Composition studies based on the muon content of air showers measured by underground muon detectors

Zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

von der KIT-Fakultät für Physik des
Karlsruher Instituts für Technologie (KIT)
und
vom Instituto de Tecnología "Prof. Jorge A. Sábato" der
Universidad Nacional de San Martín (UNSAM)

genehmigte Dissertation

von

MSc. Varada Varma Kizakke Covilakam
geboren in Manjeri

Tag der mündlichen Prüfung: 08/05/2026

Referent: Prof. Dr. Alberto Daniel Supanitsky

Korreferent: Prof. Dr. Ralph Engel

Betreuer: Dr. David Schmidt

Composition studies based on the muon content of air showers measured by underground muon detectors

For the attainment of the academic degree of

Doctorate in Science

from
the Faculty of Physics of the
Karlsruher Instituts für Technologie (KIT)
and
Instituto de Tecnología "Prof. Jorge A. Sábato" of the
Universidad Nacional de San Martín (UNSAM)

accepted dissertation

of

MSc. Varada Varma Kizakke Covilakam
born in Manjeri

Day of the oral examination: 08/05/2026

Referee: Prof. Dr. Alberto Daniel Supanitsky

Co-referee: Prof. Dr. Ralph Engel

Supervisor: Dr. David Schmidt

Composition studies based on the muon content of air showers measured by underground muon detectors

Para optar al título de

Doctor en Astrofísica

del Instituto de Tecnología "Prof. Jorge A. Sábato" de la
Universidad Nacional de San Martín (UNSAM)
y
la facultad de física del
Karlsruher Instituts für Technologie (KIT)

tesis aceptada

de

MSc. Varada Varma Kizakke Covilakam

nacido en Manjeri

Fecha de la defensa oral: 08/05/2026

Director: Prof. Dr. Alberto Daniel Supanitsky

Co-director: Prof. Dr. Ralph Engel

Colaborador: Dr. David Schmidt

Abstract

Ultra-high-energy cosmic rays (UHECR) are among the most energetic particles observed in nature, yet their origin and mass composition remain poorly understood. At energies greater than approximately 10^{15} eV, direct detection becomes impractical, and primary particles are studied through the extensive air showers (EAS) they generate in the atmosphere. The muonic component of these showers provides a particularly sensitive probe of both primary mass and hadronic interactions at energies beyond the reach of terrestrial accelerators. This thesis investigates the muon lateral distribution (MLD) and muon-related observables measured with the Underground Muon Detector (UMD) of the Pierre Auger Observatory to enhance composition sensitivity and assess the consistency between measured data and simulations of EAS based on different high-energy hadronic interaction models.

The shape of the MLD is first characterised by direct comparisons between measured radial muon densities and full detector-level Monte Carlo (MC) simulations of the EAS, constructed as weighted mixtures of primary nuclei (proton, helium, nitrogen, and iron) with fractions defined by composition models. For the first time, complete air-shower and detector simulations are compared with UMD data, including the radial dependence of the muon distribution. A bin-by-bin analysis framework is developed, in which the comparison between measurements and simulations is performed independently in each energy bin, to test whether the discrepancy can be described by a single energy-dependent multiplicative scale factor applied to a reference mixed-composition model (AugerMix), constructed as a weighted combination of primary nuclei based on composition fractions inferred from measurements of the depth of the shower maximum X_{max}. The study is based on an extended UMD data set comprising data collected in 2018–2021 (Phase 1), together with data recorded in 2023–2024 (Phase 2) following a large-scale upgrade of the detector system, including improved front-end electronics. It includes, for the first time in this context, measurements obtained with the upgraded Phase 2 electronics. The comparison reveals a systematic energy-dependent muon deficit in simulations relative to data. This conclusion is found to be robust in all contemporary high-energy hadronic interaction models, and the radial dependence of the discrepancy remains consistent, within systematic uncertainties, with a global rescaling of the simulated muon densities. This supports the interpretation of the data-simulation mismatch as a predominantly normalisation effect and provides a consistent basis for the development and interpretation of $V_{b_{\mathrm{max}}}$ observable in subsequent correlation analysis.

Using reconstructed muon densities, mass-sensitive observables based on the spatial structure of the muon distribution are developed and optimised through merit-factor studies. The statistical correlation between one of the optimised observables $V_{b_{\mathrm{max}}}$ and the depth of the shower maximum, X_{max}, is investigated as a composition probe, exploiting the robustness of correlation-based methods against calibration and scaling uncertainties. The sensitivity of this approach is quantified for both fluorescence-detector measurements and high-statistics surface-detector scenarios, including realistic detector-resolution effects. Among the tested estimators, Kendall's correlation coefficient provides the highest stability and discrimination power. The correlation-based framework is shown to constrain both the mean and the dispersion of the primary mass distribution, enabling discrimination between pure and mixed composition scenarios. Using the measured cosmic ray (CR) flux reported by the Pierre Auger Observatory and the statistical properties of Kendall's coefficient, the minimum number of events required to reject the hypothesis of pure composition is derived.

To further strengthen the reliability of the muon-based observables and to reduce potential biases associated with detector response, improved reconstruction techniques for the MLDFs are developed, based on both integrated signal measurements (continuous detector response) and hybrid combinations with particle-counting modes. Their performance is evaluated with MC simulations and compared with existing reconstruction approaches. Although these improved reconstruction methods are not propagated in the primary composition analysis presented in this thesis, they provide a foundation for future high-precision muon measurements and refined correlation-based studies.

Together, these studies establish a comprehensive framework linking muon measurements, reconstruction methodology, and correlation-based observables to constrain the mass composition of UHECR and to test hadronic interaction models at ultra-high energies.

Resumen

Los rayos cósmicos de ultra-alta energía (UHECR, por sus siglas en inglés) se encuentran entre las partículas más energéticas observadas en la naturaleza; sin embargo, su origen y composición en masa aún no se comprenden completamente. A energías mayores que aproximadamente 10^{15} eV, la detección directa se vuelve impracticable, por lo que las partículas primarias se estudian a través de las cascadas atmosféricas extensas (EAS, por sus siglas en inglés) que se generan en la atmósfera. La componente muónica de estas cascadas constituye un parámetro particularmente sensible tanto a la masa primaria como a las interacciones hadrónicas a energías que superan el alcance de los aceleradores terrestres. Esta tesis investiga la distribución lateral de muones (MLD, por sus siglas en inglés) y observables relacionados con muones medidos con el Detector de Muones Subterráneo (UMD, por sus siglas en inglés) del Observatorio Pierre Auger, con el objetivo de mejorar la sensibilidad en composición y evaluar la consistencia entre los datos medidos y simulaciones de EAS basadas en diferentes modelos hadrónicos de altas energías.

La forma de la MLD se caracteriza inicialmente mediante comparaciones directas entre densidades radiales de muones medidas y simulaciones de Monte Carlo (MC) completas a nivel de los detectores de EAS, construidas como combinaciones lineales de densidades de muones correspondientes a núcleos primarios (protón, helio, nitrógeno y hierro) con fracciones de masa definidas por modelos de composición. Por primera vez, se comparan simulaciones completas de cascadas atmosféricas y de los detectores con datos del UMD, incluyendo la dependencia radial de la distribución de muones. Se desarrolla un marco de análisis bin a bin, en el que la comparación entre mediciones y simulaciones se realiza de forma independiente en cada bin de energía, para evaluar si la discrepancia puede describirse mediante un único factor de escala multiplicativo dependiente de la energía, aplicado a un modelo de composición de referencia (AugerMix), construido como una combinación pesada de núcleos primarios basada en fracciones de composición inferidas a partir de mediciones de la profundidad del máximo de la cascada X_{max}. El estudio se basa en un conjunto de datos extendido del UMD que comprende datos recolectados en 2018–2021 (Fase 1), junto con datos registrados en 2023–2024 (Fase 2) tras una actualización a gran escala del sistema de detección, incluyendo la electrónica de front-end mejorada. Este estudio incluye, por primera vez en este contexto, mediciones obtenidas con la electrónica actualizada de la Fase 2. La comparación revela un déficit sistemático de muones en las simulaciones con respecto a los datos dependiente de la energía. Esta conclusión resulta robusta frente a todos los modelos contemporáneos de interacciones hadrónicas de alta energía, y la dependencia radial de la discrepancia es consistente, dentro de las incertidumbres sistemáticas, con un cambio de escala global de las densidades de muones simuladas. Esto respalda la interpretación de la discrepancia entre datos y simulaciones como un efecto predominantemente

de normalización y proporciona una base coherente para el desarrollo e interpretación del observable $V_{b_{\max}}$en el análisis de correlación subsiguiente.

A partir de las densidades de muones reconstruidas, se desarrollan observables sensibles a la masa basados en la estructura espacial de la distribución de muones, optimizados mediante estudios del factor de mérito. Entre ellos, se selecciona el observable $V_{b_{\max}}$,para investigar su correlación estadística con la profundidad del máximo de la cascada, $X_{\max}$, con el objetivo de realizar estudios de composición, aprovechando la robustez de los métodos basados en correlaciones frente a incertidumbres de calibración y a incertidumbres sistemáticas que se manifiestan como factores de escala. La sensibilidad de este enfoque se cuantifica tanto para mediciones de $X_{\max}$ con el detector de fluorescencia como para mediciones obtenidas a partir del detector de superficie con alta estadística, incluyendo efectos realistas de resolución instrumental. Entre los estimadores evaluados, el coeficiente de correlación de Kendall presenta la mayor estabilidad y poder de discriminación. El marco basado en correlaciones permite restringir tanto la media como la dispersión de la distribución de masas primarias, posibilitando la discriminación entre escenarios de composición pura y mixta. Utilizando el flujo de rayos cósmicos medido por el Observatorio Pierre Auger y las propiedades estadísticas del coeficiente de Kendall, se determina el número mínimo de eventos necesario para rechazar la hipótesis de composición pura.

Para reforzar la fiabilidad de los observables basados en muones y reducir posibles sesgos asociados a la respuesta del detector, se desarrollan técnicas mejoradas de reconstrucción de las MLDFs, basadas tanto en mediciones de señal integrada (respuesta continua del detector) como en combinaciones híbridas con modos de conteo de partículas. Su desempeño se evalúa mediante simulaciones MC y se compara con métodos de reconstrucción existentes. Aunque estos métodos mejorados no se incorporan en el análisis principal de composición presentado en esta tesis, constituyen una base para futuras mediciones de la componente muónica de alta precisión y estudios correlacionales más refinados.

En conjunto, estos estudios establecen un marco integral que vincula mediciones de la componente muónica de las lluvias, metodología de reconstrucción y observables basados en correlaciones para restringir la composición en masa de los UHECR y evaluar modelos de interacción hadrónica a energías ultra-altas.

Zusammenfassung

Ultrahochenergetische kosmische Strahlen (UHECR) gehören zu den energiereichsten in der Natur beobachteten Teilchen, doch ihr Ursprung und ihre Massenzusammensetzung sind weiterhin nur unzureichend verstanden. Bei Energien oberhalb von etwa 10^{15} eV ist eine direkte Detektion nicht praktikabel, sodass die Primärteilchen über die von ihnen erzeugten Schauerkaskaden, sogenannte ausgedehnte Luftschauer (EAS, engl. extended air showers), untersucht werden. Die myonische Komponente dieser Schauer bietet einen besonders sensitiven Zugang sowohl zur Primärmasse als auch zu hadronischen Wechselwirkungen bei Energien jenseits derjenigen, welche in terrestrischen Beschleunigern erreicht werden können. Diese Dissertation untersucht die myonische laterale Verteilung (MLD, von engl. muon lateral distribution) sowie andere myonenabhängige Observablen, die mit dem Underground Muon Detector (UMD) des Pierre-Auger-Observatoriums gemessen werden, um die Empfindlichkeit gegenüber der Massenkomposition zu erhöhen und die Konsistenz zwischen gemessenen Daten und Simulationen von EAS auf Basis verschiedener hadronischer Wechselwirkungsmodelle zu bewerten.

Die Form der MLD wird zunächst durch direkte Vergleiche zwischen gemessenen radialen Myondichten und vollständigen Monte-Carlo-(MC)-Simulationen von EAS auf Detektorebene charakterisiert. Dabei wird die Massenkomposition der Simulationen aus einer gewichteten Mischung von primärer Kerne (Proton, Helium, Stickstoff und Eisen) konstruiert, welche den Anteilen von Kompositionsmodellen entspricht. Erstmals werden die radialen Abhängigkeiten der Myonenverteilung von Luftschauer- und Detektorsimulationen mit denen von UMD-Messungen verglichen. Es wird ein binweiser Analyseansatz entwickelt, bei dem der Vergleich zwischen Messungen und Simulationen unabhängig in jedem Energiebereich erfolgt, um zu testen, ob die Diskrepanz durch einen einzelnen energieabhängigen multiplikativen Skalierungsfaktor beschrieben werden kann. Dieser Skalierungsfaktor wird auf ein Referenz-Kompositionsmodell (AugerMix) angewandt, das als gewichtete Kombination primärer Kerne auf Basis von aus der Tiefe des Schauermaximums X_{max}-Messungen abgeleiteter Kompositionsanteile konstruiert ist. Die Studie basiert auf einem erweiterten UMD-Datensatz, der Daten aus den Jahren 2018–2021 (Phase 1) sowie Daten aus 2023–2024 (Phase 2) umfasst, die nach einer groß angelegten Erweiterung und Modernisierung des Detektorsystems, einschließlich verbesserter Front-End-Elektronik, aufgenommen wurden. Sie beinhaltet erstmals in diesem Zusammenhang Messungen mit der aufgerüsteten Phase 2-Elektronik. Der Vergleich zeigt ein systematisches, energieabhängiges Myonendefizit in den Simulationen relativ zu den Daten. Dieses Ergebnis ändert sich nicht, wenn andere, aktuelle, hochernergetische Hadronwechselwirkungsmodelle benutzt werden. Die radiale Abhängigkeit der Diskrepanz ist innerhalb der systematischen Unsicherheiten mit einer globalen Reskalierung der simulierten Myondichten vereinbar. Dies unterstützt die Interpretation, dass die Diskrepanz zwischen Daten und Simulationen überwiegend auf einen Normierungseffekt zurückgeführt werden kann, und liefert eine konsistente Grundlage für die Entwicklung und Interpretation der Observablen $V_{b_{\max}}$ in der nachfolgenden Korrelationsanalyse.

Auf Grundlage rekonstruierter Myondichten werden massensensitive Observablen entwickelt, die auf der räumlichen Verteilung der gemessenen Myonen basieren und mittels Merit-Faktor-Studien optimiert werden. Aus diesem Satz wird das Observable $V_{b_{\max}}$ ausgewählt, dessen statistische Korrelation mit der Tiefe des Schauermaximums $X_{\max}$ als Zusammensetzungsindikator untersucht wird, wobei die Robustheit korrelationsbasierter Methoden gegenüber Kalibrierungs- und Skalierungsunsicherheiten ausgenutzt wird. Die Sensitivität dieses Ansatzes wird sowohl für Messungen mit dem Fluoreszenzdetektor als auch für hochstatistische Messungen des Oberflächendetektors quantifiziert, einschließlich realistischer Detektorauflösungseffekte. Unter den getesteten Korrelationsschätzern weist Kendalls Korrelationskoeffizient die höchste Stabilität und Diskriminationsfähigkeit auf. Der korrelationsbasierte Ansatz ermöglicht es, sowohl den Mittelwert als auch die Streuung der Primärmassenverteilung einzuschränken und zwischen reinen und gemischten Zusammensetzungsszenarien zu unterscheiden. Unter Verwendung des vom Pierre-Auger-Observatorium gemessenen Flusses der kosmischen Strahlen sowie der statistischen Eigenschaften von Kendalls Koeffizienten wird die minimale Ereigniszahl bestimmt, die erforderlich ist, um die Hypothese einer reinen Zusammensetzung zu verwerfen.

Um die Zuverlässigkeit der muonbasierten Observablen weiter zu erhöhen und mögliche Verzerrungen durch die Detektorantwort zu reduzieren, werden verbesserte Rekonstruktionsmethoden für die MLDFs entwickelt, die sowohl auf integrierten Signalmessungen (kontinuierliche Detektorantwort) als auch auf hybriden Kombinationen mit Teilchenzählmodi basieren. Ihre Leistungsfähigkeit wird mittels MC-Simulationen bewertet und mit bestehenden Rekonstruktionsansätzen verglichen. Obwohl diese verbesserten Rekonstruktionsmethoden nicht in die primäre Zusammensetzungsanalyse dieser Dissertation einfließen, bilden sie eine Grundlage für zukünftige hochpräzise Myonmessungen und weiterentwickelte korrelationsbasierte Studien.

Die vorgestellten Studien erlauben es, die Myonenmessungen, Rekonstruktionsmethodik und korrelationsbasierte Observablen miteinander zu verknüpfen. Dadurch kann die Massenzusammensetzung von ultrahochenergetischer kosmischer Strahlung nicht nur eingeschränkt werden, sondern es können auch hadronische Wechselwirkungs modelle bei höheren Energien getestet werden.

Acronyms

The list of acronyms used in this work is alphabetically sorted according to the short version.

Contents

CHAPTER 1

Introduction

Cosmic rays (CR[1]) constitute a continuous flux of highly energetic particles and nuclei of extraterrestrial origin that impinge on the Earth's atmosphere. They consist predominantly of fully ionised atomic nuclei, with a smaller fraction of electrons, photons, and neutrinos, and span more than eleven orders of magnitude in energy, from approximately 10^9 eV to beyond 10^{20} eV [1, 2]. At the highest energies, individual particles carry macroscopic amounts of energy, far exceeding those attainable in human-made accelerators. UHECR[2] therefore provide a natural laboratory for studying particle interactions at centre-of-mass energies beyond the reach of the Large Hadron Collider, while simultaneously probing the most extreme astrophysical environments [3] in the Universe, where such particles must be accelerated.

The CR energy spectrum follows an approximate power-law behaviour, with distinct spectral features such as the knee, the ankle, and flux suppression at the highest energies [2, 4]. Because flux decreases steeply with increasing energy, UHECR are extremely rare. At energies around 10^{17} eV, the flux is of the order of 10 particles per km^2 per day, while near 10^{20} eV it drops to roughly one particle per km^2 per century [2]. Above approximately 10^{15} eV, direct detection using satellite or balloon-borne instruments becomes impractical due to the rapidly decreasing flux. Instead, UHECR are studied indirectly through the EAS[3] they produce when interacting with nuclei in the atmosphere [1]. These cascades of secondary particles develop through successive hadronic and electromagnetic interactions and can spread over areas of tens of square kilometres at the highest energies, enabling their detection with large ground-based observatories.

Despite decades of experimental and theoretical progress, fundamental questions about the origin, acceleration mechanisms, and mass composition of UHECR remain unresolved [3]. Their existence implies acceleration in the most energetic astrophysical environments, potentially including active galactic nuclei, gamma-ray bursts, or other extreme sources. Determining the composition of the nuclear mass as a function of energy is essential for constraining source models, acceleration scenarios, propagation effects, and the transition from galactic to extragalactic CR [5]. However, at ultra-high energies, the composition cannot be measured directly and instead must be inferred from air-shower observables, rendering the interpretation inherently dependent on high-energy hadronic interaction models.

[1] cosmic rays

[2] ultra-high energy cosmic rays

[3] extensive air showers

The development of an extensive air shower depends both on the nature of the primary particle and on hadronic interactions at energies far beyond those accessible in laboratory experiments. Two observables are particularly sensitive to the primary mass: the atmospheric depth of shower maximum, X_{max}, and the muonic component of the shower at ground level [5, 6]. The observable X_{max} characterises the longitudinal development of the predominantly electromagnetic cascade. Lighter primaries, such as protons, penetrate deeper into the atmosphere before reaching maximum development, whereas heavier nuclei interact earlier and exhibit shallower X_{max} distributions. Measurements of X_{max} are performed using FD[4], which record the atmospheric fluorescence light emitted along the shower trajectory [7]. At Pierre Auger Observatory, the FD provides high-resolution measurements of the longitudinal profile but operates only during clear, moonless nights, resulting in a duty cycle of approximately 15% [14] and therefore limited statistics.

Complementary information is provided by the muonic component of air showers. Muons originate predominantly from the decay of charged mesons produced in the hadronic cascade and therefore trace the hadronic backbone of the shower development [1]. Because heavier primaries undergo more nucleon-nucleon interactions and distribute their energy among more secondary particles, they generate a larger number of muons at the ground for a given primary energy. The spatial distribution of muons on the ground, described by the MLD[5], encodes information about both the total muon content and the transverse structure of the hadronic cascade. Pierre Auger Observatory is equipped with dedicated muon-sensitive detectors, including the UMD[6], designed to provide high-statistics measurements of the muonic component [10].

However, interpreting muon measurements is considerably more challenging than interpreting X_{max}. Since muons originate from hadronic processes, their production is strongly dependent on hadronic interaction models, which must extrapolate accelerator data over many orders of magnitude in energy [12, 16, 17]. The systematic uncertainties associated with these models, therefore, propagate directly into composition inference. A long-standing tension between observations and simulations has emerged from this dependence of the model. Several experiments, including Pierre Auger Observatory, have reported that the muon content measured in air showers exceeds that predicted by state-of-the-art simulations [19, 20]. This discrepancy, commonly referred to as the *muon deficit* or *muon puzzle*, implies that the current high-energy hadronic interaction models may underestimate muon production at ultra-high energies. Composition estimates derived from muon measurements often suggest a heavier primary composition than those inferred from X_{max}, indicating that these observables are not consistently reproduced within a common interaction framework. At energies around 10^{19} eV, reported discrepancies can reach tens of percent. The origin of this mismatch remains unclear and could reflect deficiencies in hadronic modelling or previously unaccounted-for physical processes.

Several key aspects of the muon puzzle remain unresolved. It is not yet fully established at which energy the discrepancy emerges, how it evolves with energy, and whether it can be described purely as a global normalisation shift or whether it involves structural modifications of the muon lateral distribution. Addressing these questions requires systematic comparisons between measurements and full detector-level simulations under realistic composition assumptions. Furthermore, not all experiments report identical behaviour, highlighting the need for precise and systematic measurements. Improving our understanding of muon production in air showers is therefore essential not only for reliable composition

[4]Fluorescence detector
[5]muon lateral distribution
[6]Underground Muon detector

inference but also for testing hadronic interactions at energies far beyond those accessible in laboratory experiments.

In this context, this thesis investigates the muonic component of EAS measured with the UMD of Pierre Auger Observatory. The analysis focuses on the energy range $E/\mathrm{eV} \in (10^{17.3}, 10^{18.7})$ eV, where high-statistics measurements enable detailed comparisons between data and MC[7] simulations. The work addresses three interconnected objectives:

1. Quantify the muon deficit through direct, bin-by-bin comparisons between measured radial muon densities and detector-level simulations under mixed-composition scenarios.

2. To construct and evaluate correlation-based composition observables that combine muon-related measurements with estimates of $X_{\max}$, and to assess their statistical discrimination power between pure and mixed composition hypotheses.

3. To develop and validate improved reconstruction techniques for MLDF[8]s using ADC[9] information and hybrid binary– ADC detector modes. This is important to improve muon reconstruction and enhance composition analysis.

A central methodological aspect of this work is the preservation of radial information through bin-by-bin comparisons of measured and simulated muon densities. Instead of reducing the muon content to a single integrated observable, the radial profile is analysed independently in distance intervals from the shower core. This strategy maintains sensitivity to genuine shape differences in the MLD and allows normalisation effects to be disentangled from structural deviations. By extracting energy-dependent multiplicative scale factors relative to the AugerMix expectation and examining residuals as a function of radial distance, the analysis tests whether the discrepancy between data and simulations can be described by a global rescaling within systematic uncertainties or whether additional distortions of the distribution are required. The robustness of these conclusions is evaluated across contemporary high-energy hadronic interaction models.

In parallel, correlation-based techniques are investigated as composition probes that are intrinsically less sensitive to absolute normalisation uncertainties. Statistical correlation coefficients between an observable related to muons and $X_{\max}$ are explored as mass-sensitive indicators. Rank-based estimators, in particular Kendall's correlation coefficient [21, 74], are invariant under monotonic transformations and insensitive to global multiplicative shifts. This property makes them especially suitable in the presence of the muon deficit, where absolute muon predictions may be systematically biased. The statistical properties of different correlation estimators are quantified using MC simulations, and their sensitivity to distinguish pure from mixed composition scenarios is evaluated. Using the measured CR flux and the variance of Kendall's coefficient, the minimum number of events required to reject the hypothesis of pure composition is derived at a given confidence level.

The thesis is organised as follows:

- **Chapter 4** presents a detailed analysis of the average muon lateral distribution. Measured radial muon density profiles are compared with detector-level simulations under mixed-composition assumptions. A global multiplicative scale factor is extracted to test for residual shape differences.
- **Chapter 5** develops mass-sensitive observables derived from the muon density distribution and investigates the reconstruction of $X_{\max}$ using SD[10] information.

[7] Monte Carlo
[8] muon lateral distribution function
[9] Analog-to-digital
[10] Surface detector

- **Chapter 6** explores the correlation-based composition analysis using the statistical relationship between a muon-related observable and X_{max} measured by the FD. Different correlation estimators are evaluated, and the minimum number of events required for composition discrimination is derived.
- **Chapter 7** extends the correlation-based approach to a SD context by introducing a controlled smearing of X_{max} to emulate realistic SD resolution. The robustness and sensitivity of Kendall's rank correlation coefficient under detector-resolution effects are assessed for pure- and mixed composition scenarios.
- **Chapter 8** introduces novel reconstruction methods for MLDFs using ADC-only and hybrid binary- ADC information. Their performance is evaluated with MC simulations and compared with previously established likelihood-based approaches.

The work presented here establishes a coherent framework linking muon measurements, reconstruction methodology, and correlation-based observables to constrain the mass composition of UHECR and to test high-energy hadronic interaction models at energies beyond the reach of the accelerator.

CHAPTER 2

Physics of cosmic rays and extensive air showers

Cosmic rays (CR) are highly energetic particles of extraterrestrial origin that continuously bombard the Earth from all directions. Their study provides a powerful probe of both astrophysical particle acceleration and fundamental particle interactions at energies far beyond those accessible in terrestrial accelerators [1, 3, 25]. At the highest energies, CR constitute the most energetic particles known in nature, reaching energies exceeding 10^{20} eV [25]. Such extreme energies imply acceleration processes operating in the most violent astrophysical environments and provide a natural laboratory for studying particle physics under conditions unattainable in human-made experiments [3].

At energies greater than approximately 10^{14} eV, the flux of CR becomes too low for direct detection using balloon or satellite instruments [1, 94]. Instead, these particles are studied indirectly through the EAS they produce when they interact with nuclei in the atmosphere of the Earth [1, 94]. The atmosphere thus acts simultaneously as a target, calorimeter, and amplifier, converting a single primary particle into a cascade of billions of secondaries distributed over large areas at ground level [1, 92].

This chapter presents an overview of the physics underlying CR and EAS. The discussion includes the global properties of the CR spectrum, the mechanisms responsible for the acceleration and propagation of particles, the development of air showers in the atmosphere, and the experimental techniques used to detect them [1, 3, 94]. Particular emphasis is placed on observables that are sensitive to the mass composition of the primary CR, such as the depth of the shower maximum and the muon content of the shower, which are central to the analyses presented in subsequent chapters [6, 19, 78, 83].

2.1 Historical Development and Global Properties of Cosmic Rays

The discovery of CRs dates back to the balloon experiments of Victor Hess in 1912, which demonstrated that the level of ionising radiation increases with altitude, establishing its extraterrestrial origin [88, 89]. Early investigations of cosmic radiation were instrumental in the development of particle physics, leading to the discovery of particles such as the positron and muon [2, 90]. Today, CR research is a central pillar of astroparticle physics, connecting high-energy particle interactions, plasma astrophysics, and the physics of astrophysical accelerators [1, 3].

Cosmic rays are predominantly fully ionised atomic nuclei that span a wide range of masses, from hydrogen to iron, with protons and helium nuclei constituting the dominant components at most energies [1, 2]. Heavier nuclei beyond iron are comparatively rare because they are less stable and tend to fragment during their propagation through interstellar and intergalactic space [1, 25]. In addition to nuclei, CR also includes electrons, photons, and neutrinos, although their relative abundances are much lower [1, 2]. The differential energy spectrum spans more than eleven decades in energy, from below 10^9 eV to beyond 10^{20} eV, and nearly thirty orders of magnitude in flux, making it one of the broadest known energy distributions in nature [1, 23].

Due to scattering in Galactic and extragalactic magnetic fields, CRs arrive at Earth with an almost isotropic distribution [3]. However, small anisotropies emerge at the highest energies, where magnetic deflections become less effective [8]. These anisotropies provide important constraints on the source distributions and on the structure of large-scale magnetic fields [3, 8].

The differential energy spectrum of CR can be described, to a good approximation, by a broken power law,

$$\frac{d\Phi}{dE} \propto E^{-\gamma}, \tag{2.1}$$

where the spectral index γ varies slowly with energy E, typically between 2.6 and 3.3 depending on the energy range [1, 23]. Remarkably, the spectral index changes only modestly over many decades in energy, indicating that similar physical processes govern CR acceleration and propagation over a wide energy interval [1, 3].

Because the flux decreases rapidly with energy, direct detection of individual primary particles becomes impractical at approximately 10^{14} eV [1, 94]. At these energies, CR are studied indirectly through the EAS they generate in the atmosphere [1, 94]. These cascades contain billions of secondary particles and can be observed using fluorescence telescopes, Cherenkov detectors, or large SD arrays [14, 35]. Modern observatories such as Pierre Auger Observatory and the Telescope Array (TA) employ hybrid detection techniques that combine these methods [7, 14, 35].

Figure 2.1 shows the spectrum of CR. The measurements of the intensity of charged and neutral CRs, multiplied by the squared kinetic energy, are shown. Below 10^4 GeV, the all-particle spectrum is the sum of spline fits of the most important nuclear species. The Energy-integrated intensities are indicated by the various diagonal lines. Although the overall spectrum follows a power-law behaviour, several prominent features appear at specific energies where the spectral index changes. These features encode information about the origin, acceleration, and propagation of CR and are commonly interpreted in terms of changes in mass composition or transitions between source populations [22, 23, 25].

The first major spectral feature, known as the "knee", occurs at approximately $E_{\text{knee}} \approx (3\text{–}5) \times 10^{15}$ eV [1, 94]. At this energy, the spectrum steepens, indicating a reduction in flux relative to the extrapolated power law [22, 23]. This feature is often interpreted as the maximum energy attainable by Galactic accelerators such as supernova remnants through diffusive shock acceleration [3, 25], or as the onset of more efficient escape of particles from the Galaxy [1]. In rigidity-dependent acceleration scenarios, heavier nuclei reach higher maximum energies proportional to their charge, producing a sequence of knees for different mass groups [94]. A further steepening of the spectrum occurs near $E_{\text{2nd knee}} \sim 10^{17}$ eV, commonly referred to as the "second knee" [1, 94]. This feature is frequently associated with the gradual cut-off of the heaviest Galactic CR components or with the onset of the transition from Galactic to extragalactic CR [22, 27]. At higher energies, the spectrum exhibits a hardening known as the "ankle", $E_{\text{ankle}} \approx (3\text{–}5) \times 10^{18}$ eV, where the spectral slope becomes less steep [5, 23]. An interpretation is that the ankle marks the transition from a predominantly Galactic population to an extragalactic one [22, 27]. Alternatively, it may arise from propagation effects, such as energy losses of ultra-high-energy protons due to electron–positron pair production in interactions with cosmic microwave background

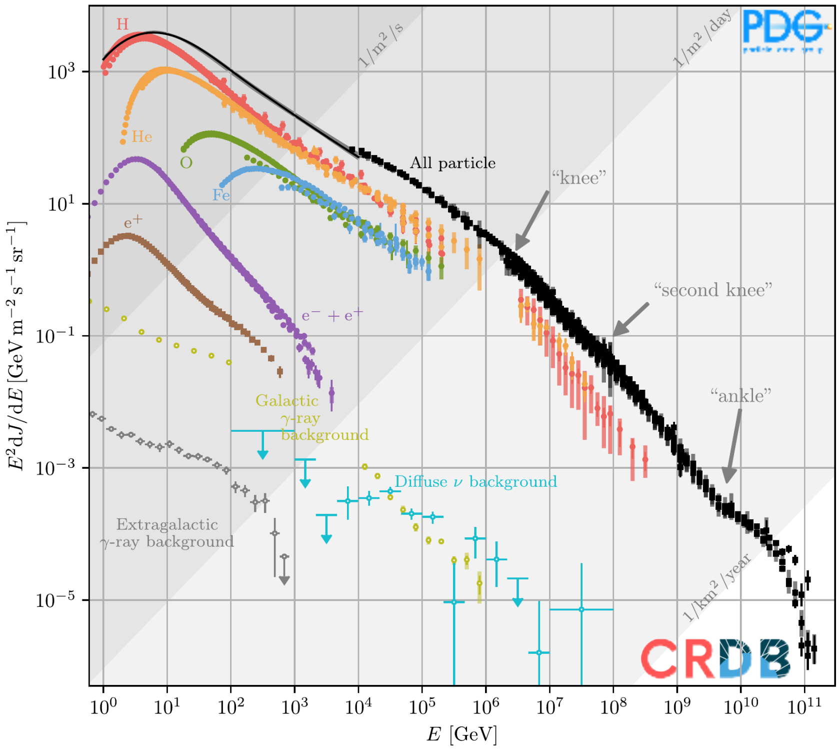

Figure 2.1: The spectrum of CR. Shown are measurements of the intensity of charged and neutral CRs, multiplied by kinetic energy squared. The figure is taken from [2].

photons [24], $p + \gamma_{\mathrm{CMB}} \to p + e^+ + e^-$. At energies greater than roughly $E \gtrsim 4 \times 10^{19}$ eV, the flux decreases sharply, a feature known as high-energy suppression [4, 5]. This behaviour is consistent with the energy losses incurred by ultra-high-energy particles interacting with cosmic background radiation [25, 28, 29]. For protons, the dominant process is photo-pion production,

$$p + \gamma_{\mathrm{CMB}} \to \Delta^+ \to \begin{cases} p + \pi^0, \\ n + \pi^+, \end{cases} \tag{2.2}$$

commonly referred to as the Greisen–Zatsepin–Kuzmin (GZK) effect [28, 29].

CRs are broadly classified into a lower-energy Galactic component and a higher-energy extragalactic component [3, 25]. Galactic CRs are widely believed to be accelerated primarily in supernova remnants through first-order Fermi acceleration, whereas those above the ankle originate in extragalactic environments [3, 25]. The transition between these populations likely occurs somewhere between 10^{15} eV and $10^{18.5}$ eV, although the precise energy range and mechanism remain subjects of active research [22, 27, 94].

Understanding the structure of the CR spectrum therefore provides direct insight into astrophysical acceleration mechanisms, propagation through magnetic fields and radiation backgrounds, and the relative contributions of Galactic and extragalactic sources [1, 3, 25].

2.2 Extensive air showers

For primary energies above $\sim 10^{14}$ eV, the CR flux becomes too small for direct particle-by-particle detection, and the reconstruction of the primary properties relies on observation of the cascades of secondary particles produced in the atmosphere, known as extensive air showers (EAS) [1, 94]. An EAS develops when a primary CR interacts hadronically with an atmospheric nucleus, initiating a multi-generation cascade that grows longitudinally along the shower axis while spreading laterally due to multiple scattering and transverse momenta [1, 92].

The particles in an EAS can be grouped into three main components: (i) the *hadronic* component, consisting primarily of pions, nucleons, anti-nucleons and neutrons; (ii) the *electromagnetic* (EM) component, consisting of electrons, positrons and photons and carrying most of the shower energy; and (iii) the *muonic* component, consisting of $\mu^{\pm}$ predominantly produced in decays of charged hadrons (mostly $\pi^{\pm}$, and also $K^{\pm}$) [1, 92]. Muons are relatively penetrating, and many reach the ground before decaying, making the muonic signal a sensitive tracer of the hadronic cascade [20, 92]. A small fraction of muons can also be produced via photo-nuclear interactions in the electromagnetic cascade, but this contribution is subdominant [1, 94]. The schematic development of an EAS is illustrated in Fig. 2.2.

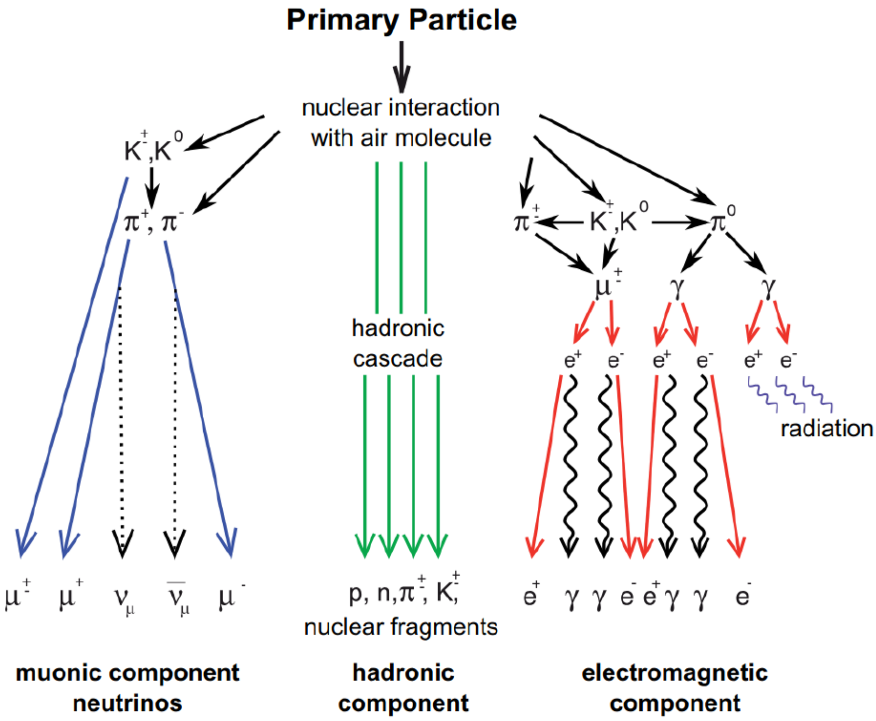

Figure 2.2: Schematic development of an EAS with its different components based on the Heitler–Matthews model. Figure taken from Ref. [30].

The first interaction of a primary particle of energy E_0 occurs at an atmospheric depth X_1 distributed approximately exponentially,

$$P(X_1)\,\mathrm{d}X_1 = \frac{1}{\lambda_{\mathrm{int}}} \exp\left(-\frac{X_1}{\lambda_{\mathrm{int}}}\right) \mathrm{d}X_1, \tag{2.3}$$

where λ_{int} is the interaction length in $\mathrm{g\,cm^{-2}}$, related to the inelastic cross section (σ_{inel}) by $\lambda_{\mathrm{int}} \propto 1/\sigma_{\mathrm{inel}}$ [1, 94]. The hadronic products of the first interaction generate secondary mesons and baryons that re-interact or decay, continuously feeding the EM and muonic components [1, 92].

The dominant interactions governing the EM component are pair production $\gamma \rightarrow e^+ + e^-$ and bremsstrahlung $e \rightarrow e + \gamma$, while photo-nuclear interactions are subdominant and are neglected in the simplest analytical treatment [1, 92]. In the Heitler model, the

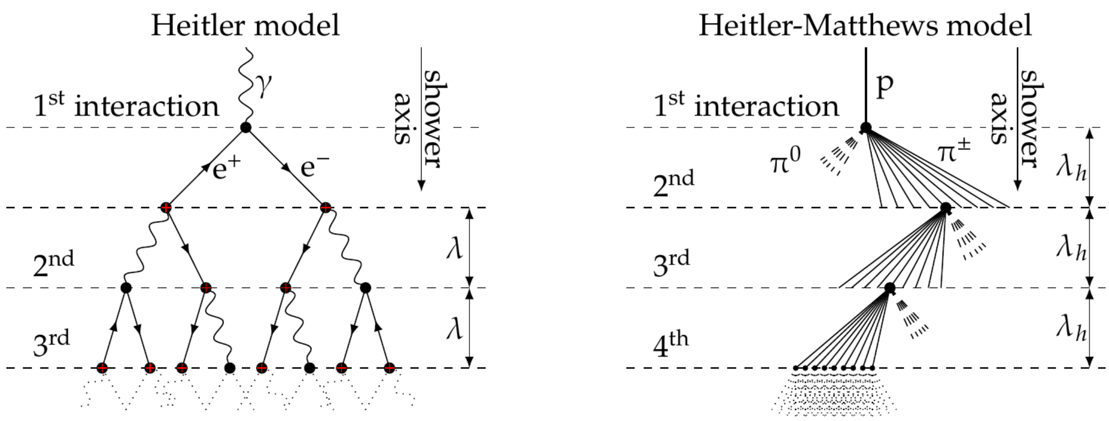

Figure 2.3: Illustrations of shower development according to the Heitler model [91] (left) and the Heitler-Matthews model [92](right). Figure taken from [18].

shower proceeds through a binary splitting process that occurs after a characteristic splitting length. In the air, this splitting length can be written as

$$\ell_{\mathrm{em}} = (\ln 2)\, X_0, \tag{2.4}$$

where $X_0 \simeq 37\,\mathrm{g\,cm^{-2}}$ is the radiation length in air [1, 91]. After n generations, the number of particles, energy per particle, and depth traversed are

$$N_n = 2^n, \qquad E_n = \frac{E_0}{2^n}, \qquad X_n = n\,\ell_{\mathrm{em}}, \tag{2.5}$$

consistent with standard cascade scaling arguments [1, 91]. The left panel of Fig. 2.3 illustrates the schematics of a purely electromagnetic shower according to the Heitler model.

Multiplication stops when E_n falls below the critical energy ξ_{em} at which ionisation losses dominate. Using $\xi_{\mathrm{em}} \simeq (84\text{–}85)\,\mathrm{MeV}$ in air [1, 2], the generation at maximum is,

$$n_c = \frac{\ln(E_0/\xi_{\mathrm{em}})}{\ln 2}. \tag{2.6}$$

At the shower maximum,

$$N_{\mathrm{max}} \simeq \frac{E_0}{\xi_{\mathrm{em}}}, \qquad X^{\mathrm{em}}_{\mathrm{max}} \simeq X_0 \ln\left(\frac{E_0}{\xi_{\mathrm{em}}}\right), \tag{2.7}$$

which captures the robust scaling $N_{\mathrm{max}} \propto E_0$ and $X_{\mathrm{max}} \propto \ln E_0$ [1, 91, 94]. Although this simple model does not describe the relative amounts of photons and electrons correctly (there are many more photons than electrons in reality), the linearity between N_{max} and E_0 holds very well.

A useful derived quantity is the elongation rate, the increase of the mean shower maximum per decade of energy,

$$\Lambda \equiv \frac{\mathrm{d}X_{\mathrm{max}}}{\mathrm{d}\log_{10} E_0}, \tag{2.8}$$

which, in the data, is a primary observable for the evolution of composition [6, 78, 94].

The Heitler–Matthews model, illustrated in the right panel of Fig. 2.3, extends the Heitler picture to hadronic showers by describing the hadronic cascade in discrete interaction steps [92, 94]. In a simplified implementation, hadronic interactions occur after a characteristic interaction length λ_π (pion interaction length in air), and each interaction produces a multiplicity N_{mult} of pions split into charged and neutral states [1, 92]. Neutral pions promptly decay as $\pi^0 \to \gamma + \gamma$, thereby transferring energy into the EM cascade, while charged pions continue to interact until their energy falls below a critical energy ξ at which decay dominates over interaction [1, 92].

Assuming equal energy sharing among secondaries, the energy per particle after n hadronic generations is $E_n = E_0/N_{\text{mult}}^n$. The generation at which $E_n = \xi$ is

$$n_{\text{dec}} = \frac{\ln(E_0/\xi)}{\ln N_{\text{mult}}}. \tag{2.9}$$

At this point, all charged pions are assumed to decay, yielding a muon number scaling.

$$N_\mu \simeq \left(\frac{E_0}{\xi}\right)^\beta, \qquad \beta = \frac{\ln(\alpha N_{\text{mult}})}{\ln N_{\text{mult}}}. \tag{2.10}$$

In the Heitler–Matthews model, $\beta \simeq 0.85$ [92], while modern air-shower simulations typically prefer $\beta \simeq 0.87$–0.93, depending on the hadronic interaction model and the definition of the analysis [12, 16, 17, 94].

The longitudinal development of an EAS is commonly described by the number of particles (or energy deposit) as a function of the depth of the atmosphere X. A widely used empirical form is the Gaisser–Hillas function [9],

$$N(X) = N_{\text{max}} \left(\frac{X - X_0}{X_{\text{max}} - X_0}\right)^{\frac{X_{\text{max}} - X_0}{\lambda}} \exp\left(\frac{X_{\text{max}} - X}{\lambda}\right), \tag{2.11}$$

where X_0 is an effective starting depth, X_{max} is the maximum depth of the shower, and λ controls the attenuation of the shower.

Complementary information is provided by the lateral distributions of the shower particles in the ground. The particle density $\rho(r)$, defined as the number of particles per unit area at a distance r from the shower axis, is often described for the electromagnetic component using the Nishimura–Kamata–Greisen (NKG) form [50, 51],

$$\rho(r) = N_e\, C(s) \left(\frac{r}{r_M}\right)^{s-2} \left(1 + \frac{r}{r_M}\right)^{s-4.5}, \tag{2.12}$$

with r_M the Molière radius, s an age parameter and $C(s)$ a normalisation. The lateral distributions of muons are typically described by empirical parameterisations; in composition studies, reference densities such as ρ_{450} can serve as practical proxies for the muon content of the shower [10, 42, 46].

Although the Heitler and Heitler–Matthews models provide conceptual insight and useful scaling relations, quantitative interpretation of EAS data requires complete MC simulations [94, 95]. The widely used air-shower codes include CORSIKA [52], AIRES [96], and CONEX [57]. These packages implement hadronic interactions using high-energy models such as EPOS [12], QGSJet [16], and Sibyll [17], and low-energy models such as FLUKA [53] and UrQMD [54]. Because the muon content of the shower is particularly sensitive to hadronic particle production, comparisons between measured muon observables and simulations provide stringent tests of hadronic interaction models at energies beyond those accessible in terrestrial accelerators [19, 83, 94].

2.3 Cosmic ray detection

Direct-detection experiments typically use spectrometers and/or calorimeters, together with tracking detectors, to measure the charge, rigidity, and energy of the primary particle [2].

Representative examples include AMS-02 on the International Space Station [86], balloon-borne missions, and satellite experiments such as PAMELA [84] and DAMPE [85].

For energies $\gtrsim 10^{15}$ eV, CR are detected indirectly using the atmosphere as an effective calorimeter and measuring either (i) the secondary particles that reach the ground or (ii) the radiation emitted as the shower develops in the atmosphere [14, 94]. Ground-based arrays sample the footprint of the shower on the surface. The two most common detector technologies are WCD[1] and scintillators [14, 35]. Examples include the Akeno Giant Air Shower Array (AGASA) [32–34], the Telescope Array (TA) [35], and Pierre Auger Observatory [14]. A key strength of surface arrays is their near-continuous operation, which yields a large exposure over many years [14].

