.1016/0010-4655(88)90004-5.
- [37] T. Dorigo et al. “Toward the end-to-end optimization of particle physics instruments with differentiable programming”. In: *Reviews in Physics* 10 (2023), p. 100085. ISSN: 2405-4283. DOI: 10.1016/j.revip.2023.100085.
- [38] L. I. Dorman. “Cosmic rays : variations and space explorations”. In: 1974. URL: <https://api.semanticscholar.org/CorpusID:120880781>.
- [39] M. Lagrange F. Bury. “Scattering-Based Machine Learning Algorithms for Momentum Estimation in Muon Tomography”. In: *Particles* 8.2 (2025). ISSN: 2571-712X. DOI: 10.3390/particles8020043.
- [40] W. Falcon and The PyTorch Lightning team. “PyTorch Lightning”. In: 2019. DOI: 10.5281/zenodo.3828935. URL: <https://github.com/Lightning-AI/lightning>.
- [41] Felipe Delestro Matos et al. *QIM3D*. Version 1.2.0. June 2025. URL: <https://platform.qim.dk/qim3d/>.
- [42] Matthias Fey and Jan E. Lenssen. “Fast Graph Representation Learning with PyTorch Geometric”. In: *ICLR Workshop on Representation Learning on Graphs and Manifolds*. 2019.
- [43] L. Frazão et al. “Discrimination of high-Z materials in concrete-filled containers using muon scattering tomography”. In: *Journal of Instrumentation* 11.07 (July 2016), P07020. DOI: 10.1088/1748-0221/11/07/P07020.
- [44] Thomas K. Gaisser. “Cosmic Rays and Particle Physics”. In: Cambridge University Press, 1990. Chap. Muons, p. 69.

- [45] Thomas K. Gaisser and A. M. Hillas. “Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers”. In: 1977. URL: <https://api.semanticscholar.org/CorpusID:117809229>.
- [46] A. Giammanco et al. “First Results, and Experimental Status of the MURAVES Experiment 2024”. In: *Journal of Advanced Instrumentation in Science* 2024.1 (Mar. 2024). DOI: 10.31526/jais.2024.501.
- [47] Andrea Giammanco et al. “Toward using cosmic rays to image cultural heritage objects”. In: *iScience* 28 (2025), p. 2589. DOI: 10.1016/j.isci.2025.112094. arXiv: 2405.10417 [physics.soc-ph].
- [48] D. Gibert et al. “Observation of the Dynamics of Hydrothermal Activity in La Soufrière of Guadeloupe Volcano with Joint Muography, Gravimetry, Electrical Resistivity Tomography, Seismic and Temperature Monitoring”. In: *Muography*. American Geophysical Union (AGU), 2022. Chap. 5, pp. 55–73. ISBN: 9781119722748. DOI: 10.1002/9781119722748.ch5.
- [49] Mengyun Guan et al. *A parametrization of the cosmic-ray muon flux at sea-level*. 2015. arXiv: 1509.06176 [hep-ex].
- [50] J. Hadamard. *Mémoire sur le problème d’analyse relatif à l’équilibre des plaques élastiques encastrées*. Mémoires présentés par divers savants à l’Académie des sciences de l’Institut de France: Extrait. Imprimerie nationale, 1908. URL: <http://books.google.com.au/books?id=BTEPAAAAIAAJ>.
- [51] C. Hagmann, D. Lange, and D. Wright. “Cosmic-ray shower generator (CRY) for Monte Carlo transport codes”. In: *2007 IEEE Nuclear Science Symposium Conference Record*. Vol. 2. IEEE, 2007, pp. 1143–1146.
- [52] W. L. Hamilton, R. Ying, and J. Leskovec. “Inductive Representation Learning on Large Graphs”. In: *arXiv e-prints* (June 2017). DOI: 10.48550/arXiv.1706.02216.
- [53] Z. S. Hartwig and P. Gumplinger. “Simulating response functions and pulse shape discrimination for organic scintillation detectors with Geant4”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Asso-*