FD observe the faint UV[2] fluorescence light emitted when atmospheric nitrogen molecules are excited by charged particles in the shower and subsequently de-excite [7, 94]. Because the fluorescence yield is approximately proportional to the local energy deposit $\mathrm{d}E/\mathrm{d}X$, this technique enables a nearly calorimetric measurement of the shower energy and provides direct access to longitudinal development, including the depth of shower maximum X_{max} [6, 78]. The main limitation of fluorescence detection is its low duty cycle, since observations require sufficiently dark and clear nights [7].

Modern UHECR observatories combine complementary techniques in a hybrid approach to improve reconstruction accuracy and reduce systematic uncertainties [14, 35]. In particular, combining a surface array with fluorescence telescopes enables (i) a robust geometric reconstruction, (ii) a calorimetric energy scale from FD data, and (iii) a SD energy estimator calibrated using high-quality hybrid events [7, 14]. Pierre Auger Observatory exemplifies this design, operating a $\sim 3000\,\mathrm{km}^2$ surface array together with fluorescence telescopes [14]. An illustration of the indirect detection methods for CR is shown in Fig. 2.4. Surface and underground particle detectors measure electromagnetic particles and muons. Imaging (IACT) and non-imaging (NIAC) atmospheric-Cherenkov detectors, as well as radio antennas, provide a measurement of the electromagnetic shower component when located in the footprint of the shower. In contrast, fluorescence light detectors can observe the shower development from the side.

2.4 Acceleration, propagation, and potential sources of ultra-high energy cosmic rays

The most widely supported class of mechanisms capable of accelerating CR to non-thermal energies is based on the Fermi process, in which charged particles gain energy through repeated scattering in a magnetised plasma [3, 25]. In such a stochastic acceleration, the fractional energy gain per interaction is approximately constant, $\Delta E \propto E$, such that after n encounters, a particle injected at energy E_0 reaches,

$$E_n = E_0(1+\xi)^n, \tag{2.13}$$

where ξ parametrizes the average fractional gain per encounter [3]. In the first-order Fermi mechanism, a shock front (like the ones produced by supernova explosions) moves at a velocity v, with which the particle interacts multiple times and is accelerated as

$$E_n = E_0\left(1+\frac{4}{3}\beta\right)^n, \tag{2.14}$$

[1] water Cherenkov detectors

[2] ultraviolet

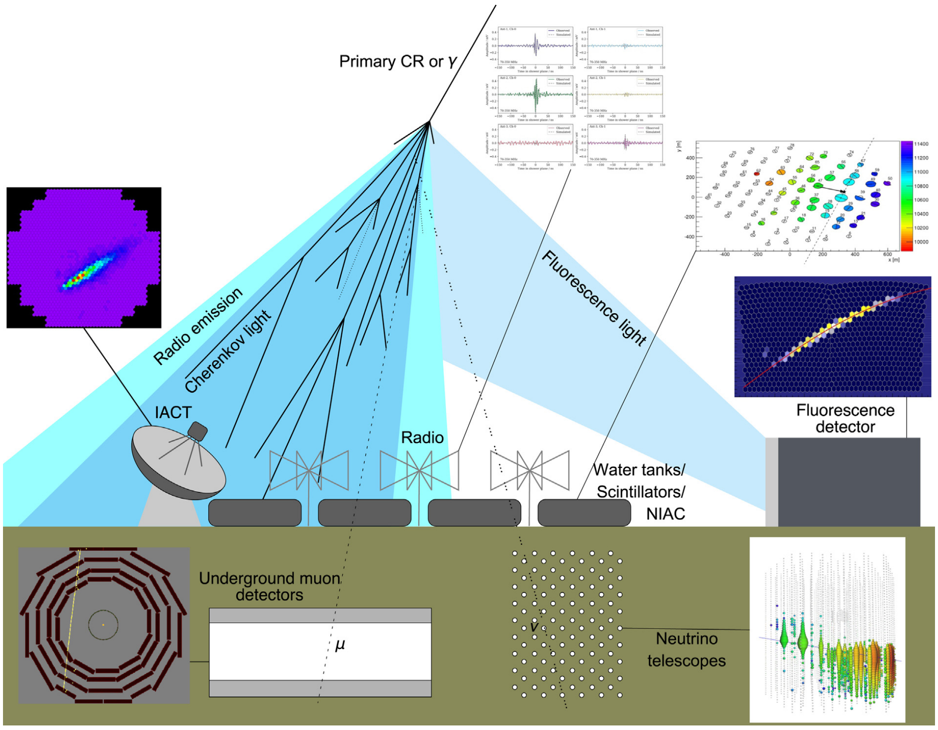

Figure 2.4: An illustration of the indirect detection methods for CR. Source: Image from Ref. [15]

where $\beta = v/c$. In the second-order Fermi mechanism, the shock front is replaced by a magnetised plasma cloud and the energy after n interactions is the same as that in Eq.(2.14), but β is replaced with β^2. The crucial difference between the two mechanisms is that, in interactions with the shock front, the particle always exits the shock upstream, and there is always an energy gain. Because this gain is systematically positive when interacting with a shock, first-order Fermi acceleration is substantially more efficient than the second-order process and is commonly accepted as the dominant mechanism for Galactic CRs up to the knee [3, 25].

Identifying astrophysical accelerators capable of producing UHECR requires, at a minimum, that the particles remain confined in the acceleration region. Using the Larmor radius $r_L = p/(ZeB)$, where Z is the charge number of the accelerated particle, e is the absolute value of the electron charge, and B is the magnetic field of the source. Requiring $r_L \lesssim R$, where R is the characteristic size of the accelerator, yields the Hillas condition [25, 31],

$$E_{\max} \propto \beta\, Ze\, B\, R, \tag{2.15}$$

This is a necessary (but not sufficient) condition that a possible source has to fulfil. Figure 2.5 shows a Hillas plot, a diagram of possible CR sources according to their magnetic field and characteristic distance. It displays upper limits on the achievable CR energy as a function of the size of the acceleration region and the magnetic field strength. The red lines indicate the upper limits due to the loss of confinement in the acceleration region for CRs at the knee, ankle, and the Greisen-Zatsepin-Kuzmin (GZK) cut-off. Any object above and to the right of the red dashed lines can confine a proton in the acceleration region up to the corresponding marked energy. The dotted grey line corresponds to a second upper limit that arises from synchrotron losses in the sources and interactions in the cosmic photon background. From Fig. 2.5, it is clear that there is not one clear and unique candidate as a source of UHECR.

At ultra-high energies, the dominant energy losses arise from interactions with background photon fields, in particular the cosmic microwave background (CMB) [25, 28, 29].

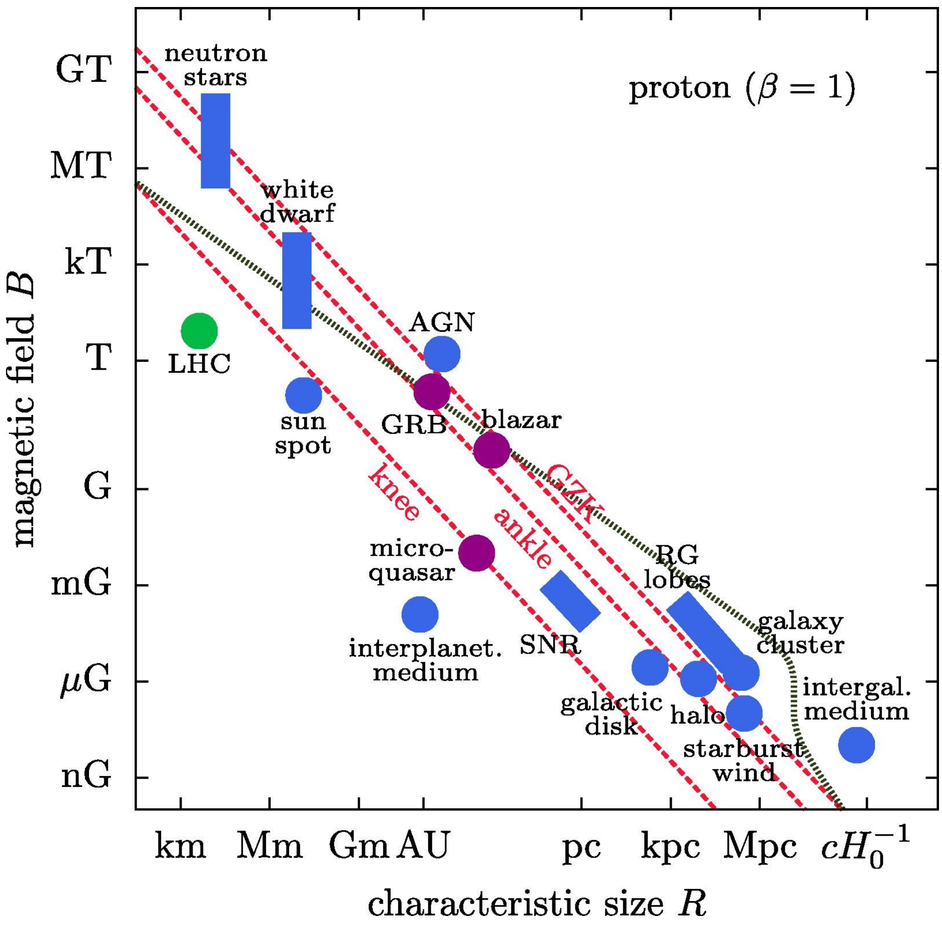

Figure 2.5: A modern adaptation of the so-called "Hillas plot". The red lines indicate the upper limits due to the loss of confinement in the acceleration region for CRs at the knee, ankle, and the Greisen-Zatsepin-Kuzmin (GZK) cut-off. The dotted grey line corresponds to a second upper limit arising from synchrotron losses in the sources and interactions with the cosmic photon background. Figure taken from Ref. [93].

For protons, the main processes are photo-pion production and Bethe–Heitler pair production [24, 28, 29]. The rapid flux suppression observed above $\sim (5\text{–}6) \times 10^{19}$ eV can be explained by these propagation losses (the GZK suppression) and/or by a maximum acceleration energy at the sources; distinguishing these scenarios remains an open question [3, 25].

For nuclei with $A > 1$, the dominant process is typically photo-disintegration, in which an absorbed photon induces the emission of nucleons or fragments [3, 25, 27],

$$^{A}X + \gamma_{\mathrm{CMB}} \rightarrow {}^{A-k}X + kN \tag{2.16}$$

where k is the number of emitted nucleons, and N denotes a nucleon (proton or neutron).

2.5 Mass composition and muon deficit in air-shower simulations

The mass (or chemical) composition of CR as a function of energy is a key observable to understand their origin, acceleration, and propagation [3, 25, 94]. At ultra-high energies, the primary particle cannot be identified directly, and the composition must instead be inferred statistically from observables of the EAS [94]. Composition measurements provide a crucial discriminator between competing interpretations of the ankle and the suppression of high-energy flux [22, 25, 27].

Above the ankle, two commonly discussed classes of scenarios are (i) a *maximum-rigidity* scenario and (ii) a *photo-disintegration* scenario [3, 27]. In a maximum-rigidity scenario, the observed suppression at the highest energies is interpreted primarily as a source effect, with a maximum energy scaling with charge.

$$E_{\max}(Z) \propto Z\, E_{\max}(Z{=}1), \tag{2.17}$$

whereas photo-disintegration scenarios emphasise composition evolution during propagation [3, 25]. In practice, both mechanisms can contribute simultaneously, and combined fits of the energy spectrum and $X_{\max}$ distributions from Pierre Auger Observatory favour a predominantly rigidity-limited picture with a mixed composition in the relevant energy range [68, 78].

The conversion of EAS measurements into composition parameters such as the mean logarithmic mass $\langle \ln A \rangle$ is based on detailed MC simulations of air showers and thus on hadronic interaction models [94, 95]. These models extrapolate accelerator constraints into unexplored phase space at energies above the LHC scale and in the forward region [12, 16, 17]. As a result, the inferred composition is intrinsically model-dependent, and composition studies simultaneously act as tests of hadronic interaction models at extreme energies [19, 83]. In Fig. 2.6, the mass composition of CR is quantified by $\langle \ln A \rangle$ as a function of CR energy E. The model predictions are indicated as markers and lines, while the data are represented as bands. The vertical arrows display the instrumental error of the leading experiments at low and high energies.

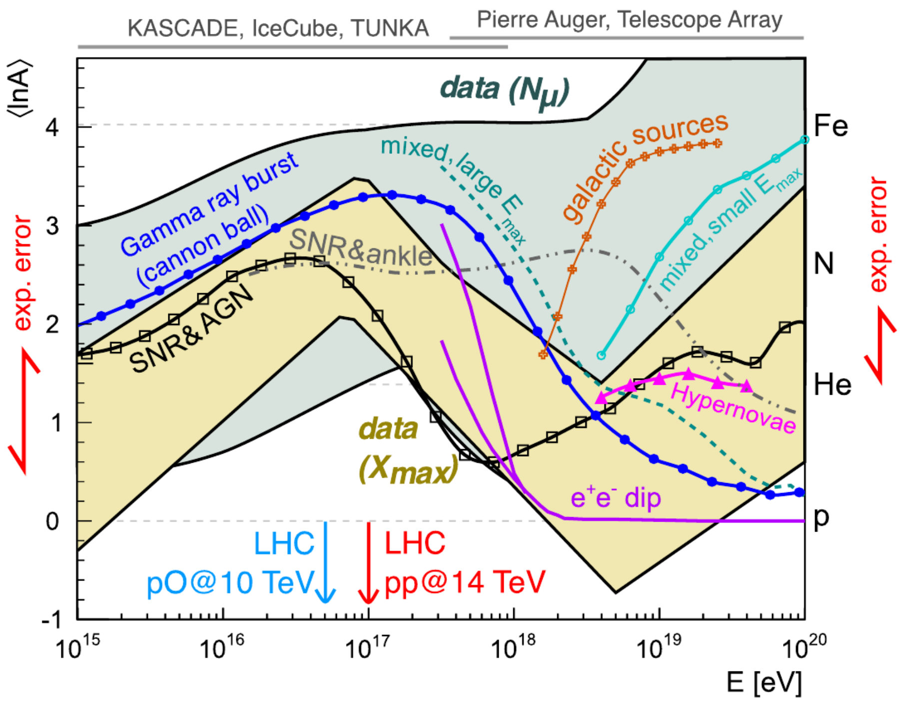

Figure 2.6: Mass composition of CR quantified by $\langle \ln A \rangle$ as a function of CR energy E. The model predictions are indicated as markers and lines, while the data is represented as bands. At the sides, vertical arrows display the instrumental error of the leading experiments at low and high energies. Figure extracted from Ref. [117].

The two most composition-sensitive air-shower observables are the depth of shower maximum $X_{\max}$ and the muon content N_μ (or experimental proxy such as a muon density at a reference distance) [6, 10, 19, 78, 94]. The observable $X_{\max}$ is measured with fluorescence telescopes (or related optical techniques), implying reduced statistics due to the limited duty cycle [7]. Observables related to muons provide complementary composition information because muons trace the hadronic cascade [10, 20, 42, 46]. Persistent tension arises when comparing the composition implications derived from $X_{\max}$ with those derived from muon measurements. Multiple experiments report that the muon content inferred from the data corresponds to a composition systematically heavier than that implied by $X_{\max}$, when interpreted within the same hadronic interaction model [19, 83, 94]. Equivalently, simulations tend to under-predict the number (or density) of muons at the ground compared to observations, giving rise to the so-called *muon deficit* [19, 83].

A practical approach is to express measurements on a normalised scale that reduces detector- and site-specific effects. One example is a "z-scale" defined by comparing the measured muon content to proton and iron expectations from full detector simulations,

$$z \equiv \frac{\ln N_{\mu,\mathrm{det}}^{\mathrm{data}} - \ln N_{\mu,\mathrm{det}}^{p}}{\ln N_{\mu,\mathrm{det}}^{\mathrm{Fe}} - \ln N_{\mu,\mathrm{det}}^{p}}, \tag{2.18}$$

so that, to first order, z is mapped onto $\langle \ln A \rangle / \ln 56$ [19, 83]. Resolving or constraining the muon deficit is essential both for reliable composition inference and for improving our understanding of hadronic interactions at ultra-high energies [19, 83, 94]. The z-scale of different experiments, obtained after applying the cross-calibration of the energy scale, is shown in Fig. 2.7 for various high-energy hadronic interaction models. The muon measurements seem to be consistent with the simulations up to about 1×10^{16} eV. However, at higher energies, a growing muon deficit is observed in the simulations.

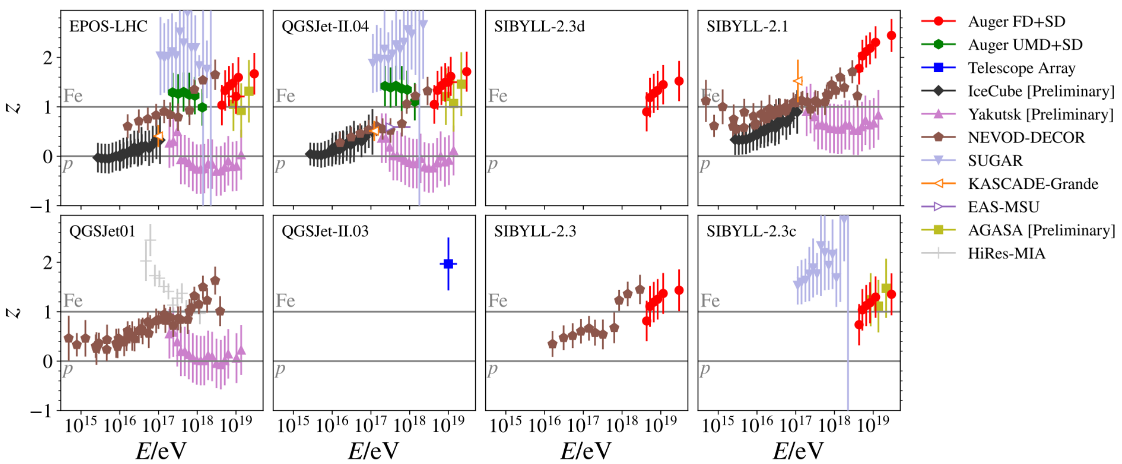

Figure 2.7: Muon density measurements converted to the z-scale, as defined in Eq. (2.18), for different hadronic interaction models. Error bars show statistical and systematic uncertainties added in quadrature. Figure taken from Ref. [97].

2.6 Summary

CRs provide a unique natural laboratory for studying particle interactions and astrophysical processes at energies far beyond those accessible to terrestrial accelerators [1, 3, 25]. Their differential energy spectrum, which extends over many orders of magnitude, exhibits distinct structures – including the knee, ankle, and high-energy suppression, which encode information about acceleration mechanisms, source populations, and propagation through cosmic radiation and magnetic fields [4, 5, 22, 23]. At energies greater than $\sim 10^{14}$ eV, the low flux of the primary particles requires indirect detection through EAS, whose development reflects both the primary energy and the mass [94]. The longitudinal and lateral structure of air showers, together with their muon content, therefore provide the principal observables to infer composition and investigate hadronic interactions at extreme energies [6, 19, 78, 83].

Modern hybrid observatories, combining SD with optical techniques such as fluorescence measurements, have enabled precise and complementary observations of the energy spectrum, anisotropies, and mass composition of UHECR [7, 8, 14, 35]. These measurements support a picture in which CR are accelerated in astrophysical environments through stochastic processes such as diffusive shock acceleration and subsequently modified by propagation effects, including magnetic deflections and interactions with background photon fields [3, 25]. However, the interpretation of air-shower observables remains limited by uncertainties in hadronic interaction models [12, 16, 17, 94].

In particular, the determination of the mass composition, commonly inferred from the depth of the shower maximum and the muon content, remains subject to systematic uncertainties and unresolved tensions. Measurements of muons in air showers often exceed model predictions, leading to the so-called muon deficit in simulations [19, 83]. Resolving this discrepancy is essential both for reliable composition inference and for improving our understanding of hadronic interactions at ultra-high energies [19, 83, 94].

Therefore, a comprehensive understanding of CR acceleration, propagation, and air-shower development is essential to interpret experimental data. The framework outlined in this chapter provides the physical and experimental basis for the analyses presented in the following chapters, with particular emphasis on the muonic component of air showers and the information encoded in the muon lateral distribution [10, 42, 46].

CHAPTER 3

Detecting extensive air showers with the Pierre Auger Observatory

The Pierre Auger Observatory, located in the southern hemisphere [14, 87], includes detectors capable of investigating CR showers at all three levels: during their development in the atmosphere, on the surface, and underground. Consequently, these showers are reconstructed to examine the three principal observables of the primary particles: (i) the energy spectrum, (ii) the arrival direction, and (iii) the chemical composition. Pierre Auger Observatory is a multi-hybrid observatory dedicated to studying UHECR through the observation of EAS. It is the world's largest observatory with an area of ≈ 3000 km^2 [14]. A large collecting area is required to measure the extremely low flux of CR at the highest energies, especially above 10^{18} eV with sufficient statistics. The observatory is located in Malargüe in the Province of Mendoza in Argentina. The observatory operates as a hybrid detection system, combining a large SD array of WCD that continuously samples secondary particles at ground level with a FD that observes the longitudinal development of air showers through ultraviolet fluorescence light emitted by atmospheric nitrogen [7, 14]. This hybrid configuration enables complementary measurements of the same shower, reducing systematic uncertainties and improving reconstruction accuracy. The SD array measures the footprint of secondary particles at ground level with a duty cycle of nearly 100% and provides large statistics, measuring the time structure and lateral distribution of shower particles [98, 99], while the FD observes the longitudinal development of the shower through atmospheric fluorescence with a duty cycle of approximately 15% [7]. The FD provides a quasi-calorimetric and nearly model-independent measurement of the shower energy, which is used to calibrate the SD energy estimator [4]. Cross-calibration between the SD and FD enables precise energy reconstruction over the entire data set [4]. The FD also provides a direct measurement of the depth of the shower maximum X_{max}, which is one of the most sensitive observables to the primary mass [77, 78]. Together, the two detection systems allow for precise determination of the primary energy of CR and the arrival direction, and mass-sensitive observables such as X_{max}.

In its second phase of operation, the Observatory has undergone a major upgrade known as AugerPrime, designed to improve sensitivity to the composition of the primary mass at the highest energies [87]. The upgrade incorporates additional detector systems with complementary particle sensitivities, notably the SSD[1], the RD[2] and the UMD. In this thesis, particular emphasis is placed on the UMD, which provides a direct measurement of the

[1]Surface Scintillator Detector

[2]Radio detector

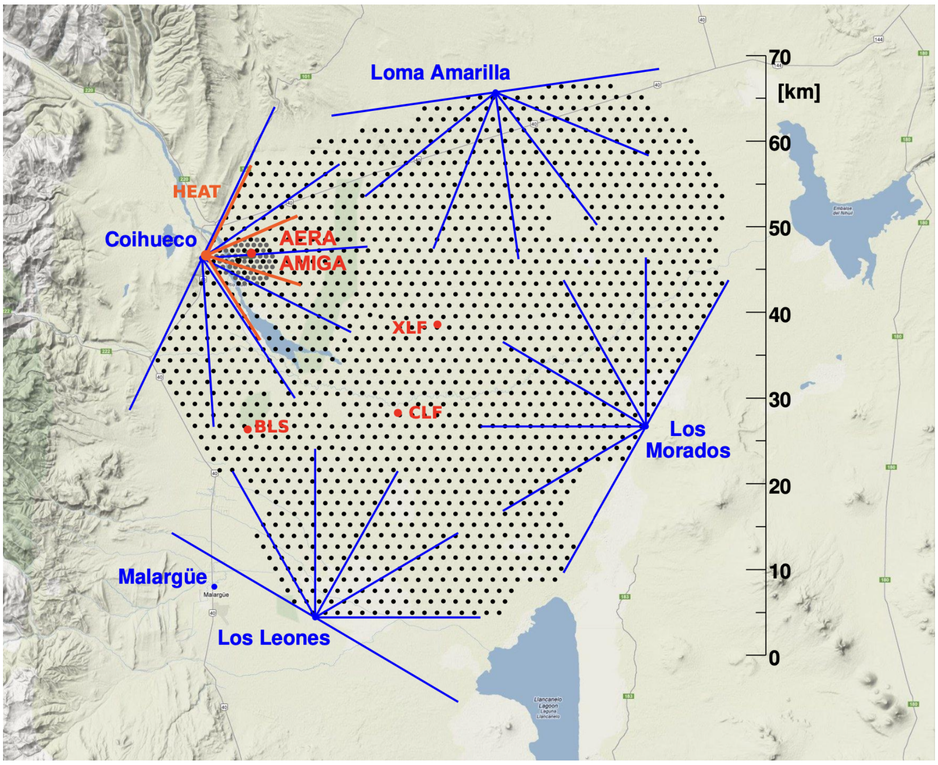

Figure 3.1: Layout of Pierre Auger Observatory. On the right, a scale in kilometres is included. The SD stations (black dots) cover an area of $\approx 3000\,\mathrm{km}^2$. The fields of view of the fluorescence telescopes (blue lines) overlap with said area. The position of other detection systems, including the UMD, can also be seen (red dots and lines). Image taken from Ref. [62].

muonic component of air showers. A schematic depiction of the various detectors in the observatory is shown in Fig. 3.1.

3.1 The surface detector

The SD of Pierre Auger Observatory is designed to operate with a duty cycle that is nearly 100%, to measure the time structure of EAS, and to provide an aperture that becomes effectively energy-independent above the full trigger efficiency [14]. The array consists of more than 1600 WCD deployed on an area of approximately $3000\,\mathrm{km}^2$. The stations are arranged in an isometric triangular grid with a spacing of 1500 m (SD–1500), providing sensitivity to air showers above $\sim 10^{18.5}$ eV [110]. For zenith angles below 60°, the corresponding geometric aperture reaches $7350\,\mathrm{km}^2$ sr [99]. The stations are located at altitudes between 1300 and 1600 m above sea level [100]. To extend the sensitivity of the Observatory towards lower energies, two denser sub-arrays are embedded within the SD–1500 array. The SD–750 covers approximately $28\,\mathrm{km}^2$ with a spacing of 750 m and reaches an energy threshold of $\sim 10^{17.5}$ eV, while the SD–433 covers $\sim 1\,\mathrm{km}^2$ with a spacing of 433 m and is sensitive down to $\sim 10^{16.5}$ eV [100, 110].

A picture of an SD-station can be appreciated in the left panel of Fig. 3.2, showing its main components. Each SD station consists of a cylindrical polyethylene tank of 3.6 m diameter and 1.2 m height, containing approximately 12,000 litres of ultra-pure water enclosed in a multi-layer reflective liner. Three 9-inch PMT[3]s are mounted on top of the tank, allowing the water volume to be seen through clear windows. The charged secondary particles that

[3] photo-multiplier tube

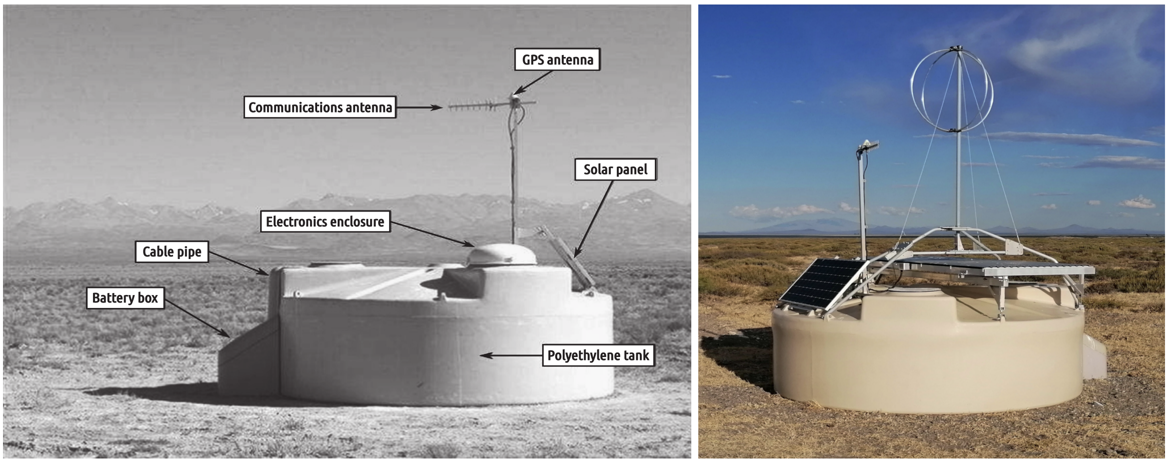

Figure 3.2: A SD station with the main components marked. Image taken from Ref. [106] (Right) Photograph of one of the stations of the SD, after its upgrade. Image taken from Ref. [11].

traverse the water emit Cherenkov light, which is detected by the PMTs [100]. The detector geometry also allows the indirect detection of high-energy photons via electron–positron pairs produced inside the tank.

Each PMT provides two output signals: a low-gain signal taken from the anode and a high-gain signal taken from the last dynode and amplified by a factor of approximately 32, thus extending the dynamic range [99, 100]. The signals are digitised at 40 MHz with 10-bit Flash ADCs (FADCs) and time-stamped using GPS with the GPS antenna with an absolute accuracy of about 10-12 ns [14, 99]. The station electronics evaluate local trigger conditions, buffer the traces, and communicate with the central data acquisition system (CDAS). Each station is powered by solar panels and batteries that produce an average power of approximately 10 W [14]. The SD employs a hierarchical trigger system [99]. The first two levels, T1 and T2, are formed locally at each station. Two complementary T1 triggers are implemented: a Threshold trigger (T1-Th), primarily sensitive to muons, and a Time-over-Threshold trigger (T1-ToT), more sensitive to the electromagnetic component. Selected T1 triggers are promoted to T2 and transmitted to the CDAS. Spatial-temporal coincidences among at least three neighbouring stations consistent with a shower front propagating at the speed of light generate a T3 array-level trigger. The next set of triggers is applied offline to the stored T3 data. The fourth-level trigger (T4) serves as a physics trigger and selects genuine air showers, while the fiducial 6T5 condition ensures the confinement of the shower core within the array [99]. This fiducial trigger guarantees a reliable reconstruction of the impact point of the shower on the ground. This trigger requires that the station with the largest signal be surrounded by a set of working (i.e., non-broken) stations.

The reconstruction procedure determines the station signals, the shower geometry and the shower size [100]. The signal in each station is obtained by integrating the FADC traces within a signal window. If the HG[4] channel is saturated, the LG[5] trace is used. The integrated charge is converted to units of VEM[6], defined as the mean charge deposited by a vertical muon, obtained from continuous calibration using background muons [100].

The geometry of the shower is reconstructed iteratively. An initial estimate of the core position $\mathbf{x}_c$ is obtained from the signal-weighted barycenter of the participating stations. The arrival direction $\hat{a}$ is first estimated by assuming a plane shower front propagating at the speed of light and is subsequently refined by fitting a curved shower front model to the

[4]high gain
[5]low gain
[6]vertical-equivalent muon

station start times through a χ^2 minimisation that accounts for timing uncertainties [100]. Accidental stations are rejected on the basis of timing and geometric consistency. The shower size is defined as the expected signal at an optimal distance r_{opt} from the shower axis, chosen to minimise systematic uncertainties associated with the unknown LDF[7] shape. For the SD–1500 array, $r_{\text{opt}} = 1000\,\text{m}$, while for the SD–750 array, $r_{\text{opt}} = 450\,\text{m}$ [110]. The LDF used in $\overline{\text{Off}}\underline{\text{line}}$ is based on a modified Nishimura–Kamata–Greisen (NKG) form [100],

$$S(r) = S(r_{\text{opt}}) \left(\frac{r_{\text{opt}}}{r}\right)^{\beta} \left(\frac{r_{\text{opt}} + r_0}{r + r_0}\right)^{\beta+\gamma}, \tag{3.1}$$

where $r_0 = 700\,\text{m}$, and β and γ are fitted or fixed to parameters derived from well-sampled events. The shower size $S(r_{\text{opt}})$ and the position of the core are obtained by maximising the logarithmic likelihood that accounts for small signals (Poissonian regime), large signals (Gaussian approximation), saturated stations and non-triggered stations [100]. Thus, the shower size estimators are referred to as S (1000) (S(450)) for SD–1500 (SD–750) arrays. In Fig. 3.3, an example LDF fit for an event is shown.

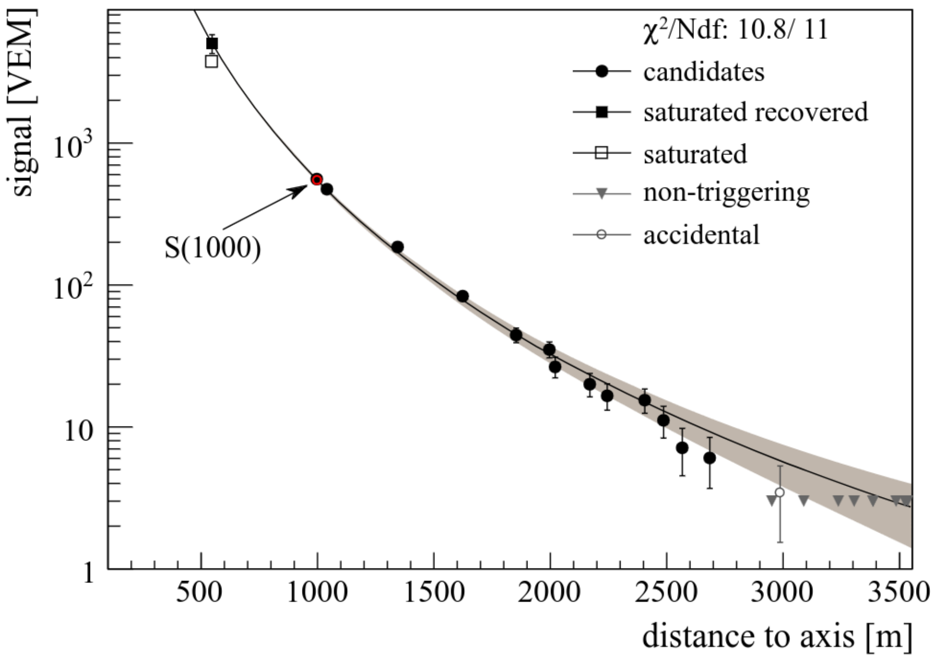

Figure 3.3: Example of an LDF fit of the SD. Taken from Ref. [14].

Since the measured shower size depends on the zenith angle due to atmospheric attenuation, a correction is applied using the CIC[8] method [110]. The attenuation-corrected size is given by

$$S_{\theta_{\text{ref}}} = \frac{S(r_{\text{opt}})}{f_{\text{att}}(\theta)}, \tag{3.2}$$

where the reference angle corresponds to the median of the zenith-angle distribution, (e.g. 38° for SD–1500 (covering a range of $\theta \in [0°, 60°]$) and 35° for SD–750 (ranging from $\theta \in [0°, 55°]$). The corrected observables are denoted S_{38} and S_{35} for the SD–1500 and SD–750 arrays, respectively [110].

[7] lateral distribution function

[8] Constant Intensity Cut

3.2 The fluorescence detector

The FD measures the longitudinal development of EAS by detecting UV fluorescence light emitted by atmospheric nitrogen molecules excited by charged particles in the cascade. Because the fluorescence photon yield is (to the first order) proportional to the energy deposited in air, the FD provides a quasi-calorimetric energy measurement and thereby sets the absolute energy scale of the Observatory [4, 7]. The FD is optimised to detect showers with energies above 3×10^{18} eV during clear, moonless nights, corresponding to an overall duty cycle of $\sim 15\%$ [7]. The FD consists of 24 telescopes distributed across four sites - Los Leones (a picture of the FD building can be seen in the left panel of Fig. 3.4), Los Morados, Loma Amarilla, and Coihueco - overlooking the SD array [7, 14]. Each site hosts six independent telescopes; each telescope covers an elevation range from 1.5° to 30° and 30° in azimuth; together, the six telescopes provide a 180° azimuthal FoV[9] facing the array interior [7]. The three high-elevation fluorescence telescopes, collectively referred to as HEAT[10], operate at the Coihueco site [13] and are shown in the right panel of Fig. 3.4. HEAT enables the observation of air showers that develop at higher altitudes and correspond to lower primary energies by extending the FoV to elevation angles between 30° and 58°. In combination with the SD-750 array, this configuration extends the range of high-quality hybrid measurements down to energies of approximately $10^{17.2}$ eV.

Figure 3.4: (Left) FD building at Los Leones during the day. (Right) Photo of the three HEAT telescopes in tilted mode. The container for DAQ, slow control, and calibration hardware is behind the enclosure of the second telescope. Images taken from Ref [14].

The different parts of a fluorescence telescope can be seen on the left panel of Fig. 3.5. Each FD telescope is a Schmidt optical system composed of a circular aperture, a $\sim 13\,\mathrm{m}^2$ segmented mirror, and a PMT camera located on the focal surface [7]. The camera comprises 440 PMTs arranged in a 22×20 matrix, providing a camera FoV of approximately 30° in azimuth and $\approx$ 28° in elevation [7]. A UV bandpass filter transmits most relevant fluorescence light while suppressing the visible background, improving the signal-to-noise ratio [7]. A Schmidt corrector ring increases the effective aperture while maintaining a lightweight design [7]. The telescope buildings are climate-controlled and equipped with an automatic shutter system to protect optics and electronics from adverse weather and bright light sources [7].

The FD electronics digitise and record the PMT signals and employ a hierarchical trigger system including pixel-level threshold triggers, pattern triggers to identify track-like signatures, and higher-level software rejection of non-shower events [7]. Events that pass the final trigger are assembled locally at each site; those that meet air-shower quality criteria can

[9] field of view

[10] High Elevation Auger Telescopes

generate a hybrid trigger communicated to the Central Data Acquisition System (CDAS), which then requests the corresponding SD data from stations expected to be hit by the same shower [7, 14]. In the FD camera, an air shower appears as a track of activated pixels with a clear time sequence. The shower geometry is reconstructed using pixel pointing directions and timing, which define the shower-detector plane and constrain the shower axis via a timing fit. When SD stations participate (hybrid events), the shower core and axis are further constrained, significantly improving geometrical resolution [77]. Typical hybrid resolutions are of the order 50 m in the core position and $\sim 0.6^\circ$ in the arrival direction [77]. The depth of shower maximum, $X_{\rm max}$, is reconstructed from the longitudinal development measured by the FD [77, 78].

Given the reconstructed geometry, atmospheric conditions, and the fluorescence yield (photons emitted per unit energy loss), the collected light profile can be converted into an energy-deposit profile as a function of slant depth, $\mathrm{d}E/\mathrm{d}X(X)$ [4, 7]. The longitudinal profile is then fitted with a Gaisser–Hillas parametrisation in Eq. (2.11). An example of the fit to the shower profile is shown in Fig. 3.5. The calorimetric energy is obtained by integrating the fitted profile, which captures the dominant electromagnetic energy deposit in the atmosphere [7]. The total primary energy is obtained by applying a correction for invisible energy carried by neutrinos and high-energy muons [4]. The resulting FD energy resolution is $\lesssim 10\%$ [7]. Although the FD operates only under stringent conditions and thus has limited exposure, its quasi-calorimetric nature makes it the reference for the absolute energy scale. High-quality hybrid events reconstructed by both SD and FD are therefore used to calibrate the SD energy estimator, allowing the SD data set - collected with nearly 100% duty cycle - to be placed on the FD energy scale [4].

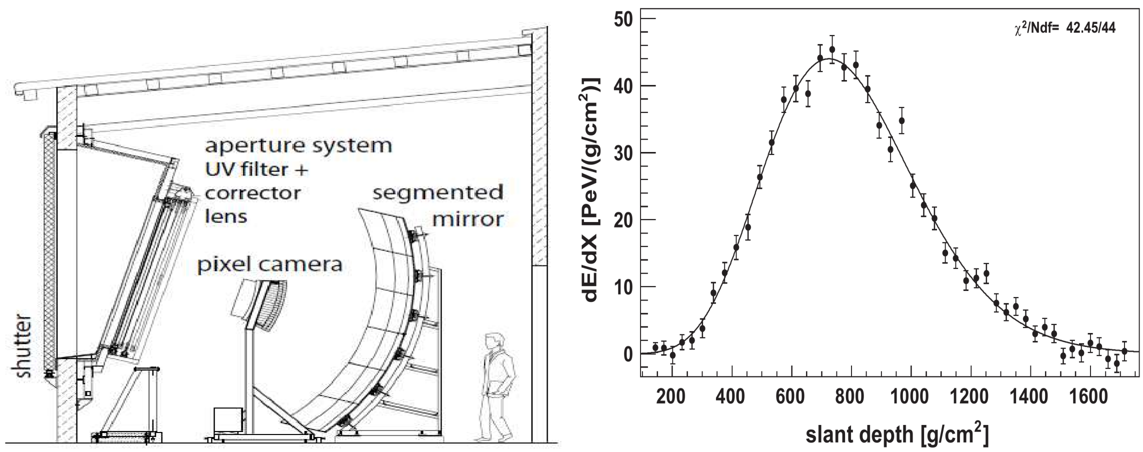

Figure 3.5: (Left) Schematic view of a fluorescence telescope with a description of its main components. Image taken from Ref. [7]. (Right) Energy deposit per slant depth as a function of the slant depth for an example event. The measured data points are in markers, and the fit of a Gaisser-Hillas function (see Eq. (2.11)) can be seen as a solid line. Figure taken from Ref. [14].

Energy calibration

The SD energy estimator is calibrated against the quasi-calorimetric FD energy using high-quality hybrid events [111]. The calibration curve is expressed as,

$$E_{\rm FD} = A \left(\frac{S_x}{\rm VEM} \right)^B , \tag{3.3}$$

where A and B are determined using a likelihood-based fit and S_x is the shower size corrected for attenuation. This procedure transfers the FD energy scale to the SD, enabling the reconstruction of the shower energy for the full SD data set with high statistical precision [111]. The correlation between the SD energy estimators (S_x) and the FD energy, along with the corresponding fits to Eq. (3.3), is displayed in Fig. 3.6.

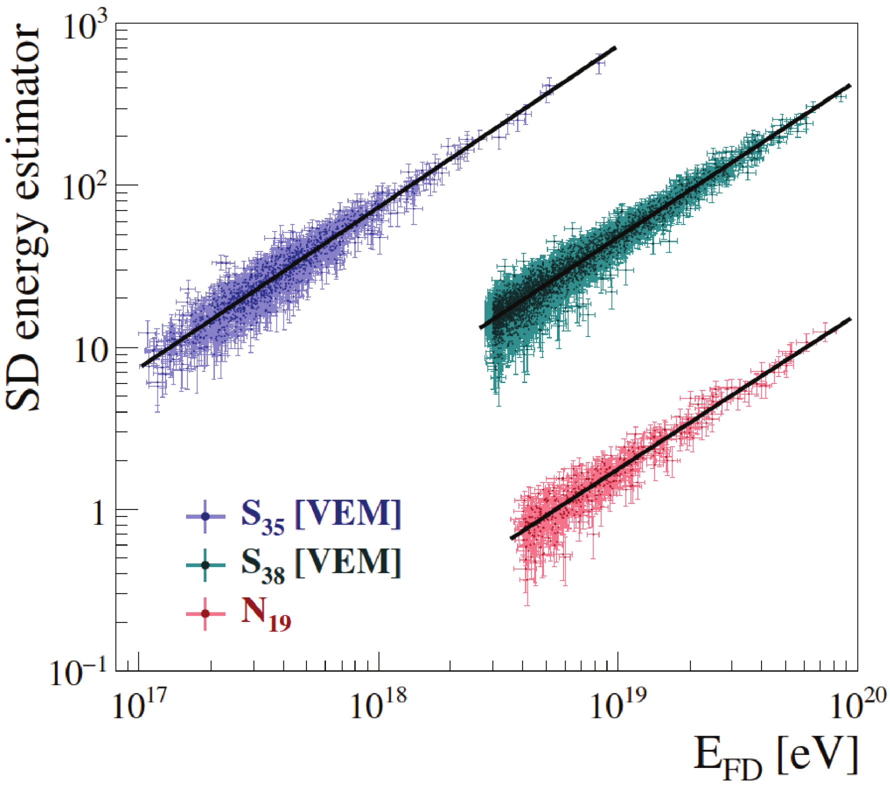

Figure 3.6: Correlation between the SD energy estimates, (S_x where $x = 38, 35$ for SD–1500 and SD–750 arrays respectively), and the FD energy. The correlation with the SD estimate N_{19} used for inclined air showers ($\theta > 60°$) and the FD energy is also shown. Lines indicate fits to Eq.(3.3). Figure taken from Ref. [87].

3.3 AugerPrime

Despite the success of the hybrid design, precisely determining the primary mass composition at the highest energies remains a challenge. The AugerPrime upgrade was implemented to improve sensitivity to the electromagnetic and muonic components of air showers and thus improve composition measurements [87]. The upgrade introduces detectors with complementary particle sensitivities, improved electronics, and extended dynamic range. By combining measurements from different detector types, AugerPrime enables a more direct determination of the muon content of air showers, which is strongly correlated with primary mass and provides a stringent test of hadronic interaction models [19, 20].

The principal components of AugerPrime include the SSD, RD systems, UUB[11] electronics, small-PMT, and UMD [87]. An image of an upgraded SD station is shown in the right panel of Fig. 3.2 with the SSD and the radio antenna visible on top.

[11] upgraded unified board

3.3.1 The surface scintillator detector

As part of the AugerPrime upgrade, each SD station has been equipped with a SSD installed above the existing WCD, allowing both detectors to sample the same region of the shower front. The SSD was designed to enhance the capability of the Observatory to separate the electromagnetic and muonic components of EAS, thereby improving sensitivity to the mass composition of primary CR [87, 103].

Each SSD consists of plastic scintillator modules with a total active detection area of $\sim 3.8\,\mathrm{m}^2$. The detector is constructed from extruded scintillator strips read out via WLS[12] fibres coupled to a PMT, enabling efficient light collection and timing [103]. The SSD and WCD have complementary responses to the components of air showers. The SSD is comparatively more sensitive to the electromagnetic component, while the WCD exhibits enhanced sensitivity to muons due to the larger Cherenkov signal in water. Combining both signals enables a disentangling of electromagnetic and muonic contributions at ground level and provides improved mass sensitivity [87, 103].

3.3.2 Radio detectors

EAS also produce coherent radio emission as they propagate through the atmosphere. This emission arises primarily from the geomagnetic deflection of electrons and positrons in the Earth's magnetic field, with a secondary contribution from the time-varying negative charge excess in the shower front. Since the radio signal is generated by the electromagnetic component of the shower, it provides a complementary probe of shower development, including sensitivity to X_{max} and to the calorimetric energy [104]. This technique was demonstrated at Pierre Auger Observatory by AERA[13], which validated radio measurements for energy and shower-development studies and provided the technical basis for radio extensions [104, 105].

As part of AugerPrime, a dedicated RD has been deployed to extend the sensitivity to mass composition, particularly for inclined air showers, while operating with an almost 100% duty cycle [87]. The RD is installed at SD stations and is read out synchronously with SD triggers to ensure coincident measurements of the same air-shower event [87].

3.3.3 The Underground Muon Detector

The UMD is the muon-measurement component of the Auger Muons and Infill for the Ground Array (AMIGA) low-energy enhancement [112]. It consists of plastic scintillator muon counters buried 2.30 m underground in the SD–750 and SD–433 infill arrays [10, 112]. The soil overburden absorbs most of the electromagnetic component of the shower and corresponds to $\sim 540\,\mathrm{g\,cm^{-2}}$, imposing a vertical muon energy threshold of approximately 1 GeV [10]. By measuring penetrating muons in the ground, UMD provides a direct handle on the muonic component of EAS in the energy range $\sim 10^{16.5}$ to 10^{19} eV [10]. A picture of a UMD station during deployment (left panel) and a UMD module during assembly (right panel) is shown in Fig. 3.7.

The UMD operates in slave mode with respect to the WCD. The SD station provides the trigger for data acquisition, and the UMD records the corresponding underground response [112]. The complete detector will comprise 219 scintillation modules installed at 73 positions ("UMD counters"), with three modules per counter [10]. Of these, 61 counters are located in the SD–750 array and 12 in the SD–433 array [10]. The three modules are deployed within ~ 20 m of their paired SD station to sample approximately the same shower particle

[12] wavelength-shifting

[13] Auger Engineering Radio Array

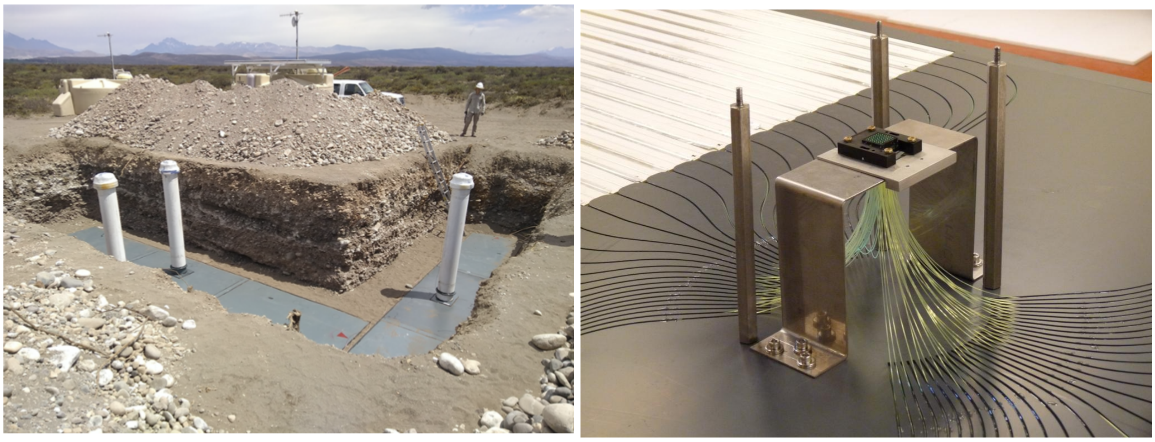

Figure 3.7: (Left) Deployment of a UMD station. (Right) The scintillator module manifold is used to route the fibres to the optical connector. Images taken from Ref. [108].

densities as the SD [10]. This multi-hybrid configuration enables cross-checks of the muon content inferred from other detector components [10]. A scheme of an SD station together with its three UMD modules can be seen in Fig. 3.8.

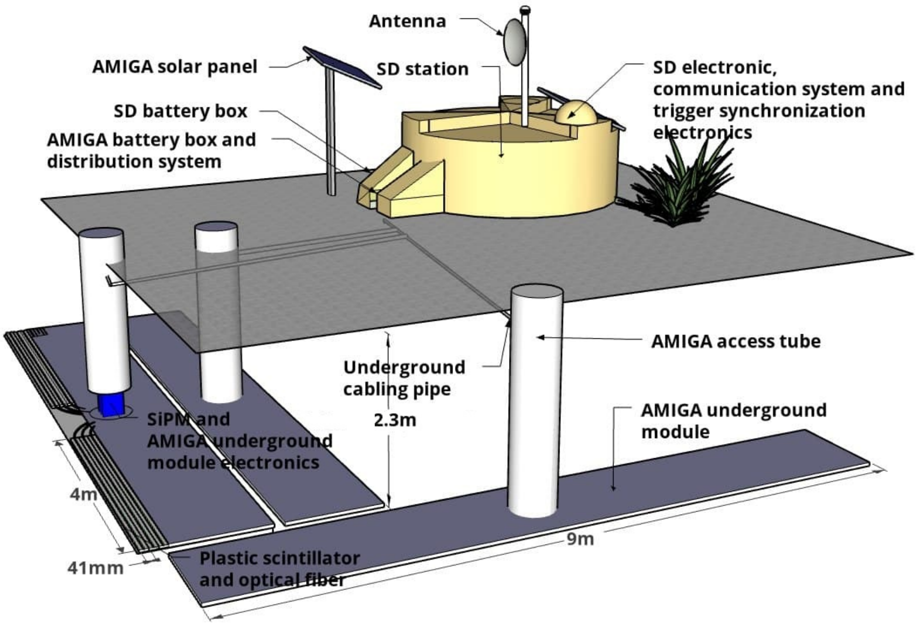

Figure 3.8: General overview of a station. The figure shows a counter and SD. Three modules are shown buried 2.3 m deep, with a detection area of 10 m^2. Module electronics are accessed for field installation and service via a plastic tube. In addition, the solar panel and the additional battery box are shown. Image taken from Ref. [39].

Each UMD module has an active area of 10 m^2, yielding 30 m^2 per counter [10]. A module is segmented into 64 plastic scintillator strips arranged in two identical panels of 32 strips. The strips are approximately 4 m long, $\sim$ 4 cm wide, and $\sim$ 1 cm thick, and are covered by a reflective titanium dioxide (TiO_2) coating to reduce photon losses. Each strip contains an embedded wavelength-shifting (WLS) optical fibre that transports the scintillation light to the photo-sensor and readout electronics housed in a central "dome". For environmental protection, modules are enclosed in a PVC container before burial [112]. Muon detection relies on fluorescence and wavelength shifting in the scintillator-fibre system [40, 113].

The scintillator is based on polystyrene (Dow Styron 663W) doped first with PPO (2,5-diphenyloxazole), which emits other UV photons with a larger attenuation length. Since the plastic is not transparent for these photons, a second dopant POPOP (1,4-bis(5-phenyloxazol-2-yl) benzene) is introduced, which convert the initially produced ultraviolet photons into blue light at $\sim 420\,\mathrm{nm}$ with an attenuation length of order 5–25 cm in the scintillator [40]. Since the strips are 4 m long, WLS fibres are required to transport light efficiently to the photo-sensors. The Saint-Gobain BCF-99-29AMC multi-clad WLS fibres absorb the blue photons and re-emit green photons, a fraction of which are guided by total internal reflection to the photo-sensor over meter-scale distances [40, 112]. Fig. 3.9 shows a scheme of the scintillation strip with the WLS optical fibre (left panel) and a scheme of the fibre (right panel). In the latter, the refractive indices of the core and cladding are indicated.

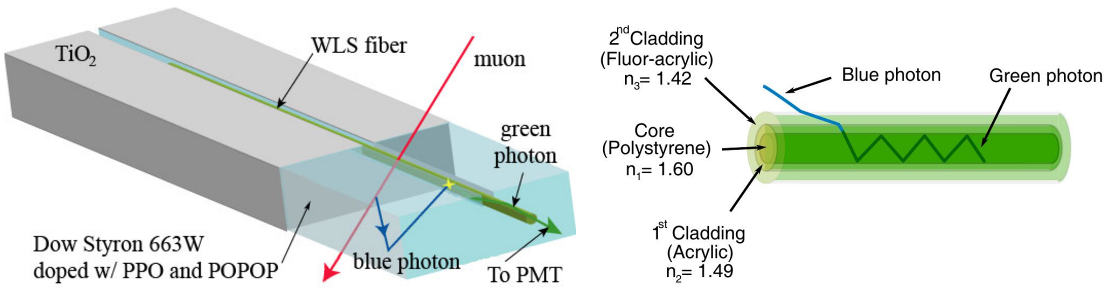

Figure 3.9: (Left) A scintillator bar being excited by a muon. Highlighted are the trajectories of the impinging cosmic particle (or muon, in red) and the photons produced within the bar (blue photon generated by fluorescence in the bar) and in the fibre (green photon generated by fluorescence in the fibre). (Right) Scheme of the WLS optical fibre. Images taken from Ref. [108] and Ref. [40] respectively.

The photo-sensors are silicon photomultipliers (SiPM[14]). The UMD uses the Hamamatsu S13361-2050NE-08 device, which contains 1584 micro-cells operated in Geiger mode [40, 112]. Incident photons trigger avalanches in individual cells, producing quantised signals; the number of triggered cells is referred to as the number of PE[15] [40]. Earlier engineering-array designs employed multi-anode photomultipliers, but the final design adopts SiPMs as described above [10, 112].

To achieve a broad dynamic range, two complementary acquisition modes operate simultaneously in each module of the UMD: a binary (counter) mode and an ADC (integrator) mode [112]. The binary mode is optimised for low muon densities and constitutes the primary reconstruction strategy used throughout this thesis (except Chapter 8, where the ADC mode is employed and discussed in detail). In binary mode, the 64 SiPM channels are processed independently using two 32-channel CITIROC application-specific integrated circuits (ASICs) (manufactured by WEEROC), each providing a pre-amplifier, fast shaper, and discriminator stage [40, 112]. The discriminator output of each channel is digitised by a Field-Programmable Gate Array (FPGA) at a sampling frequency of 320 MHz (corresponding to 3.125 ns), producing binary traces of 2048 time bins (6.4 μs) per channel. Each time bin records a logical "1" or "0" depending on whether the signal exceeds a threshold of approximately 2.5 PE, which is chosen to suppress most of the SiPM dark noise, including contributions from optical crosstalk [10, 114]. The detector calibration procedure ensures a uniform response across all channels [114].

Muon identification in binary mode relies on the characteristic temporal structure of the signal. Laboratory measurements show that a single muon typically produces a signal

[14]silicon photo-multipliers
[15]photo-electron

extending over multiple consecutive time bins, with a minimum duration of approximately 12.5 ns above threshold [114]. Consequently, muons are identified by requiring sequences of at least four consecutive positive samples ("1111"), which discriminates against short-lived noise fluctuations [61]. Once such a sequence is detected, an inhibition window is applied to avoid multiple counting of the same particle. An optimal window of 12 samples (corresponding to $\sim$ 37.5 ns) has been established, capturing approximately 99% of the muon signal while balancing under-counting and over-counting effects [114].

The ADC mode, in contrast, is designed for high muon densities where signal pile-up prevents reliable counting. In this mode, the SiPM signals are summed and processed through parallel low- and high-gain amplification chains, followed by digitisation at 6.25 ns sampling, producing two traces of 1024 samples each [40]. The muon content is then reconstructed from the integrated signal charge, normalised by the average single-muon response. This mode is not used in the main analyses of this thesis but is introduced and exploited in Chapter 8.

Using the shower geometry provided by the SD, the reconstructed muon counts are converted into muon densities by dividing by the effective detector area and assigning each module a distance to the shower axis in the shower plane [40]. Each event thus provides a sampling of the MLDF, which is typically fitted and evaluated at a fixed reference distance (e.g. $r = 450$ m for SD–750) to define an event-level estimator of the muonic shower size [40].

A schematic representation of the electronics chain for both acquisition modes is shown in Fig. 3.10.

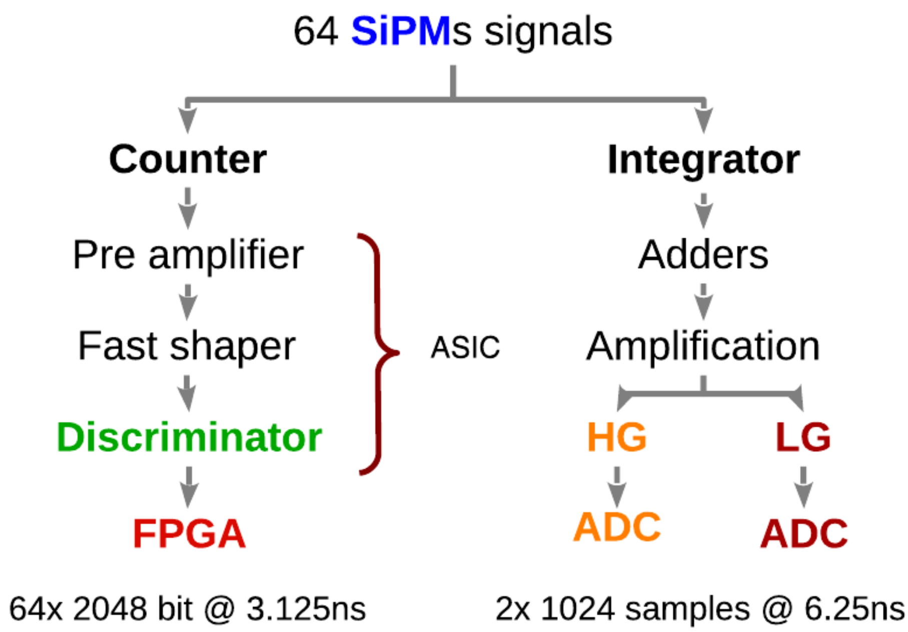

Figure 3.10: Schematics of the electronics of the two acquisition modes of the UMD. Image taken from Ref. [40].

Different reconstruction strategies can be defined depending on how the inhibition window is implemented (e.g. one-bin and infinite-window approaches [62]). The reconstruction in binary mode is subject to several sources of bias inherent to segmented detectors. These include pile-up, where multiple muons crossing the same scintillator bar within the detector time resolution are counted as a single particle, leading to under-counting, and corner-clipping muons, where a single inclined muon traverses adjacent bars and produces multiple signals, resulting in over-counting [114]. The latter effect depends on the detector geometry and increases with zenith angle and the relative azimuthal orientation between the detector and the shower axis. Additional effects such as detector inefficiency (missed muons) and noise (spurious signals from scintillation or electronics) are subdominant after calibration and threshold [114]. Dedicated correction and reconstruction procedures have been developed to mitigate these effects [40, 114]. These detector effects motivate the adoption of robust reconstruction strategies in binary mode. In this work, the infinite-window approach is used together with the corner-clipping corrections described in Ref. [61].

The infinite-window approach, introduced in Ref. [46], reconstructs the muon content without exploiting the detailed time structure of the binary trace. Instead, once a muon is identified, the detector is considered inhibited for the entire duration of the trace (i.e. all 2048 time bins). This ensures that each scintillator strip can contribute at most one counted signal per event, making the method robust against multiple triggering from wide signals or detector effects.

Within this framework, the response of a segmented detector can be described probabilistically. For a detector segmented into n_s independent strips, the probability of observing k fired strips given a true number of muons N_μ is given by [114]

$$P(k \mid N_\mu, n_s) = \binom{n_s}{k} S(N_\mu, k) \frac{k!}{n_s^{N_\mu}} \quad (1 \leq k \leq n_s), \tag{3.4}$$

where $S(N_\mu, k)$ are the Stirling numbers of the second kind, accounting for the number of ways of distributing N_μ indistinguishable muons into k non-empty strips.

An estimator for the number of muons can be derived from the mean response of the detector. In the limit of large n_s, this leads to the approximate relation [114]

$$\hat{N}_\mu \simeq \frac{\ln(1 - k/n_s)}{\ln(1 - 1/n_s)}, \tag{3.5}$$

which provides an unbiased reconstruction of the muon number under the assumption of a uniform distribution of muons over the detector surface.

3.4 $\overline{\text{Off}}\underline{\text{line}}$ analysis framework

The $\overline{\text{Off}}\underline{\text{line}}$ framework is the official simulation and reconstruction environment of Pierre Auger Observatory [101]. It provides a unified and modular platform for processing EAS data, allowing both the simulation of detector responses and the reconstruction of high-level observables from experimental measurements. Its primary role is the standardised reconstruction of shower quantities such as the arrival direction, core position, and primary energy.

The $\overline{\text{Off}}\underline{\text{line}}$ framework is a modular software system written predominantly in C++, organised around three main components: a processing pipeline, an event data model, and a detector description system. The processing pipeline consists of independent modules that perform specific tasks (e.g., calibration, geometry reconstruction, signal processing). These modules are executed sequentially and configured via XML steering files (bootstraps), allowing flexible yet reproducible workflows. Each module interacts with a central event data structure, which is progressively populated throughout the reconstruction chain. The event data model provides a unified container for both simulated and reconstructed quantities. It ensures a consistent interface between different stages of the analysis and allows intermediate results to be stored and accessed for validation and systematic studies. The detector description system supplies time-dependent information about the detector and environment, including calibration constants, detector status, and atmospheric conditions.

This ensures that both simulated and real events are processed under realistic and consistent conditions.

A dedicated geometry package manages coordinate transformations between different reference frames (e.g., shower frame, ground frame, detector coordinates), which is essential for lateral distribution analyses and accurate determination of core distances.

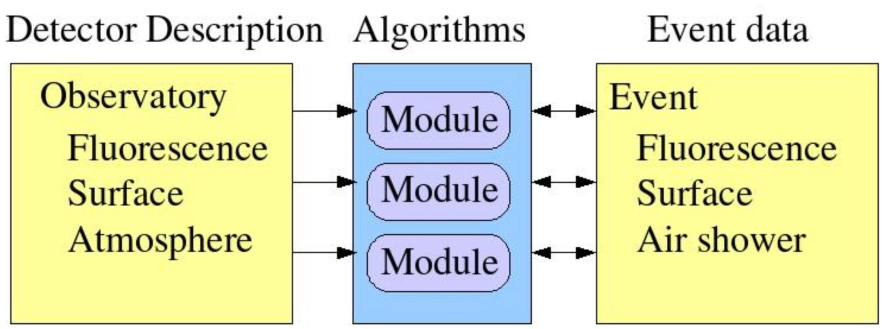

Figure 3.11: Schematic structure of the Offline framework, illustrating the modular processing pipeline, event data model, and detector description system.

A key feature of the framework is its integration with the air-shower simulation program CORSIKA (COsmic Ray SImulations for KAscade) [52]. CORSIKA simulates the development of EAS in the atmosphere, providing the list of secondary particles reaching ground level. These particles are then processed within Offline to simulate detector responses, completing the full chain from primary CR to recorded detector signal. This coupling ensures that simulated events are reconstructed using the same algorithms applied to experimental data, enabling direct and consistent comparisons. The development of the air shower depends on the chosen high-energy hadronic interaction model, which directly affects observables such as the muon content and shower fluctuations. These dependencies constitute an important source of systematic uncertainty in composition studies.

Event data within Offline are stored in the ROOT format [82] using Auger Data Summary Trees (ADST), which implement a hierarchical structure tailored to event-based measurements. Most observables used in this thesis are directly derived from ADST variables, following a class–member naming convention. Additional quantities are computed from these base variables or extracted directly from simulation outputs.

3.4.1 Simulation of the Underground Muon Detector

The UMD simulation implemented in Offline provides a detailed modelling of the muonic component of air showers. This is particularly relevant for mass composition studies, as muons originate primarily from hadronic interactions in the cascade. Muons generated in the air shower must traverse an overburden of soil before reaching the detector. During propagation, they lose energy primarily through ionisation processes. As a result, only muons above a certain threshold energy contribute to the UMD signal, $E_\mu \gtrsim 1$ GeV. This threshold acts as a natural filter, removing low-energy muons and making the UMD sensitive predominantly to the high-energy component of the muon spectrum. Within the scintillator modules, muons produce scintillation light, which is transported via optical fibres and detected by SiPM. The detector response simulation includes photon production, light attenuation, and the electronic response of the SiPM, providing a realistic representation of the measured signals.

All air-shower events used in this thesis were reconstructed using a recent development version of Offline, ensuring consistency with the reconstruction procedures applied to experimental data. The accurate modelling of detector responses, particularly for the muonic component, is essential for interpreting observables such as lateral density distributions and their correlation with shower development parameters. In particular, the muon content influences the shape and normalisation of lateral distribution functions, and therefore directly impacts composition-sensitive observables. The UMD simulation provides an independent constraint on the hadronic component of the shower, which is critical for reducing systematic uncertainties in composition analyses.

3.5 Summary

In this chapter, the detection systems of Pierre Auger Observatory were described, emphasising their hybrid design and complementarity [14]. The Observatory combines a nearly 100% duty-cycle SD with a fluorescence-based calorimetric measurement (FD), enabling high-statistics measurements with controlled systematics.

The SD reconstructs shower geometry and size from the lateral distribution of signals in WCD. Using a modified NKG lateral distribution function and a likelihood-based reconstruction, the optimal signal $S(r_{\text{opt}})$ is obtained [100]. Atmospheric attenuation effects are corrected via the CIC method [102], and the energy scale is established through cross-calibration with high-quality hybrid events reconstructed by the FD [4].

The FD measures the longitudinal energy-deposit profile of air showers by detecting nitrogen fluorescence light [7]. The profile is fitted with a Gaisser–Hillas function [9], yielding both the calorimetric energy and the depth of shower maximum X_{max}, the most powerful composition-sensitive observable at ultra-high energies [6]. Corrections for invisible energy carried by neutrinos and high-energy muons are applied to obtain the total primary energy [4]. Continuous atmospheric monitoring ensures control of systematic uncertainties in light transmission and fluorescence yield [7].

AugerPrime enhances mass sensitivity through additional detector components [87]. The SSDs enable a decomposition of electromagnetic and muonic signals at ground level by exploiting the different responses of scintillator and WCD [103]. The UUB electronics increase sampling frequency and dynamic range, and the small-PMT extends the measurable signal range near the shower core. The RD provides nearly continuous sensitivity to the electromagnetic component and X_{max} through coherent radio emission in the 30–80 MHz band [104, 105].

Finally, the UMD provides a direct measurement of the muonic component over the energy range $10^{16.5}$–10^{19} eV [10, 39]. By reconstructing muon densities and time profiles, it offers an independent and composition-sensitive probe that is essential for testing hadronic interaction models and for addressing the muon deficit problem [19, 83]. Together, these systems form a multi-hybrid observatory capable of disentangling electromagnetic and hadronic shower components and of constraining the mass composition of ultra-high-energy CR with unprecedented precision.

CHAPTER 4

Analysis and characterisation of the muon lateral distribution

The radial distribution of muons in EAS provides a sensitive probe of both the primary CR mass and the modelling of hadronic interactions in the atmosphere across a wide energy range, with particular sensitivity to the low-energy interactions governing meson decay into muons. At fixed energy, heavier primaries generate a larger number of muons and therefore produce systematically higher muon densities at ground level compared to lighter primaries. The MLD, which describes the mean muon density as a function of distance from the shower axis, is consequently a key observable for composition studies and for evaluating the consistency among shower simulations, detector response modelling, and reconstructed data.

Beyond its overall normalisation, the MLD shape encodes information on the transverse development of the hadronic cascade. The lateral spread of muons at ground level is governed by the transverse momentum of the parent mesons (predominantly pions and kaons) and by the geometry of muon production and propagation in the atmosphere. Larger typical transverse momentum produces a broader muon footprint and a flatter radial dependence, whereas smaller transverse momentum yields a more concentrated distribution and a steeper MLD. Because these features are controlled by hadronic-particle production and subsequent decay kinematics, the MLD shape provides complementary sensitivity to both the primary mass and high-energy hadronic-interaction modelling.

In this chapter, the average MLD measured with the UMD detector is evaluated in bins of reconstructed energy and zenith angle and compared with detector-level predictions from simulated air showers under realistic mixed-composition scenarios. The comparison is performed directly at the level of profiled muon densities computed in radial bins and confronted with composition-weighted MC predictions, without imposing an explicit parametric form for the MLD. This bin-by-bin strategy reduces reliance on functional parametrisations and preserves sensitivity to genuine shape differences between data and simulation. In each energy bin, a single global scale factor is extracted as a first-order test of compatibility with a predominantly multiplicative muon deficit.

The remainder of this chapter is organised as follows: Section 4.1 introduces the simulation framework and experimental dataset, including the air-shower library, detector-level reconstruction, and the construction of mixed-composition reference samples based on externally constrained mass fractions. Section 4.2 presents the empirical calibration that aligns simulated energies with the experimental energy scale and evaluates its performance through bias and resolution studies. Section 4.3 describes the spectral re-weighting applied to the simulated events to reproduce the measured CR flux, ensuring that averaged observables reflect physically realistic energy distributions. Section 4.4 defines the event-selection strategy used to construct reliable radial profiles. This includes determining the minimum and maximum usable radial distances based on reconstruction stability and trigger efficiency, and an outlier-rejection procedure that suppresses pathological measurements while preserving the physical distribution. Section 4.6 then describes the construction of profiled muon density distributions and introduces unbiased estimators that correct for missing or non-physical entries, ensuring statistically consistent radial averages. The central analysis is presented in Section 4.7, where measured and simulated average MLD profiles are compared through a global multiplicative scale fit. This section details the extraction of the scale factor and the interpretation of residuals as indicators of normalisation and shape agreement. In this section, systematic uncertainties are also described. Finally, Section 4.8 summarises the main results of the analysis and discusses their implications for the interpretation of the muon content of EAS and for the modelling of hadronic interactions at ultra-high energies.

4.1 Simulations and data

4.1.1 Data

The experimental data set used in this analysis consists of reconstructed events recorded by the UMD modules instrumented with SiPMs. Data collected between January 2018 and December 2021 is referred to as Phase 1, while data collected between March 2023 and September 2024 is referred to as Phase 2. All event reconstruction was performed using the Offline framework (git version 9a9cf5811), explained in Section 3.4. This reconstruction chain provides calibrated detector observables and reconstructed shower parameters that are directly comparable to the detector-level simulation outputs. The infinite-window approach, together with corner-clipping corrections [61] (see Section 3.3.3), was used to reconstruct the number of muons measured by the binary mode of the UMD. From December 2021 to February 2023, during the transition between the old and new electronics of the WCD in the SD–750 array, no SD hexagons shared the same electronics. Consequently, no event fulfils the 6T5 (see Section 3.1) and the usage of old electronics simultaneously. After that, when the entire array was equipped with new electronics, this period of data acquisition is termed Phase 2.

Pierre Auger Observatory produces an official dataset for each of the biannual ICRC[1]. All standard quality cuts used for the production of official Pierre Auger Observatory SD data for ICRC 2025 were applied to the two datasets, which include a physical T4 trigger (see Section 3.1), a successful fit of the SD LDF, and the rejection of known bad data acquisition periods or lighting problems. In addition to the standard SD quality cuts, the UMD requires additional conditions for events. These include the requirement that the SD station with the largest signal in the event has a non-rejected UMD and a successful fit of the MLDF.

4.1.2 Simulations

Air shower library

A MC air-shower library was used as the basis for all simulation studies in this work. The library contains simulated showers initiated by proton, helium, nitrogen, and iron primaries and was originally produced for the thesis work reported in Ref. [62]. The showers were generated using COsmic Ray SImulations for KAscade (CORSIKA) v7.740 [52]. High-energy hadronic interactions were modelled using three alternative interaction models, EPOS-LHC [12], QGSJet-II.04 [55], and Sibyll-2.3c [56]. Low-energy hadronic interactions were simulated

[1] International Cosmic Ray Conference

using the Ultra-relativistic Quantum Molecular Dynamics model (UrQMD) [54]. This combination of interaction models provides a consistent description of particle production across the full energy range of the air shower. The simulated arrival directions were isotropically distributed with zenith angles in the range $\theta \in [0°, 48°]$. Primary energies were generated with a uniform distribution in $\log_{10}(E/\mathrm{eV})$ over the interval $\log_{10}(E/\mathrm{eV}) \in [16.8, 18.7]$. For each primary species and for each high-energy hadronic interaction model, 3000 independent air showers were simulated. This provides comparable statistical samples across primary masses and interaction model assumptions.

Detector simulations with $\overline{\text{Off}}\underline{\text{line}}$

The air-shower library in Section 4.1.2 was propagated through a full detector simulation using the $\overline{\text{Off}}\underline{\text{line}}$ framework to produce simulated detector responses as explained in Section 3.4. The MdSdInfillSimRec module was used to simulate the response of both the UMD and the SD array in the infill region of Pierre Auger Observatory. Similar to Phase 1 and Phase 2 data, the infinite-window approach together with the corner-clipping corrections was used to reconstruct the number of muons simulated by the binary mode of the UMD.

To enhance the effective detector-level statistics, each simulated air shower was reused five times by repositioning its core at different locations across the detector array before reconstruction. This procedure generates multiple detector-level realisations of the same underlying shower, differing only in their geometry relative to the array layout, while preserving the intrinsic shower development. The potential bias introduced by reusing input showers was explicitly evaluated. In agreement with the findings reported in Ref. [115], the effect was found to be negligible within the statistical precision of the analysis, indicating that this approach does not introduce a significant systematic distortion in the reconstructed observables.

4.1.3 Mixed composition scenarios

To model realistic primary CR compositions, mixed-composition simulation datasets were constructed. These datasets, collectively referred to as AugerMix, are designed to reproduce the energy-dependent mass composition inferred from experimental measurements. The AugerMix samples were generated by combining simulated air showers from the proton, helium, nitrogen, and iron primaries according to composition fractions obtained from fits to the observed X_{max} distributions measured by Pierre Auger Observatory [59]. These fits provide the relative abundance of each nuclear species as a function of energy and are derived by comparing the measured X_{max} distributions with air-shower predictions from different high-energy hadronic interaction models.

The composition fractions used in this work correspond to those obtained for the high-energy hadronic interaction models, EPOS-LHC, QGSJet-II.04, and Sibyll-2.3c. For each model, the predicted X_{max} distributions differ due to model-dependent shower physics; consequently, the inferred mass fractions are also model-dependent and were therefore constructed for each interaction model to preserve internal consistency between shower development and composition inference. The mixed datasets were generated by sampling simulated events from each primary species in proportions that corresponded to the corresponding energy-dependent mass fractions. In this way, AugerMix represents a physically motivated approximation of the true CR composition, incorporating both the experimental constraints of the X_{max} measurements and the theoretical framework of the chosen high-energy hadronic interaction model.

The mass fractions obtained using EPOS-LHC, QGSJet-II.04, and Sibyll-2.3c high-energy hadronic interaction models are shown in Fig. 4.1. The error bars indicate the statistical uncertainties associated with the fitted composition fractions. These uncertainties reflect the finite size of the experimental data set and propagate into the mixed-composition simulations through the sampling procedure. The trend of the He and N fractions as a function of energy has a strong dependence on the particular high-energy hadronic interaction model used. The same is also true for proton primaries at energies $\log_{10}(E/\mathrm{eV}) > 10^{18}$.

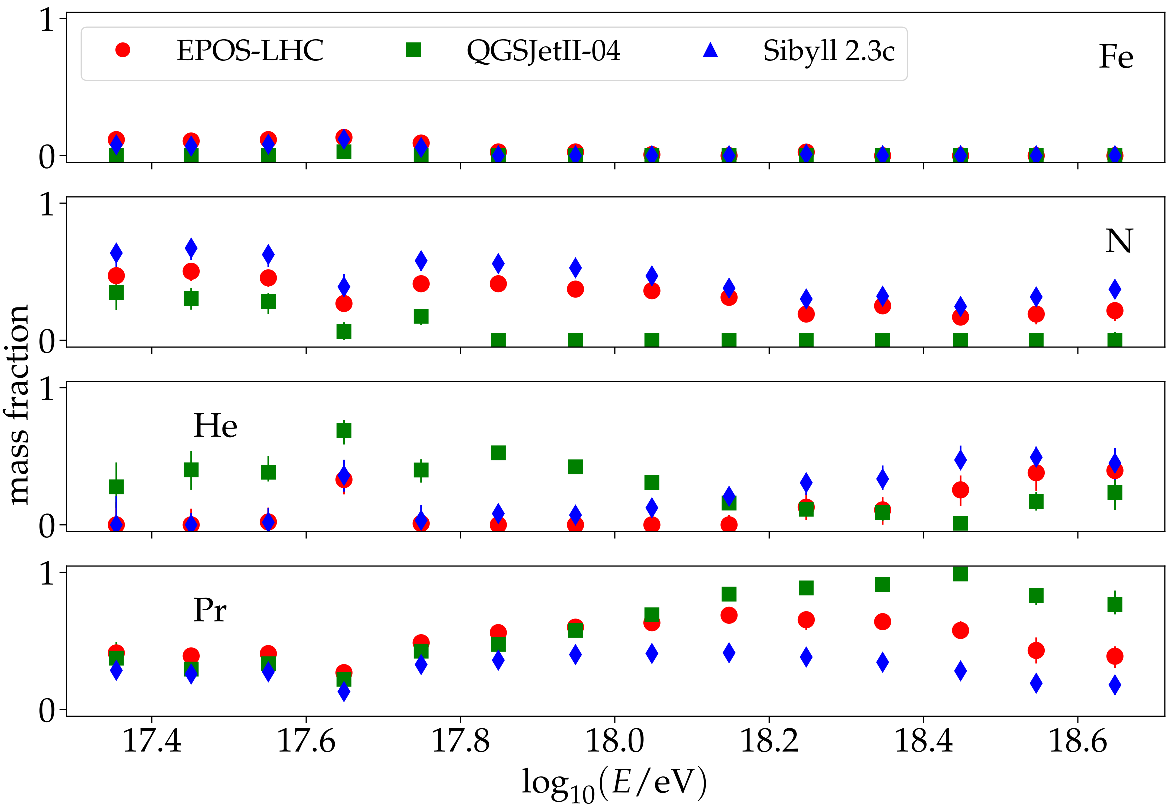

Figure 4.1: Mass fraction fits to the experimental $X_{\max}$ distributions as a function of the logarithm of the primary energy for EPOS-LHC, QGSJet-II.04 and Sibyll-2.3c high-energy hadronic interaction models [59]. The error bars indicate the statistical errors.

4.2 Energy calibration

To compare simulated events directly with the data, the MC energies are mapped onto the Pierre Auger Observatory reconstructed energy scale. This empirical calibration accounts for the systematic differences between the true energy assigned to simulated events, E_{MC}, and the energy that would be reconstructed using the Pierre Auger Observatory calibration procedure. The reconstructed energy is parametrised as a power-law function of the MC energy,

$$E_{\mathrm{rec}} = A\, E_{\mathrm{MC}}^{B} \tag{4.1}$$

This functional form captures both a global rescaling and a possible non-linear distortion of the energy scale. The calibration parameters are obtained from a MC dataset generated using Auger mass fractions reported at ICRC 2017, described in Section 4.1.3. This ensures that the calibration is derived under realistic composition assumptions consistent with experimental observations.

For numerical stability during the fitting procedure, both energies are normalised by a reference energy $E_{\mathrm{ref}} = 10^{18}$ eV. The scaled fitting relation is therefore written as,

$$\frac{E_{\mathrm{rec}}}{E_{\mathrm{ref}}} = A_{\mathrm{scaled}} \left(\frac{E_{\mathrm{MC}}}{E_{\mathrm{ref}}} \right)^{B} .$$

This normalisation reduces correlations between fit parameters and improves stability when fitting over a wide dynamic energy range. After determining the fit parameters, the physical normalisation constant is recovered by restoring the reference-energy scaling,

$$\log_{10} E_{\mathrm{rec}} = \log_{10} A_{\mathrm{phys}} + B \log_{10} E_{\mathrm{MC}},$$

where $\log_{10} A_{\mathrm{phys}} = \log_{10} A_{\mathrm{scaled}} + (1 - B) \log_{10} E_{\mathrm{ref}}$. This linear relation in log-space provides a direct mapping between the simulated and reconstructed energy scales. Using the fitted calibration parameters, the corrected MC energy is obtained by inverting the above relation,

$$\log_{10} E_{\mathrm{corr}} = \frac{\log_{10} E_{\mathrm{rec}} - \log_{10} A_{\mathrm{phys}}}{B}.$$

This transformation maps simulated energies onto the reconstructed Auger energy scale, enabling direct comparison between simulation and data across the full energy range of the analysis. The applied calibration corrects the dominant, first-order systematic offset in the energy scale between simulation and experiment. Residual higher-order or model-dependent systematic effects may remain, but are sub-leading compared to the primary scale correction addressed here.

The left panel of Fig. 4.2 illustrates the energy calibration obtained for EPOS-LHC. The reconstructed energy derived from the AugerMix dataset is shown as a function of the MC energy, together with the fitted calibration relation. The identity line is also shown for reference. The systematic deviation from the identity line reflects the initial mismatch between the simulation and Auger energy scales, while the fitted relation quantifies the empirical correction required to align them. The error bars represent statistical uncertainties associated with the reconstructed energy estimates.

The performance of the energy calibration is evaluated by quantifying both the systematic bias and the statistical resolution of the corrected energy estimate. This is done independently within bins of true MC energy. For each true energy bin, the fractional deviation between the corrected and true energies is defined as,

$$E_{\mathrm{ratio}} = \frac{E_{\mathrm{corr}}}{E_{\mathrm{MC}}} - 1. \tag{4.2}$$

This quantity measures the relative deviation of the corrected energy from the true simulated energy. A value of zero corresponds to perfect reconstruction, positive values indicate overestimation, and negative values indicate underestimation. The distribution of E_{ratio} within each energy bin is modelled using a Gaussian function,

$$f(E_{\mathrm{ratio}}) = A \exp\left[-\frac{1}{2} \left(\frac{E_{\mathrm{ratio}} - \mu}{\sigma} \right)^{2} \right],$$

where μ and σ characterise the first two moments of the fractional deviation distribution. The parameter μ represents the bias, defined as the average fractional deviation between corrected and true energies ($E_{\mathrm{corr}}/E_{\mathrm{MC}} - 1$). It quantifies any systematic shift in the reconstructed energy scale after calibration. A value of μ close to zero indicates that the calibration successfully aligns the reconstructed energy with the true MC energy on average. The

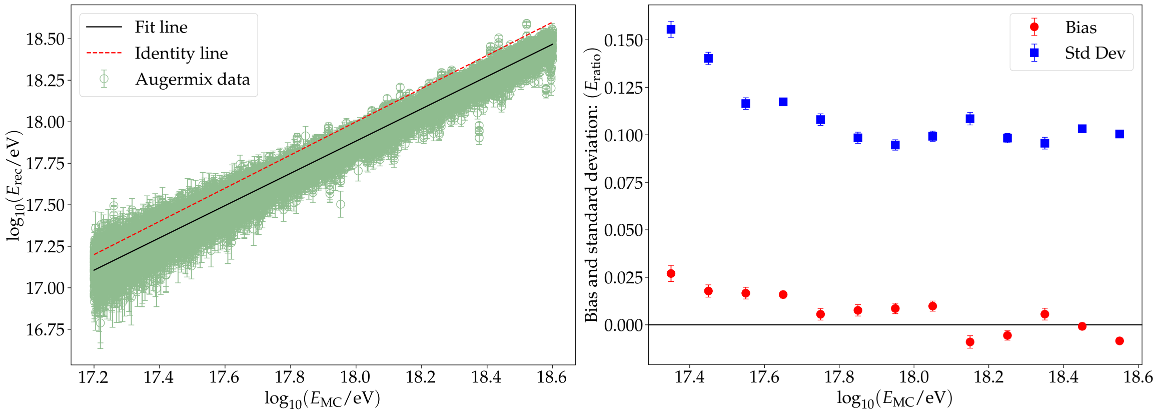

Figure 4.2: (Left) Reconstructed energy as a function of MC energy for the AugerMix dataset generated with EPOS-LHC high-energy hadronic interaction model. The black solid line shows the fitted calibration relation in Eq. (4.1), while the red dashed line indicates the identity relation corresponding to perfect agreement between reconstructed and simulated energies. Error bars represent statistical uncertainties on the reconstructed energy. (Right) Bias (red circles) and resolution (blue squares) of the calibrated energy as a function of true MC energy.

parameter σ represents the standard deviation of the fractional deviation distribution and is interpreted as the energy resolution. It characterises the statistical spread of reconstructed energies for showers with the same true energy.

The right panel of Fig. 4.2 shows the bias and energy resolution as functions of true MC energy. The residual bias is small across the full energy range, with an average value close to zero, indicating that the dominant systematic scale mismatch has been effectively corrected. Small fluctuations around zero are consistent with statistical uncertainties in individual energy bins. The energy resolution improves with increasing energy due to both a reduction in relative shower-to-shower fluctuations and enhanced detector sampling. Higher-energy primaries generate more particles at ground level and trigger a larger number of surface detector stations, increasing the sampling of the lateral distribution. This improved sampling reduces statistical uncertainties in reconstructed observables. Overall, the calibration yields an approximately unbiased energy estimate with an energy resolution that improves toward higher energies, consistent with the expected performance of the detector and reconstruction chain. The procedure was repeated by fitting the AugerMix composition model for QGSJet-II.04 and Sibyll-2.3c high-energy hadronic interaction models, and a similar energy correction was applied.

4.3 Spectrum

The simulated events are generated with a flat energy spectrum, whereas the measured CR flux follows the energy-dependent spectrum reported by Pierre Auger Observatory. To enable a physically meaningful comparison between simulations and data, the simulated events must therefore be re-weighted to reproduce the measured Auger spectrum. This re-weighting is necessary because the event distribution produced by the simulation does not reflect the true relative occurrence of different primary energies. In particular, a flat-generated spectrum overpopulates high-energy events relative to the physical flux. If uncorrected, this mismatch would bias any energy-dependent observable derived from the simulations. For example, averaged quantities such as the MLD would be systematically skewed toward the properties of higher-energy showers, since these would be over-represented relative to their true abundance. Applying spectral weights corrects this sampling imbalance and ensures that each energy interval contributes according to the measured CR flux. Each simulated event is therefore assigned a weight defined by the ratio between the physical flux and the generated flux,

$$w(E) = \frac{J_{\text{Auger}}(E)}{J_{\text{gen}}(E)} = \frac{J_{\text{Auger}}(E)}{E^{-1}} = E\, J_{\text{Auger}}(E), \tag{4.3}$$

where $J_{\text{gen}}(E) \propto E^{-1}$ describes the flat spectrum used in the simulation, and $J_{\text{Auger}}(E)$ is the parametrised Auger energy spectrum [58]. The latter is modelled as,

$$J_{\text{Auger}}(E) = J_0 \left(\frac{E}{10^{17}\text{eV}}\right)^{-\gamma_0} \prod_{i=0}^{1} \left[1 + \left(\frac{E}{E_{ij}}\right)^{\frac{1}{\omega_{ij}}}\right]^{(\gamma_i - \gamma_j)\omega_{ij}} . \tag{4.4}$$

with $j = i + 1$. The normalisation factor J_0 , the three spectral indices γ_i, and the transition parameter ω_{01} constitute the free parameters. The transition parameter ω_{12}, constrained with much more sensitivity using data from the SD–1500, is fixed. The best-fit values of the spectral parameters are shown in Table 4.1.

Table 4.1: Best-fit values of the spectral parameters. Obtained from Ref. [58].

Free Parameter	Values
$E_{12}/(\text{ eV})$	3.9×10^{18}
γ_1	3.34
γ_2	2.6
ω_{01}	0.49
$J_0\, /(\text{km}^2\,\text{yr sr eV})$	1.09×10^{-13}
Fixed Parameter	**Values**
γ_0	2.64
$E_{01}/(\text{ eV})$	1.24×10^{17}
ω_{12}	0.05

These break energies mark changes in the spectral slope and reproduce the observed structure of the measured CR flux relevant for the SD–750 array. The smooth transition terms ensure continuity of the spectrum across these regions.

After computing the event weights, a normalisation is applied so that the total effective number of weighted simulated events matches the desired sample size. The normalised weight assigned to the i-th event is

$$w_i^{(\text{norm})} = \frac{w(E_i)}{\sum_j w(E_j)}\, N, \tag{4.5}$$

where N is the total number of events. This normalisation preserves the relative spectral weighting while maintaining a controlled statistical scale for subsequent analyses. Through this procedure, the simulated data set is effectively reshaped to follow the measured Auger energy spectrum, ensuring that comparisons between simulation and data are unbiased across the full energy range.

The left panel in Fig. 4.3 illustrates the effect of the re-weighting procedure. The weighted simulated spectrum reproduces the shape of the Phase 1 data spectrum over the analysis range. This agreement provides direct validation that the weighting procedure corrects spectral sampling and prevents energy-dependent biases in averaged observables such as the MLD.

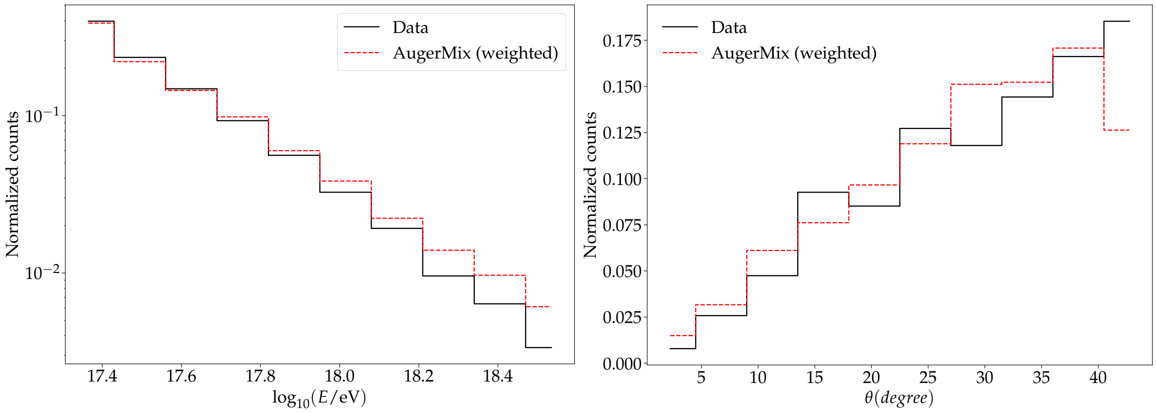

Figure 4.3: (Left) Comparison of the normalised reconstructed energy distributions for Phase 1 data and AugerMix simulations. The dashed histogram shows the AugerMix sample after applying spectral weights to reproduce the Auger flux. The black step histogram shows the reconstructed data distribution. (Right) Zenith angle distribution at $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$. Normalised event counts are shown for the measured data and the weighted AugerMix using EPOS-LHC high-energy hadronic interaction model.

The right panel in Fig. 4.3 compares the distribution of the zenith angle of the events at $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$ for EPOS-LHC high-energy hadronic interaction model between the data and the weighted AugerMix reference sample. Normalised counts are shown as a function of the zenith angle in the range $[0°, 45°]$. The distributions exhibit close agreement across the full angular range, indicating that the weighting procedure applied to the AugerMix sample successfully reproduces the observed zenith distribution of the data. No significant angular-dependent discrepancies are visible, suggesting that differences in event composition or detector acceptance are adequately accounted for in the reference model at this energy. This agreement is essential for a meaningful comparison of muon densities, as the muon content depends strongly on zenith angle through atmospheric slant depth and shower evolution. Any mismatch in the zenith distributions would introduce a bias in the comparison, making it impossible to disentangle genuine differences in muon content from purely geometrical or acceptance effects. Similar results are obtained for other energy bins and using other high-energy hadronic interaction models.