- ciated Equipment* 737 (2014), pp. 155–162. ISSN: 0168-9002. DOI: 10.1016/j.nima.2013.11.027.
- [54] Dieter Heck, Tanguy Pierog, and Johannes Knapp. *CORSIKA: An Air Shower Simulation Program*. Astrophysics Source Code Library, record ascl:1202.006. Feb. 2012.
 - [55] R. Hess. *Blender Foundations: The Essential Guide to Learning Blender 2.6*. Focal Press, 2010. ISBN: 0240814304, 9780240814308.
 - [56] Victor F. Hess. “Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten”. In: *Physikalische Zeitschrift* 13 (1912), pp. 1084–1091. URL: <https://api.semanticscholar.org/CorpusID:125608176>.
 - [57] *Identification of Jets Containing b-Hadrons with Recurrent Neural Networks at the ATLAS Experiment*. Tech. rep. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2255226>.
 - [58] International Atomic Energy Agency. *Muon imaging: Present Status and Emerging Applications*. IAEA TECDOC 2012. IAEA, Vienna, 2022. URL: <https://www.iaea.org/publications/15182/muon-imaging>.
 - [59] D. Jiang et al. “Denoising of 3D magnetic resonance images with multi-channel residual learning of convolutional neural network”. In: *Japanese Journal of Radiology* 36 (2018), pp. 566–574. DOI: 10.1007/s11604-018-0758-8.
 - [60] K. Jourde et al. “Experimental detection of upward going cosmic particles and consequences for correction of density radiography of volcanoes”. In: *Geophysical Research Letters* 40.24 (2013), pp. 6334–6339. DOI: 10.1002/2013GL058357.
 - [61] G. Kasieczka et al. “Deep-learning Top Taggers or The End of QCD?”. In: *JHEP* 05 (2017), p. 006. DOI: 10.1007/JHEP05(2017)006.
 - [62] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization”. In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. 2015. URL: <http://arxiv.org/abs/1412.6980>.

- [63] B. Lefevre et al. “3D Reconstruction of a Nuclear Reactor by Muon Tomography: Structure Validation and Anomaly Detection”. In: *PRX Energy* 4 (1 Jan. 2025), p. 013002. DOI: 10.1103/PRXEnergy.4.013002.
- [64] G. Li et al. “DeepGCNs: Can GCNs Go as Deep as CNNs?” In: *arXiv e-prints* (Apr. 2019). DOI: 10.48550/arXiv.1904.03751.
- [65] L. Liu et al. “On the Variance of the Adaptive Learning Rate and Beyond”. In: *arXiv e-prints* (2019). DOI: 10.48550/arXiv.1908.03265.
- [66] Z. Lu and Y. Chen. “Single Image Super Resolution based on a Modified U-net with Mixed Gradient Loss”. In: *Signal, Image and Video Processing* 16 (2022). URL: <https://link.springer.com/article/10.1007/s11760-021-02063-5>.
- [67] P. Martínez et al. “Process monitoring: measurement of the metal–slag interface in furnace ladles”. In: *IAEA TECDOC SERIES. Muon Imaging. Present Status and Emerging Applications 2012* (2022), pp. 34–35. URL: <https://www.iaea.org/publications/15182/muon-imaging>.
- [68] B. Masschaele et al. “HECTOR: A 240kV micro-CT setup optimized for research”. In: *Journal of Physics: Conference Series* 463.1 (2013), p. 012012.
- [69] D. Maurin et al. “A cosmic-ray database update: CRDB v4.1”. In: *The European Physical Journal C* 83.10 (Oct. 2023). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-12092-8.
- [70] P. K. Mohanty et al. “Transient Weakening of Earth’s Magnetic Shield Probed by a Cosmic Ray Burst”. In: *Phys. Rev. Lett.* 117 (17 Oct. 2016), p. 171101. DOI: 10.1103/PhysRevLett.117.171101.
- [71] Spandan Mondal and Luca Mastrolorenzo. “Machine learning in high energy physics: a review of heavy-flavor jet tagging at the LHC”. In: *Eur. Phys. J. ST* 233.15-16 (2024), pp. 2657–2686. DOI: 10.1140/epjs/s11734-024-01234-y.
- [72] M. P. Morigi et al. “Application of X-ray computed tomography to cultural heritage diagnostics”. In: *Applied Physics A* 100.3 (2010), pp. 653–661.