4.4 Data selection

4.4.1 Minimum distance cut: core reconstruction

A common procedure in the reconstruction of the MLD of the UMD is to keep the shower core position fixed to the value obtained from the SD reconstruction. While this approach simplifies the reconstruction, it does not account for the uncertainty in the SD core determination. Any uncertainty in the reconstructed core position propagates directly into the shower-plane distances assigned to the UMD detector stations. Since MLD is evaluated as a function of this distance, the uncertainty of the core position introduces a systematic distortion in the inferred radial profile. This effect is most significant at small radial distances from the shower axis. In this region, the MLD rises steeply, so even a small error in the reconstructed distance can lead to a large change in the inferred muon density. Consequently, fixing the core position without accounting for its uncertainty can bias the reconstructed MLD, particularly in the inner region where the radial gradient is largest.

The deviation between reconstructed and true shower-plane distances is illustrated in the left panel of Fig. 4.4, which shows the relationship between reconstructed distance and true distance for simulated proton and iron primaries. The simulations corresponding to EPOS-LHC high-energy hadronic interaction model in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$. The corresponding impact on the inferred muon density is shown in the right panel of the same figure for the proton primary. The distortion in density arises solely from the misassignment of radial distance due to core position uncertainty.

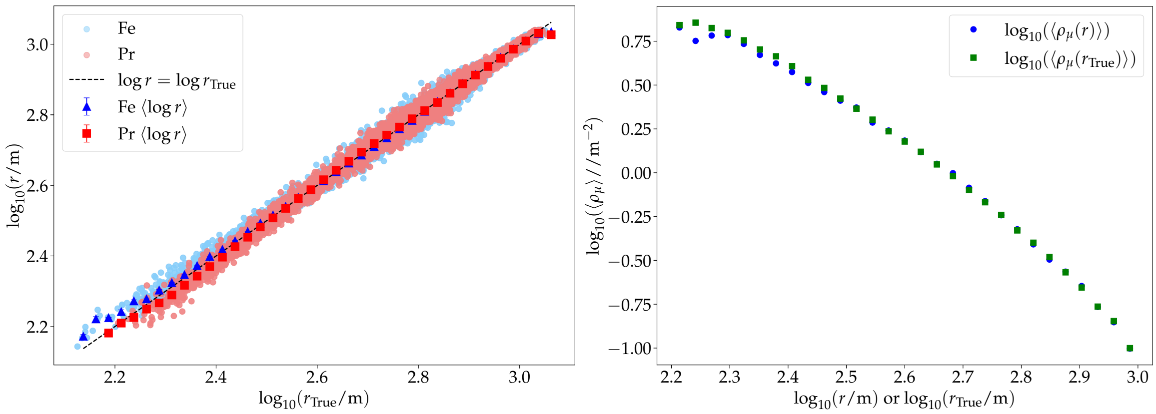

Figure 4.4: (Left) Reconstructed distance as a function of true distance for proton and iron primaries, with mean trends shown by binned averages. (Right) Comparison of the mean muon density for simulated proton primaries evaluated at reconstructed and true radial positions. Both figures corresponds to simulations with EPOS-LHC high-energy hadronic interaction model in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$.

To quantify the impact of this effect, two observables are studied as functions of radial distance,

$$\Delta\rho_\mu = \frac{\rho_\mu(r) - \rho_\mu(r_{\mathrm{True}})}{\rho_\mu(r_{\mathrm{True}})},$$

which measures the fractional deviation of the reconstructed density from the density evaluated at the true radial position, and

$$\Delta r = \frac{r}{r_{\mathrm{True}}} - 1,$$

which quantifies the relative error in reconstructed shower-plane distance.

The minimum distance threshold is derived using proton primaries because they exhibit the largest intrinsic shower-to-shower fluctuations. Proton-induced showers typically trigger fewer stations and have a lower muon content at ground, leading to larger uncertainties in the reconstructed core position. This makes them a conservative choice for defining the region affected by core-position uncertainty. Since the instability addressed by the cut is primarily geometric and reconstruction-driven, heavier primaries show equal or smaller deviations in the retained radial range. The proton-based threshold, therefore, defines a composition-independent reference region where radial coordinates and inferred densities are stable. The radial dependence of the parameter $\Delta\rho_\mu$ for proton simulations using EPOS-LHC high-energy hadronic interaction model in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$ and zenith range $\theta \in [0°, 45°]$, is shown in the left panel of Fig. 4.5. At small distances from the shower axis, the density deviation exhibits large fluctuations and systematic offsets. This behaviour reflects the strong sensitivity of the MLD to core position errors in the inner region. A reasonable minimum reconstructed core distance, above which the density deviation becomes sufficiently small, is inferred from the figure, and the procedure is repeated for the other energy bins. The radial dependence of the parameter Δr is shown in the right panel of Fig. 4.5 for proton and iron simulations using EPOS-LHC high-energy hadronic interaction model in the same energy bin and zenith range. This radial bias also exhibits large fluctuations at small distances. At larger distances, where the MLD varies more gradually with radius and the condition $\sigma_{\mathrm{core}} \ll r_i$ is satisfied, the effect of core-position uncertainty is reduced, and both observables approach zero.

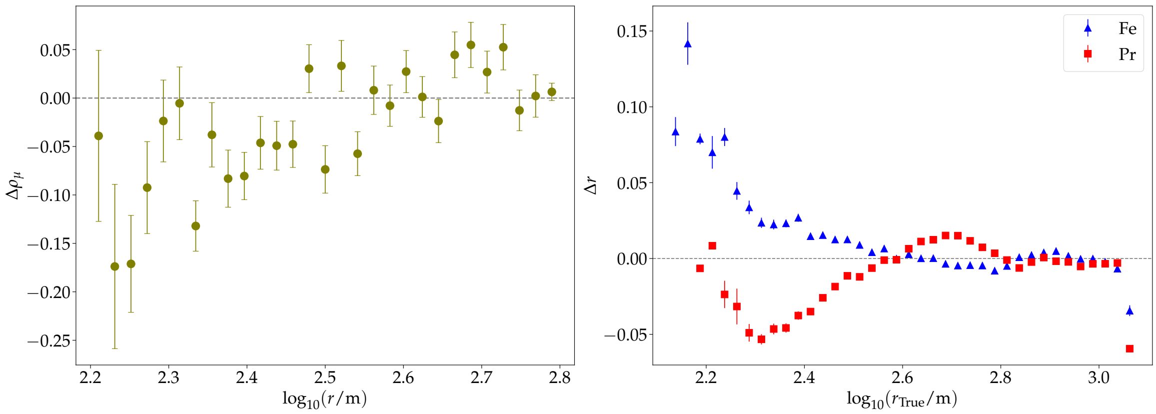

Figure 4.5: (Left) Parameter $\Delta\rho_\mu$ as a function of $\log_{10}(r/\mathrm{m})$ for proton primaries. (Right) Radial dependence of parameter Δr for proton and iron primaries. Both figures corresponds to simulations using EPOS-LHC high-energy hadronic interaction model in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$ and zenith range $\theta \in [0°, 45°]$.

A minimum distance cut is therefore introduced to exclude the region where the reconstructed radial coordinate is strongly affected by core-position uncertainty. The threshold is chosen to remove the unstable inner region while preserving the physically meaningful part of the MLD. Importantly, the selected cut maintains the overall shape of the MLD, ensuring that composition-sensitive features are not artificially suppressed. The minimum distance threshold is determined independently in each energy bin by identifying the radial region where both the density deviation and the radial bias become small and approximately stable. The resulting energy-dependent thresholds (r_{min}) are shown in Fig. 4.6. The threshold fluctuates moderately with energy. This selection ensures that the MLD is constructed using detector stations for which the radial coordinate is reliably determined, thereby reducing systematic distortions arising from fixed-core reconstruction while preserving the physical information encoded in the lateral distribution. The cut is therefore applied in all subsequent analyses.

4.4.2 Maximum distance cut: Lateral trigger probability of the surface detector

At large distances from the shower axis, the probability that a SD station triggers decreases due to the rapid drop in particle density. In this regime, only upward statistical fluctuations of the detector signal can produce a trigger. As a consequence, measurements at large core distances are subject to a selection bias: triggered stations preferentially correspond to upward fluctuations, which leads to a systematic overestimation of the measured muon density. To suppress this bias, a maximum usable distance from the shower axis is defined based on LTP[2] of the SD. The LTP describes the probability that a detector station located at a distance r from the shower axis triggers, given the shower energy and the zenith angle. Since particle density increases with primary energy and decreases with atmospheric depth, the

[2] lateral trigger probability

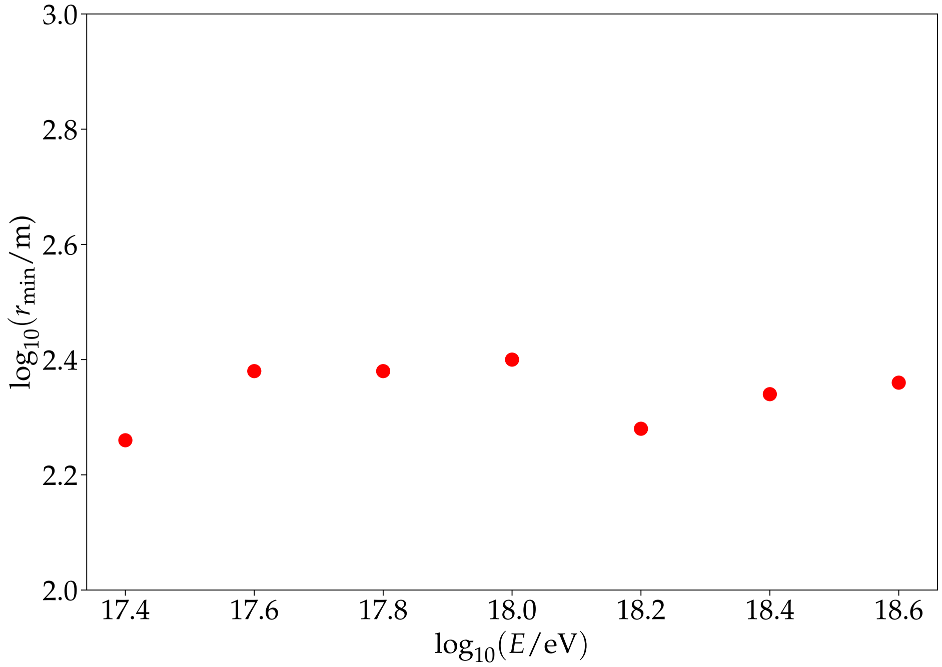

Figure 4.6: Energy dependence of the minimum radial distance cut derived from stability criteria on radial bias and density deviation.

LTP depends on both the energy and the zenith angle of the shower. The LTP is parametrised as a smooth function of distance, with a transition around a characteristic scale R_0 [60],

$$LTP(r) = \begin{cases} \dfrac{1}{1+\exp\left(\dfrac{r-R_0}{\Delta R}\right)} & \text{for } r < R_0, \\ \dfrac{1}{2\exp\left(\dfrac{r-R_0}{2\Delta R}\right)} & \text{for } r > R_0 \end{cases} \tag{4.6}$$

where R_0 and ΔR are free fit parameters with R_0 being the distance where the LTP is equal to 0.5 and ΔR controlling the steepness of the LTP. The parametrisations are obtained from fits to hybrid data measured with the FD and SD–750 array, as reported in Ref. [60]. This functional form describes a gradual transition from the near-unity trigger probability at small distances to the exponentially suppressed probability at large distances. The transition scale R_0 and the width ΔR depend on the shower energy and the zenith angle, reflecting the spatial change in the particle distribution at the ground level. The left panel of Fig. 4.7 shows representative LTP curves as functions of distance for different energies at fixed zenith angle $\sin^2\theta = 0.25$, and the right panel shows different zenith angles at fixed energy $\log_{10}(E/\mathrm{eV}) = 18$. The curves shift outward with increasing energy and inward with increasing zenith angle, consistent with the expected scaling of the shower footprint on the ground. To ensure that only regions of high and well-characterised trigger efficiency, close to unity, are used in the analysis, a maximum radial distance is defined by requiring the LTP to exceed a fixed threshold. In this work, the maximum distance is taken as the value of r at which $LTP(r) = 0.95$.

This criterion defines an energy- and zenith-dependent radial cut that excludes regions where trigger inefficiency would introduce selection biases. The resulting maximum allowed distances are shown in Fig. 4.8 as functions of energy for different zenith angles.

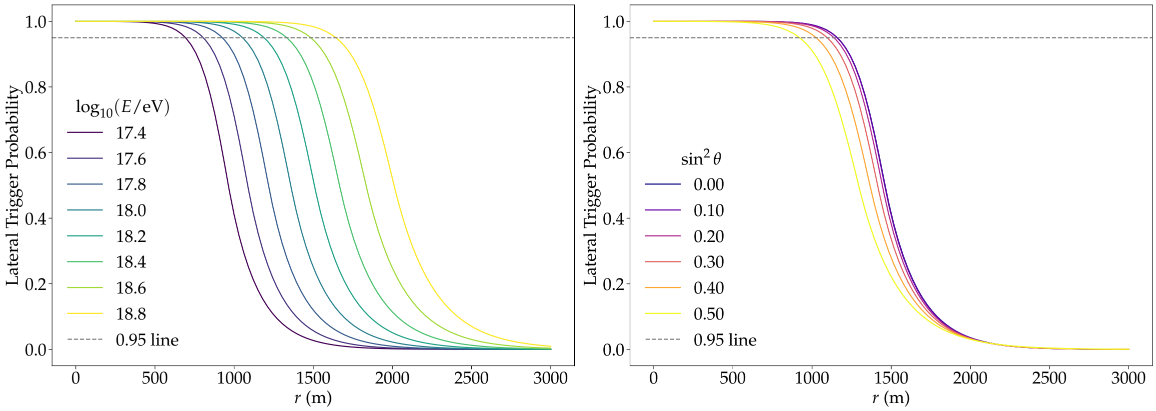

Figure 4.7: (Left) LTP as a function of the distance to the shower axis, as parametrised in Ref. [60] for different energies at fixed $\sin^2\theta = 0.25$. (Right) LTP as a function of the distance to the shower axis for different zenith angles at $\log_{10}(E/\mathrm{eV}) = 18$. Grey dashed lines mark 95% LTP.

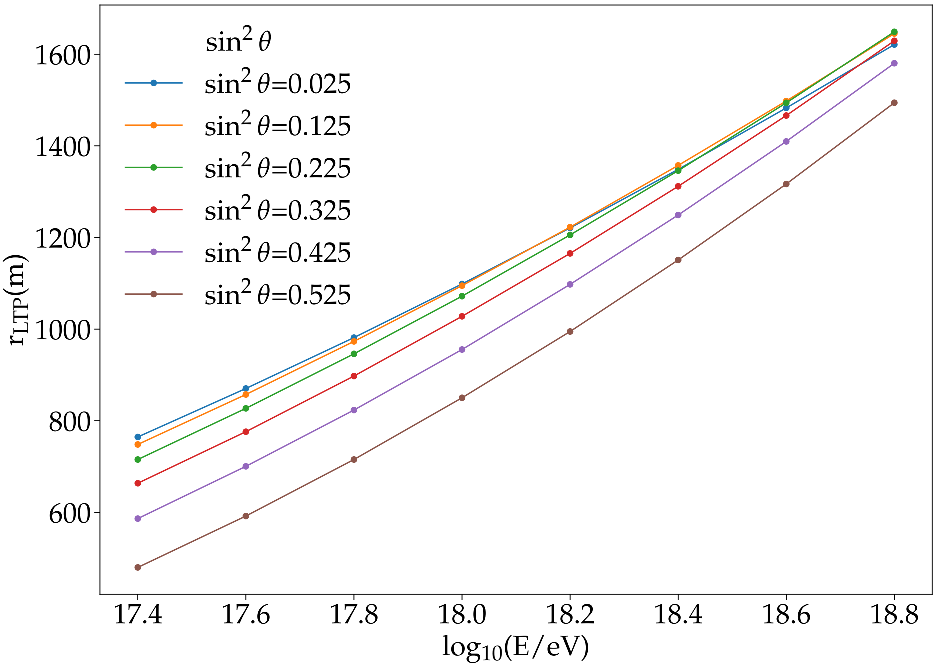

Figure 4.8: Maximum distance cuts corresponding to an LTP of 95%, as a function of the logarithmic energy, for different zenith angles (indicated by different colours).

This LTP-based cut ensures that the muon density measurements are restricted to regions where the detector response is effectively unbiased and the trigger efficiency is well controlled and is applied in the subsequent analyses.

4.4.3 Outlier detection

To suppress pathological density measurements that can distort the profiled MLD, an outlier rejection procedure is applied independently within the radial distance bins. The density distribution at the station level in each $\log_{10}(r/\mathrm{m})$ bin is characterised by robust statistics based on quantiles rather than the mean, to reduce sensitivity to outliers arising from occasional failures of the detection system. For each radial bin with width $\Delta\log_{10}(r/\mathrm{m}) = 0.025$ and energy bin width $\Delta\log_{10}(E/\mathrm{eV}) = 0.2$, in the zenith range $\theta \in [0°, 45°]$, the median density is calculated together with the 16th and 84th percentiles. A robust estimate of the local dispersion is then defined as

$$\sigma(r) = \frac{q_{84}(r) - q_{16}(r)}{2},$$

which corresponds to the standard deviation for a Gaussian distribution but remains stable in the presence of asymmetric or heavy-tailed data. Measurements are classified as outliers if they lie more than $5\sigma(r)$ from the median value in their corresponding radial bin,

$$\rho_\mu(r) < \text{median}(r) - 5\sigma(r) \quad \text{or} \quad \rho_\mu(r) > \text{median}(r) + 5\sigma(r).$$

This threshold is intentionally conservative and removes only extreme deviations that are unlikely to arise from physical shower fluctuations or statistical measurement noise. The procedure is applied separately for each energy bin and data period. Figure 4.9 illustrates the method for Phase 1 data in the energy interval $\log_{10}(E/\text{eV}) \in [17.9, 18.1)$. In the figure, the muon density is shown as a function of radial distance. The grey points represent all measurements, while the red stars indicate entries identified as outliers. The data are grouped into logarithmic radial bins. For each bin, the black diamonds show the median muon density. The downward arrow indicates the upper limits corresponding to the $\rho_\mu + 5\sigma(r)$ cut. When the lower bound of the interval falls below zero, the error bar is truncated to respect the physical constraint $\rho_\mu(r) \geq 0$.

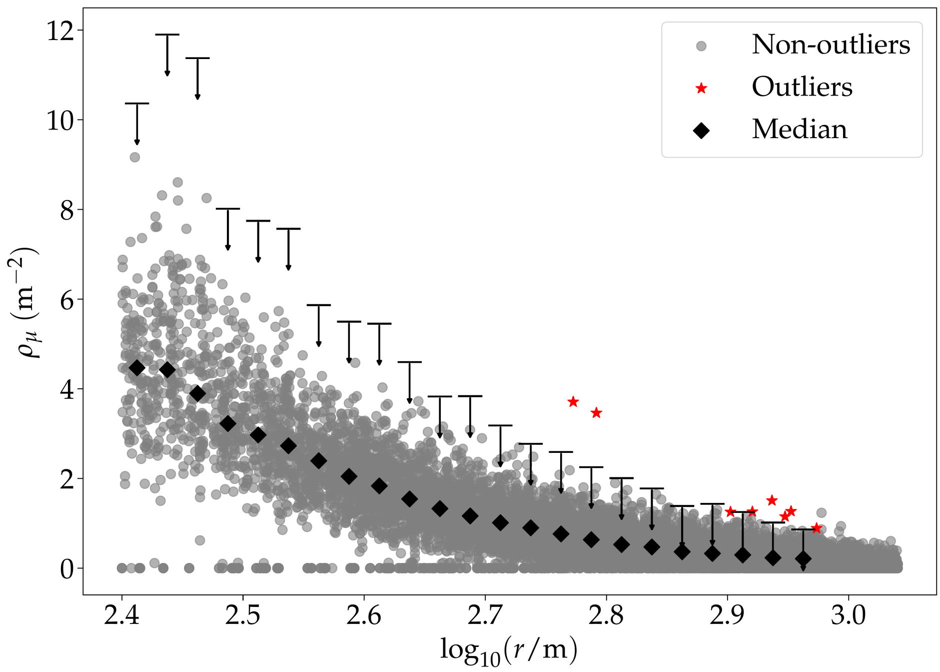

Figure 4.9: Muon density as a function of radial distance for Phase 1 data in the energy interval $\log_{10}(E/\text{eV}) \in [17.9, 18.1)$. Grey points show all measurements, while red stars mark entries identified as outliers. For each radial bin, the black diamonds indicate the median density and the downward arrow indicates the upper limit corresponding to the $\rho_\mu + 5\sigma(r)$ cut, where σ is estimated from the 16th and 84th percentiles.

Only a very small fraction of station measurements are identified as outliers by the rejection procedure. For Phase 1 data, 212 measurements were flagged from a total of 1,664,178, corresponding to less than 0.02% of all entries. Comparable fractions are observed for the Phase 2 data. This confirms that the adopted criterion 5σ removes only extreme measurements and does not significantly reduce the statistical power of the data set.

Figure 4.10 shows the distribution of the reconstructed events for both Phase 1 and Phase 2 combined as a function of the reconstructed energy and the zenith angle after the outliers are removed. Data are binned in logarithmic energy, $\log_{10}(E/\mathrm{eV})$, and in $\sin^2\theta$. The colour scale represents the number of events in each bin, and the corresponding counts are also displayed numerically. The distribution reflects the steeply falling CR spectrum, with event statistics rapidly decreasing toward higher energies.

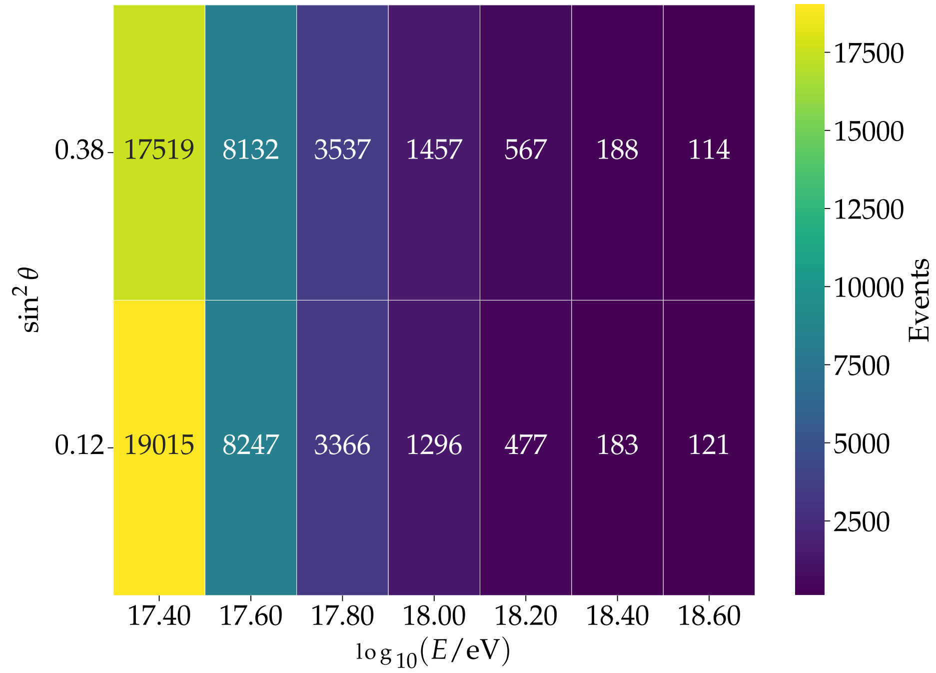

Figure 4.10: Distribution of reconstructed events as a function of the logarithmic reconstructed energy $\log_{10}(E/\mathrm{eV})$, and $\sin^2\theta$. The colour scale indicates the number of events in each bin, with explicit numerical counts.

4.5 Phase 1–Phase 2 data comparison

After the quality cuts mentioned in 4.4, the corresponding energy distributions of the Phase 1–Phase 2 datasets in the zenith range $\theta \in [0°, 45°]$ are shown in the left panel of Fig. 4.11. The distributions exhibit similar shapes across the full energy range, with comparable population levels in each bin. The final dataset comprises 32314 events from Phase 1 and 31905 events from Phase 2.

The consistency between the Phase 1 and Phase 2 datasets is evaluated through a comparison of their average muon density profiles as a function of the radial distance from the shower axis in the zenith range $\theta \in [0°, 45°]$. For each event, the muon density is expressed as a function of $\log_{10}(r/\mathrm{m})$, and the data are grouped in bins of reconstructed energy.

Within a given energy interval, the average muon density profile is constructed by dividing the events by $\log_{10}(r/\mathrm{m})$ and computing, in each radial bin, the arithmetic mean of the measured muon densities. The statistical uncertainty associated with each bin is estimated as the standard error of the mean, $\sigma/\sqrt{N}$, where σ is the standard deviation of the densities in the bin and N is the number of contributing events. This procedure yields, for each energy bin and for each phase, a profile $\langle\rho_\mu(r)\rangle$ that characterises the radial distribution of muons. To assess the agreement between the two data-taking periods, the ratio of the mean muon densities is then computed bin-by-bin in $\log_{10}(r/\mathrm{m})$:$\langle\rho_{\mu,1}(r)\rangle/\langle\rho_{\mu,2}(r)\rangle$.

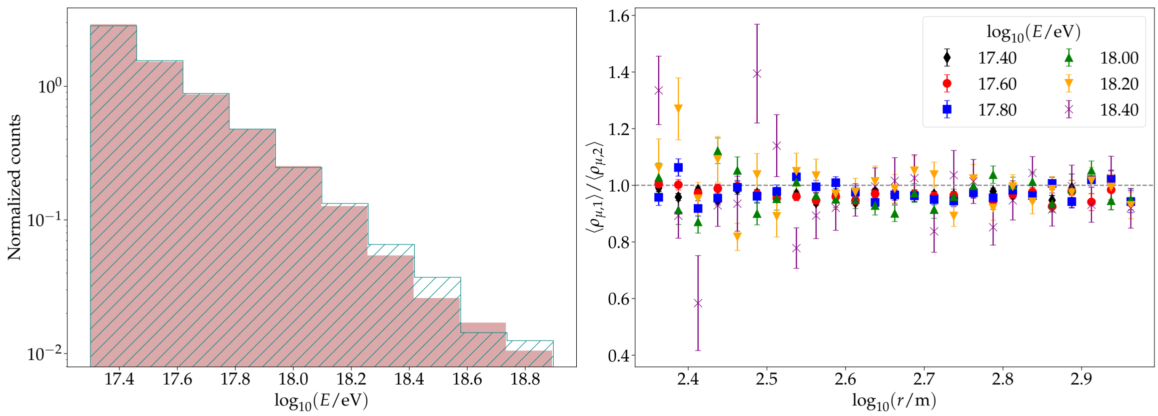

Figure 4.11: (Left) Energy distribution of two phases (Phase 1 and Phase 2) of data. (Right) Ratio of the average muon density of the two datasets compared for different energy bins as a function of $\log_{10}(r/\mathrm{m})$.

As shown in the right panel of Fig. 4.11, the ratio of the average muon densities remains close to unity over the full radial range considered. No systematic radial dependence is observed, and the ratios fluctuate around 1 within statistical uncertainties for all energy bins. The scatter increases slightly at the highest energies, consistent with reduced event statistics in those bins rather than with a systematic detector effect. Overall, the absence of coherent offsets or distance-dependent trends indicates that the measured muon density profiles are statistically compatible between Phase 1 and Phase 2.

This level of agreement demonstrates that the detector performance remained stable across the two operational periods and that the reconstruction procedure yields consistent muon density estimates. No evidence is observed for time-dependent calibration shifts, variations in detector efficiency, or reconstruction-induced distortions affecting the MLD measurement. Furthermore, the upgrade of the SD electronics has no measurable impact on the UMD data. Given this demonstrated compatibility, the Phase 1 and Phase 2 datasets can be combined without introducing additional systematic uncertainty, thereby increasing the statistical power of the analysis.

4.6 Calculation of the muon density

The reconstructed muon densities from the data are compared with MC simulations to assess the agreement between models and measurements and to evaluate the composition sensitivity of the data. The analysis is restricted to energy $\log_{10}(E/\mathrm{eV}) \in [17.3, 18.6)$ and zenith angle $\theta \in [0°, 45°]$, as the reconstruction of more inclined events becomes increasingly challenging, leading to larger systematic uncertainties. Events are selected within the reconstructed energy interval. This selection ensures that the compared showers share similar stages of development and detector acceptance. Prior data cleaning removes saturated detector signals and previously identified outliers, ensuring that only physically meaningful density measurements contribute to the radial profiles.

Simulated events generated with the same selection criteria are profiled using the same binning procedure as applied to the data. This ensures that differences between data and simulations reflect physical or modelling effects rather than analysis artefacts.

Figure 4.12 shows the mean muon density as a function of the radial distance for simulated proton and iron primaries using EPOS-LHC high-energy hadronic interaction model, together with measurements from the data from Phase 1 and Phase 2. Data and simulations are shown for the energy range $\log_{10}(E/\mathrm{eV}) \in [17.4, 17.6)$ and the zenith angle range

$\theta \in [0°, 45°]$. The muon density decreases approximately quadratically in the log − log space, consistent with the expected power-law behaviour of the MLD. At all radial distances, iron primaries produce systematically higher muon densities than proton primaries, reflecting the larger hadronic multiplicity and more advanced shower development associated with heavier nuclei. The reconstructed data from both detector phases lie between the proton and iron expectations across the full radial range. The two data phases exhibit mutually consistent profiles, indicating stable detector performance and reconstruction across data collection periods.

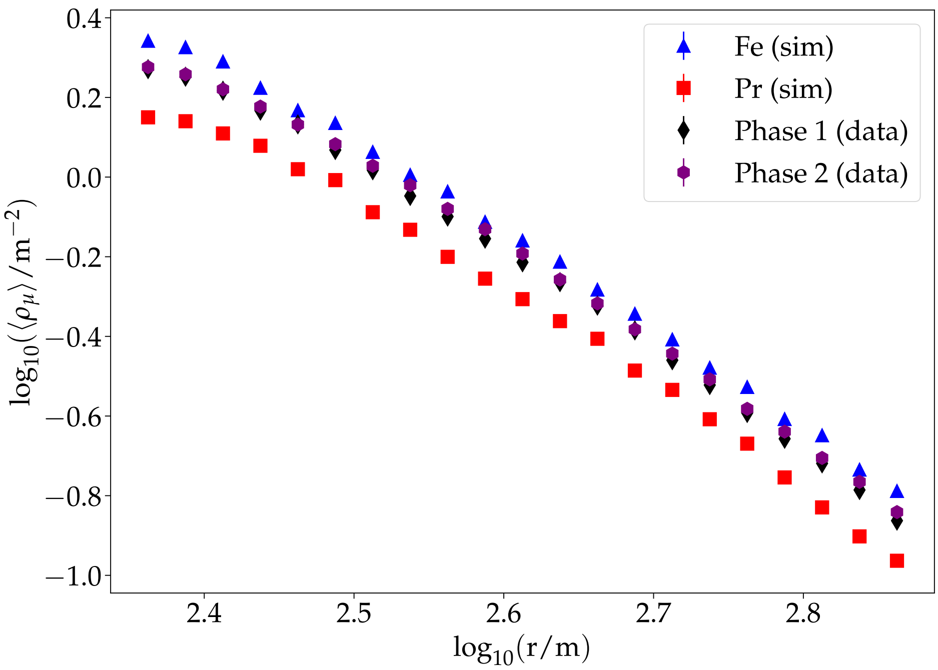

Figure 4.12: Mean muon density as a function of radial distance from the shower axis. Symbols show binned averages for simulated proton and iron primaries using EPOS-LHC high-energy hadronic interaction model and reconstructed data from Phase 1 and Phase 2. Error bars represent statistical uncertainties on the mean density. The data and simulations are shown for energy range $\log_{10}(E/\mathrm{eV}) \in [17.4, 17.6)$ and zenith angle range $\theta \in [0°, 45°]$.

No strong radius-dependent deviation is observed between the data and the simulation, suggesting that the model reproduces the shape of the MLD. The dominant separation between proton and iron profiles appears primarily as an overall normalisation difference rather than a change in slope. This confirms that the MLD retains a strong sensitivity to primary mass while remaining robust against moderate variations in radial scale.

Overall, the comparison demonstrates good agreement between the reconstructed data and simulations. The measured muon densities are compatible with expectations for a mixed composition, show no evidence of significant distortion in the radial profile shape, and confirm their suitability as an input observable for subsequent composition analyses.

4.6.1 Pinocchios and unbiased estimators

In the station-level muon density sample, a small fraction of measurements corresponds to non-physical pinocchio entries, hereafter referred to as pinocchios. These arise when the muon trace is empty or when no muons are recorded, which leads the observable to be effectively absent but stored as a zero-valued entry [63]. Such entries do not represent physical muon densities and therefore introduce a systematic dilution of statistical estimators if treated as ordinary data. The occurrence of pinocchios is modelled as a Bernoulli process in this work. In each trial, with probability p, the measurement is a Pinocchio and takes the value $\rho_\mu = 0$, while with probability $1 - p$ the measurement is drawn from the underlying physical distribution of ρ_μ. In this work, the probability of Pinocchio is taken to be constant at $p = 0.014$ [63].

Let $\{\rho_{\mu,i}\}_{i=1}^{m}$ denote the m observed values in a given $\log_{10}(r/\mathrm{m})$ bin, including pinocchio entries. The usual sample mean,

$$\bar{\rho}_\mu \equiv \frac{1}{m}\sum_{i=1}^{m}\rho_{\mu,i}, \tag{4.7}$$

is biased with respect to the mean of the underlying physical distribution, since the presence of pinocchios reduces the expectation value of the mean of the sample by a factor $(1 - p)$. An unbiased estimator of the true physical mean is therefore obtained by correcting for this dilution,

$$\bar{\rho}_{\mu,\mathrm{ub}} = \frac{\bar{\rho}_\mu}{1-p} = \frac{1}{m(1-p)}\sum_{i=1}^{m}x_i. \tag{4.8}$$

To consistently propagate statistical uncertainties, an unbiased estimator of the variance of the physical distribution must also be constructed. Since the presence of Pinocchio entries modifies both the first and second moments of the observed sample, the standard sample variance is no longer unbiased. The variance estimator must therefore correct for the Bernoulli dilution and the finite sample size. The resulting unbiased estimator is

$$S^2_{\mathrm{ub}} = \frac{m - mp + p}{m(m-1)(1-p)^2}\sum_{i=1}^{m}\rho^2_{\mu,i} - \frac{m}{(m-1)(1-p)^2}\,\bar{\rho}^2_\mu, \tag{4.9}$$

This estimator has the property that its expectation value equals the variance of the underlying physical muon density distribution. In the limit $p \to 0$, it reduces to the standard unbiased sample variance. The statistical uncertainty on the unbiased mean is estimated as follows.

$$\sigma(\bar{\rho}_{\mu,\mathrm{ub}}) = \frac{S_{\mathrm{ub}}}{\sqrt{m(1-p)}}, \tag{4.10}$$

where $m(1 - p)$ is the expected number of non-Pinocchio entries in the bin.

These estimators are applied independently in each radial bin when constructing profiled muon density distributions $\rho_\mu(r)$, ensuring that both the mean and its uncertainty reflect the properties of the underlying physical signal rather than the mixture of physical and missing measurements. Note that this effect constitutes only a small correction, as inferred from the low Pinocchio probability.

4.7 Global Scale Comparison of Muon Density Profiles

The objective of this analysis is to compare the measured average muon density profile with the expectation from air-shower simulations, without introducing an explicit functional model for the lateral distribution, and thus no additional shape parameters. Rather than fitting a parametric form of the MLD, the comparison is carried out between the average muon density measured in the radial bins and that obtained from the simulations. This approach avoids potential biases associated with functional parametrisations and instead focuses on a direct, bin-by-bin comparison between data and simulations.

The analysis is performed in discrete bins of primary energy and zenith angle. For each energy bin, events are selected within a narrow interval in the reconstructed $\log_{10}(E/\mathrm{eV})$ where, for simulations, the energy is calibrated as described in Section 4.2 and restricted to predefined zenith ranges. The observable of interest is the radial muon density $\rho_\mu(r)$, expressed as a function of the logarithmic distance from the core $\log_{10}(r/\mathrm{m})$. Radial profiles are constructed by dividing events into bins in $\log_{10}(r/\mathrm{m})$ and computing the mean density in each bin. To correct for the bias associated with the Pinocchio thinning procedure, a bias-corrected estimator of the mean and its statistical uncertainty is used (see Section 4.6.1). This produces the data profile $y_D(r) = \langle \rho_\mu(r) \rangle_{\text{data}}$ and the corresponding statistical uncertainties $\sigma_{\text{stat}}[y_D(r)]$ for each energy-zenith bin.

MC simulations are generated separately for different primary mass groups. For each hadronic interaction model and energy bin, these primaries are combined into a composition mixture using externally determined fractions for AugerMix. Events are weighted by the assumed CR flux and the true simulated energy to reproduce the expected physical spectrum. The weighted radial profiles are then calculated for each primary and linearly combined to produce the predicted mixed profile $y_M(r) = \sum_A f_A y_M^A(r)$ with $y_M^A(r) = \langle \rho_\mu^A(r) \rangle$, where the sum runs over the primary species and f_A corresponds to the fraction of the assumed mass composition. The statistical uncertainty of the mixture $\sigma_{\text{stat}}[y_M(r)]$ is obtained by propagating the weighted MC errors of each component.

4.7.1 Evaluation of systematic uncertainties

The systematic uncertainties of the data are evaluated bin-by-bin in radius and energy bin width $\Delta \log_{10}(E/\mathrm{eV}) = 0.2$, in the zenith range $\theta \in [0°, 45°]$ at the level of the mean muon density profile. Each source of systematic uncertainty represents a potential modification of the measured profile arising from detector calibration, environmental effects, or reconstruction biases. The main sources of systematic uncertainties are discussed in Ref. [61]. The following subsections further explain each source in more detail.

Energy scale systematic

The absolute energy calibration of the detector is subject to systematic uncertainty, which affects the assignment of events to the reconstructed energy bins. Since the muon density profile is constructed within finite energy bins, it is sensitive to uncertainties in the absolute energy scale. The systematic uncertainty is evaluated by applying a shift to the reconstructed energies by $\pm 14\%$ [14], corresponding to the uncertainty of the energy scale. For each shifted scale, the event selection is repeated using the same nominal energy bin boundaries. This effectively changes the set of events entering each bin. The muon density profile is then recomputed for the shifted selections. The difference between the shifted and nominal profiles, evaluated at each radial distance, is taken as the energy-scale systematic uncertainty. The systematic uncertainty on the energy scale is asymmetric and is indicated as $\sigma_{\text{ES}}^{\uparrow}(r)$ and $\sigma_{\text{ES}}^{\downarrow}(r)$.

Muon pattern systematic (MP)

Due to the differences in the UMD modules (SiPM gain, scintillators, fibres, electronics, etc.), certain detectors in the UMD array can produce wider signals than others. The typical binary muon signal is $\approx 7 - 8$ "1s" wide. In each bar in the module, the muon signal is identified if there is at least one sequence of four consecutive "1s" in its binary trace. In Ref. [61], the impact of selecting four consecutive "1s" as the muon pattern is studied by repeating the reconstruction procedure with a shorter pattern and a wider pattern. The shorter pattern resulted in a shift of $+5\%$, while the wider pattern resulted in a shift of -7%. These values were taken directly from Ref. [61] and correspond to the systematic uncertainties due to muon patterns. The uncertainty is implemented as an asymmetric fractional variation of the mean muon density profile and is denoted as $\sigma^{\uparrow}_{\mathrm{MP}}(r)$ and $\sigma^{\downarrow}_{\mathrm{MP}}(r)$.

Corner-clipping systematic (CC)

The corner-clipping systematic accounts for the geometric response of the muon detectors, which depends on the angle of incidence of the shower and the azimuthal orientation. These arise from inclined muons that traverse two neighbouring scintillator bars, leading to double counting and consequently to an overestimation of the reconstructed muon density. This effect is described by a parametric correction factor $P_{cc}(\theta, \Delta\phi)$ defined in Ref. [61], which modifies the reconstructed muon density through a multiplicative response term.

$$P_{\mathrm{CC}}(\theta, \Delta\phi) = (m_0 + m_1 t)\, s + b_0 + b_1 t, \tag{4.11}$$

where $t = \sec\theta - 1.2$ and $s = |\sin\Delta\phi| - 0.5$. Here, θ is the shower zenith angle and $\Delta\phi$ is the azimuthal angle relative to the detector orientation. The parameters $m_0 = 0.03, m_1 = 0.14, b_0 = 0.035, b_1 = 0.076$ are obtained from calibration studies and are taken directly from Ref. [61]. The corrected muon density for each event is given by

$$\rho_{\mu,\mathrm{corr}} = \frac{\rho_\mu}{1 + P_{\mathrm{CC}}(\theta, \Delta\phi)}, \tag{4.12}$$

The mean corrected density in a given radial bin is, therefore, obtained by averaging the corrected event-level estimates,

$$\langle \rho_{\mu,\mathrm{corr}} \rangle = \left\langle \frac{\rho_\mu}{1 + P_{\mathrm{CC}}} \right\rangle. \tag{4.13}$$

The systematic uncertainty arises from the finite precision of the parameters that define the corner-clipping model and is obtained by propagating the uncertainties of the model parameters $p_i \in \{m_0, m_1, b_0, b_1\}$ to the mean corrected density. The variance of the mean density is computed by propagating the parameter uncertainties (including their correlations) through the correction using standard error-propagation formulas

$$\sigma^2_{\mathrm{CC}} = \left(\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial m_0}\right)^2 \sigma^2_{m_0} + \left(\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial m_1}\right)^2 \sigma^2_{m_1} + 2\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial m_0}\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial m_1}\mathrm{cov}(m_0, m_1) \tag{4.14}$$

$$+ \left(\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial b_0}\right)^2 \sigma^2_{b_0} + \left(\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial b_1}\right)^2 \sigma^2_{b_1} + 2\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial b_0}\frac{\partial\langle\rho_{\mu,\mathrm{corr}}\rangle}{\partial b_1}\mathrm{cov}(b_0, b_1) \tag{4.15}$$

where, $\sigma_{m_0} = 0.002, \sigma_{m_1} = 0.01, \sigma_{b_0} = 0.001$, and $\sigma_{b_1} = 0.004$. The covariances are $\mathrm{cov}(m_0, m_1) = 8.9 \times 10^{-6}$ and $\mathrm{cov}(b_0, b_1) = 7.6 \times 10^{-7}$. All these values were obtained from the covariance matrix of the fits reported in Ref. [61] for the measured data.

In short, in this analysis, the correction factor is evaluated event-by-event as a function of zenith angle and azimuthal orientation, using the calibrated model parameters. For each radial bin, the sensitivity of the mean muon density to each model parameter is computed, and the corresponding uncertainties and covariances are combined to obtain the systematic uncertainty on the bin-averaged density. This produces a distance-dependent CC uncertainty σ_{CC}.

Long-term stability systematic (LTS)

The detector response may vary over time due to gradual calibration drift, environmental conditions, or hardware changes. To quantify these effects, the data set is divided into distinct time periods, and the muon density profile is independently computed for each period. Differences between the period-wise profiles reflect time-dependent variations in the detector response. The spread of these profiles relative to the profile obtained from the entire dataset defines the long-term stability systematics σ_{LTS}. This uncertainty represents the potential bias introduced by temporal variations in detector performance.

Total systematic contributions

For each radial bin, the individual systematic contributions are combined in quadrature to obtain asymmetric total systematic uncertainties $\sigma^{\uparrow}_{\rm tot}(r)$ and $\sigma^{\downarrow}_{\rm tot}(r)$,

$$[\sigma^{\uparrow}_{\rm tot}(r)]^2 = [\sigma^{\uparrow}_{\rm ES}(r)]^2 + [\sigma^{\uparrow}_{\rm MP}(r)]^2 + [\sigma_{\rm CC}(r)]^2 + [\sigma_{\rm LTS}(r)]^2, \tag{4.16}$$

$$[\sigma^{\downarrow}_{\rm tot}(r)]^2 = [\sigma^{\downarrow}_{\rm ES}(r)]^2 + [\sigma^{\downarrow}_{\rm MP}(r)]^2 + [\sigma_{\rm CC}(r)]^2 + [\sigma_{\rm LTS}(r)]^2, \tag{4.17}$$

Figure 4.13 shows the systematic fractional uncertainties affecting the mean muon density $\rho_\mu(r)$ at $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$ in the zenith angle range $[0°, 45°]$. The uncertainties are displayed as functions of the radial distance, expressed in $\log_{10}(r/\mathrm{m})$, and are decomposed into individual contributions associated with different sources of systematic uncertainty. These include positive and negative energy-scale variations, long-term stability (LTS) effects, calibration-related contributions, and model-dependent components, along with the corresponding total systematic uncertainty obtained from their combination. The magnitude of the uncertainties varies with radial distance, reflecting the non-uniform sensitivity of the measurement to different systematic sources across the shower footprint. The total systematic uncertainty is dominated by the contribution from the energy-scale uncertainty and defines the overall uncertainty envelope used in the comparison between data and model predictions.

For each energy interval, individual systematic components are first evaluated in absolute units of the mean muon density profile. To reduce bin-to-bin fluctuations, each component is then smoothed as a function of $\log_{10}(r/\mathrm{m})$ by fitting $\log_{10}\sigma$ with a second-degree polynomial. The resulting fitted systematic uncertainties are then expressed as fractional uncertainties by normalising to the measured data profile $y_D(r)$, without applying any additional smoothing to the data. These fractional uncertainties are presented as functions of the radial distance. Similar behaviour is observed for all other energy intervals.