- [73] K. Morishima et al. “Discovery of a big void in Khufu’s Pyramid by observation of cosmic-ray muons”. In: *Nature* 552.7685 (Dec. 2017), pp. 386–390. ISSN: 1476-4687. DOI: 10.1038/nature24647.
- [74] C. L. Morris et al. “Horizontal cosmic ray muon radiography for imaging nuclear threats”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 330 (2014), pp. 42–46. ISSN: 0168-583X. DOI: 10.1016/j.nimb.2014.03.017.
- [75] *Muon systems*. 2015. URL: <https://muon.systems/en/muoncargo> (visited on 06/02/2025).
- [76] S. Navas et al. “Review of particle physics”. In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.
- [77] E. Niederleithinger et al. *Muon tomography of a reinforced concrete block – first experimental proof of concept*. 2020. URL: <https://arxiv.org/abs/2008.07251>.
- [78] W. O’Donnell et al. “Muographic Image Upsampling with Machine Learning for Built Infrastructure Applications”. In: *Particles* 8.1 (2025). ISSN: 2571-712X. DOI: 10.3390/particles8010033.
- [79] L. Oláh et al. “Development of Machine Learning-Assisted Spectra Analyzer for the NEWCUT Muon Spectrometer”. In: *Journal of Advanced Instrumentation in Science* 2022 (Apr. 2022). DOI: 10.31526/jais.2022.264.
- [80] L. Oláh et al. “High-definition and low-noise muography of the Sakurajima volcano with gaseous tracking detectors”. In: *Scientific Reports* 8.1 (Feb. 2018), p. 3207. ISSN: 2045-2322. DOI: 10.1038/s41598-018-21423-9.
- [81] “Muography: Exploring Earth’s Subsurface with Elementary Particles”. In: ed. by László Oláh, Hiroyuki K. M. Tanaka, and Dezső Varga. American Geophysical Union, 2022. Chap. Tomographic Imaging of Volcanic Structures with Cosmic-Ray Muons. DOI: 10.1002/9781119722748.
- [82] D. Pagano et al. “EcoMug: An Efficient COsmic MUon Generator for cosmic-ray muon applications”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1014 (2021), p. 165732. ISSN: 0168-9002. DOI: 10.1016/j.nima.2021.165732.

- [83] Particle Data Group. *Atomic and Nuclear Properties of Materials*. <https://pdg.lbl.gov/2015/AtomicNuclearProperties/>. Part of the Review of Particle Physics. 2015.
- [84] A. Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [85] D. Poulson et al. “Application of muon tomography to fuel cask monitoring”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377.2137 (2019), p. 20180052. DOI: 10.1098/rsta.2018.0052.
- [86] D. Lo Presti et al. “Three Years of Muography at Mount Etna, Italy”. In: *Muography*. American Geophysical Union (AGU), 2022. Chap. 7, pp. 93–108. ISBN: 9781119722748. DOI: 10.1002/9781119722748.ch7.
- [87] W. C. Priedhorsky et al. “Detection of high-Z objects using multiple scattering of cosmic ray muons”. In: *Review of Scientific Instruments* 74.10 (Oct. 2003), pp. 4294–4297. ISSN: 0034-6748. DOI: 10.1063/1.1606536.
- [88] H. Qu and L. Gouskos. “ParticleNet: Jet Tagging via Particle Clouds”. In: *Phys. Rev. D* 101.5 (2020), p. 056019. DOI: 10.1103/PhysRevD.101.056019.
- [89] H. Qu, C. Li, and S. Qian. “Particle Transformer for Jet Tagging”. In: (2022). arXiv: 2202.03772 [hep-ph].
- [90] E. T. Rand et al. “Nonparametric Dense-Object Detection Algorithm for Applications of Cosmic-Ray Muon Tomography”. In: *Phys. Rev. Appl.* 14 (6 2020), p. 064032. DOI: 10.1103/PhysRevApplied.14.064032.
- [91] B. C. Rastin. “An accurate measurement of the sea-level muon spectrum within the range 4 to 3000 GeV/c”. In: *Journal of Physics G: Nuclear Physics* 10.11 (Nov. 1984), p. 1609. DOI: 10.1088/0305-4616/10/11/017.