4.7.2 Scale-factor fit between measured data and simulations

The agreement between the measured and predicted mean radial distributions is quantified through a single multiplicative scale factor f, defined so that the model prediction becomes $f\,y_M(r)$,

$$y_D(r) = f\,y_M(r). \tag{4.18}$$

The best-fit value of f is obtained by minimising a function χ^2 that includes statistical uncertainties in the data and statistical uncertainties in MC. By restricting the fit to a single normalisation parameter, the method deliberately avoids absorbing shape differences into additional free parameters. The use of a χ^2 estimator is derived directly from the statistical model for the measurements. For each radial bin centred at r_i, the measured profile value is assumed to fluctuate with respect to the prediction of the scaled model according to Gaussian statistics. This assumption is justified by the central limit theorem, as both the measured

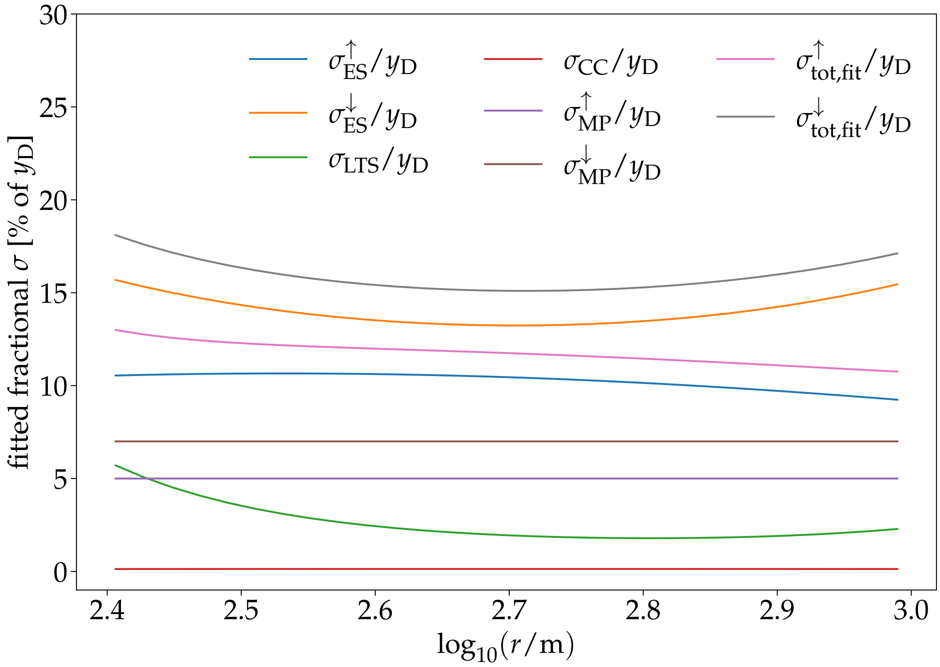

Figure 4.13: Fractional systematic uncertainties on the mean muon density $\rho_\mu(r)$ plotted as functions of radial distance in $\log_{10}(r/\mathrm{m})$ at $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$ in the zenith angle range $[0°, 45°]$. Individual contributions from energy-scale variations, long-term stability, calibration effects, and model-related uncertainties are shown, together with the corresponding total systematic uncertainty.

data profile $y_D(r)$ and the prediction of the simulated model $y_M(r)$ are obtained from the averages of many events within each bin, allowing their fluctuations to be well approximated by a normal distribution.

$$\epsilon_i = y_D(r_i) - f\, y_M(r_i), \qquad \epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \tag{4.19}$$

where $\sigma^2_{\mathrm{stat,i}} = \sigma^2_{\mathrm{stat}}[y_D(r_i)] + f^2\, \sigma^2_{\mathrm{stat}}[y_M(r_i)]$ is the variance of the residual between the data and the prediction. Under this assumption, the likelihood of observing the set of measurements $\{y_D(r_i)\}$ for a given scale factor f is

$$\mathcal{L}(f) = \prod_{i=1}^{N_{\mathrm{bins}}} \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left[-\frac{\epsilon_i^2}{2\sigma_i^2}\right]. \tag{4.20}$$

Maximising the likelihood in Eq. (4.20) is equivalent to minimising the negative logarithmic likelihood.

$$-2\ln\mathcal{L} = \sum_{i=1}^{N_{\mathrm{bins}}} \frac{\epsilon_i^2}{\sigma^2_{\mathrm{stat,i}}} + \ln(\sigma^2_{\mathrm{stat,i}}). \tag{4.21}$$

The logarithmic term depends only weakly on the fit parameter f compared to the quadratic residual term, and therefore contributes negligibly to the position of the minimum. Neglecting this term, one obtains, up to an additive constant independent of f, the χ^2 function,

$$\chi^2(f) = \sum_{i=1}^{N_{\mathrm{bins}}} \frac{\epsilon_i^2}{\sigma^2_{\mathrm{stat,i}}}. \tag{4.22}$$

where the statistical uncertainties between data and simulations are assumed to be independent. Thus, under the assumption of Gaussian-distributed bin uncertainties, minimising χ^2

is equivalent to estimating the maximum likelihood of the normalisation parameter f. The best-fit scale factor is $f = \arg\min \chi^2$, where the fit is performed using statistical uncertainties only (data statistical uncertainty and MC statistical uncertainty).

To probe the radial dependence, the fit is performed independently in several distance intervals to test possible shape mismatches between the data and the simulation. The full-distance fit includes all radial bins allowed by the minimum usable distance $r_{\min}(E)$, as described in Section 4.4.1, and an upper bound given by r_{LTP}, as described in Section 4.4.2. This range provides maximum sensitivity to deviations in the overall shape of the MLD, and the resulting scale factor is denoted by f_{full}. Two complementary subranges are also considered: from $r_{\min}$ to 500 m and from 500 m to r_{LTP}. The corresponding scale factors are denoted by f_1 and f_2, respectively. In addition, a localised fit is performed around $\log_{10}(450) + 0.025$, corresponding to the reference distance of the UMD-750 detector, and the resulting scale factor is indicated by f_{450}. In this narrow region, the MLDs predicted by different models tend to intersect, reducing the sensitivity to slope variations. The comparison in this region therefore isolates normalisation differences more directly.

Figure 4.14 shows the comparison between the measured average muon density profile and the corresponding fit at $\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$, in the zenith-angle range $[0°, 45°]$ for EPOS-LHC high-energy hadronic interaction model, together with the residuals of the fit. The upper panel presents the radial dependence of the mean muon density $\langle \rho_\mu(r) \rangle$ as a function of $\log_{10}(r/\mathrm{m})$, comparing the data points with the AugerMix reference distribution after fitting different definitions of the scale factors. The grey band represents the systematic uncertainty of the detector. The multiple fitted curves correspond to scale factors obtained from different radial fit ranges or definitions, including f_{full}, f_1, f_2, and f_{450}. The fitted scale factors differ slightly depending on the chosen range, with values clustering around $f \approx 1.2$, indicating a consistent overall rescaling between definitions.

The residuals plotted in the bottom panel of Fig. 4.14 assess the radial agreement between the measured and predicted muon density profiles beyond the log-scaled visualisation MLD. The residual is defined as the fractional deviation of the data from the fitted model prediction,

$$R(r) = \frac{y_D(r) - f\, y_M(r)}{f\, y_M(r)}. \tag{4.23}$$

Statistical uncertainties are displayed as error bars on $R(r)$, calculated by propagating the statistical uncertainty on the measured mean. In addition, the total asymmetric systematic uncertainty in the measurement is propagated in the residual space and shown as a shaded band, corresponding to $\pm\sigma_{\mathrm{tot}}^{\uparrow\downarrow}(r)/(f\, y_M(r))$. Horizontal reference lines at $\pm 2\sigma$ (with σ_{stat} taken as a representative median statistical uncertainty across bins) provide a visual guide to the typical significance of deviations: residuals extending beyond these levels indicate tension that is unlikely to be explained by statistical fluctuations alone, while agreement within the systematic band indicates consistency once systematic effects are accounted for.

Most residuals remain within the statistical confidence interval and lie within the systematic band, indicating good agreement between the scaled model and the measurements throughout the full radial range. No strong radial trend is observed in the residuals, suggesting that the global scale factor provides an adequate description of the data within uncertainties. Similar plots were generated for different energy ranges and zenith ranges, as well as different high-energy hadronic interaction models, and a similar trend was observed.

Additional plots for other energy bins and high-energy hadronic interaction models are shown in Appendix A.

The physical interpretation of the fitted scale factor depends on the distance range over which it is determined. At approximately the reference distance of 450 m, where the sensitivity to slope variations is minimised, a single scale factor provides an accurate

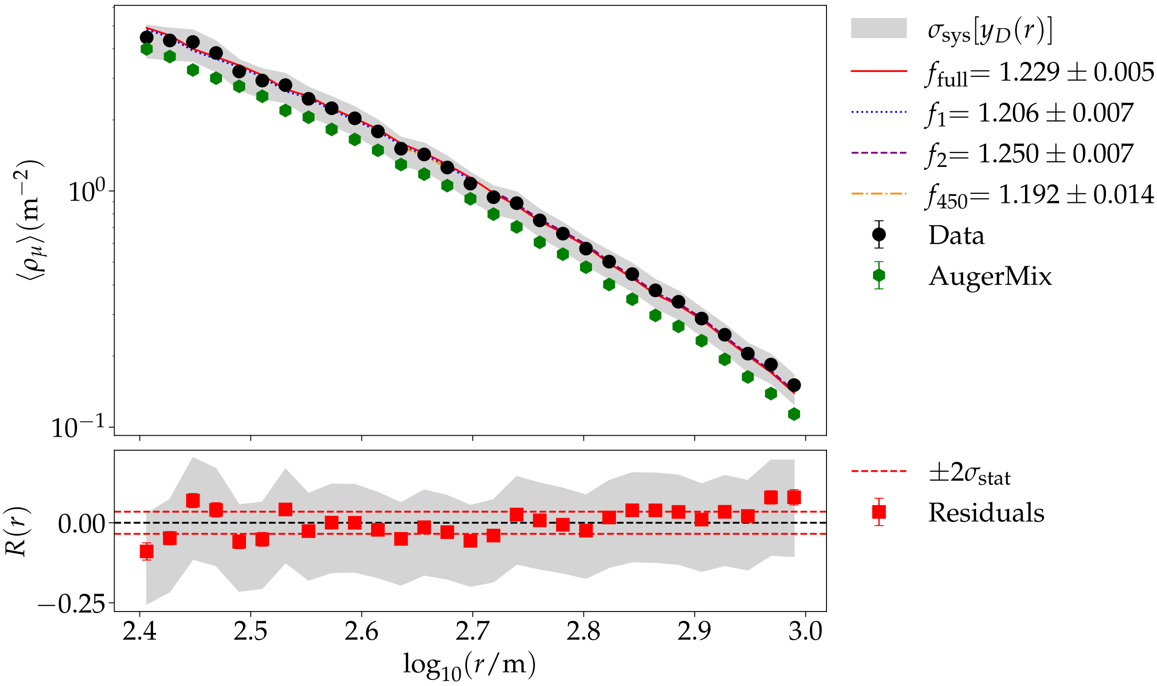

Figure 4.14: (Top) Muon density profile as a function of radial distance compared with the scaled AugerMix reference model. Different curves correspond to scale factors derived from various fit windows, including the full radial range and restricted intervals. The multiple fitted curves correspond to scale factors obtained from different radial fit ranges ($f_{\rm full}$, f_1, f_2, and f_{450}). (Bottom) Residuals between data and the scaled model (see Eq.(4.23)) are shown as a function of radial distance. The red dashed lines correspond to the $\pm 2\sigma$ statistical confidence limits. In both figures, the grey band indicates the detector systematic uncertainty. Both figures corresponds to $\log_{10}(E/\rm eV) \in [17.9, 18.1)$ in the zenith angle range $[0°, 45°]$ for EPOS-LHC high-energy hadronic interaction model.

description of the relative normalisation between data and the simulation. In this regime, the fitted value of f can be interpreted as a local normalisation difference that is largely independent of residual shape mismatches.

In contrast, when the full radial range is included, a single scale factor must simultaneously reconcile differences across multiple regions of the MLD. If the data profile is systematically steeper or flatter than the simulation, the inner and outer radial bins will favour different normalisation adjustments. The fitted scale factor, therefore, becomes an effective compromise that partially absorbs the average slope mismatch. In this case, f no longer represents a purely local normalisation but instead reflects a global adjustment that depends on the detailed radial structure of the discrepancy.

Different distance selections, therefore, probe different aspects of the data–simulation comparison. The reference-distance fit isolates normalisation differences with minimal slope sensitivity, while the full-distance fit tests the global validity of the assumption that the data and simulation differ only by a scale factor. The value obtained near the reference distance is generally the most physically interpretable, whereas the full-range result provides a more stringent consistency check.

The upper panel of Fig. 4.15 compares four definitions of the scale factor f obtained with the EPOS-LHC high-energy hadronic interaction model for showers in the zenith range $[0°, 45°]$, as a function of primary energy. All four estimators, $f_{\rm full}$, f_1, f_2, and f_{450}, exhibit an overall increasing trend with $\log_{10}(E/\rm eV)$, indicating a progressively larger rescaling of the simulated signal relative to observations at higher energies. Although the general energy dependence is consistent across definitions, systematic offsets are visible between estimators at fixed energy. In particular, f_{450} tends to yield the largest values at high energy, whereas f_2 and f_1 remain slightly lower, and f_{full} generally lies between them depending on energy. All four estimators agree within systematic uncertainties. The error bars reflect statistical uncertainties, which increase toward the highest energies due to reduced event statistics. The agreement in the overall trend suggests robustness of the inferred scaling behaviour, while the spread between estimators quantifies the sensitivity of the result to the fitting definition or radial range used.

Only events with reconstructed energies $E \geq 10^{17.5}$ eV are considered suitable for analysis for the SD–750 array. This requirement ensures that the SD and UMD in the array operate with full trigger efficiency and avoids potential composition-dependent biases that may arise at lower energies [61].

The lower panel of Fig. 4.15 shows the energy dependence of the scale factor f_{full} obtained with EPOS-LHC high-energy hadronic interaction model for three different zenith angle intervals, $\theta \in [0°, 30°]$, $\theta \in [30°, 45°]$ and $\theta \in [0°, 45°]$, together with results from previous work in Ref. [61] for comparison.

The muon-rescaling factor f introduced in Ref. [61] quantifies the mismatch between the muon density measured with the UMD and the muon density expected from the mass composition inferred via $\langle X_{\text{max}} \rangle$. For each hadronic interaction model and fixed energy, simulations of different primary masses define an approximately linear relation between $\langle X_{\text{max}} \rangle$ and $\langle \ln \rho_{35} \rangle$. The observable ρ_{35} is defined as the muon density at a distance of 450 m from the shower core, reconstructed with the SD–750 array and corrected for the attenuation of the zenith angle to a reference angle of 35° using the CIC method (see Section3.1). This correction allows for a direct comparison of events recorded at different zenith angles. The measured $\langle X_{\text{max}} \rangle$ is then used to infer the expected $\langle \ln \rho_{35} \rangle_{\text{sim}}$ and the discrepancy with the measured $\langle \ln \rho_{35} \rangle_{\text{data}}$ is parametrised as

$$f = \exp(\langle \ln \rho_{35} \rangle_{\text{data}} - \langle \ln \rho_{35} \rangle_{\text{sim}}.$$

Thus, $f = 1$ corresponds to consistency between the muon- and X_{max}-based composition inferences, while $f > 1$ indicates that the hadronic model underpredicts the muon content.

All zenith selections exhibit a general increase of f with primary energy, indicating a progressively larger rescaling of the simulated signals relative to observations at higher energies. The differences between the zenith ranges are evident. The $[30°, 45°]$ interval tends to yield larger scale factors at high energy, while the $[0°, 45°]$ combined sample lies between the narrower bins. The previous measurements in Ref. [61] are systematically higher than the current estimates at lower energies. Uncertainties grow toward the highest energies due to reduced event statistics, especially for the restricted zenith bins.

The increase of f with energy suggests that the discrepancy between the measured and simulated signals becomes progressively larger at higher primary energies. This behaviour may reflect limitations in the modelling of extensive air showers, including the description of hadronic interactions and the predicted muon content, which are known to become increasingly relevant at ultra-high energies. The dependence on zenith angle is comparatively modest, indicating that the inferred rescaling is largely robust against the choice of angular range, although the systematic differences between bins may point to residual angular effects or detector-related systematics that deserve further investigation. The offsets with respect to Ref. [61] likely arise from differences in the reconstruction, event selection, calibration, or analysis methodology adopted in the present work. A detailed comparison would be required to identify the dominant source of these differences. Similar energy-dependent behaviour is observed for the other high-energy hadronic interaction models and for the

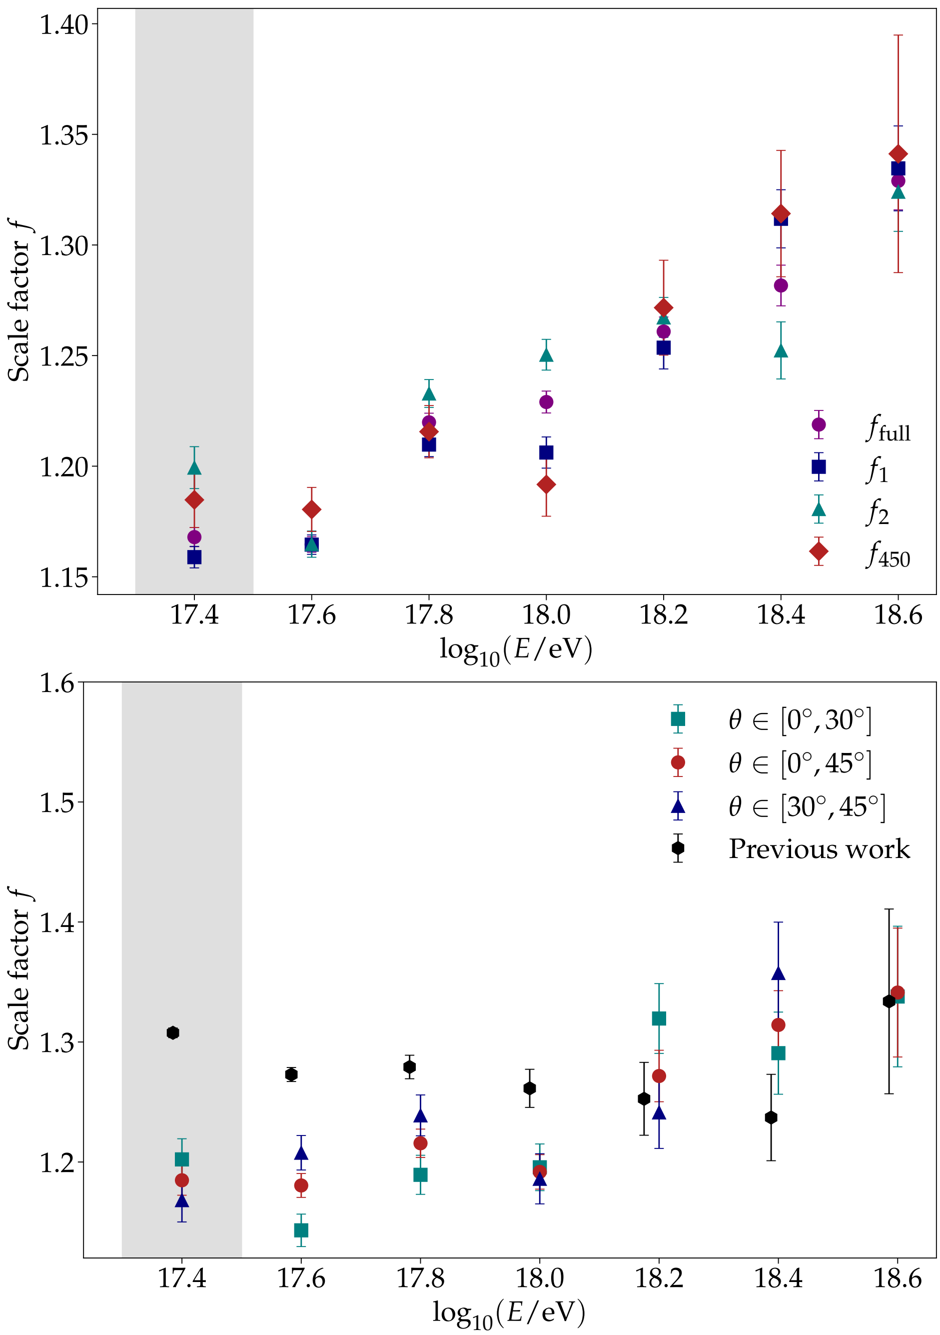

Figure 4.15: (Upper) Scale factor f as a function of primary energy for EPOS-LHC simulations in the zenith range $[0°, 45°]$. Four estimators are shown: f_{full}, f_1, f_2, and f_{450}. (Lower) Scale factor f_{full} as a function of primary energy for EPOS-LHC simulations in different zenith angle ranges. Results are shown for $\theta \in [0°, 30°]$, $\theta \in [30°, 45°]$, and $\theta \in [0°, 45°]$, together with values from previous work in Ref [61]. Error bars represent statistical uncertainties in both figures.

alternative estimators (f_1, f_2, and f_{450}), suggesting that the trend is not specific to a single observable.

Scale factor f for different high-energy hadronic interaction models

Figure 4.16 shows the energy dependence of the full-window scale factor f_{full} obtained using three different high-energy hadronic interaction models: EPOS-LHC, Sibyll-2.3c, and

QGSJet-II.04 in the zenith angle interval $\theta \in [0°, 45°]$. All models exhibit a clear increase in f_{full} with the primary energy, indicating that the rescaling required to match the simulated and observed signals increases at higher energies. However, a systematic separation between models is evident across the entire energy range. QGSJet-II.04 consistently yields the largest scale factors, while Sibyll-2.3c produces the lowest values, with EPOS-LHC lying between them and close to Sibyll-2.3c. The relative ordering is stable with energy, suggesting that model-dependent differences in predicted shower development or muon content persist across the full energy range considered. The uncertainties remain modest overall but grow slightly at the highest energies, reflecting reduced statistics. These results highlight the strong dependence of the inferred scale factor on the choice of high-energy hadronic interaction model.

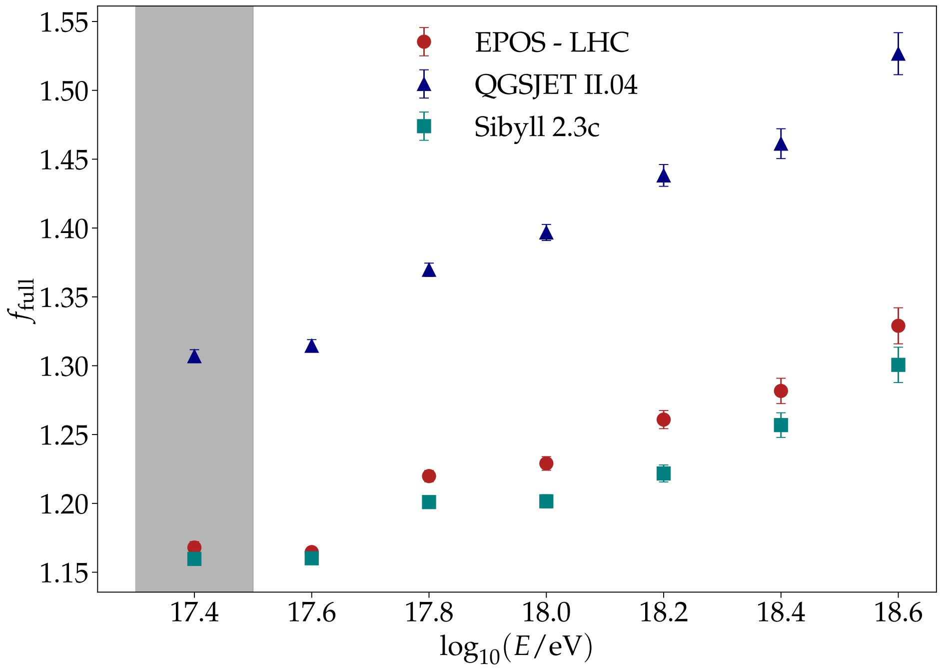

Figure 4.16: Full-window scale factor f_{full} as a function of primary energy for three hadronic interaction models. Results are shown for EPOS-LHC, Sibyll-2.3c, and QGSJet-II.04 in the zenith angle interval $\theta \in [0°, 45°]$. Error bars represent statistical uncertainties.

4.7.3 Systematic bracketing of the scale factor

To propagate systematic effects into the inferred scale factor, systematic bracketing fits are performed. The measured profile is shifted upward and downward by the total systematic envelope, producing modified profiles

$$y_D^+(r) = y_D(r) + \sigma_{\text{tot}}^{\uparrow}(r), \qquad y_D^-(r) = \max\left[y_D(r) - \sigma_{\text{tot}}^{\downarrow}(r), 0\right]. \tag{4.24}$$

The scale factor is then refitted using these shifted profiles. The resulting parameters $(f_{\text{full}}^+, f_{\text{full}}^-)$, (f_1^+, f_1^-), $(f_,^+ f_2^-)$, and (f_{450}^+, f_{450}^-) corresponding to the different distance ranges define the systematic range of the scale factor consistent with the total uncertainty envelope of the data.

Figure 4.17 compares the energy dependence of the scale factor f_{450} in the zenith angle interval $\theta \in [0°, 45°]$ obtained in this work with values reported in previous work in Ref. [61]. The central markers show the best-fit scale factors as a function of primary energy, while the bracketed symbols of matching colour indicate the corresponding systematic uncertainty

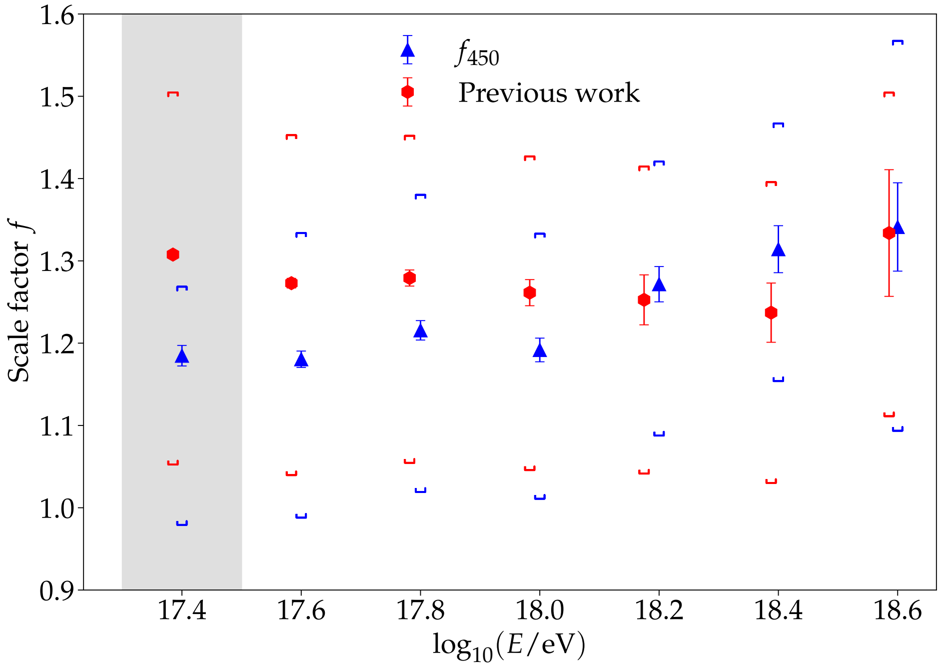

Figure 4.17: Comparison of the scale factor f_{450} with previous results reported in Ref. [61] as a function of primary energy in the zenith angle interval $\theta \in [0°, 45°]$. Central markers show the best-fit values for this work and for previous work. Brackets of matching colour indicate systematic uncertainty ranges: red brackets correspond to the systematic interval reported in the previous analysis, while blue brackets denote the f_{450}^{+} and f_{450}^{-} variations derived in this work.

intervals. For the analysis in this chapter, the blue brackets represent the variations f_{450}^{+} and f_{450}^{-} derived from the systematic uncertainty evaluation, while the red brackets show the systematic uncertainty band reported in Ref. [61]. Overall, while the central scale factors remain comparable, both estimators are compatible within the total uncertainties.

4.8 Summary and conclusions

In this chapter, the average muon lateral distribution measured with the UMD detector was characterised in bins of reconstructed energy and zenith angle and compared to detector-level expectations from air-shower simulations. The comparison was performed directly at the level of profiled muon densities in radial bins using mixed-composition (AugerMix) reference samples constructed from externally constrained mass fractions. This bin-by-bin strategy avoids imposing an explicit MLD parametrisation and, therefore, preserves sensitivity to genuine shape differences between data and simulation. After validating the compatibility of the Phase 1 and Phase 2 datasets, the combined data sample was used throughout, improving statistical precision without introducing evidence for time-dependent detector effects.

A set of selection criteria was introduced to ensure that the reconstructed MLD is built from geometrically and instrumentally stable measurements. At small core distances, a minimum-distance cut was motivated by the sensitivity of the steep inner MLD to SD core-position uncertainties in fixed-core reconstruction. At large distances, an energy- and zenith-dependent maximum distance was defined using the SD LTP threshold to suppress trigger-induced selection bias in the low-density regime. In addition, an outlier rejection procedure based on robust quantiles was applied per radial and energy bin, in the complete zenith range considered, removing only a negligible fraction of measurements and ensuring that the profiled densities are not distorted by rare detector malfunction. Together with the treatment of pinocchios through unbiased estimators, these steps produce stable radial profiles and statistically significant uncertainties for subsequent comparisons.

The core compatibility test was formulated as a single-parameter scale fit between the measured and predicted profiles, $y_D(r) = f\, y_M(r)$, with f obtained by minimising a χ^2 function that includes both the data and MC statistical uncertainties. By construction, this fit cannot absorb radial-shape mismatches into additional degrees of freedom. Consequently, the fitted scale factor quantifies the dominant multiplicative deficit in the simulated muon normalisation. Across energy, inferred scale factors show a systematic increase with $\log_{10}(E/\mathrm{eV})$, indicating that the relative normalisation offset between data and simulation increases with energy. This behaviour is robust across the different scale estimators considered (f_{full}, f_1, f_2, and f_{450}), although their absolute values differ, reflecting sensitivity to the radial interval and therefore to shape effects. The analysis further shows that, despite differences in the fit window, all fitted values of f are compatible within the total uncertainties once systematic effects are included.

Systematic uncertainties were evaluated at the level of the mean muon density in each radial bin, including asymmetric contributions from the energy scale and muon-pattern normalisation, together with distance-dependent corner-clipping and long-term stability uncertainties. The resulting systematic uncertainties are generally larger and more asymmetric than those reported in earlier work in Ref. [61], reflecting both the AugerMix composition adopted here and the explicit propagation of distance-dependent systematic effects into the profile-level comparison. The comparison with previous results shows compatibility within the total uncertainties.

Finally, the dependence on the hadronic interaction model was assessed using EPOS-LHC, Sibyll-2.3c, and QGSJet-II.04. Although all models require a scale factor $f > 1$ and show an increase with energy, stable model ordering is observed, indicating that model-dependent differences in predicted muon content persist throughout the entire energy range considered. Together, the results support the existence of a global muon deficit in simulations relative to the data. However, its precise quantification is conditional on the radial sensitivity of the estimator.

CHAPTER 5

Mass sensitive observables

As discussed in Chapter 3, the detectors at Pierre Auger Observatory observe EAS using complementary observational techniques. The FD provides the most direct and robust measurement of the composition of the primary mass through the longitudinal development of the electromagnetic component. Differences in primary mass modify the depth of the first interaction and the subsequent evolution of the shower, which are reflected in the average depth of shower maximum, X_{max}, and in fluctuations of its distribution, $\mathrm{Var}[X_{\mathrm{max}}]$. Measurements indicate a transition towards a lighter composition between 10^{17} eV and $10^{18.3}$ eV, followed by a gradual increase in average mass at higher energies [78].

Despite its reliability, the fluorescence technique operates with a limited duty cycle. SD observables therefore provide an essential complementary approach, offering nearly continuous operation and high statistics. In particular, the muonic component of air showers is sensitive to primary mass, with heavier nuclei yielding more muons as a consequence of the superposition principle. However, the interpretation of muon measurements is complicated by significant uncertainties in high-energy hadronic interaction models. The long-standing muon deficit problem detailed in Ref. [83] highlights discrepancies between simulations and data, indicating that model extrapolations to ultra-high energies remain incomplete.

Motivated by these considerations, this chapter investigates two other approaches to mass sensitivity. First, two new muon density–based observables, V_b and V_b^*, are introduced to capture both the magnitude and the radial structure of the muon distribution without assuming a specific analytical form of the MLDF. Second, the possibility of estimating X_{max} from SD risetime measurements using the $\langle\Delta\rangle$ asymmetry method is examined. The analysis uses simulations and data described in Section 4.1, with the corresponding energy calibration summarised in Section 5.1. The study is carried out within the EPOS-LHC high-energy hadronic interaction model. The properties and discrimination power of V_b and V_b^* are presented in Section 5.2, and a suitable observable is selected. The reconstruction of X_{max} using $\langle\Delta\rangle$ method is discussed in Section 5.3, and the chapter concludes in Section 5.4.

5.1 Energy calibration

In this analysis, both SD-only events and hybrid events (detected simultaneously by the SD and the FD) are used. The reconstruction of the primary CR energy differs for these two event classes, as explained in Chapter 3. Hybrid events benefit from the nearly calorimetric energy measurement provided by the FD, while SD events only require an indirect calibration that relates SD observables to the shower energy. Since the detector simulations used in this work do not incorporate a full FD response within the adopted simulation setup, an approximate procedure is required to estimate a reconstructed energy similar to FD for the simulated events.

5.1.1 Fluorescence detector

The operation of the FD is described in detail in Section 3.2. As discussed in Section 4.1, the simulation configuration adopted in this analysis does not propagate events through a full FD response and reconstruction chain. Consequently, the energy reconstructed by the FD must be approximated using the MC energy.

The FD provides a nearly calorimetric measurement of the shower energy by observing the fluorescence light emitted when charged particles deposit energy in the atmosphere, primarily through ionisation losses. Since the fluorescence yield is proportional to the energy deposited, the integral of the longitudinal shower profile measured by the FD corresponds to the calorimetric energy of the air shower,

$$E_{\rm cal} = \int \mathrm{d}X \frac{\mathrm{d}E}{\mathrm{d}X}. \tag{5.1}$$

A small fraction of the primary energy is carried away by particles that do not produce fluorescence light, such as neutrinos and high-energy muons. This component, referred to as the invisible energy ($E_{\rm inv}$), is taken into account through a correction factor applied to the calorimetric energy. The total energy resolution of the FD is estimated to be approximately 8%, approximately independent of energy [64].

Figure 5.1 shows the contributions to the FD energy resolution arising from the detector (including reconstruction), the atmosphere and the invisible energy correction.

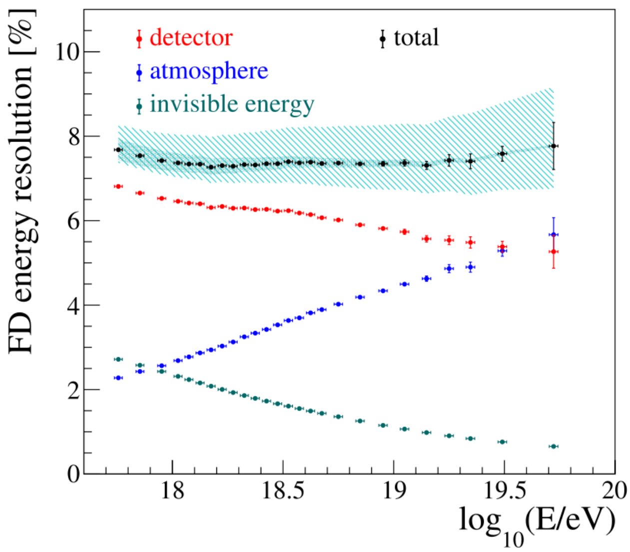

Figure 5.1: The FD energy resolution, showing the contributions from the detector (including reconstruction), the atmosphere, and the invisible energy correction. The shaded blue area represents the total systematic uncertainty on the resolution, dominated by the detector/atmosphere contributions but including the effect of $E_{\rm inv}$ (darker shading). Figure taken from [64].

To approximate the FD-reconstructed energy ($E_{\rm FD}$) for simulated events, a Gaussian smearing was applied to the MC energy ($E_{\rm MC}$), with a mean $\mu = 0$ and standard deviation

$\sigma = 0.08$. The value $\sigma = 0.08$ reflects an energy resolution of Pierre Auger Observatory FD of approximately 8%.

5.1.2 Surface detector

The energy is calibrated by assuming a power-law relationship between the reconstructed energy and the MC energy, following Section 4.2.

5.2 Mass sensitive observables from underground muon detectors

The parameter V_b combines the muon density and its spatial distribution without requiring an explicit analytical parametrisation of the MLDF. The definition of V_b is inspired by the observable S_b defined in Ref. [65], given by the sum of the signals S_i measured with the WCD weighted by $(r_i/r_0)^b$, and by the parameter M_b defined in Ref. [66], which corresponds to the logarithm of the sum of the weighted muon densities measured with the UMD, using the same radial weighting $(r_i/r_0)^b$. The M_b parameter was originally introduced in the context of photon searches.

For a given air-shower event with N triggered counters, V_b is defined as the sum of the muon density ρ_i for each counter i, weighted by the distance r_i between the counter and the shower axis,

$$V_b = \sum_{i=1}^{N} \rho_i \times \left(\frac{r_i}{r_0}\right)^b \text{ in m}^{-2}. \tag{5.2}$$

The muon density ρ is defined for N_μ muons impinging on a detector with an area A at a zenith angle θ as $\rho = N_\mu/(A\cos\theta)$. For the SD–750 array of the UMD at Pierre Auger Observatory, the reference distance is $r_0 = 450$ m, which corresponds to the optimal distance for an array with a detector spacing of 750 m. This choice ensures that V_b is evaluated in a region where the impact of shower-to-shower fluctuations and detector effects is minimised. The exponent b is a free parameter chosen to maximise the separation between primary particles. In the limiting case, $b = 0$, call counters contribute equally, and the observable reduces to a purely density-based quantity, independent of the spatial distribution of muons. For increasing values of b, a greater weight is assigned to detectors located at larger core distances, thereby incorporating information on the shape of the muon lateral distribution.

In addition to the observable V_b, a closely related variant is introduced, denoted V_b^*, to disentangle the impact of energy scaling on discrimination power. V_b^* is defined as the sum of the weighted muon densities, multiplied by an explicit energy-dependent normalisation factor,

$$V_b^* = \sum_{i=1}^{N} \left(\frac{10^{17}}{E}\right) \rho_i \times \left(\frac{r_i}{r_0}\right)^b \text{ in m}^{-2}, \tag{5.3}$$

Each event is rescaled by the ratio $10^{17}/E$, where E is the reconstructed energy of the shower, such that the observable is approximately normalised to a fixed reference energy. For this analysis, the energy measured by the FD as described in section 5.1.1 is considered. This construction reduces the explicit dependence of the observable on the primary energy and allows showers of different energies within the same bin to be compared on a common scale. By construction, V_b^* emphasises differences in muon content and topology that are not trivially driven by the overall energy scaling of the shower.

Although both observables V_b and V_b^* share the same underlying weighting scheme controlled by the exponent b, they differ in how the effects of energy are treated. The observable V_b retains both the intrinsic energy dependence of the shower and the radial structure of the muon distribution, providing a formulation referencing a fixed core distance that enables direct comparison with conventional observables evaluated at specific lateral distances. In contrast, V_b^* suppresses the dominant energy scaling through explicit normalisation, thereby emphasising differences in the muon distribution topology.

5.2.1 Discrimination power of observables

Taking into account the protons and iron primaries, the discrimination power of the parameter V_b is quantified by evaluating the merit factor as a function of the exponent b. The merit factor is a standard figure of merit in classification problems and provides a measure of the separation between the bulk of two distributions in units of their combined spread. In this context, b determines the relative weight assigned to each detector as a function of its distance to the shower axis and is selected to maximise the discrimination between proton- and iron-induced air showers. The overlap between the V_b distributions of the proton and iron primaries is characterised by the merit factor, defined as the absolute difference between the mean values of the two distributions divided by the square root of the sum of their variances. Here, the mean and variance are evaluated separately for proton and iron events. The merit factor (MF) is defined as

$$MF = \frac{\mid E[V_{b,\text{fe}}] - E[V_{b,\text{pr}}] \mid}{\sqrt{Var[V_{\text{fe}}] + Var[V_{b,\text{pr}}]}}, \tag{5.4}$$

where E[$V_{b,A}$] and Var[$V_{b,A}$] are the mean and variance, respectively, of the distribution function of the parameter V_b where $A = \text{pr}, \text{fe}$.

The merit factor thus provides a dimensionless and normalisation-independent measure of separation that is primarily sensitive to differences in the central values of the distributions relative to their intrinsic shower-to-shower fluctuations and reconstruction effects. Although it does not explicitly account for non-Gaussian tails, it is well-suited to assess the relative discrimination power of V_b in different choices of b and to compare the impact of different high-energy hadronic interaction models.

The merit factor was evaluated for exponent values b ranging from 0 to 8 in steps of 0.5, in different energy bins in the range $\log_{10}(E/\text{eV}) \in [17.3, 18.7]$. The muon density used for this evaluation was obtained by reconstructing the showers described in Section 4.1.2. The resulting merit factors as a function of b are shown in the left panel of Fig. 5.2 for the observable V_b and in the right panel of Fig. 5.3 for V_b^*.

In both cases, the merit factor exhibits a clear maximum whose location depends on energy. Although the absolute value of the merit factor varies across the energy range, the position of the maximum changes only moderately, indicating a weak but systematic energy dependence of the optimal exponent $b_{\text{max}}\big(\log_{10}(E/\text{eV})\big)$. This behaviour is shown in the right panels of Figs. 5.2 and 5.3 for the parameters V_b and V_b^*, respectively.

To describe this energy dependence, the values of $b_{\text{max}}\big(\log_{10}(E/\text{eV})\big)$ obtained in each energy bin were parametrised as a smooth function of energy. Several functional forms were tested. A cubic polynomial provides a purely descriptive fit, but lacks physical interpretability and exhibits strong parameter correlations. A sigmoid function was also considered, but the available energy range does not fully sample both asymptotic regimes, leading to poorly constrained fit parameters. For these reasons, a saturating exponential model was adopted. This form provides a stable fit while naturally describing a monotonic increase that approaches an asymptotic high-energy value. The optimal exponent is therefore parametrised as,

$$b_{\text{max}}\big(\log_{10}(E/\text{eV})\big) = b_\infty - Ae^{-k(\log_{10}(E)-\log_{10}(E_0))}, \tag{5.5}$$

where $\log_{10}(E_0) = 17.35$. The parameter b_∞ represents the asymptotic value of b_{max} at high energy, A sets the amplitude of energy evolution, and k controls the rate at which the function

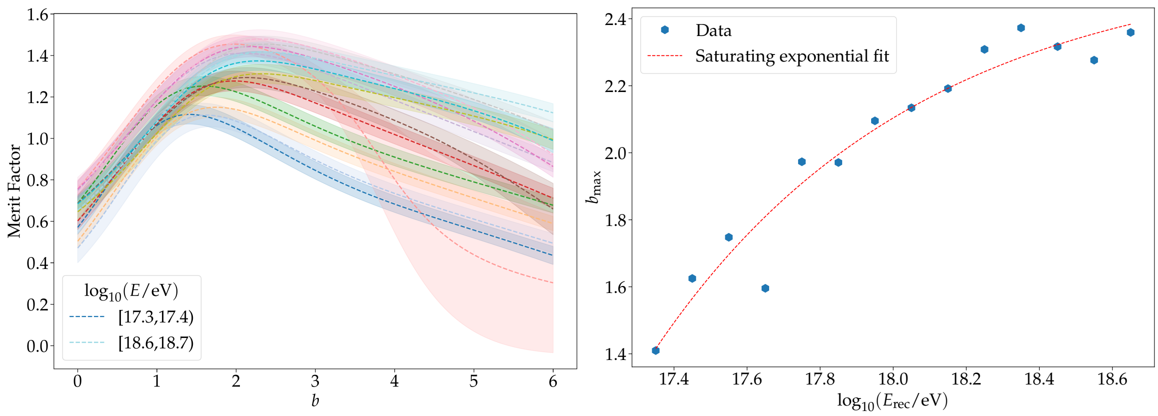

Figure 5.2: (Left) Merit factor as a function of the parameter b for the observable V_b, evaluated for different energy bins $\log_{10}(E/\mathrm{eV}) \in [17.3, 18.7]$. (Right) The optimal b values as a function of energy are fitted with a saturating exponential function.

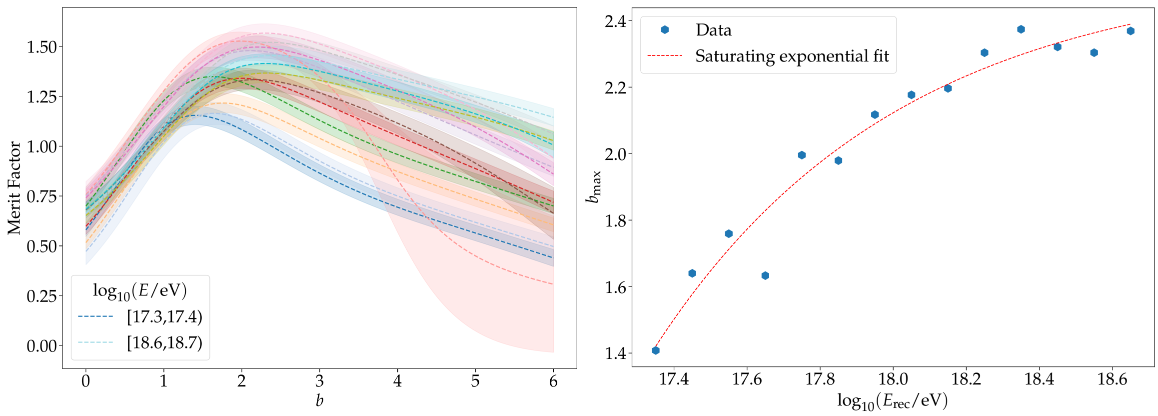

Figure 5.3: (Left) Merit factor as a function of the parameter b for the observable V_b^*, evaluated for different energy bins $\log_{10}(E/\mathrm{eV}) \in [17.3, 18.7]$. (Right) The optimal b values as a function of energy are fitted with a saturating exponential function.

approaches saturation. The coefficients of this parametrisation were obtained by fitting the values of b_{max} extracted in each energy bin and are summarised in Table 5.1. The resulting function reproduces the observed energy dependence in the explored range for both V_b and V_b^*.

Table 5.1: The coefficients of the saturating exponential function were obtained by fitting the values of $b_{\mathrm{max}}(\log_{10}(E/\mathrm{eV}))$ extracted in each energy bin.

Coefficient	V_b **Value** $\pm$ **Error**	V_b^* **Value** $\pm$ **Error**
b_∞	2.57 ± 0.18	2.55 ± 0.14
A	1.16 ± 0.16	1.13 ± 0.12
k	1.38 ± 0.46	1.5 ± 0.42

The asymptotic value $b_\infty \cong 2.5$ is comparable to the effective radial fall-off of the muon lateral distribution in the distance range relevant to the analysis. This indicates that the optimal weighting partially compensates for the radial decrease of the muon signal, enhancing sensitivity to composition across the sampled shower footprint. Nevertheless, b_{max}

should not be interpreted as a direct measurement of the LDF slope, since it also incorporates detector and reconstruction effects.

Although the observable V_b^* was introduced to reduce the explicit energy dependence of the muon signal by applying an event-by-event energy normalisation, its use does not lead to a significant improvement in the discrimination power between proton and iron primaries. As shown by the merit-factor analysis, both V_b and V_b^* exhibit similar trends as a function of the exponent b, with comparable maximum values and nearly identical optimal regions.

The energy dependence of the optimal exponent $b_{\max}\big(\log_{10}(E/\mathrm{eV})\big)$ is weak and can be adequately described by saturated exponential parametrisation in both cases. This indicates that the dominant contribution to the separation power arises from the radial distribution of muons rather than from the residual energy scaling within the considered energy bins. Consequently, the explicit energy scaling implemented in V_b^* does not provide additional discriminatory information beyond what is already captured by V_b. For these reasons and given the absence of a measurable performance gain, the remainder of this analysis is based exclusively on the observable V_b, which provides a simpler, more transparent, and equally effective measure of mass sensitivity. Using the coefficients in Table 5.1, an optimised version of the observable, denoted $V_{b_{\max}}$, is constructed by evaluating the observable with the energy-dependent exponent $b_{\max}\big(\log_{10}(E/\mathrm{eV})\big)$. This optimised observable is used in the following sections to assess the discrimination power between proton and iron primaries. To study the maximum degree of separation between the proton and iron distributions without energy bias, the same analysis was repeated for V_b with fixed energies, and the results are reported in Appendix B.1.

5.2.2 General features of optimised muon observable

The normalised distributions of $V_{b_{\max}}$ for the iron and proton primaries are shown in the left panel of Fig. 5.4 for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$. A clear separation between the two distributions is observed, with iron-induced showers exhibiting systematically larger values of $V_{b_{\max}}$ than proton-induced showers. This behaviour reflects the larger muon content and the flatter lateral distribution of muons expected for heavier primaries. The right panel shows the merit factor obtained using the optimised observable $V_{b_{\max}}$ as a function of the reconstructed energy. The merit factor increases with energy at low energies, evolving from values close to unity and reaching a plateau of $\sim 1.4 - 1.5$ above $\log_{10}(E/\mathrm{eV}) \gtrsim 17.8$. This trend indicates that the discrimination power between the proton and iron primaries improves with energy and saturates at the highest energies for the optimal choice of the exponent b.

The relatively smooth evolution of the merit factor with energy demonstrates that the saturating exponential parametrisation of $b_{\max}\big(\log_{10}(E/\mathrm{eV})\big)$ successfully captures the energy dependence of the optimal weighting without introducing artificial discontinuities in the discrimination performance. Residual fluctuations from bin to bin are compatible with statistical uncertainties and shower-to-shower fluctuations. This result confirms that the optimised observable $V_{b_{\max}}$ provides a stable and effective mass-sensitive parameter throughout the considered energy range, justifying its use in the subsequent analysis.

As shown in the left panel of Fig. 5.5, the mean value of the optimised observable $\log_{10}(\langle V_{b_{\max}}\rangle)$ exhibits an approximately linear increase with $\log_{10}(E/\mathrm{eV})$ for both the proton and iron primaries. This behaviour reflects the expected scaling of the muon content with the primary energy. It indicates that $V_{b_{\max}}$ preserves the intrinsic energy dependence of the shower while maintaining a clear separation between light and heavy primaries throughout the entire energy range considered. The right panel of Fig. 5.5 shows $\log_{10}(\langle V_{b_{\max}}\rangle)$ as a function of $\sin^2\theta$ for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$. Within statistical uncertainties, a

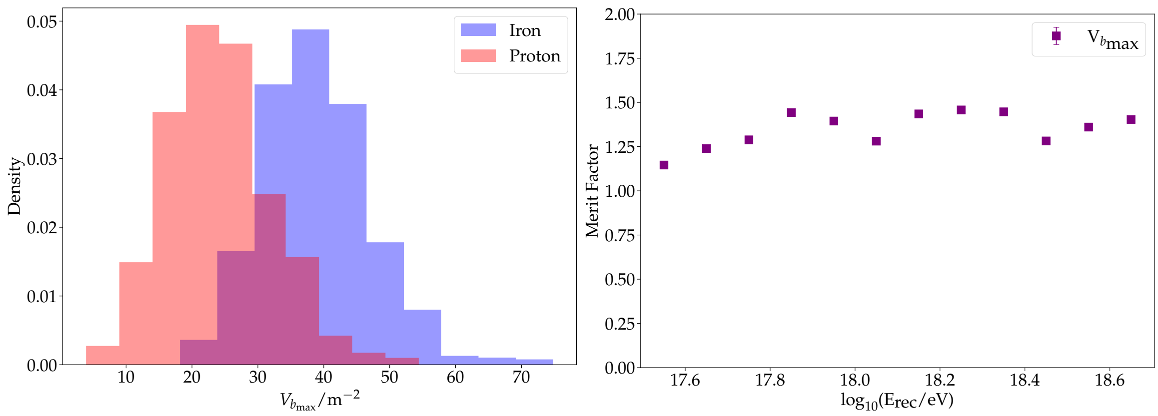

Figure 5.4: (Left)The normalised histogram distributions of $V_{b_{\max}}$ for iron and proton primaries for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$. (Right) Dependence of merit factor on different energy bins $\log_{10}(E/\mathrm{eV}) \in [17.3, 18.7]$.

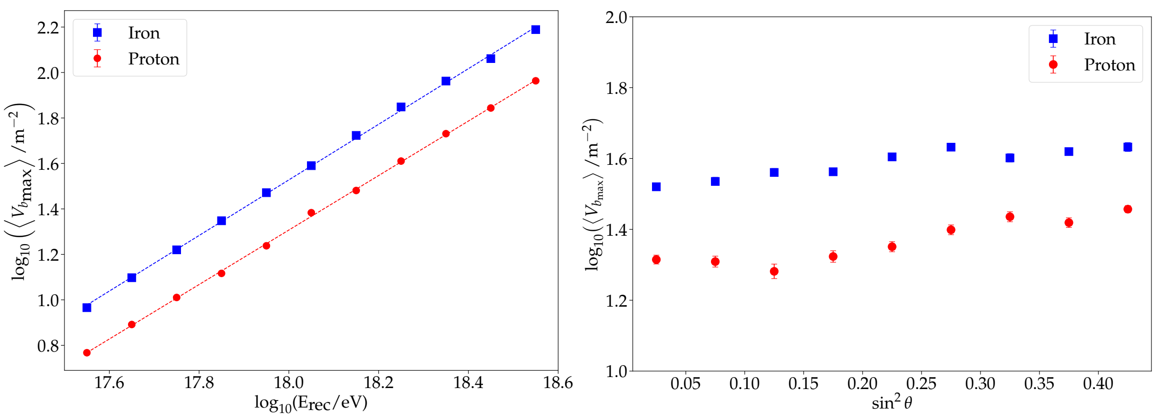

Figure 5.5: (Left) Variation of $\log_{10}(\langle V_{b_{\max}} \rangle)$ with $\log_{10}(E/\mathrm{eV})$ for proton and iron primaries. (Right) The behaviour of $\log_{10}(\langle V_{b_{\max}} \rangle)$ as a function of $\sin^2\theta$ is plotted for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$.

very weak dependence on the zenith angle is observed for both proton and iron primaries. This indicates that $V_{b_{\max}}$ does not have residual geometrical effects related to the slope of the shower. Similar behaviour is observed in the other energy bins, demonstrating that $V_{b_{\max}}$ is both energy-scalable and weakly dependent on the zenith angle. These properties make $V_{b_{\max}}$ a robustly mass-sensitive observable, suitable for composition studies without requiring additional corrections for shower geometry.

In Fig. 5.6, for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$, the merit factor defined in Eq. (5.4) is evaluated as a function of b varied from 0 to 8 in steps of 0.5 for three different high-energy hadronic interaction models (EPOS-LHC, QGSJet-II.04 and Sibyll-2.3c). The results indicate that all models exhibit a maximum merit factor around ≈ 2 for the selected energy bin and are largely independent of the hadronic interaction model. The peak values and their locations differ only marginally between models, particularly in the region of the largest merit factor where the proton–iron separation is strongest. For higher values of b, the merit factor decreases steadily for all models, with a more pronounced degradation observed for QGSJet-II.04, reflecting an increased sensitivity to shower-to-shower fluctuations when higher weights are assigned to distant stations. The agreement in the value of the maximum further suggests the robustness of the observable against both model uncertainties and moderate variations in the weighting scheme.

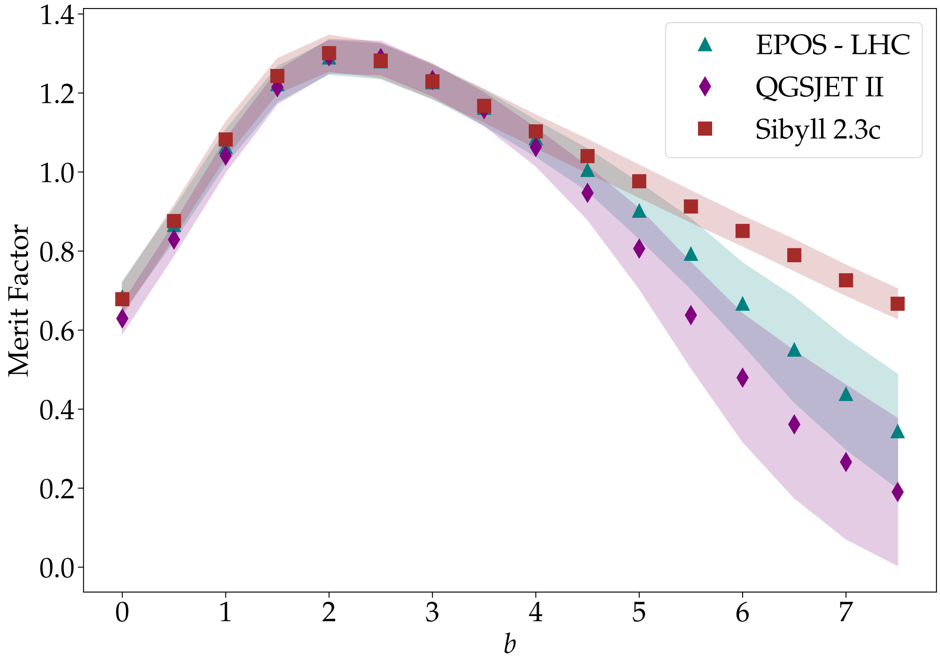

Figure 5.6: Merit factor as a function of the parameter b for the observable V_b, evaluated for different energy bins $\log_{10}(E/\mathrm{eV}) \in [17.3, 18.7]$ in steps of $\log_{10}(E/\mathrm{eV}) = 0.1$.

5.3 Depth of the shower maximum from surface detectors

The time profiles of the signals recorded by the WCD of Pierre Auger Observatory SD arise from a combination of the muonic and electromagnetic components of EAS. The temporal structure of these signals, characterised by the signal risetime, carries information about the longitudinal development of the shower. From the measured risetimes, the parameter $\langle\Delta\rangle$ can be constructed and used as an estimate of shower development.

This section describes the $\langle\Delta\rangle$ method introduced in Ref. [68], and applies the method in the SD–750 energy range. The observable $\langle\Delta\rangle$ is calibrated against X_{max} measured by the FD, allowing the derivation of an estimator based on SD $X^{\Delta}_{\mathrm{max}}$. The bias and resolution of the reconstructed $X^{\Delta}_{\mathrm{max}}$ are subsequently evaluated.

5.3.1 Risetime

The distribution of particle arrival times at the ground level encodes information on the longitudinal development of UHECR air showers. Particles produced at different stages of the cascade originate at different atmospheric depths and therefore traverse different path lengths before reaching the detector. As a result, particles generated deeper in the atmosphere arrive later than those produced higher up, even when propagating at comparable velocities. This temporal spread, as measured by the SD, reflects the longitudinal structure of the shower, as it maps the depth of particle production onto their arrival times. In particular, showers that develop deeper in the atmosphere exhibit a broader distribution of arrival times, consistent with a more extended production profile.

The temporal spread of particle arrivals increases with distance from the shower axis, providing additional information on the lateral development of the air shower. For a given primary energy, showers initiated by lighter primary particles reach maximum development at greater atmospheric depths, resulting in a wider spread of arrival times at ground detectors.

This relationship between primary particle composition, atmospheric depth of shower maximum and temporal spread is critical to understanding the mass composition of CR.

The arrival-time distribution of the particles is recorded in the FADC traces of the SD. However, mass composition analyses typically do not take advantage of the full temporal structure of the signal. Instead, different segments of the pulse carry distinct physical information. The early part of the signal is dominated by muons, which arrive in a relatively narrow time window and are closely related to the hadronic component of the shower. In contrast, the latter part of the pulse is dominated by the electromagnetic component, which is delayed due to scattering and attenuation in the atmosphere. Although the early signal is more sensitive to shower development and primary mass, it is also more difficult to reconstruct accurately due to its rapid onset, limited time resolution, low particle statistics, and sensitivity to detector response and trigger uncertainties. In contrast, the late signal, although more stable, carries a reduced sensitivity to the longitudinal development of the shower. Therefore, an optimal time window is selected to balance the reliability and sensitivity of the measurement with the underlying physics of the shower.

The risetime ($t_{1/2}$) is defined using the cumulative (integrated) signal recorded by the FADC trace. The integrated signal is constructed as the running sum of the trace amplitudes, and the total signal corresponds to its final value. The risetime is then defined as the time difference between the points at which the integrated signal reaches 10% to 50% of this total. This definition avoids the unstable leading edge of the signal and the late-time tail, providing a robust and composition-sensitive measure of the temporal structure of the shower front.

In this work, the risetime for each SD station is obtained directly from $\overline{\text{Off}}\underline{\text{line}}$, where a linear interpolation algorithm is applied between the edges of the bin to estimate the precise time required for the signal to rise from 10% to 50% of its final magnitude [69]. Given that the UB traces have a bin width of 25 ns, this interpolation enables a more accurate determination of the risetime. The analysis is done with Phase 1 data of the SD–750 array mentioned in Section 4.1.1.

Due to the spatial separation between detectors, even when positioned at the same axial distance from the shower core, the risetime of the signal varies depending on the sequence in which the detectors are struck during the shower's passage across the array. Detectors impacted earlier in this sequence exhibit a slower risetime compared to those struck later. This asymmetry arises from a combination of factors: the attenuation of the electromagnetic component and the variation in the angular distribution of muons sampled at different positions within the array. For small zenith angles, geometrical effects play a significant role, while for $\theta > 30^{\circ}$ the attenuation of the electromagnetic component is the dominant effect.

The azimuthal asymmetry in risetime refers to the variation in risetime as a function of the azimuthal angle on the shower plane (ζ). This angle is specific to each triggered SD in an event. This asymmetry arises from geometrical and shower-development effects that lead to systematic differences in the risetime measured by stations located at different azimuthal angles with respect to the shower axis, particularly between early and late regions of the shower front.

For the present analysis, the azimuthal asymmetry of the risetime is corrected outside the $\overline{\text{Off}}\underline{\text{line}}$ framework using the parametrisation developed for the SD–750 array in Ref. [68]. The corrected risetime is obtained from the measured risetime according to the relation,

$$t_{1/2}^{\text{corrected}} = t_{1/2}^{\text{measured}} - g(r,\theta)\cos\zeta \;\text{ in ns}, \tag{5.6}$$

where the function $g(r,\theta)$ is defined as $g(r,\theta) = m(\theta)r^2$ with $m = (a\sec\theta + b\sec^3\theta + c)\sqrt{\sec\theta - 1}$. The parameters a, b, c have the values taken from Ref. [68] and are specific to the SD–750 array,

$$a = 0.00029\,\text{ns m}^{-2}, \quad b = -0.00007\,\text{ns m}^{-2}, \quad c = -0.00013\,\text{ns m}^{-2}.$$

Figure 5.3.1 illustrates the effect of this correction on the risetime as a function of ζ. The asymmetry uncorrected risetime (red points) exhibits a pronounced azimuthal dependence. In contrast, the risetime values corrected for asymmetry (blue points) show a significantly reduced variation, indicating that the dominant asymmetry has been effectively removed. The asymmetry-corrected risetime is therefore used for the remainder of the analysis, ensuring that the observable is insensitive to the azimuthal position of the detector and can be reliably compared across events.

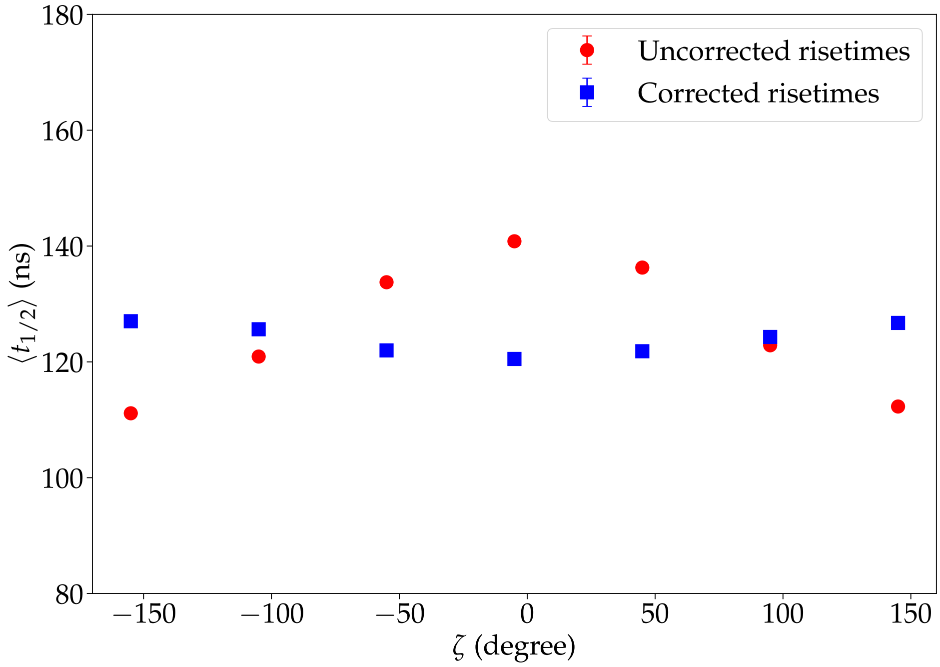

Figure 5.7: Effect of the asymmetry correction on data of SD–750 array for one zenith bin ($1.2 < \sec\theta < 1.3$) for all energies. The uncertainty on the mean is included in the figure but is negligibly small.

5.3.2 Quality cuts

A set of quality cuts is applied at both the event and station levels to ensure reliable reconstruction and to minimise systematic effects associated with detector response, geometry, and low-signal fluctuations. These selections restrict the analysis to a phase space where the risetime is well defined and exhibits stable behaviour, while preserving sufficient statistics for a meaningful study.

- Event level cuts:
  - $17.5 < \log_{10}(E/\mathrm{eV}) < 18.5$
  - $\sec\theta < 1.3$ (or $\theta \lesssim 40°$)
  - Requires a minimum of 3 detectors in the event.
- Station-level cuts:
  - Signal > 6 VEM

    Some negative risetime values remain after the asymmetry correction using Eq. (5.6) and must be removed to reduce trigger biases. Although Ref. [68] applies a signal threshold of > 3 VEM, this cut still leaves some negative risetime values in the present analysis. Therefore, a higher signal threshold of > 6 VEM is chosen for this study.

    Although a hard cut on the signal may introduce a selection bias by preferentially removing stations with low particle densities, this selection is necessary to ensure a reliable determination of the risetime observable. At low signal levels, timing information is strongly affected by fluctuations and detector effects, leading to unphysical or poorly reconstructed values. The chosen threshold therefore represents a compromise between minimising bias and ensuring measurement stability. The impact of this selection is further mitigated by the applied core-distance and energy cuts, which restrict the analysis to a region where the detector response is well behaved.

- Core-distance selection: 350 m $< r <$ 850 m
  To ensure a stable and physically meaningful definition of the risetime observable, a core-distance selection is applied. This range is determined by examining the behaviour of $t_{1/2}/r$ as a function of the distance from the shower axis. Figure 5.8 shows $\langle t_{1/2}/r \rangle$ as a function of r for two different energy bins of the Phase 1 data. Within the intermediate distance range, the observable exhibits an approximately linear behaviour, indicating a well-defined and stable scaling of the risetime with core distance. At smaller distances ($r \leq 350$ m), deviations from linearity are observed, likely due to saturation effects and the strong electromagnetic component near the shower core. At larger distances ($r \geq 850$ m), increased statistical fluctuations and reduced signal strength lead to larger uncertainties and departures from the linear trend. The selected interval 350 m $< r <$ 850 m therefore represents the region where the risetime observable behaves in a controlled and approximately linear manner throughout the studied energy range. This choice minimises systematic biases while preserving sufficient statistics for the analysis.

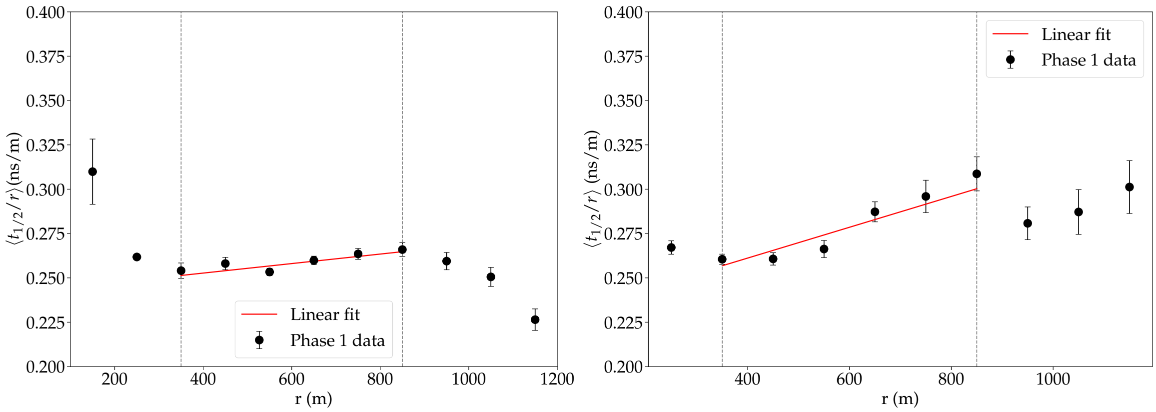

Figure 5.8: Risetime over core distance as a function of core distance for two energy bins, $\log_{10}(E/\mathrm{eV}) \in [17.8, 17.9)$ (left) and $\log_{10}(E/\mathrm{eV}) \in [18.4, 18.5)$ (right), is shown. The core distance region (between the dashed vertical lines) shows approximately linear behaviour and is hence chosen for the present analysis.

5.3.3 Risetime uncertainty

Uncertainty in risetime arises from several sources, including intrinsic shower-to-shower fluctuations, finite sampling of the shower front, electronic response, digitisation effects, and the reconstruction procedure. In this work, the parametrisation of uncertainty using twins and pairs of stations developed in Ref. [68] is adopted. The total uncertainty on the risetime is expressed as

$$\sigma_{1/2} = \sqrt{\left(\frac{J(r,\theta)}{\sqrt{S}}\right)^2 + \left(\sqrt{2}\frac{25}{\sqrt{12}}\right)^2} \text{ in ns,} \tag{5.7}$$

where S is the total signal in units of VEM. The first term accounts for statistical fluctuations that scale inversely with the square root of the signal, whereas the second term represents the contribution from the digitisation of the FADC traces, assuming a 25 ns sampling window. The function $J(r, \theta)$ encodes the dependence of the statistical component of the uncertainty on the distance from the core and the angle of zenith and is parametrised as $J(r, \theta) = p_0(\theta) + p_1(\theta)r$. The coefficients p_0 and p_1 are defined separately for $r \leq 650$ m and $r > 650$ m, reflecting the change in behaviour of the risetime fluctuations across these distance regimes as

$$p_0(\theta) = \begin{cases} (-340 \pm 30) + (186 \pm 21)\sec\theta \text{ if } r \leq 650\text{m}, \\ (-447 \pm 30) + (224 \pm 22)\sec\theta \text{ if } r > 650\text{m}, \end{cases}$$

$$p_1(\theta) = \begin{cases} (0.94 \pm 0.03) + (-0.44 \pm 0.01)\sec\theta \text{ if } r \leq 650\text{m} \\ (1.12 \pm 0.03) + (-0.51 \pm 0.02)\sec\theta \text{ if } r > 650\text{m} \end{cases}$$

where units of p_0 are (ns VEM$^{1/2}$) and units of p_1 are (ns VEM$^{1/2}$ m^{-1}).

5.3.4 Benchmark

The benchmark function characterises the average behaviour of the asymmetry-corrected risetime as a function of core distance and zenith angle for a fixed energy bin. Because the present analysis uses slightly different quality cuts from Ref. [68], the benchmark must be recalculated to ensure internal consistency. For each SD detector, two time traces are recorded on the LG and HG channels. In the case where no saturation occurs in either channel, the risetime is obtained from the trace corresponding to the HG channel. If this channel is saturated, the trace from the LG channel is used to compute the risetime. If the LG is saturated as well, which can occur for stations close to the core in high-energy events, that station is not selected for this analysis. The risetimes measured for a station on the two channels are not identical; hence, this analysis treats the risetimes obtained from different readout channels differently. This difference in the measured risetime is a result of the threshold imposed to remove very small signals during the digitisation process, which affects the LG traces much more compared to the HG traces Ref. [68]. To define the benchmark function, the data in the energy bin $17.7 < \log_{10}(E/\text{eV}) < 17.8$ are used, providing a balanced population of stations over the core-distance range covered by both HG and LG channels.

The HG risetime is modelled using the functional form,

$$t_{1/2}^{\text{HG}} = 40 + \sqrt{A_{HG}(\theta)^2 + B_{HG}(\theta)r^2} - A_{HG}(\theta) \text{ in ns}, \tag{5.8}$$

which provides a smooth interpolation between small and large core distances and reproduces the approximately linear growth of risetime at intermediate distances. The additive constant of 40 ns represents the intrinsic minimum risetime scale of the detector response.

Using the values $A_{HG}(\theta)$ and $B_{HG}(\theta)$ obtained from the above function, the risetime values measured by the LG channels are fitted with the function below.

$$t_{1/2}^{\text{LG}} = 40 + N(\theta)\left(\sqrt{A_{HG}(\theta)^2 + B_{HG}(\theta)r^2} - A_{HG}(\theta)\right) \text{ in ns}, \tag{5.9}$$

where the factor $N(\theta)$ accounts for the systematic shift between LG and HG risetimes. This multiplicative correction preserves the radial dependence of the HG benchmark while allowing channel-dependent scaling.

Figure 5.9 shows the risetime as a function of the core distance together with the corresponding fits for the benchmark energy bin and $\sec\theta \in [1, 1.05)$. The risetimes are modelled according to Eqs. (5.8) and (5.9) and are taken from Ref. [68]. The risetime uncertainties calculated from Section 5.3.3 are included in the fits but not shown in the figure. The same is

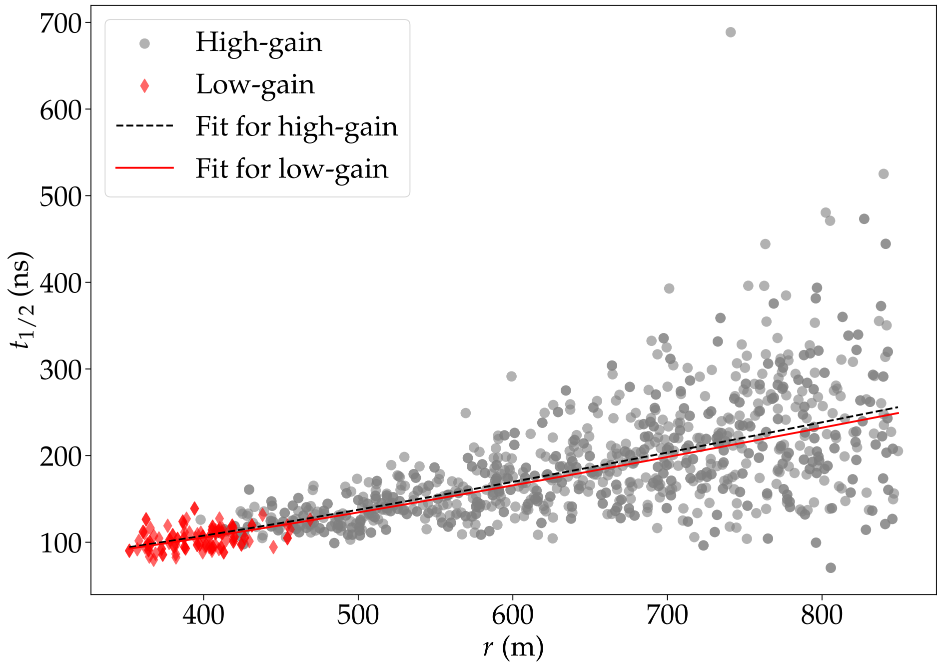

Figure 5.9: Risetimes as a function of core distance with their respective fits for data of the SD–750 array for the benchmark energy bin and $\sec\theta \in [1, 1.05)$. The risetimes are modelled according to Eqs. (5.8) and (5.9) and are taken from Ref. [68]. The risetime uncertainties calculated from Section 5.3.3 are included in the fits but not shown in the figure.

repeated for other $\sec\theta$ bins. The extracted zenith-angle dependencies of $A_{HG}(\theta), B_{HG}(\theta)$ and $N(\theta)$ are shown in Fig. 5.10 that have the functional form,

$$A_{sat}(\theta) = a_0 + a_1 \sec^{-4}\theta, \quad B_{sat}(\theta) = b_0 + b_1 \sec^{-4}\theta, \quad N(\theta) = n_0 + n_1 \sec^2\theta + n_2 \left(e^{\sec\theta}\right),$$

where the values of the parameters are given in Table 5.2 and the function form is obtained from Ref. [68].

Table 5.2: Values of the parameters with statistical uncertainties to define the benchmark.

Parameter	Value ± Error
a_0	-63.42 ± 21.41
a_1	233.72 ± 41.2
b_0	-0.037 ± 0.014
b_1	0.20 ± 0.028
n_0	8.74 ± 0.92
n_1	7.61 ± 0.95
n_2	-5.66 ± 0.69

These parametrisations define the benchmark function used in the subsequent analysis.

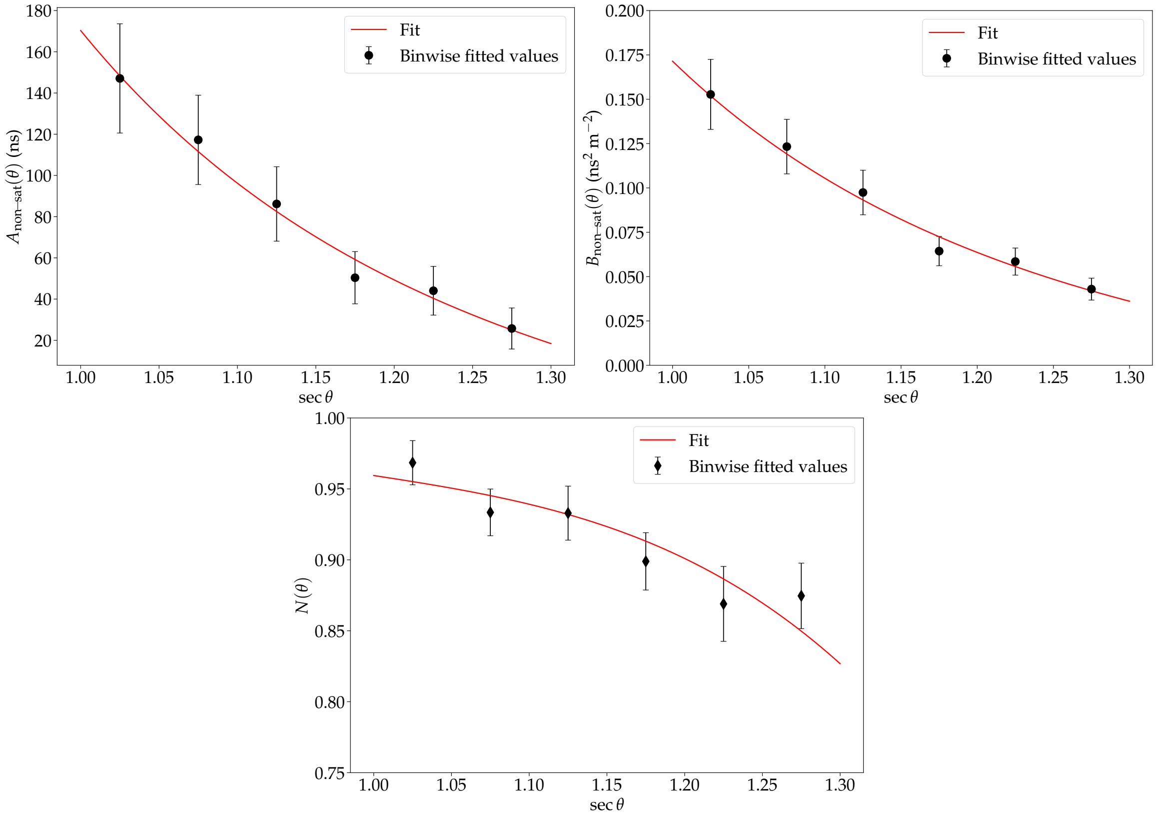

Figure 5.10: Functional form of parameters in the benchmark definition for the SD–750 array.

5.3.5 $\langle\Delta\rangle$ method

The parameter $\langle\Delta\rangle$ combines the risetime values of individual SD stations into a single observable that characterises the whole event. It is the average deviation from the benchmark of the risetimes of an event expressed in terms of the risetime uncertainty,

$$\Delta_i = \frac{t_{1/2,i} - t_{1/2}^{\text{benchmark}}(r_i, \theta_i)}{\sigma_{1/2,i}}, \tag{5.10}$$

where $t_{1/2,i}$ is the risetime measured by each detector, $t_{1/2}^{\text{benchmark}}(r_i, \theta_i)$ is the benchmark parametrisation evaluated at the core distance of each detector, and the zenith angle of the shower (i.e., the benchmark calculated from the data and explained in Section 5.3.4), and $\sigma_{1/2,i}$ is the risetime uncertainty. The dimensionless parameter $\langle\Delta\rangle$ is defined as

$$\langle\Delta\rangle = \frac{1}{N}\sum_{i=1}^{N}\Delta_i \tag{5.11}$$

If the benchmark accurately describes the average risetime behaviour in the reference energy bin, then the expectation value of the normalised deviation is zero by construction. Since the benchmark is derived directly from the data, it is expected that the distributions of Δ_i and $\langle\Delta\rangle$ for the data sample are centred on zero. This behaviour is observed in the benchmark energy bin, where the distribution of Δ_i for Phase 1 data is shown in the left panel of Fig. 5.11, and the distribution of $\langle\Delta\rangle$ is shown in the right panel of the same figure.

The corresponding distributions of Δ_i and $\langle\Delta\rangle$ obtained from the proton and iron simulations are shown in Fig. 5.12. In this case, the mean values exhibit a small residual bias with respect to zero. This is expected because the benchmark parameterisation was obtained from

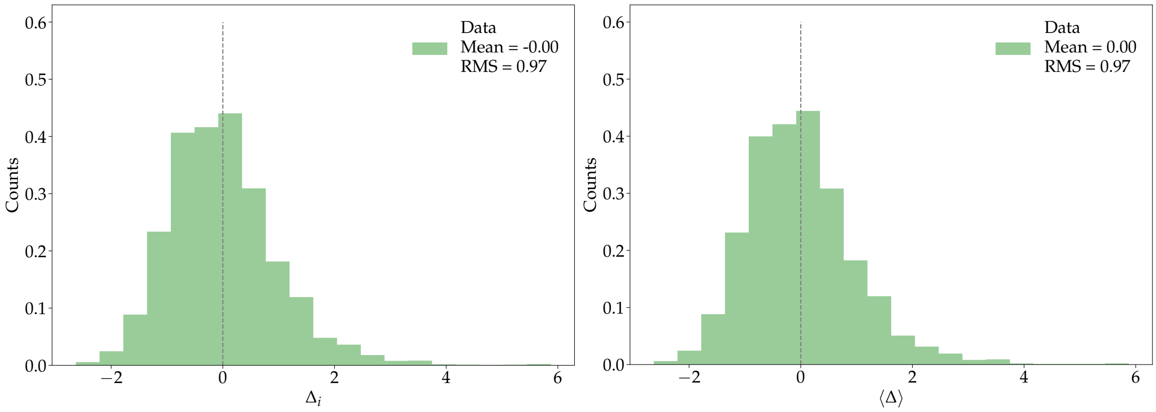

Figure 5.11: (Left) Δ_i distribution for the benchmark energy bin for Phase 1 data. (Right) $\langle\Delta\rangle$ distribution in the same energy bin for Phase 1 data.

data rather than from simulations. Nevertheless, the distributions for simulated showers remain close to zero within statistical uncertainties, indicating that the data-derived benchmark provides a reasonable description of the risetime behaviour in the simulations as well. Examining the distributions for simulated showers is important to verify that the benchmark constructed from the data can be consistently applied to simulated events. This validation ensures that comparisons between data and simulations using the observables Δ_i and $\langle\Delta\rangle$ are meaningful and that any systematic deviation from zero observed in other energy bins can be interpreted as a change in shower development, potentially reflecting differences in the primary mass composition.

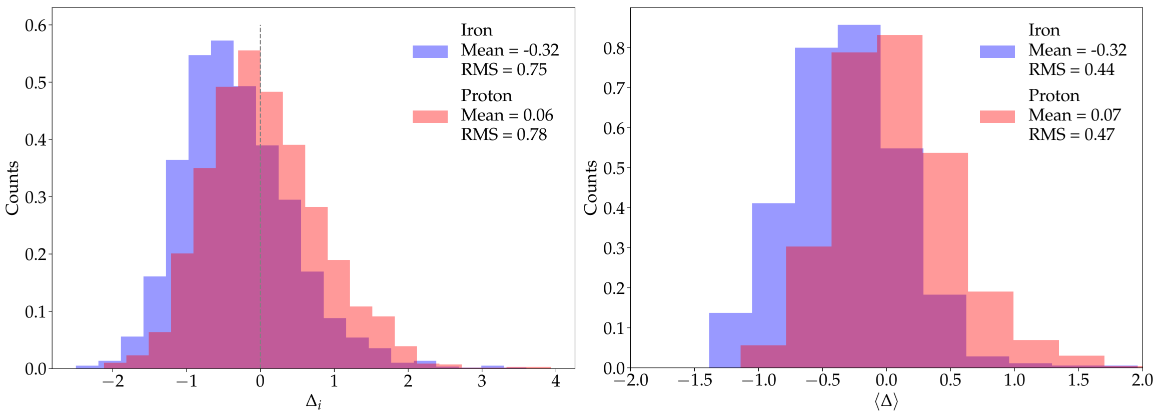

Figure 5.12: (Left) Δ_i distribution for the benchmark energy bin. (Right) $\langle\Delta\rangle$ distribution in the same energy bin.

The uncertainty of each Δ_i has contributions from two sources: the uncertainty of the risetime and the uncertainty coming from the parametrisations used in the analysis (one to describe the benchmark and the other to describe the risetime uncertainty). Assuming that the measured risetime uncertainty dominates the uncertainty of Δ_i, the uncertainty arising from the parametrisations can be neglected. Using propagation of errors, the uncertainty in $\langle\Delta\rangle$ is

$$\sigma_{\langle\Delta\rangle} = \frac{1}{\sqrt{N}}, \tag{5.12}$$

where N is the number of SD used in the calculation of $\langle\Delta\rangle$.

To verify that the benchmark parametrisation correctly captures the zenith-angle dependence of the risetime, the mean values of the event-level observable $\langle\Delta\rangle$ are evaluated and plotted in Fig. 5.13 as a function of $\sec\theta$ in the benchmark energy bin. Since the benchmark was constructed from data in this energy interval, no residual zenith dependence is expected for the data sample if the parametrisation accurately models the angular behaviour. The mean values of $\langle\Delta\rangle$ for the SD–750 array data are consistent with a flat distribution around zero within uncertainties across the explored $\sec\theta$ ranges. This confirms that the benchmark function successfully removes the zenith-angle dependence of the risetime in the reference energy bin. In contrast, the simulated predictions for the pure proton exhibit a residual dependence on $\sec\theta$, while iron remains systematically negative and approximately flat within uncertainties. The AugerMix prediction lies between the two pure compositions and shows a mild angular trend. The absence of a significant angular dependence in the data confirms that the benchmark correction effectively captures the zenith scaling of the risetime for the reference energy bin. The residual trends observed in simulations therefore reflect differences in the modelled shower time structure rather than deficiencies in the benchmark parametrisation.

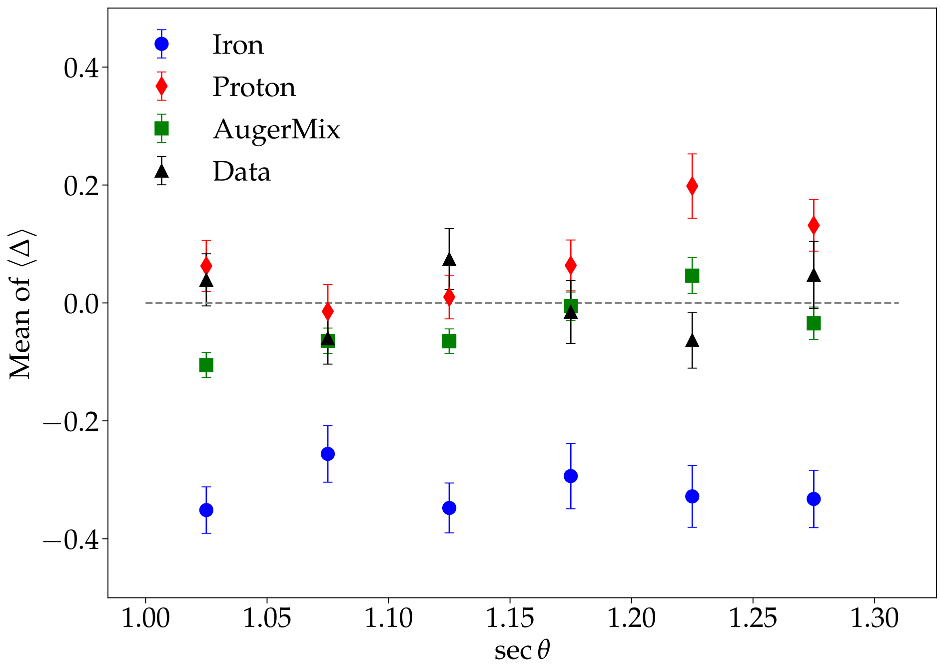

Figure 5.13: Mean values of $\langle\Delta\rangle$ as a function of $\sec\theta$ in the benchmark energy bin.

The mean values of the observable $\langle\Delta\rangle$ as a function of energy are shown in Fig. 5.14 for data collected with the SD–750 array, together with simulations of proton, iron and the AugerMix composition model. A clear and approximately monotonic increase of $\langle\Delta\rangle$ with energy is observed for all compositions, as well as for the measured data. Across the full energy range, the $\langle\Delta\rangle$ for the proton yields systematically larger $\langle\Delta\rangle$ values than iron, with AugerMix lying between the two, as expected. The data points follow a trend closer to the $\langle\Delta\rangle$ values of proton than to pure iron, particularly at higher energies, although they remain compatible within uncertainties with a mixed-composition scenario. To quantify the energy evolution, straight-line fits were performed for each composition hypothesis. The resulting slopes are summarised in Table 5.3. Although proton and AugerMix exhibit comparable slopes within uncertainties, the data and iron prediction show a significantly weaker energy dependence. As is also evident in Fig. 5.14, a purely linear description does not fully capture the behaviour of $\langle\Delta\rangle$, especially at lower energies, i.e., $\log_{10}(E/\mathrm{eV}) < 17.7$.

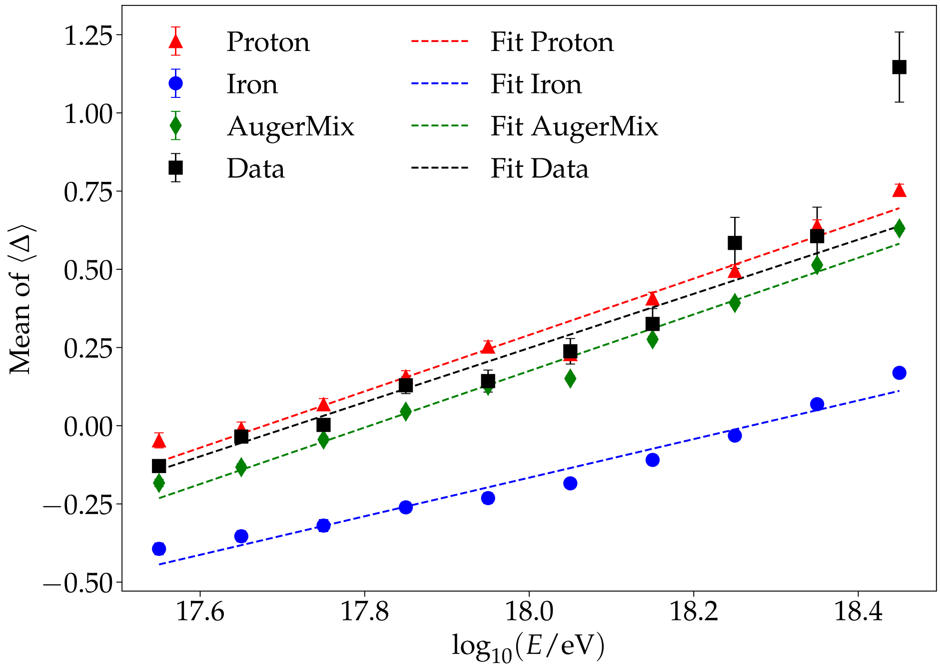

Figure 5.14: Mean values of $\langle\Delta\rangle$ as a function of energy for data gathered with the SD–750 array compared against simulations of proton, iron, and AugerMix.

Table 5.3: Slopes of $\langle\Delta\rangle$ as a function of the energy for the different compositions shown in Fig. 5.14.

Mass Composition	**Slope$\pm$Error**
p	0.9 ± 0.02
Fe	0.61 ± 0.02
AugerMix	0.9 ± 0.01
Data	0.87 ± 0.05

5.3.6 Converting $\langle\Delta\rangle$ to $\langle X^{\Delta}_{\max}\rangle$

The function that converts $\langle\Delta\rangle$ to $X_{\max}$ is found by calibrating with Golden Hybrid events. Since the $X_{\max}$ measurements from FD are unknown in this analysis, three methods were considered and are explained in the subsequent subsections.

Application of the calibration coefficients from Ref. [68]

As a first step, the calibration function and coefficients reported in Ref. [68] were directly applied to the present dataset. This provides a useful reference to assess how the previously published calibration performs when applied to the benchmark parametrisation and quality cuts adopted in this analysis. In Ref. [68], it is assumed that the depth of the shower maximum depends on the observable $\langle\Delta\rangle$ according to

$$X_{\max} = a + b\langle\Delta\rangle + c\log_{10}(E_{SD}/\mathrm{eV}) \tag{5.13}$$

The coefficients $a = 636 \pm 20$ g cm^{-2}, $b = 96 \pm 10$ g cm^{-2}, and $c = 2.9 \pm 1.2$ g cm^{-2} were used in conjunction with Eq. (5.13) to calculate the $X^{\Delta}_{\max}$ values based on the $\langle\Delta\rangle$ parameter determined in this analysis. The behaviour of $\langle X^{\Delta}_{\max}\rangle$ as a function of the logarithm of energy is shown in Fig. 5.15. The markers represent the $\langle X^{\Delta}_{\max}\rangle$ values obtained using Eq. (5.13), while the solid lines indicate the MC $\langle X_{\max} \rangle$ values for the simulations. The solid black line represents the measured $\langle X_{\max} \rangle$ data reported in Ref. [78]. As seen in the figure, the separation between the $\langle X^{\Delta}_{\max} \rangle$ values for the proton and iron primary is less than 50 g cm^{-2}, significantly smaller than the intrinsic MC separation. This indicates a reduced mass discrimination power when the published calibration coefficients are directly applied to the present benchmark parametrisation. The difference in the slope in the figure between the calculated $\langle X^{\Delta}_{\max} \rangle$ (markers) and the MC $\langle X_{\max} \rangle$ values (lines) highlights a mass-dependent bias in $\langle X^{\Delta}_{\max} \rangle$.