- [92] J. Ren et al. *SCA3D: Enhancing Cross-modal 3D Retrieval via 3D Shape and Caption Paired Data Augmentation*. 2025. arXiv: 2502.19128 [cs.CV]. URL: <https://arxiv.org/abs/2502.19128>.
- [93] J. Resnick. “The radon transforms and some of its applications”. In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 33.1 (1985), pp. 338–339. DOI: 10.1109/TASSP.1985.1164533.
- [94] Francesco Riggi. *Messengers from the Cosmos*. Springer Cham, 2024. DOI: 10.1007/978-3-031-24762-0.
- [95] A. Bueno Rodriguez et al. “Advanced Material Identification in Muon Tomography”. In: *European Workshop on Maritime Systems Resilience and Security 2024 (MARESEC 2024)*. MARESEC 2024 Proceedings, Nov. 2024. URL: <https://elib.dlr.de/209907/>.
- [96] A.B. Rodriguez et al. “B2G4: A synthetic data pipeline for the integration of Blender models in Geant4 simulation toolkit”. In: *arXiv preprint arXiv:2311.06327* (2023).
- [97] B. P. Roe et al. “Boosted decision trees as an alternative to artificial neural networks for particle identification”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 543.2 (2005), pp. 577–584. ISSN: 0168-9002. DOI: 10.1016/j.nima.2004.12.018.
- [98] G. Van Rossum and F.L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. ISBN: 1441412697.
- [99] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. “Learning representations by back-propagating errors”. In: *Nature* 323 (Oct. 1986), pp. 533–536. DOI: 10.1038/323533a0.
- [100] H. Sato, T. Kin, and A. Giammanco. “Measurement of energy differential spectrum of cosmic-ray muons below 400 MeV”. In: *Journal of Instrumentation* 17.08 (2022), P08009.
- [101] L. J. Schultz et al. “Statistical Reconstruction for Cosmic Ray Muon Tomography”. In: *IEEE Transactions on Image Processing* 16.8 (2007), pp. 1985–1993. DOI: 10.1109/TIP.2007.901239.

- [102] L.J. Schultz et al. “Image reconstruction and material Z discrimination via cosmic ray muon radiography”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 519.3 (2004), pp. 687–694. ISSN: 0168-9002. DOI: 10.1016/j.nima.2003.11.035.
- [103] S J Sciutto. “AIRES A system for air shower simulations. User’s guide and reference manual”. en. In: (2019). DOI: 10.13140/RG.2.2.12566.40002.
- [104] Raman Sehgal et al. “Application of Machine learning in Muon Tomography”. In: 2020. URL: <https://api.semanticscholar.org/CorpusID:221969336>.
- [105] G. Shen et al. “Measurement of multiple scattering at 50 to 200 GeV/c”. In: *Phys. Rev. D* 20 (7 Oct. 1979), pp. 1584–1588. DOI: 10.1103/PhysRevD.20.1584.
- [106] Prashant Shukla and Sundaresh Sankrith. “Energy and angular distributions of atmospheric muons at the Earth”. In: *Int. J. Mod. Phys. A* 33.30 (2018), p. 1850175. DOI: 10.1142/S0217751X18501750.
- [107] *Silent Border*. URL: <https://silentborder.eu/> (visited on 06/02/2025).
- [108] Todor Stanev. “Cosmic Ray Showers”. In: *High Energy Cosmic Rays*. Cham: Springer International Publishing, 2021, pp. 179–229. ISBN: 978-3-030-71567-0. DOI: 10.1007/978-3-030-71567-0_8.
- [109] M. Stapleton et al. “Angle Statistics Reconstruction: a robust reconstruction algorithm for Muon Scattering Tomography”. In: *Journal of Instrumentation* 9.11 (Nov. 2014), P11019. DOI: 10.1088/1748-0221/9/11/P11019.
- [110] T. J. Stocki et al. “Machine learning for the cosmic ray inspection and passive tomography project (CRIPT)”. In: *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*. 2012, pp. 91–94. DOI: 10.1109/NSSMIC.2012.6551067.
- [111] G. C. Strong et al. “TomOpt: differential optimisation for task- and constraint-aware design of particle detectors in the context of muon tomography”. In: *Machine Learning: Science and Technology* 5.3 (July 2024), p. 035002. DOI: 10.1088/2632-2153/ad52e7.