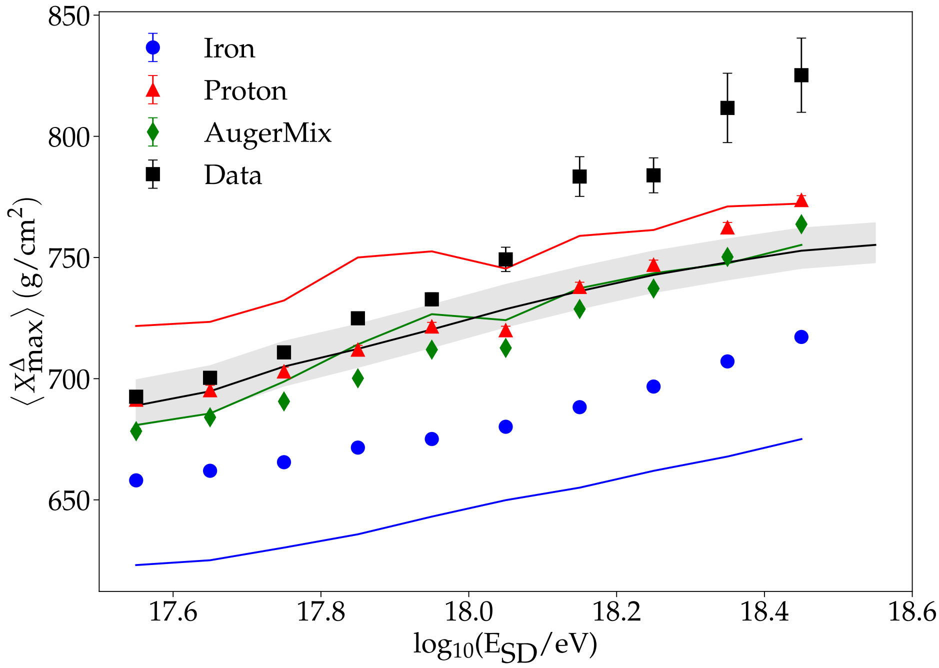

Figure 5.15: $\langle X^{\Delta}_{\max} \rangle$ as a function of energy for data gathered with the SD–750 array compared against simulations of proton, iron, and AugerMix (markers). $X^{\Delta}_{\max}$ is calculated using the calibration coefficients from Ref. [68] on Eq. (5.13). The solid lines indicate the MC $\langle X_{\max} \rangle$ values for the simulations. The solid black line represents the measured $\langle X_{\max} \rangle$ data reported in Ref. [78].

The performance of this calibration can be further evaluated by examining the residuals between the reconstructed and expected values of $X_{\max}$. Since the calibration coefficients are not re-fitted in this step, the residuals are computed with respect to the MC truth, i.e., $X^{\Delta}_{\max} - X^{\mathrm{MC}}_{\max}$. The distribution of these residuals for the AugerMix composition is shown in the left panel of Fig. 5.16. The mean value of the distribution is approximately ~ -5 g cm^{-2} with an RMS value of $\sim$ 56 g cm^{-2}. The small mean indicates a modest overall bias in the reconstructed $X^{\Delta}_{\max}$, while the RMS quantifies the dispersion of the reconstruction. The energy dependence of the bias is shown in the right panel of Fig. 5.16. The bias exhibits a clear dependence on the primary mass. Proton showers show a negative bias except in the highest energy bin, whereas iron showers display a positive bias throughout the explored energy range. The AugerMix composition lies between these two extreme cases and remains closer to the proton prediction. In addition, the magnitude of the bias decreases rapidly with increasing energy for the proton and AugerMix compositions at energies > 10^{18} eV. These results demonstrate that the direct application of the published calibration coefficients introduces a composition-dependent bias when used with the benchmark parametrisation adopted in this work.

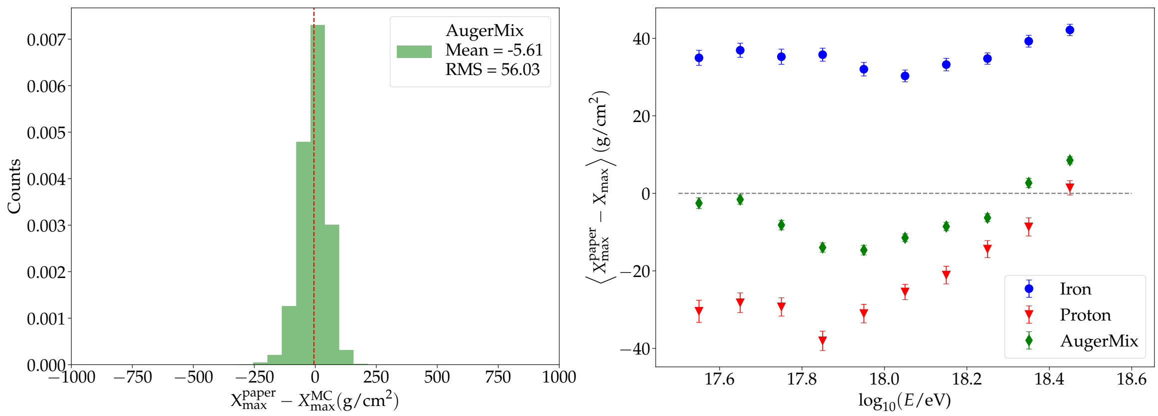

Figure 5.16: (Left) Differences between X^{Δ}_{max} and MC X_{max} values for the AugerMix composition. X^{Δ}_{max} is calculated using the calibration coefficients from Ref. [68] on Eq. (5.13). (Right) The bias of the X^{Δ}_{max} obtained for this analysis as a function of energy. X^{Δ}_{max} is calculated using the calibration coefficients from Ref. [68] on Eq. (5.13).

Using X_{max} measurements from FD

The results presented above indicate that the direct application of the calibration coefficients from Ref. [68] does not fully reproduce the behaviour of the observables obtained with the benchmark parametrisation used in this work. The reduced separation between the proton and iron simulations and the presence of a composition-dependent bias suggest that the relationship between $\langle \Delta \rangle$ and X_{max} differs from that assumed in the original calibration. This is expected since the present analysis employs updated quality cuts and a different benchmark parametrisation for the risetime observable. Consequently, the calibration coefficients must be re-evaluated using the golden hybrid dataset corresponding to the current analysis. In the following subsection, the parameters of Eq. (5.13) are therefore determined directly from the hybrid events used in this study. The calibration was repeated using golden hybrid events from the ICRC 2025 production, for which both FD and SD information are simultaneously available. X_{max}, measured by the FD, was used to calibrate the risetime parameter $\langle \Delta \rangle$ obtained from the SD.

Figure 5.17 shows the correlation between $\langle \Delta \rangle$ and $\langle X_{max} \rangle$ for hybrid data (black markers). The X_{max} from simulations (proton in red and iron in blue) is also shown in the figure. To quantify the relationship between the mean shower depth and the risetime observable, a linear model was adopted and is given as

$$X_{max} = a + b\langle \Delta \rangle. \tag{5.14}$$

The parameters of the linear calibration relation are obtained using a weighted least-squares fit to the binned means, where the weights correspond to the inverse squared uncertainties of the means. Under the assumption of Gaussian uncertainties, this procedure is equivalent to a maximum-likelihood estimation. The fit yields

$$a = 688.9 \pm 0.5\ \mathrm{g\,cm^{-2}}, \qquad b = 15.57 \pm 0.514\ \mathrm{g\,cm^{-2}},$$

with uncertainties obtained from the regression covariance matrix. The fitted calibration curve is overlaid (blue line) in Fig. 5.17. The positive slope confirms that larger values of the asymmetry parameter correspond, on average, to a deeper shower development in the atmosphere.

Using this calibration, $\langle X^{\Delta}_{max} \rangle$ values were reconstructed for all simulated compositions and data. The behaviour of $\langle X^{\Delta}_{max} \rangle$ as a function of the energy logarithm is shown in Fig. 5.18.

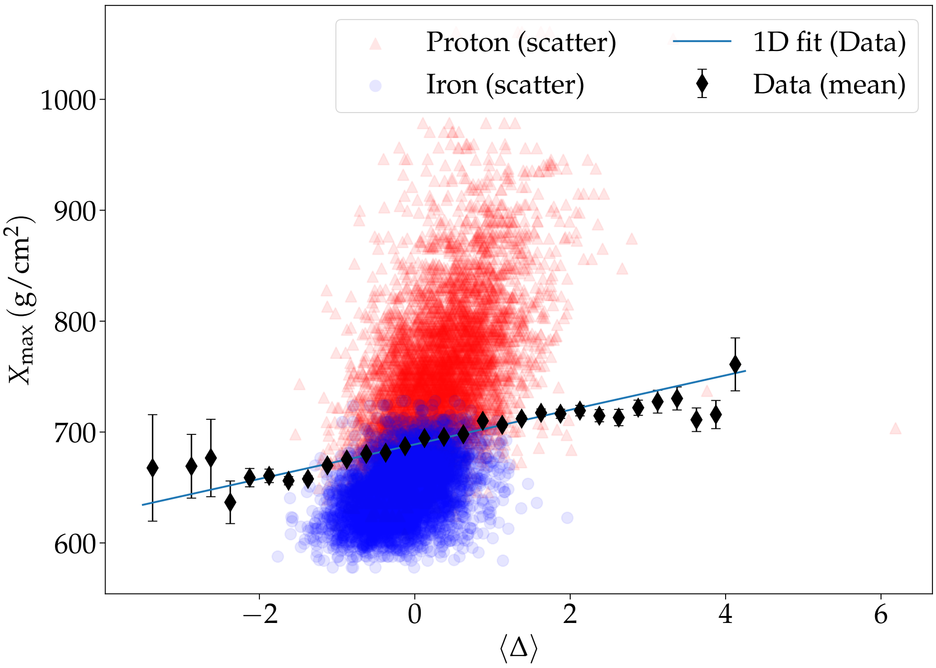

Figure 5.17: $\langle X_{max} \rangle$ from the golden hybrid events (in black) and X_{max} from simulations (proton in red and iron in blue) as a function of $\langle \Delta \rangle$. The blue line is the calibration curve obtained with a weighted least-squares fit to the binned means, where the weights correspond to the inverse squared uncertainties of the means.

Markers represent the $\langle X^{\Delta}_{max} \rangle$ values obtained using Eq. (5.14), while solid lines indicate the MC X_{max} values for the simulations. The solid black line represents the published $\langle X_{max} \rangle$ data reported in Ref. [78] and the grey band denotes the systematic uncertainty of the mean, both of which are obtained from measurements of X_{max} with Pierre Auger Observatory FD. Although the global energy trend is reproduced, the separation between the reconstructed proton and iron curves is significantly reduced compared to the true MC separation. From the figure, it is clear that the separation in $\langle X^{\Delta}_{max} \rangle$ between proton and iron is $\lesssim 25$ g cm^{-2} at higher energies and even smaller at lower energies, whereas the corresponding separation in the MC X_{max} is much larger, at roughly ~ 100 g cm^{-2}. This compression indicates that the linear calibration partially washes out intrinsic mass sensitivity. Moreover, a residual, mass-dependent bias is evident across the energy range.

The calibration performance can be assessed by studying the residual distribution between the reconstructed depth of the shower maximum obtained from the SD observable and the value measured directly by the FD. In this case, the residual is defined as the difference between the reconstructed value X^{Δ}_{max} and the FD measurement X^{FD}_{max}. The distribution of this quantity for the golden hybrid events used in the calibration is shown in the left panel of Fig. 5.19. For the hybrid data sample, the distribution has a mean value of approximately 3 g cm^{-2} and an RMS of approximately ≈ 88 g cm^{-2}. The mean of the residual distribution quantifies the global bias of the reconstruction, i.e., whether the method systematically overestimates or underestimates the true value of X_{max}. A mean value close to zero indicates that the calibration procedure does not introduce a significant systematic shift between the reconstructed and measured quantities. The small positive value obtained here implies a slight tendency of the method to underestimate X_{max} by only a few g cm^{-2}, which is negligible compared to the typical fluctuations of air showers. The RMS of the residual distribution characterises the dispersion of the reconstruction, which arises from several sources. These include intrinsic shower-to-shower fluctuations, measurement uncertainties

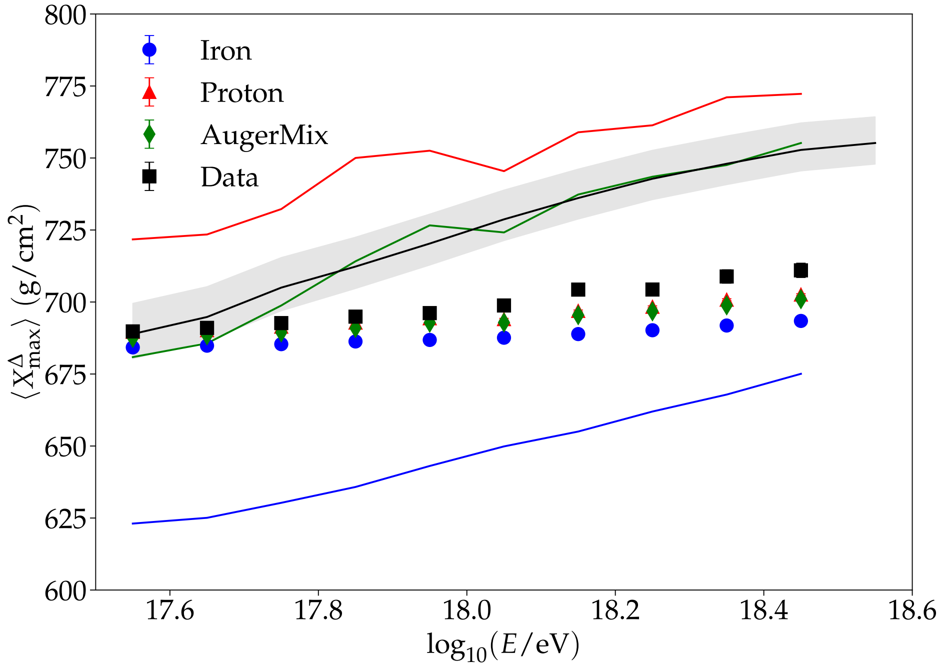

Figure 5.18: $\langle \Delta \rangle$ vs $\langle X_{\max} \rangle$ for the golden hybrid events (in black) and simulations (proton in red and iron in blue). The red line is the calibration curve obtained with the maximum likelihood method. (Right) $\langle X_{\max}^{\Delta} \rangle$ as a function of energy for data gathered with the SD–750 array compared against simulations of proton, iron, and AugerMix (markers). The solid lines indicate the MC $X_{\max}$ values for the simulations. The solid black line represents the measured $\langle X_{\max} \rangle$ data reported in Ref. [78]. The grey band corresponds to the systematic error of the mean.

in the SD signals, and the finite resolution of the FD measurement used as reference. Because the natural fluctuations of $X_{\max}$ in EAS are large (typically of the order of 50–80 g cm^{-2} depending on the primary mass), the observed RMS reflects both the intrinsic physical fluctuations of the showers and the resolution of the reconstruction method. Therefore, the width of the distribution provides a measure of the effective reconstruction resolution of $X_{\max}^{\Delta}$ when applied to real data.

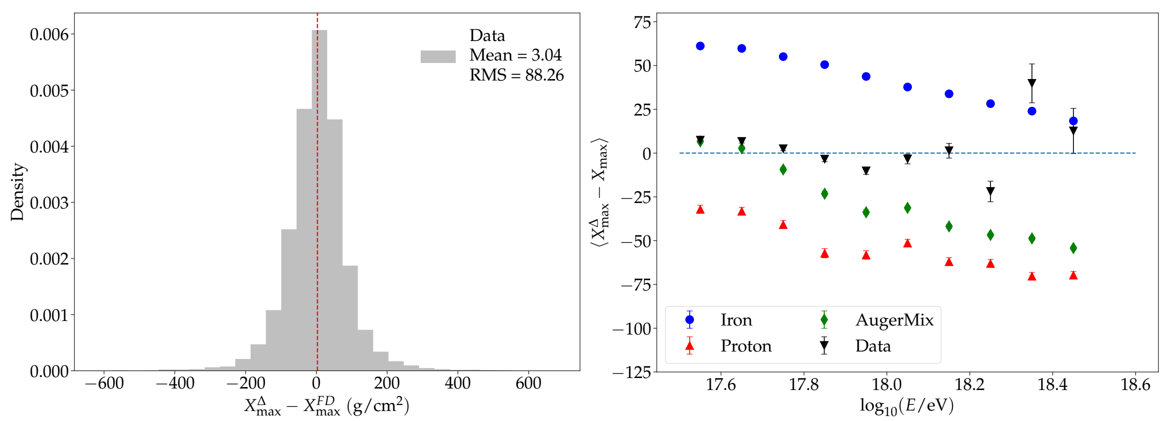

Figure 5.19: (Left) Differences between $X_{\max}^{\Delta}$ and $X_{\max}^{\text{FD}}$ values for the golden hybrid events. $X_{\max}^{\Delta}$ is calculated using Eq. (5.14), where the calibration coefficients were determined through a least-squares minimisation applied to binned hybrid data. (Right) The bias of the $X_{\max}^{\Delta}$ obtained for this analysis as a function of energy.

To further evaluate the performance of the calibration, the bias of the reconstructed $X_{\max}^{\Delta}$ is examined as a function of energy. The bias is defined as the mean value of the residual distribution $\langle X_{\max}^{\Delta} - X_{\max}^{\mathrm{MC}} \rangle$ for simulations and $\langle X_{\max}^{\Delta} - X_{\max}^{\mathrm{FD}} \rangle$ for golden hybrid data. The right panel of Fig. 5.19 shows this quantity for simulations of proton and iron primaries, as well as for the mixed composition scenario (AugerMix) and the golden hybrid data sample as a function of energy. The simulations exhibit a clear dependence of the bias on the primary mass of the CRs. Proton showers show a negative bias, meaning that the reconstructed $X_{\max}^{\Delta}$ tends to underestimate the true depth of the shower maximum obtained from the MC truth. In contrast, iron showers show a positive bias, indicating a systematic overestimation of $X_{\max}^{\Delta}$. This behaviour is expected because the risetime observable used in this analysis is sensitive to the longitudinal development of the shower and, therefore, depends on the mass of the primary particle. Proton showers develop deeper in the atmosphere and have larger shower-to-shower fluctuations, while iron showers develop earlier and exhibit smaller fluctuations. Since the calibration is derived from the golden hybrid data sample, which contains an unknown mixture of primary masses, it reflects the average behaviour of the composition present in the data. When this calibration is applied to simulations of pure primaries, it cannot simultaneously reproduce the exact behaviour of both proton and iron showers, leading to opposite biases for these extreme mass compositions. The AugerMix composition, which represents a realistic mixture of primary masses consistent with previous Pierre Auger Observatory measurements, lies between the proton and iron curves and remains closer to the proton prediction. This behaviour is consistent with a composition dominated by light nuclei in the considered energy range. The bias observed for the golden hybrid data remains close to zero throughout the explored energy range. This indicates that the calibration derived from the hybrid sample successfully reproduces the average behaviour of the real data and does not introduce a significant systematic shift in the reconstructed $X_{\max}^{\Delta}$. The fact that the simulated compositions bracket the data behaviour provides confidence that the method correctly captures the physical dependence of the observable on the shower development while remaining unbiased for the measured events.

Using the calibration function from Ref. [68]

The calibration function from Ref. [68] given in Eq. (5.13) was fitted with golden hybrid events from the ICRC 2025 production. The calibration coefficients were determined through a least-squares minimisation applied to binned hybrid data. The resulting fit parameters are,

$$a = 117.9 \pm 27.69\ \mathrm{g\,cm^{-2}}, \qquad b = 32.22 \pm 0.39\ \mathrm{g\,cm^{-2}}, \qquad c = 12.89 \pm 2.69\ \mathrm{g\,cm^{-2}},$$

with uncertainties obtained from the regression covariance matrix. The reconstructed observable $\langle X_{\max}^{\Delta} \rangle$, obtained using these coefficients, is shown as a function of energy for events recorded by the SD–750 array in Fig. 5.20. The data were compared with simulations of the compositions of the proton, iron, and AugerMix (markers). The solid coloured curves correspond to the MC $X_{\max}$ values for the simulated primaries. The solid black line represents the measured $\langle X_{\max} \rangle$ data reported in Ref. [78], while the grey band indicates the corresponding systematic uncertainty. The reconstructed values follow the expected elongation trend and lie between the predictions for proton and iron primaries, consistent with a mixed mass composition.

As described in detail in the previous subsections, the calibration performance can be evaluated by examining the distribution of the residuals between the reconstructed and FD measurements of the maximum depth of the shower. The left panel of Fig. 5.21 shows the distribution of $X_{\max}^{\Delta,E} - X_{\max}^{\mathrm{FD}}$ for the golden hybrid events. The mean value of this distribution is ~ 3 g cm^{-2} with an RMS value of ~ 88 g cm^{-2}. The small mean indicates that

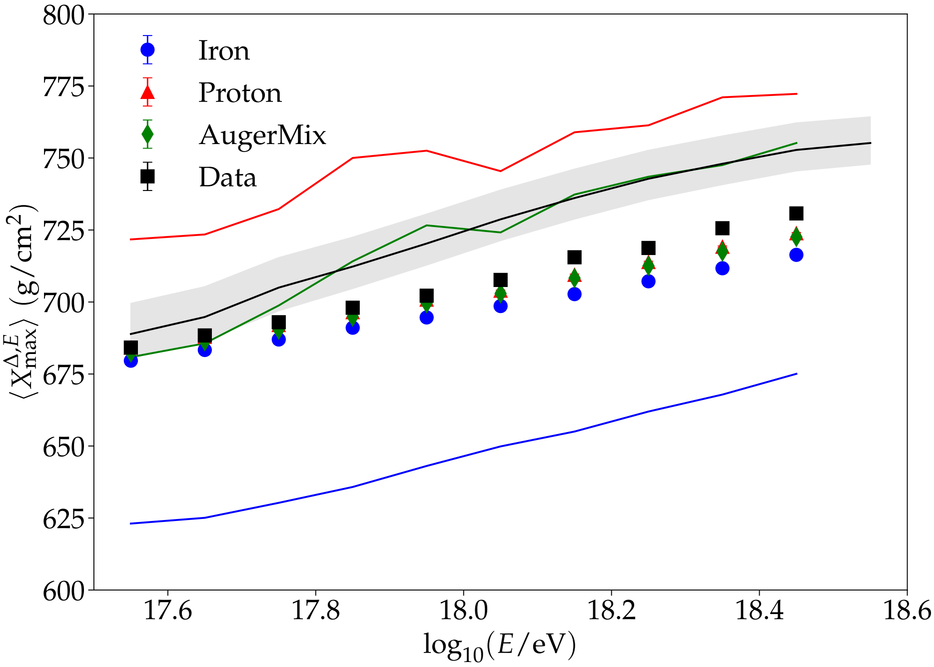

Figure 5.20: $\langle X^{\Delta}_{\rm max}\rangle$ as a function of energy for data gathered with the SD–750 array compared against simulations of proton, iron, and AugerMix (markers). The solid lines indicate the MC $X_{\rm max}$ values for the simulations. The solid black line represents the measured $\langle X_{\rm max}\rangle$ data reported in Ref. [78]. The grey band corresponds to the systematic error of the mean.

the calibration does not introduce a significant systematic bias, while the RMS reflects the dispersion arising from shower-to-shower fluctuations and the reconstruction uncertainties. The energy dependence of the bias is shown in the right panel of Fig. 5.21. As in the case of the calibration based solely on Eq. (5.14), the bias obtained for simulations exhibits a clear mass dependence. Proton showers show a negative bias, whereas iron showers display a positive bias that decreases with increasing energy. The AugerMix composition lies between these two extreme cases. The bias obtained for the golden hybrid data remains close to zero throughout the explored energy range, indicating that the calibration derived from the hybrid sample provides an unbiased estimator of $X_{\rm max}$ for the measured events. These results indicate that the dominant contribution to the calibration arises from the term proportional to $\langle\Delta\rangle$ in Eq. (5.13). The additional energy term primarily accounts for the energy dependence of the observable and ensures a smooth conversion between $\langle\Delta\rangle$ and $X_{\rm max}$, but it has only a minor impact on overall calibration.

5.4 Summary and conclusions

In this chapter, a mass-sensitive observable based on the muon content measured by the underground muon detectors was introduced and systematically optimised. The observable V_b corresponds to the weighted sum of the muon densities evaluated with the reference distance fixed to $r_0 = 450$ m and was shown to capture both the total muon content and the radial structure of the muon distribution without relying on an explicit parametrisation of the muon lateral distribution function. A merit-factor analysis demonstrated that the discrimination power between proton and iron primaries is maximised for values of the exponent b that depend weakly on energy and can be described by a saturating exponential parametrisation. The resulting optimised observable $V_{b_{\rm max}}$ exhibits an increasing separation power with energy, a stable scaling with primary energy, and no significant dependence on

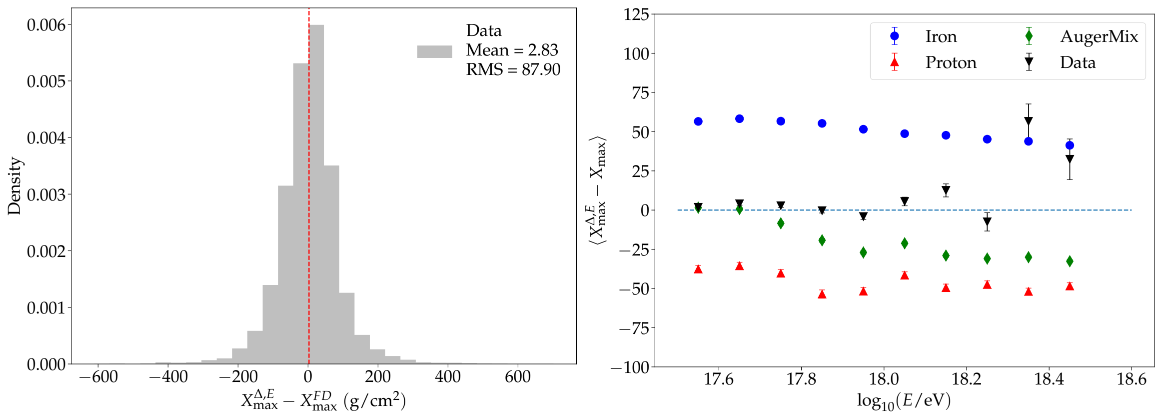

Figure 5.21: (Left) Differences between $X^{\Delta,E}_{\rm max}$ and $X^{\rm FD}_{\rm max}$ values for the golden hybrid events. $X^{\Delta,E}_{\rm max}$ is calculated using Eq. (5.13), where the calibration coefficients were determined through a least-squares minimisation applied to binned hybrid data. (Right) The bias of the $X^{\Delta,E}_{\rm max}$ obtained for this analysis as a function of energy.

the zenith angle. Although the discrimination power is not significantly improved when considering the normalised observable V_b^*, this quantity incorporates an explicit energy normalisation that reduces sensitivity to energy-scale uncertainties. In this sense, V_b^* provides a more robust observable under systematic variations, particularly those associated with energy reconstruction and calibration. Therefore, while V_b^* does not produce a substantial gain for discrimination by mass composition itself, it can represent a more stable and reliable choice in analyses where systematic effects play a dominant role. However, in the context of the present work, V_b^* is not adopted but is definitely a choice. The observable reflects the integrated muon density profile evaluated relative to the optimal reference distance of the UMD infill array, preserving the natural energy dependence of the muon signal. By construction, V_b provides a model-independent summary of both the muon content and its radial distribution, making it a robust mass-sensitive observable that can be optimised through the choice of b and systematically compared between hadronic interaction models. The merit factor of the new parameter is comparable to those obtained for $X_{\rm max}$ and the fitted muon density evaluated at a reference distance as reported in Ref. [45].

The risetime asymmetry method was implemented and validated in the infill energy range, successfully reproducing the behaviour reported in Ref. [68]. The benchmark parametrisation removes the zenith-angle dependence of the risetime observable in the data, yielding a flat behaviour of $\langle\Delta\rangle$ as a function of $\sec\theta$ in the reference energy bin. However, simulations exhibit residual angular and energy-dependent structures that are not present in the data. This indicates model-dependent differences in the predicted development of EAS and highlights the limitations of using simulations to describe the detailed behaviour of the risetime observable. When the published calibration coefficients are applied, the reconstructed $X^{\Delta}_{\rm max}$ values of the simulations exhibit a clear energy-dependent and mass-dependent bias with respect to the MC $X_{\rm max}$ values. In particular, the separation between the proton and iron showers is reduced relative to the intrinsic MC separation, and the bias varies with energy. This behaviour limits the direct use of $X^{\Delta}_{\rm max}$ as a mass-sensitive observable when the original calibration is applied to the present benchmark parametrisation. Although $\langle\Delta\rangle$ method provides a fully SD-based and data-driven estimator of $X_{\rm max}$, the calibration used to convert $\langle\Delta\rangle$ into $X^{\Delta}_{\rm max}$ described in Ref. [68] does not fully account for the composition dependence of shower evolution. Consequently, while the risetime asymmetry method retains sensitivity to the longitudinal development of air showers and their energy evolution, the reconstructed $X^{\Delta}_{\rm max}$ obtained in the current manner cannot be used as a direct substitute for the FD measurement of X_{max} in precision mass composition studies without further calibration or model-dependent corrections. In particular, the presence of a systematic, energy-dependent bias largely washes out the intrinsic mass separation, leaving only limited residual sensitivity to the primary mass.

CHAPTER 6

Composition analysis from muon content and $X_{\max}$ correlation

This chapter investigates whether the statistical correlation between the two composition-sensitive observables introduced in Chapter 5 can be used as a probe of the primary mass of UHECR. The observables $V_{b_{\max}}$ and $X_{\max}$ provide complementary information about the development of the shower. $X_{\max}$ reflects the longitudinal evolution of the electromagnetic component, while $V_{b_{\max}}$ is related to the muonic component at ground. Since the development of air showers depends strongly on the mass of the primary particle, the joint behaviour of these observables and, in particular, their statistical correlation is expected to encode composition information. This approach is closely related to previous studies by Pierre Auger Observatory, where the correlation between $X_{\max}$ and muon-sensitive observables has been used to constrain the spread of primary masses. In particular, the work presented in Ref. [71] exploits the correlation between $X_{\max}$ and the surface detector signal $S(1000)$ to demonstrate that a negative correlation is indicative of a mixed mass composition, providing a robust observable largely insensitive to systematic uncertainties and hadronic interaction models. Similarly, direct measurements of the muonic component, as reported in Ref. [70], highlight the importance of accurately characterising the muon content of air showers as a key composition-sensitive quantity. The analysis developed in this chapter follows the same underlying principle: the interplay between the electromagnetic and muonic components of air showers carries information about the primary mass. However, instead of relying on global observables, such as the total number of muons or the surface detector signal, this analysis employs the correlation between $V_{b_{\max}}$ and $X_{\max}$, providing an alternative and complementary probe of composition.

An additional advantage of using a correlation coefficient is its invariance under linear transformations of the observables. Changes or rescaling of $X_{\max}$ or $V_{b_{\max}}$, which may arise from calibration uncertainties or detector-dependent systematics, do not affect the value of the correlation coefficient. This reduces sensitivity to absolute scale uncertainties and makes correlation-based observables particularly robust for composition studies. Furthermore, correlation-based analyses are largely insensitive to the so-called muon deficit problem in air-shower simulations (see Section 2.5). Since the deficit manifests itself primarily as a global mismatch in the absolute normalisation of the muonic component, it affects $V_{b_{\max}}$ (or other muon-related observables) through an overall shift or rescaling. Such transformations leave the rank ordering and relative event-by-event fluctuations essentially unchanged, and therefore have little impact on the correlation coefficient. As a result, the inferred composition information from the correlation remains robust even in the presence of discrepancies between simulated and observed muon densities.

The primary objective of this chapter is to determine whether the correlation between $V_{b_{\max}}$ and $X_{\max}$ can discriminate between different primary composition scenarios and to quantify the statistical power required for such discrimination. To achieve this, several correlation estimators are first evaluated, and the most suitable metric is selected. The dependence of the chosen correlation coefficient on high-energy hadronic interaction models is then examined. Its sensitivity to primary mass composition is studied by using simulated pure and mixed compositions. Finally, the minimum number of events required to reject a pure-composition hypothesis at a specified significance level is derived and compared with the number of events expected to be observed by the detector.

The structure of the chapter is as follows. Section 6.1 reviews different correlation coefficients found in the literature. Section 6.2 calculates the various correlation coefficients between $X_{\max}$ and $V_{b_{\max}}$ and chooses the most suitable correlation coefficient for mass composition studies. Section 6.2 also investigates the dependence of the selected estimator on high-energy hadronic interaction models. Section 6.3 evaluates the sensitivity of the correlation to primary composition. Section 6.4 derives the minimum number of events required to reject the pure-composition hypothesis. In Section 6.5, the correlation coefficient between $X_{\max}$ and $V_{b_{\max}}$ of the data is explored. The chapter ends in Section 6.6.

6.1 Correlation coefficients

The correlation coefficients quantify the strength and direction of statistical associations between pairs of variables. They do not imply causation, nor do they identify dependent and independent variables. Instead, they measure the degree to which two quantities vary together according to a specified criterion (e.g., linear or monotonic association). Any relationship inferred from a correlation coefficient must therefore be interpreted strictly as an association rather than a causal effect. The correlation coefficients considered in this work are briefly described in the following subsections.

Pearson's product-moment correlation coefficient (r_{p})

Pearson's product-moment correlation coefficient is appropriate when both variables are approximately normally distributed, and their relationship is linear [72]. Because it depends directly on the numerical values of the observations, it is sensitive to extreme values, which may either exaggerate or suppress the apparent strength of association. Consequently, it may perform poorly when one or both variables deviate substantially from normality. For variables X and Y, Pearson's correlation coefficient is defined as

$$r_{\mathrm{p}} = \frac{\mathrm{cov}(X, Y)}{\sigma(X)\sigma(Y)}.$$

Pearson's correlation coefficient ranges from -1 to 1. A value of 0 implies that there is no linear association between two variables. A value of zero indicates the absence of a linear association. Positive values indicate that the variables tend to increase together, whereas negative values indicate that one variable tends to decrease as the other increases.

Spearman's rank correlation coefficient (r_{s})

The Spearman rank correlation coefficient is defined as the Pearson correlation applied to the ranked values of the data [73]. It is particularly useful when variables are non-normally distributed, skewed, or ordinal. Because it depends only on the relative ordering of observations, it is more robust to extreme values than Pearson's coefficient. For a sample size of N, observations X_i and Y_i are replaced by their ranks $R(X_i)$ and $R(Y_i)$, respectively, and the coefficient is defined as

$$r_s = \frac{\mathrm{cov}(R(X), R(Y))}{\sigma(R(X))\sigma(R(Y))}.$$

Like Pearson's coefficient, Spearman's coefficient ranges from -1 to 1. However, it measures the strength of a monotonic relationship rather than strictly linear dependence. Thus, it can capture associations that are consistently increasing or decreasing, even if they are non-linear.

Gideon–Hollister rank correlation coefficient (r_{GH})

The Gideon–Hollister coefficient is a non-parametric measure based on the ranks of paired observations (x_i, y_i), where $i = 1, ..., N$ [76]. Data are first ordered in ascending order of x_i so that the ranks of x_i satisfy $\mathcal{R}(x_i) = i$. The perfect correlation corresponds to the case where the ranks of y_i follow the same ordering, $\mathcal{R}(y_i) = i$, while the perfect anti-correlation corresponds to $\mathcal{R}(y_i) = N + 1 - i$. The coefficient quantifies the maximum cumulative deviation of the observed rank ordering from these ideal cases. Specifically, it evaluates how strongly the sequence of ranks $\mathcal{R}(y_i)$ departs from perfect correlation or anti-correlation by computing cumulative counts of deviations. The cumulative sum is defined as

$$d_i^- = \sum_{j=1}^{i} I(\mathcal{R}(y_i) < N + 1 - i), \quad d_i^+ = \sum_{j=1}^{i} I(\mathcal{R}(y_i) > i),$$

where $I(\dot{)}$ is the indicator function, equal to 1 if the condition is true and 0 otherwise. The correlation coefficient is then determined from the maximum cumulative deviations,

$$r_G = \frac{\max\{d_i^-\} - \max\{d_i^+\}}{N/2}.$$

Because it is based solely on rank ordering and cumulative deviations, the Gideon–Hollister coefficient is more robust to outliers than Pearson's and Spearman's coefficients.

Kendall correlation coefficient (τ)

Kendall's tau [75] measures the degree of concordance between two variables and is particularly suitable when their distributions deviate substantially from normality. It is based entirely on the relative order of observations rather than their numerical values, making it a distribution-free measure of association.

Consider paired observations (x_i, y_i) for $i = 1, ..., N$. For any pair of observations (x_i, y_i) and (x_j, y_j) with $i < j$, the pair is concordant if both variables change in the same direction (either both increase or both decrease) and discordant if they change in opposite directions. If either variable is tied (x_i = x_j or y_i = y_j), the pair is classified as tied and is excluded from concordant and discordant counts. Let n_c and n_d denote the number of concordant and discordant pairs, respectively. Let n_1 and n_2 denote the number of pairs tied in each variable, and let n_0 be the total number of possible pairs. Kendall's correlation coefficient is then defined as

$$\tau = \frac{n_c - n_d}{\sqrt{(n_0 - n_1)(n_0 - n_2)}}.$$

Kendall's coefficient is generally more robust and statistically efficient than Pearson's and Spearman's coefficients when distributions are non-normal or contain outliers. Values close to 1 indicate strong agreement in ordering, while values close to -1 indicate strong disagreement.

Table 6.1 compares the correlation coefficients considered in this work. The table summarises statistical assumptions, robustness properties, and practical interpretation of experimental data. In the following sections, these correlation coefficients between the two observables $V_{b_{\max}}$ and $X_{\max}$ are calculated.

6.2 Correlation coefficient between $X_{\max}$ and $V_{b_{\max}}$

A complete simulation of UMD–FD hybrid events using the $\overline{\text{Off}}\underline{\text{line}}$ software is computationally challenging. Therefore, in this analysis, the observable $X_{\max}$ is taken directly from the MC shower and subsequently smeared to account for detector resolution. This smearing is implemented by adding a Gaussian fluctuation with zero mean and an energy-dependent standard deviation given by $\sigma[X_{\max}] = 14.78 + 3.4 \times (\log(E/\text{eV}) - 19.6)^2$ in g/cm^2 as reported in Ref. [59]. The resulting fluctuating quantity is denoted $X_{\max}$ throughout the remainder of the analysis. The analysis is done in the energy bin $\log_{10}(E/\text{eV}) \in [17.5, 18.7]$, and the energy corresponds to the FD energy described in Section 5.1.1.

The left panel of Fig. 6.1 illustrates the effect of applying Gaussian smearing to the MC $X_{\max}$ values to emulate detector resolution in a single energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$. The unsmeared MC $X_{\max}$ is shown as a function of the smeared value for the proton and iron primaries, together with the identity line. The spread of points around the identity line reflects the magnitude of the applied Gaussian fluctuations, while the overall linear trend indicates that the smearing preserves the underlying ordering of events. This procedure mimics the finite experimental resolution of the FD while retaining the physical separation between different primary compositions.

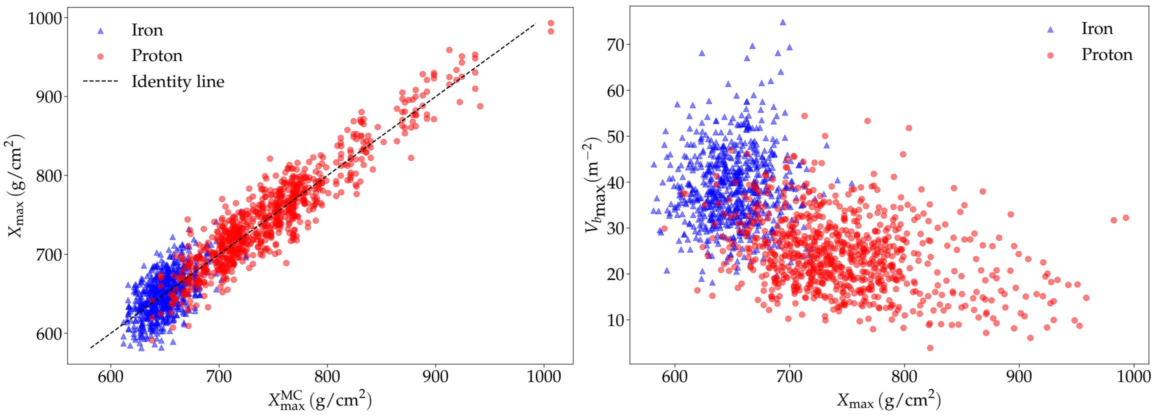

Figure 6.1: (Left) MC $X_{\max}$ values after applying Gaussian smearing to model detector resolution, shown as a function of the original (unsmeared) $X_{\max}$ for proton and iron primaries in a single energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$. The dashed line represents the identity relation. (Right)$V_{b_{\max}}$ as a function of $X_{\max}$ for proton and iron primaries in a single energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$, used for the correlation analysis. Both plots correspond to simulations using EPOS-LHC high-energy hadronic interaction model.

The joint distribution of the observables $X_{\max}$ and $V_{b_{\max}}$ for proton and iron primaries in a single energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$ is shown in the right panel of Fig. 6.1. Restricting the analysis to a narrow energy interval minimises residual energy-dependent effects, ensuring that the observed correlation structure reflects intrinsic shower-to-shower fluctuations rather than variations driven by energy. The simulations used to generate these data are described in Section 4.1.2.

A clear separation between the proton and iron populations is visible in the $X_{\max}$–$V_{b_{\max}}$ plane. Showers initiated by heavier primaries develop earlier in the atmosphere, resulting in lower values of $X_{\max}$, and produce a larger number of muons at ground level, leading to higher values of $V_{b_{\max}}$. The earlier development arises from the larger interaction cross section of heavy nuclei and, more importantly, from the smaller energy per nucleon in the superposition picture, which causes the electromagnetic cascade to reach its maximum at smaller atmospheric depths. In contrast, lighter primaries penetrate deeper before reaching the shower maximum and generate fewer muons. This opposite mass dependence of $X_{\max}$ and $V_{b_{\max}}$ gives rise to the characteristic anti-correlation between the two observables and forms the basis for using their correlation as a composition-sensitive discriminator.

Using the simulated data sets described in Section 4.1, the correlation between $X_{\max}$ and $V_{b_{\max}}$ was evaluated using the four estimators described in Section 6.1. The objective of this comparison was to identify the estimator that provides the most stable and statistically efficient measure of the sensitivity of composition. The analysis was performed within a single reconstructed energy bin, $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$, and was repeated across additional energy bins to verify that the observed behaviour is not driven by residual energy dependence.

The left panel of Fig. 6.2 shows the dependence of the correlation coefficient between $X_{\max}$ and $V_{b_{\max}}$ on the proton fraction,

$$C_p = \frac{N_p}{N_p + N_{Fe}}, \tag{6.1}$$

where N_p and N_{Fe} denote the numbers of protons and iron events in the sample, respectively. Artificial mixed compositions were constructed by randomly combining proton and iron simulations such that C_p varied from 0 to 1 in steps of 0.1. The limiting cases correspond to samples of pure iron ($C_p = 0$) and pure proton ($C_p = 1$). For each composition, the correlation coefficient was computed, and its statistical uncertainty was estimated using bootstrap resampling.

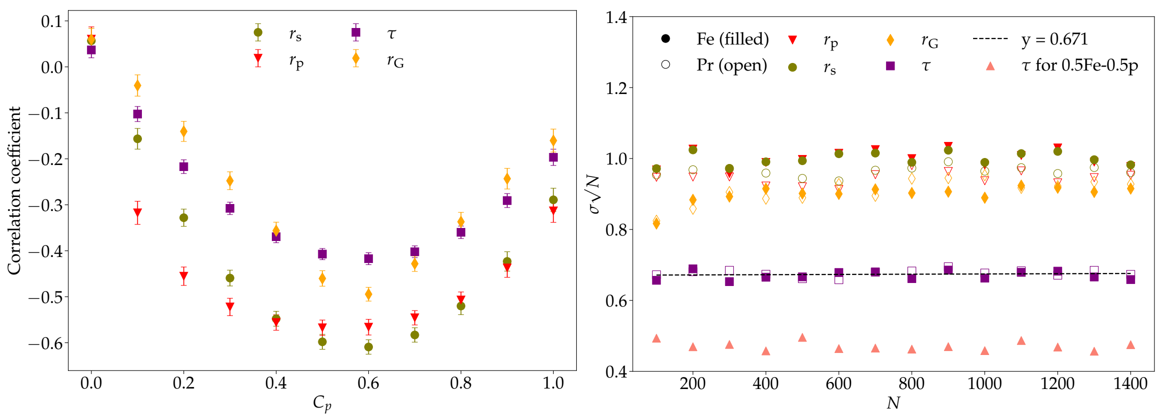

Figure 6.2: (Left) Correlation coefficients between $X_{\max}$ and $V_{b_{\max}}$ as a function of proton fraction C_p in a fixed energy bin. Different markers indicate different correlation estimators. (Right) Statistical uncertainty of the estimators, shown as a function of sample size N. Both figures correspond to primaries simulated with EPOS-LHC high-energy hadronic interaction model.

All correlation estimators exhibit a characteristic symmetric dependence on C_p, with the strongest (most negative) correlation occurring near equal mixing, $C_p \approx 0.5$. This behaviour arises because proton and iron showers populate distinct regions of the $(X_{\max}, V_b)$ phase space. When both populations are present in comparable proportions, the combined distribution acquires a pronounced global ordering between the two observables, enhancing the overall monotonic structure and therefore increasing the magnitude of the correlation. In contrast, pure compositions exhibit larger intrinsic fluctuations within a single population, which reduces the strength of the observed correlation. The qualitative dependence of the correlation coefficient on C_p was found to be consistent across energy bins, indicating that the composition sensitivity of the observable pair remains stable within the studied energy range.

To characterise the finite-sample behaviour of the correlation estimators, a bootstrap study based on case resampling was performed. For proton, iron, and a mixed (50% Fe – 50% Pr) composition, samples of size $N = 100, 200, \ldots, 1500$ were repeatedly generated by resampling events with replacement of the full simulated datasets. For each resampled dataset, the correlation coefficient was evaluated, and the procedure was repeated 1000 times to obtain the empirical distribution of the fixed estimator n. The dispersion of this distribution was quantified by its standard deviation, which was then multiplied by $\sqrt{N}$ to examine the expected $1/\sqrt{N}$ scaling of statistical fluctuations and to compare the asymptotic stability of the different correlation estimators.

The right panel of Fig. 6.2 shows $\sigma\sqrt{N}$ for different correlation estimators as a function of N. The approximately constant behaviour of $\sigma\sqrt{N}$ confirms that the statistical uncertainty follows the expected scaling $\sigma \propto 1/\sqrt{N}$. As seen in the figure, among the estimators considered, $\tau(V_{b_{\max}}, X_{\max})$ exhibits the smallest and most stable statistical uncertainty, which is well described by approximately $0.67/\sqrt{N}$. The value of $\sigma\sqrt{N}$ is similar for pure proton and pure iron compositions, while it is slightly smaller for the mixed 50% Fe - 50% Pr sample due to the stronger underlying correlation structure. The smaller statistical variance observed for the mixed composition arises because the combined proton–iron sample exhibits a stronger global ordering between $X_{\max}$ and $V_{b_{\max}}$ than either pure population alone. When two well-separated populations are mixed, a larger fraction of observation pairs becomes consistently concordant or discordant, which reduces the probability of ambiguous pair comparisons. Since Kendall's τ is based on pairwise ordering, this increased structural separation lowers the sampling variability of the estimator, leading to a smaller empirical variance in the bootstrap distribution. The $\sigma\sqrt{N}$ for different correlation estimators as a function of N was plotted for other high-energy hadronic interaction models and is included in Appendix C.1.