- [112] G. Chatham Strong et al. *TomOpt: Differential Muon Tomography Optimisation*. Version v0.1.0. Feb. 2024. DOI: 10.5281/zenodo.10673885.
- [113] H. Su et al. “Multi-view Convolutional Neural Networks for 3D Shape Recognition”. In: *arXiv e-prints* (2015). DOI: 10.48550/arXiv.1505.00880.
- [114] C. Thomay et al. “A binned clustering algorithm to detect high-Z material using cosmic muons”. In: *Journal of Instrumentation* 8.10 (Oct. 2013), P10013. DOI: 10.1088/1748-0221/8/10/P10013.
- [115] H. Touvron et al. “Going deeper with Image Transformers”. In: *arXiv e-prints* (2021). DOI: 10.48550/arXiv.2103.17239.
- [116] Yung-Su Tsai. “Pair production and bremsstrahlung of charged leptons”. In: *Rev. Mod. Phys.* 46 (4 Oct. 1974), pp. 815–851. DOI: 10.1103/RevModPhys.46.815.
- [117] S. Vanini et al. “Muography of different structures using muon scattering and absorption algorithms”. In: *Philosophical Transactions of the Royal Society A* 377 (2018). DOI: 10.1098/rsta.2018.0051.
- [118] A. Vaswani et al. “Attention Is All You Need”. In: *arXiv e-prints* (2017). DOI: 10.48550/arXiv.1706.03762.
- [119] Gaber MA Wahab, Mostafa Gouda, and Gamal Ibrahim. “Study of physical and mechanical properties for some of Eastern Desert dimension marble and granite utilized in building decoration”. In: *Ain Shams Engineering Journal* 10.4 (2019), pp. 907–915.
- [120] G. Wang, L. Schultz, and J. Qi. “Statistical Image Reconstruction for Muon Tomography Using a Gaussian Scale Mixture Model”. In: *IEEE Transactions on Nuclear Science* 56.4 (2009), pp. 2480–2486. DOI: 10.1109/TNS.2009.2023518.
- [121] Y. Wang et al. “Dynamic Graph CNN for Learning on Point Clouds”. In: *arXiv e-prints* (2018). DOI: 10.48550/arXiv.1801.07829.
- [122] Z. Wang and A. C. Bovik. “Mean squared error: Love it or leave it? A new look at Signal Fidelity Measures”. In: *IEEE Signal Processing Magazine* 26.1 (2009), pp. 98–117. DOI: 10.1109/MSP.2008.930649.

- [123] Z. Wang et al. “Image quality assessment: From error visibility to structural similarity”. In: *IEEE Transactions on Image Processing* 13 (4 Apr. 2004), pp. 600–612. ISSN: 10577149. DOI: 10.1109/TIP.2003.819861.
- [124] L. S. Waters et al. “The MCNPX Monte Carlo Radiation Transport Code”. In: *AIP Conference Proceedings* 896.1 (Mar. 2007), pp. 81–90. ISSN: 0094-243X. DOI: 10.1063/1.2720459.
- [125] M. J. Weekes et al. “Material identification in nuclear waste drums using muon scattering tomography and multivariate analysis”. In: *Journal of Instrumentation* 16.05 (May 2021), P05007. DOI: 10.1088/1748-0221/16/05/P05007.
- [126] K. Xu et al. “Representation Learning on Graphs with Jumping Knowledge Networks”. In: *arXiv e-prints* (June 2018). DOI: 10.48550/arXiv.1806.03536.
- [127] K. Zhang et al. “Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising”. In: *IEEE Transactions on Image Processing* 26.7 (2017), pp. 3142–3155. DOI: 10.1109/TIP.2017.2662206.