To remain conservative, the statistical uncertainty associated with $\tau(V_{b_{\max}}, X_{\max})$ is taken to be $0.67/\sqrt{N}$ for the remainder of this work. The bootstrap results do not show a significant dependence of this scaling on energy within the range considered. From Appendix C.1, it can also be concluded that this choice is independent of the high-energy hadronic interaction model considered.

6.2.1 Effect of outliers

Artificial outliers were introduced to assess the robustness of the correlation estimators against rare events. For each data set (p, Fe) for EPOS-LHC high-energy hadronic interaction model, the events were first restricted to the energy bin $\log_{10}(E/\text{eV}) \in [E_{\min}, E_{\max})$. A small fraction f_{out} of events (in this case $f_{\text{out}} = 0.01$, i.e., 1%) was then uniformly selected at random without replacement using a fixed random seed to ensure reproducibility. For the selected events, the observables entered in the correlation analysis, $V_{b_{\max}}$ and $X_{\max}$, were modified according to a shift contamination model, $x \to x + \Delta$, where Δ is a large random offset with a random sign and magnitude proportional to the intrinsic scale of the variable. Specifically, the offset amplitude was taken to be of the order $\sim 10\sigma$, where σ denotes the baseline standard deviation of the corresponding observable. This procedure generates extreme, physically inconsistent values that mimic rare reconstruction failures or pathological detector responses.

The left panel of Fig. 6.3 shows the dataset after the introduction of artificially generated outliers, displaying the joint distribution of the observables $X_{\max}$ and $V_{b_{\max}}$ for proton- and iron-induced showers in a single energy bin, $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$, for the EPOS-LHC high-energy hadronic interaction model model. The right panel shows the dependence of the correlation coefficient between $X_{\max}$ and $V_{b_{\max}}$ on the proton fraction defined in Eq. (6.1), for the different correlation estimators computed from the data set shown in the left panel.

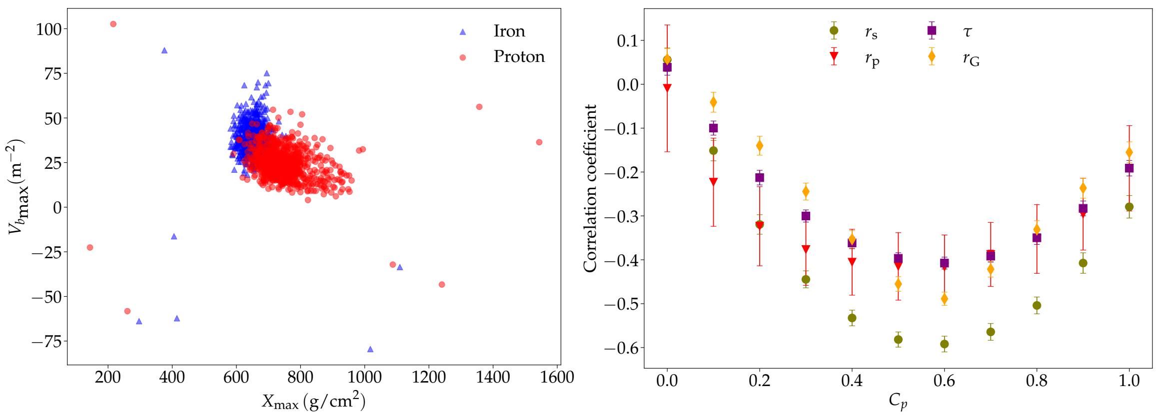

Figure 6.3: (Left) $V_{b_{\max}}$ as a function of $X_{\max}$ for proton and iron primaries in a single energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$, after introducing artificial outliers. (Right) Correlation coefficients are shown as a function of proton abundance C_p in the same energy bin, evaluated for the contaminated datasets. Different markers correspond to different correlation estimators. Both plots correspond to simulations using EPOS-LHC high-energy hadronic interaction model.

The presence of a small fraction of extreme events significantly affects estimators that depend directly on the numerical values or global rank structure of the data. Pearson's correlation is most strongly affected, showing large statistical fluctuations and reduced stability, reflecting its sensitivity to extreme deviations in the covariance. The Spearman coefficient and the Gideon-Hollister estimator are also affected by distortions of global rank ordering. In contrast, Kendall's correlation coefficient, which depends only on pairwise concordance, remains comparatively stable under contamination. This behaviour demonstrates the superior robustness of Kendall's τ in the presence of rare pathological events, further supporting its use as the primary estimator in subsequent analysis.

6.2.2 Comparison of Kendall's correlation estimator for different high-energy hadronic interaction models

Figure 6.4 shows the dependence of the correlation estimator $\tau(V_{b_{\max}}, X_{\max})$ on the proton fraction C_p for the three high-energy hadronic interaction models considered in this study. All models predict a similar qualitative trend in which the magnitude of the correlation varies smoothly with composition. However, small systematic differences are visible between the models, reflecting differences in predicted shower development and muon production. Despite these differences, the overall behaviour of the correlation remains comparable for all high-energy hadronic interaction models in the full range of compositions.

The weak dependence of the model of $\tau(V_{b_{\max}}, X_{\max})$ indicates that the correlation between $V_{b_{\max}}$ and $X_{\max}$ is mainly governed by the geometric and mass-dependent properties of air-shower development rather than details of the hadronic interaction modelling. This relative insensitivity to the choice of interaction model is advantageous for composition studies, as it reduces systematic uncertainties associated with model assumptions.

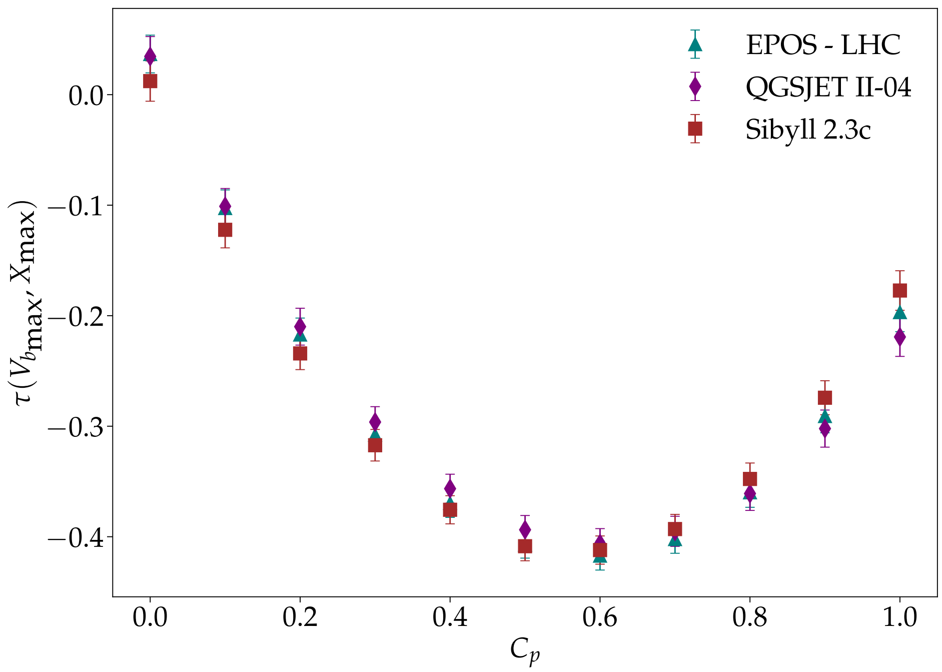

Figure 6.4: Dependence of $\tau(V_{b_{\max}}, X_{\max})$ on proton fraction C_p for three hadronic interaction models (EPOS-LHC, QGSJet-II.04, and Sibyll-2.3c).

6.2.3 Effect of energy dependence

Both observables entering the correlation estimator, $V_{b_{\max}}$ and $X_{\max}$, exhibit a dependence on the primary energy. Consequently, the measured value of Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ may be affected by the width of the energy bin used in the analysis. If a wide energy interval is considered, the intrinsic dependence of both observables on energy introduces an additional correlation component that is not directly related to the primary mass composition. To quantify this effect, the analysis was repeated using two different energy bin widths, $\Delta \log_{10}(E/\mathrm{eV}) = 0.1$ and $\Delta \log_{10}(E/\mathrm{eV}) = 0.2$. In the upper panels of Fig. 6.5, Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ is shown as a function of the logarithmic mass ($\ln A$) for pure primary compositions for the three different high-energy hadronic interaction models in the energy interval $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ (upper left), $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$ (upper right), $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.7)$ (lower left), and $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.2)$ (lower right). The values of the correlation estimator fall within the systematic uncertainties for the different high-energy hadronic interaction models. The increase in the width of the bin is found to reduce the magnitude of the correlation between $V_{b_{\max}}$ and $X_{\max}$. This behaviour arises because the broader bin mixes showers of different energies, thereby diluting the mass-dependent component of the correlation.

As a result, the ability of the estimator $\tau(V_{b_{\max}}, X_{\max})$ to distinguish between pure and mixed compositions is reduced when a bin width of 0.2 is used. This behaviour highlights the importance of using sufficiently narrow energy intervals when studying correlations between observables that both scale with primary energy. For this reason, the correlation analysis presented in this work is performed using energy intervals of width $\Delta \log_{10}(E/\mathrm{eV}) = 0.1$.

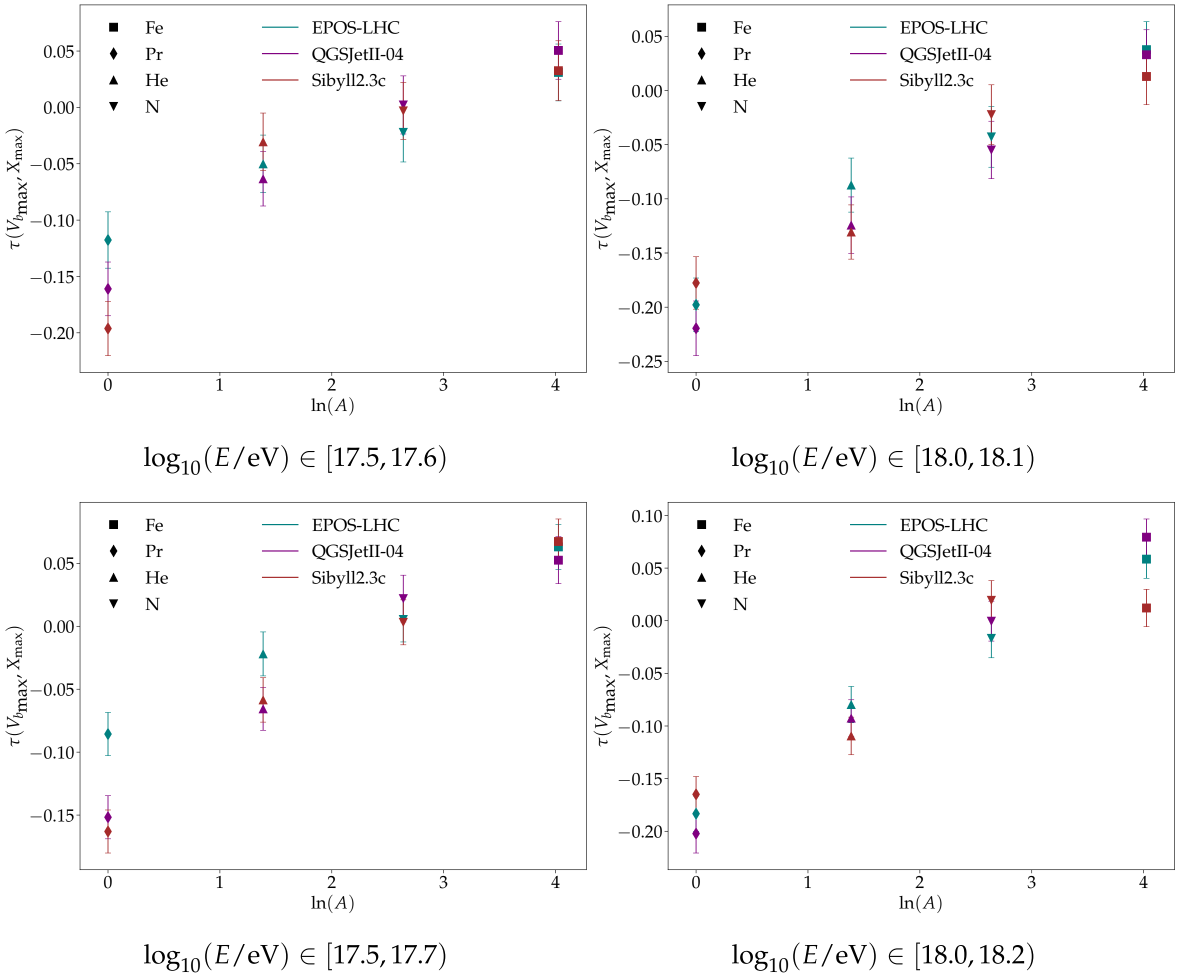

$\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$

$\log_{10}(E/\mathrm{eV}) \in [17.5, 17.7)$ $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.2)$

Figure 6.5: Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ as a function of the logarithmic mass ($\ln A$) for pure primary compositions for the three different high-energy hadronic interaction models using two different energy bin widths, $\Delta \log_{10}(E/\mathrm{eV}) = 0.1$ (upper panels) and $\Delta \log_{10}(E/\mathrm{eV}) = 0.2$ (lower panels).

6.3 Mass composition studies using Kendall's correlation coefficient

Kendall's tau ($\tau(V_{b_{\max}}, X_{\max})$) was selected as the primary correlation estimator for composition analysis based on its overall statistical and systematic performance. Although Pearson and Spearman coefficients exhibit larger absolute correlation values, they are more sensitive to fluctuations in distribution shape and to the presence of extreme events. The Gideon-Hollister coefficient, while robust, shows stronger dependence on global rank structure and lacks a direct probabilistic interpretation.

In contrast, $\tau(V_{b_{\max}}, X_{\max})$ provides a distribution-free measure of pairwise ordering that remains stable under measurement smearing, outlier contamination, and changes in the underlying hadronic interaction model. Bootstrap studies show that its statistical uncertainty scales as $1/\sqrt{N}$, with a small, energy-independent prefactor of approximately 0.67, i.e., $\sigma[\tau(V_{b_{\max}}, X_{\max})] = 0.67/\sqrt{N}$. Furthermore, its weak sensitivity to hadronic interaction modelling reduces systematic uncertainties in composition inference. Taken together, these properties make $\tau(V_{b_{\max}}, X_{\max})$ the most robust and reliable estimator for extracting composition information from the correlation between $X_{\max}$ and $V_{b_{\max}}$.

The Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ was evaluated for four pure primary compositions (proton, helium, nitrogen, and iron) as well as for the AugerMix composition, within the energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$. The resulting values are shown as a function of $\langle \ln A \rangle$ in Fig. 6.6. The mean logarithmic mass $\langle \ln A \rangle$ is defined as $\langle \ln A \rangle = \sum_i f_i \ln A_i$, where f_i denotes the fractional abundance of primaries with mass number A_i. For pure compositions $\langle \ln A \rangle = \ln A$.

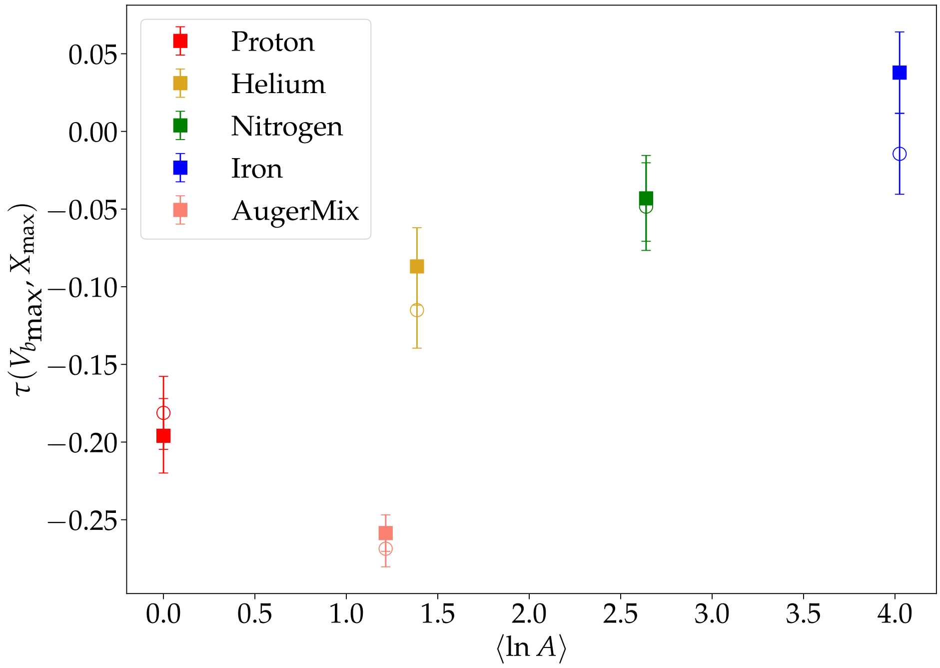

Figure 6.6: Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ as a function of the mean logarithmic mass $\langle \ln A \rangle$ for pure primary compositions and the AugerMix model in the energy bin $\log_{10}(E/\text{eV}) \in [18.0, 18.1)$ for EPOS-LHC high-energy hadronic interaction model. Filled markers correspond to reconstructed energy, while open markers show results obtained using true MC energy.

Fig. 6.6 shows a linear dependence of $\tau(V_{b_{\max}}, X_{\max})$ on the mean logarithmic mass $\langle \ln A \rangle$. Moving from light to heavy primaries, the correlation between $V_{b_{\max}}$ and $X_{\max}$ increases systematically, reflecting the mass dependence of shower development. Heavier primaries develop earlier in the atmosphere and produce larger muon content, which modifies the joint ordering of the two observables and therefore changes the value of $\tau(V_{b_{\max}}, X_{\max})$. The monotonic behaviour observed for the pure compositions indicates that $\tau(V_{b_{\max}}, X_{\max})$ is sensitive to the underlying mass scale of the primary particles. The AugerMix composition yields the most negative value of $\tau(V_{b_{\max}}, X_{\max})$, consistent with it representing a mixture of different primaries. This confirms that the correlation coefficient provides a composition-sensitive observable that can discriminate between different mass scenarios. Filled markers represent values obtained using reconstructed energy, while unfilled markers correspond to calculations performed using the true MC energy. The close agreement between the two indicates that the correlation measurement is only weakly affected by energy reconstruction uncertainties. This demonstrates that $\tau(V_{b_{\max}}, X_{\max})$ is primarily sensitive to the intrinsic shower development and composition rather than to the details of the energy reconstruction procedure. The same qualitative behaviour was observed in the remaining energy bins and is demonstrated in Fig. 6.7.

The left panel of Fig. 6.7 shows the dependence of $\tau(V_{b_{\max}}, X_{\max})$ on the primary mass scale $\ln A$ across multiple reconstructed energy bins. A consistent monotonic trend is ob-

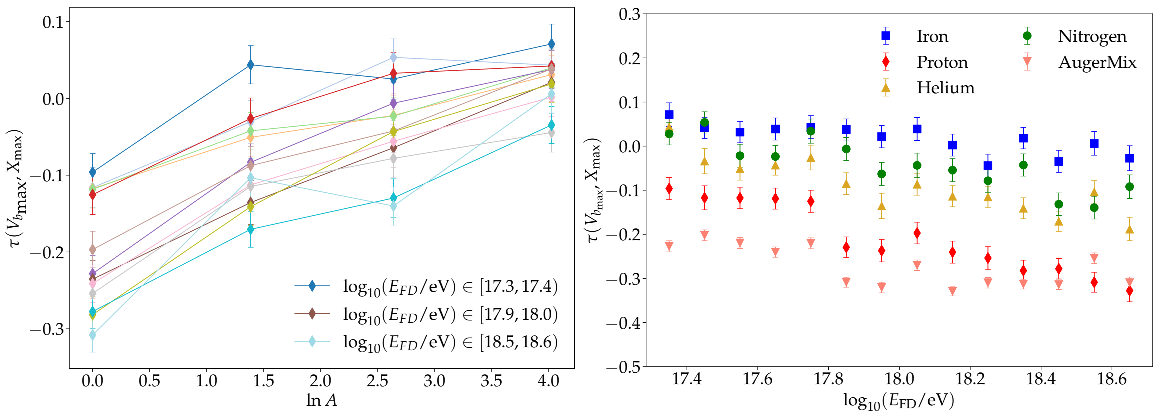

Figure 6.7: (Left) Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ as a function of the logarithmic mass $\ln A$ for different reconstructed energy bins. Each curve corresponds to a fixed energy interval, and error bars indicate statistical uncertainties. (Right) $\tau(V_{b_{\max}}, X_{\max})$ as a function of reconstructed energy for different primary compositions. Error bars represent statistical uncertainties obtained from bootstrap resampling. The simulated compositions use EPOS-LHC high-energy hadronic interaction model.

served across all energy intervals, with the correlation becoming progressively less negative as the primary mass increases. This behaviour reflects the systematic change in shower development with mass and confirms that the sensitivity of the correlation coefficient to composition is preserved over the full energy range considered. While the absolute value of τ varies slightly between energy bins, the overall ordering of primary masses and the slope of the τ–$\ln A$ relation remain similar. This indicates that the mapping between the correlation coefficient and the underlying mass scale is largely energy-independent. Consequently, $\tau(V_{b_{\max}}, X_{\max})$ provides a stable composition-sensitive observable that can be applied consistently across different energies without requiring strong energy-dependent recalibration.

Similarly, the right panel of Fig. 6.7 shows the dependence of $\tau(V_{b_{\max}}, X_{\max})$ on reconstructed energy for different primary compositions. Each composition exhibits a relatively stable correlation value across the full energy range, with only modest fluctuations within statistical uncertainties. At all energies, the ordering of the correlation coefficients follows the expected mass hierarchy, with lighter primaries showing less negative values and heavier or mixed compositions exhibiting stronger anti-correlation. The separation between compositions remains approximately constant with energy, indicating that the composition sensitivity of τ does not degrade at higher energies. This behaviour demonstrates that the correlation coefficient provides a stable discriminator of primary mass across the entire energy range considered.

Figure 6.8 shows the dependence of the correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ on the mass dispersion, quantified through $\sigma[\ln A]$, for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$ using the EPOS-LHC high-energy hadronic interaction model. Each simulated point corresponds to a specific mixture of primary nuclei (p, He, N, Fe), where the fractional abundance is varied in steps of 0.1 subject to the constraint $\sum_i f_i = 1$. As expected, pure compositions lie at $\sigma[\ln A] = 0$, while mixed compositions span a range of mass dispersions. The colour scale indicates the mean logarithmic mass $\langle \ln A \rangle$ of each mixture, with larger values corresponding to heavier average compositions. A clear monotonic trend is observed: the correlation coefficient becomes increasingly negative as the mass dispersion increases. Physically, this behaviour reflects the enhanced anti-correlation between $V_{b_{\max}}$ and $X_{\max}$ in mixed compositions, where event-by-event fluctuations in primary mass simultaneously broaden the $X_{\max}$ distribution

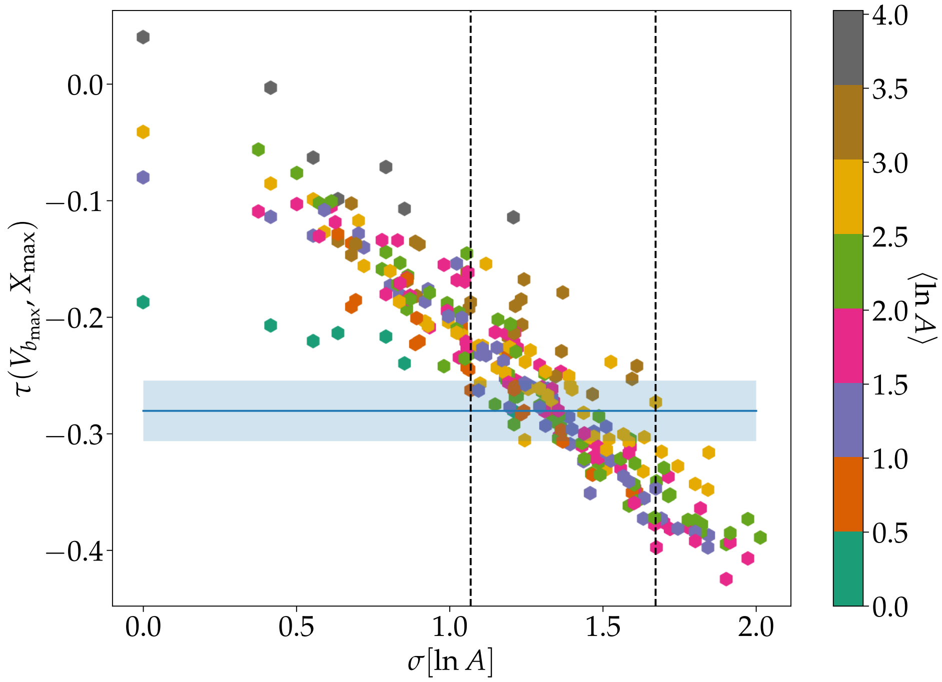

Figure 6.8: Dependence of $\tau(V_{b_{\max}}, X_{\max})$ on mass dispersion $\sigma[\ln A]$ for the energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$ using the EPOS-LHC high-energy hadronic interaction model. Each point corresponds to a simulated mixture of primary nuclei. The horizontal band indicates the predicted value for the AugerMix composition and its statistical uncertainty. The dashed black vertical lines correspond to the range of $\sigma[\ln A]$ in simulations compatible with the observed correlation in the AugerMix prediction.

and modify the muon content at ground level, strengthening the global ordering between the observables. The horizontal blue line indicates the value of $\tau(V_{b_{\max}}, X_{\max})$ obtained for the simulated AugerMix composition in this energy bin, and the shaded band represents its statistical uncertainty, taken as $0.67/\sqrt{N}$, where N is the number of events. The intersection of this band with the simulated grid defines the range of $\sigma[\ln A]$ values consistent with the measured correlation, providing a direct constraint on the allowed mass dispersion of the primary CR composition. The dashed black vertical lines correspond to the range of $\sigma[\ln A]$ in simulations compatible with the observed correlation in the AugerMix prediction. A comparison with the AugerMix prediction indicates significant mixing of primary masses. Specifically, $\sigma[\ln A] = 1.37 \pm 0.3$ with values of $\sigma[\ln A] \in [1.07, 1.67]$ for EPOS-LHC high-energy hadronic interaction model.

6.4 Determination of the minimum number of events required to reject a pure composition hypothesis

6.4.1 Statistical separation of correlation coefficients

The sensitivity of the observable under consideration is quantified through the difference between the correlation coefficient measured for a pure primary composition and that measured for a mixed composition. The estimated correlation coefficients for the pure and mixed composition scenarios, denoted $\hat{\tau}_{\mathrm{pure}}$ and $\hat{\tau}_{\mathrm{mix}}$, respectively, are obtained from finite samples of events and therefore exhibit statistical fluctuations. The sampling distributions of the correlation estimators are evaluated using bootstrap resampling. For sufficiently large samples, the bootstrap distributions are well approximated by Gaussian distributions centred on their expectation values. The Gaussian behaviour of τ within a fixed energy bin is verified in Appendix C.2. Accordingly, the estimators are modelled as,

$$\hat{\tau}_{\text{pure}} \sim \mathcal{N}(\mu_{\text{pure}}, \sigma^2_{\text{pure}}) \quad \text{and} \quad \hat{\tau}_{\text{mix}} \sim \mathcal{N}(\mu_{\text{mix}}, \sigma^2_{\text{mix}}), \tag{6.2}$$

where the means represent the expected values of the correlation coefficients and the variances represent their statistical uncertainties.

The quantity relevant for hypothesis testing is the difference between the two estimators, $\Delta\hat{\tau} = \hat{\tau}_{\text{pure}} - \hat{\tau}_{\text{mix}}$. Since the difference of two independent Gaussian random variables is also Gaussian, i.e., $\Delta\hat{\tau} \sim \mathcal{N}(\langle\Delta\hat{\tau}\rangle, \sigma_{\Delta\hat{\tau}})$, with

$$\langle\Delta\hat{\tau}\rangle = \mu_{\text{pure}} - \mu_{\text{mix}} \quad \text{and} \quad \sigma^2_{\Delta\hat{\tau}} = \sigma^2_{\text{pure}} + \sigma^2_{\text{mix}}. \tag{6.3}$$

The expected separation between the pure and mixed composition hypotheses is therefore quantified by the mean difference $\langle\Delta\hat{\tau}\rangle$. In contrast, the statistical uncertainty of this separation is given by $\sigma_{\Delta\hat{\tau}}$. The pure composition considered in this study is proton since it has the smallest value of τ.

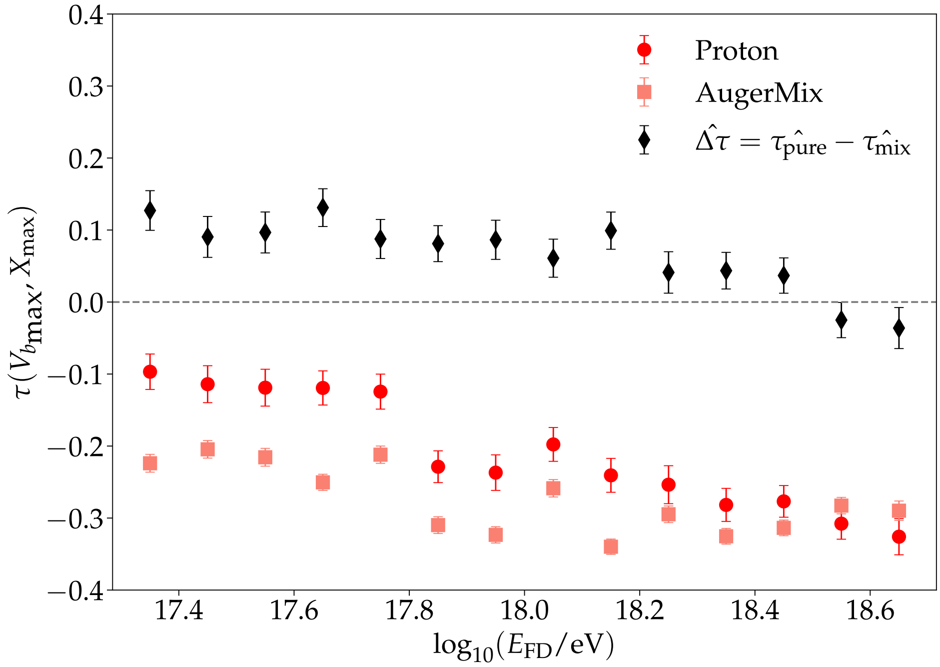

Figure 6.9: Correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ as a function of reconstructed energy for proton and AugerMix compositions, together with their difference $\Delta\hat{\tau} = \hat{\tau}_{\text{pure}} - \hat{\tau}_{\text{mix}}$. Error bars represent statistical uncertainties obtained from bootstrap resampling.

Figure 6.9 shows the energy dependence of the correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ for proton and AugerMix compositions, together with their difference $\Delta\hat{\tau} = \hat{\tau}_{\text{pure}} - \hat{\tau}_{\text{mix}}$. Over most of the energy range, the proton sample exhibits systematically less negative correlation values than the mixed composition, resulting in a positive $\Delta\hat{\tau}$. This indicates that the ordering relationship between $V_{b_{\max}}$ and $X_{\max}$ differs measurably between the two composition scenarios, providing statistical leverage to distinguish between them. The magnitude of $\Delta\hat{\tau}$ varies with energy, reflecting changes in shower development and composition sensitivity across the energy spectrum. Regions where $\Delta\hat{\tau}$ is large correspond to stronger separability and therefore require fewer events to reject the pure-composition hypothesis. At the highest energies, however, $\Delta\hat{\tau}$ decreases and eventually becomes consistent with zero or negative values within statistical uncertainties. When $\Delta\hat{\tau}$ approaches zero, the required sample size diverges, reflecting the fact that no finite dataset can reliably distinguish between the two hypotheses using this observable alone. A negative $\Delta\hat{\tau}$ implies that the correlation measured for the mixed composition is equal to or more positive (i.e., less negative) than that for the pure proton sample. In this regime, the two hypotheses are not statistically ordered in the expected direction, and the correlation observable loses discriminatory power. Consequently, the inferred separation between composition scenarios becomes small, leading to a rapid increase in the number of events required for a statistically significant test. Moreover, this regime complicates the physical interpretation, as it indicates that contributions from heavier nuclei beyond protons must be explicitly taken into account to properly describe the observed correlations. For this reason, in the present work, only the regime $\Delta\hat{\tau} > 0$ is considered, where the correlation ordering follows the expected physical behaviour and the observable retains its discriminatory power.

6.4.2 Expected number of observed events

The statistical requirement derived in the previous section must be compared with the number of events expected to be recorded by the detector. For an energy interval $[E_i, E_{i+1}]$, the number of detected events is given by,

$$N_{\mathrm{obs},i} = \varepsilon \times T \times 2\pi \times \int_0^{\theta_{\max}} \sin\theta \cos\theta \, d\theta \times \int_{E_i}^{E_{i+1}} J(E)\, dE \times A_{\mathrm{eff}}, \tag{6.4}$$

where $\theta_{\max} = 45^\circ$ is the maximum zenith angle considered. The azimuthal integral yields a factor 2π, and the angular integral accounts for the projection of the detector area and the geometric acceptance. In this work, the detector efficiency is taken as $\varepsilon = 0.1$ for the FD [67], the effective area is $A_{\mathrm{eff}} = 24\ \mathrm{km}^2$, and the exposure time is $T = 5$ or 10 years. The differential CR flux $J(E)$ is modelled in Eq. (4.4). The energy bins are chosen to be regularly sized in decimal logarithm, $\Delta \log E_i = 0.1$

6.4.3 Definition of the minimum number of events

To determine the statistical power required to reject the pure-composition hypothesis, a target significance level must be specified. In this work, a separation corresponding to five standard deviations is adopted. The statistical uncertainty of the correlation estimator is empirically found to scale with sample size as $\sigma = 0.67/\sqrt{N}$ (see the right panel of Fig. 6.2), where N is the number of events.

Requiring that the expected separation between the pure and mixed scenarios corresponds to a 5σ detection implies that the signal-to-noise ratio satisfies $\Delta\tau/\sigma_\tau = 5$. Substituting the empirical scaling of the statistical uncertainty gives,

$$\Delta\tau = \frac{A}{\sqrt{N}}. \tag{6.5}$$

where $A = 5 \times 0.67$.

Solving for the minimum number of events required to achieve this separation yields,

$$\sqrt{N_{\min}} = \frac{A}{\Delta\tau}. \tag{6.6}$$

Equation (6.6) shows that the minimum sample size required to reject the pure-composition hypothesis is inversely proportional to the expected separation between the correlation coefficients of the competing composition scenarios.

The separation $\Delta\tau$ is not known exactly since it is estimated from a finite sample of events, but it can be approximated as a Gaussian random variable. Consequently, the minimum required number of events cannot be treated as a deterministic quantity. Instead, it must also be regarded as a random variable derived from the transformation,

$$x = \sqrt{N_{\mathrm{min}}} = \frac{A}{\Delta\tau}. \tag{6.7}$$

Because this transformation involves the reciprocal of a Gaussian variable, the distribution of x is non-Gaussian and asymmetric. Proper statistical inference, therefore, requires the determination of the full probability distribution of x. The probability density function of x is obtained by applying the standard change-of-variables formula to the Gaussian distribution of $\Delta\hat{\tau}$. If $\Delta\hat{\tau} \sim \mathcal{N}(\langle\Delta\hat{\tau}\rangle, \sigma^2_{\Delta\hat{\tau}})$ has density f, then for $x \neq 0$,

$$f(x) = f\left(\frac{A}{x}\right)\left|\frac{d}{dx}\left(\frac{A}{x}\right)\right|. \tag{6.8}$$

Evaluating the derivative and substituting the Gaussian form gives,

$$f(x) = \begin{cases} \frac{A}{x^2}\frac{1}{\sqrt{2\pi}\sigma_{\Delta\hat{\tau}}}\exp\left[-\frac{\left(\frac{A}{x}-\langle\Delta\hat{\tau}\rangle\right)^2}{2\sigma^2_{\Delta\hat{\tau}}}\right] & \text{if} \quad x \neq 0 \\ 0 & \text{if} \quad x = 0 \end{cases} \tag{6.9}$$

This distribution is strongly right-skewed because small values of $\Delta\tau$ correspond to very large required event counts. In Appendix C.3, the agreement between the numerical and analytical curves, confirming that the reciprocal transformation of a normally distributed variable produces the expected asymmetric, right-skewed distribution, is verified. Consequently, symmetric error propagation is not appropriate.

To report a central estimate and uncertainty for the minimum event requirement, the cumulative distribution function of x, $F(x)$, is constructed by integrating the probability density. The integral to calculate $F(x)$ (see Appendix C.4) can be evaluated analytically in terms of the error function. The median of the distribution is obtained by solving $F(x_{50}) = 0.5$, and the statistical uncertainty is defined using the central 68.27% probability interval, $F(x_{16}) = 0.15865$ and $F(x_{84}) = 0.84135$. The minimum number of events is therefore reported as

$$\sqrt{N_{\mathrm{min}}} = k^{+(x_{84}-k)}_{-(k-x_{16})}, \tag{6.10}$$

where $k = x_{50}$. This procedure fully accounts for the non-Gaussian nature of the transformation and provides an unbiased estimate of the required statistical power. The value of $\sqrt{N_{\mathrm{min}}}$ depends on the statistical separation between the correlation coefficients predicted for pure and mixed compositions.

Equation (6.4) provides the expected number of detected events as a function of energy. Its square root is compared with the inferred value of $\sqrt{N_{\mathrm{min}}}$ in Eq. (6.10) obtained from the statistical separation analysis, as shown in Fig. 6.10. The quantity N_{min} represents the minimum number of events required in a given energy bin to reject the hypothesis of a pure primary composition at the chosen confidence level using the correlation observable $\tau(V_{b_{\mathrm{max}}}, X_{\mathrm{max}})$. If the number of observed events satisfies $N_{\mathrm{obs}} \geq N_{\mathrm{min}}$, the experiment has sufficient statistical power to discriminate between a pure composition scenario and a mixed composition. Conversely, if the available statistics fall below this threshold, the observable lacks sufficient discriminatory power in that energy interval. The blue curve in Fig. 6.10 represents a polynomial fit to the median estimate of $\sqrt{N_{\mathrm{min}}}$ obtained from the statistical analysis performed in each energy bin. The shaded blue band corresponds to the

1σ uncertainty of the fitted curve, derived from the covariance matrix of the fit parameters. This band, therefore, reflects the statistical uncertainty in the inferred minimum number of events required for composition discrimination. The orange and green curves show the expected statistics, expressed as $\sqrt{N_{\mathrm{obs}}}$, for 5 and 10 years of exposure, respectively, and are calculated from Eq. (6.4). Because the CR flux decreases rapidly with energy, the number of detected events decreases accordingly, resulting in a declining $\sqrt{N_{\mathrm{obs}}}$ with increasing energy. The energy range below the expected statistics curves and the fitted $\sqrt{N_{\mathrm{min}}}$ curve defines the energy range over which the correlation observable provides sufficient statistical power to reject the pure-composition hypothesis. The dashed vertical lines indicate the energies at which the expected event statistics become equal to the required minimum number of events. Below these energies, the experiment collects enough events to statistically distinguish between composition scenarios, whereas above these energies, the available statistics become insufficient. This comparison directly identifies the upper limit in energy below which the correlation observable provides sufficient statistical sensitivity to reject a pure primary composition, linking the theoretical discrimination power of the observable with the practical limitations imposed by the available event statistics. The results indicate that for $\log_{10}(E/\mathrm{eV}) \lesssim 17.9$, the expected statistics from 10 years of Pierre Auger Observatory data exceed the required threshold, allowing the pure-composition hypothesis to be rejected using the correlation observable. This region is represented by the grey shaded band in Fig. 6.10. At higher energies, however, the rapidly decreasing CR flux reduces the number of detected hybrid events, and the statistical uncertainty of Kendall's correlation coefficient increases. Consequently, the minimum number of events required for discrimination becomes larger than the number of events expected from the detector exposure, limiting the sensitivity of the method in this energy range. Fig. 6.10 shows that the discrimination power of the $\tau(V_{b_{\mathrm{max}}}, X_{\mathrm{max}})$ observable is not limited by intrinsic shower fluctuations in this energy range, but primarily by the available statistics of hybrid events.

Only events with reconstructed energies $E_{\mathrm{FD}} \geq 10^{17.5}$ eV are considered suitable for analysis for the SD–750 array. This requirement ensures that the SD and UMD in the array operate with full trigger efficiency and avoids potential composition-dependent biases that may arise at lower energies [61].

6.5 Kendall's correlation coefficient from data

To perform the correlation analysis using real events, the FD–SD Golden hybrid Phase 1 and Phase 2 datasets reconstructed for ICRC 2025 were matched event-by-event with the UMD reconstructed Phase 1 and Phase 2 datasets described in Section 4.1.1. This matching allows the depth of the shower maximum measured by the FD, X_{max}, to be matched directly with the corresponding $V_{b_{\mathrm{max}}}$ reconstructed at ground level using UMD. The resulting data set enables a composition-sensitive analysis between X_{max} and $V_{b_{\mathrm{max}}}$ from real data.

Standard quality cuts [77] were applied to ensure reliable event reconstruction, including fiducial area requirements, zenith angle cuts, detector status filters, and minimum X_{max} constraints. As in simulated events, using the muon density reconstructed from the data, the event-level observable $V_{b_{\mathrm{max}}}$ was calculated using an energy-dependent slope parameter $b_{\mathrm{max}}(\log_{10}(E/\mathrm{eV}))$ according to Eq. (5.5). The total $V_{b_{\mathrm{max}}}$ per event was then obtained by summing the contributions of all triggered stations.

The analysis was performed in fixed logarithmic energy intervals of width $\Delta \log_{10}(E/\mathrm{eV}) = 0.1$. For each energy bin, the distribution of $V_{b_{\mathrm{max}}}$ as a function of X_{max} in real data was compared with simulated showers generated with pure proton, helium, nitrogen, and iron primaries. Figure 6.11 shows the event-by-event distributions of $V_{b_{\mathrm{max}}}$ as a function of X_{max} for selected energy bins of measured data and proton and iron primaries obtained from

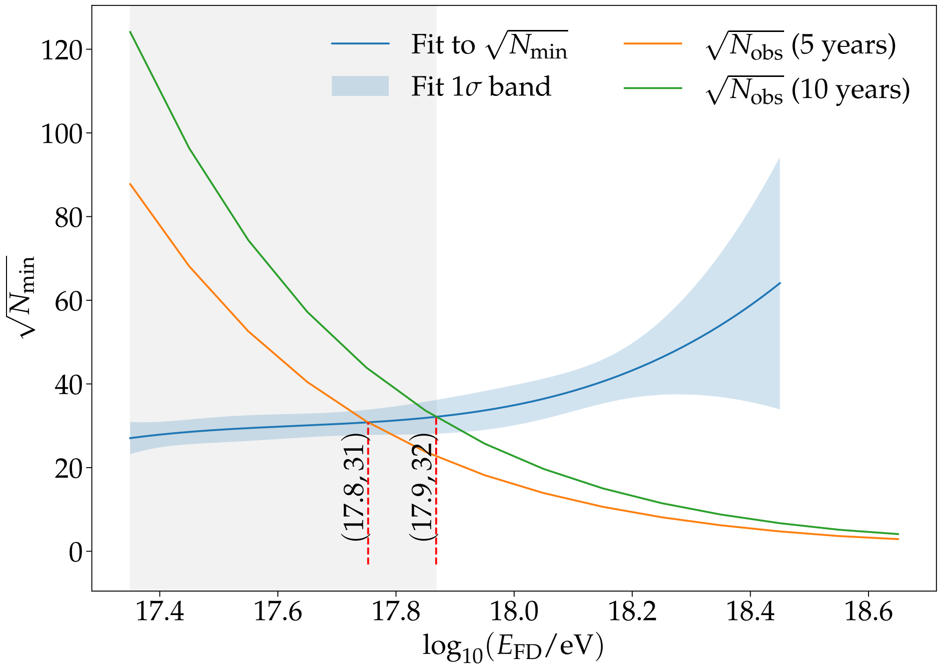

Figure 6.10: The square root of the number of hybrid events expected to be measured by the detectors of Pierre Auger Observatory plotted against the square root of the minimum number of events required to reject the hypothesis of a pure composition with 5σ significance. The blue curve represents a polynomial fit to the median estimate of $\sqrt{N_{\rm min}}$ obtained from the statistical analysis performed in each energy bin. The shaded blue band corresponds to the 1σ uncertainty of the fitted curve, derived from the covariance matrix of the fit parameters. The orange and green curves show the expected statistics, expressed as $\sqrt{N_{\rm obs}}$, for 5 and 10 years of exposure, respectively. The grey shaded band indicates the region below which the expected statistics from 10 years of Pierre Auger Observatory data exceed the required threshold, allowing the pure-composition hypothesis to be rejected using the correlation observable. $\sqrt{N_{\rm min}}$ is calculated for EPOS-LHC high-energy hadronic interaction model.

simulations using EPOS-LHC high-energy hadronic interaction model. To summarise the distribution of simulated events in the $(X_{\rm max}, V_{b_{\rm max}})$ plane, covariance ellipses were drawn for proton and iron primaries. For each simulated sample, the mean vector and the 2×2 covariance matrix were computed from the event-by-event values of $X_{\rm max}$ and $V_{b_{\rm max}}$. The covariance matrix was diagonalised to determine the principal axes of the distribution. The resulting ellipse is centred at the sample mean, oriented along the eigenvectors of the covariance matrix, and its semi-axes correspond to the square roots of the covariance eigenvalues. These contours, therefore, represent the characteristic 1σ-scale spread of the event distribution in the principal-axis basis, including the correlation between $X_{\rm max}$ and $V_{b_{\rm max}}$. The number of golden hybrid events in each energy bin is indicated in the figure labels. At higher energies, the number of available hybrid events decreases due to the steeply falling CR flux, which leads to larger statistical uncertainties in the calculated τ values. Due to the limited number of hybrid events available in individual energy bins, the resulting correlation coefficient does not provide sufficient statistical power to discriminate strongly between pure-composition scenarios (see Fig. 6.10). To mitigate this limitation, the solution would be to repeat the analysis with larger logarithmic energy bin widths. However, as described in Section 6.2.3, increasing the energy bin width is unsuitable for the correlation analysis.

Fig. 6.12 shows the dependence of the correlation estimator $\tau(V_{b_{\rm max}}, X_{\rm max})$ on the mean logarithmic mass $\langle \ln A \rangle$ for the energy bin $\log_{10}(E/{\rm eV}) \in [17.5, 17.6)$. The markers represent

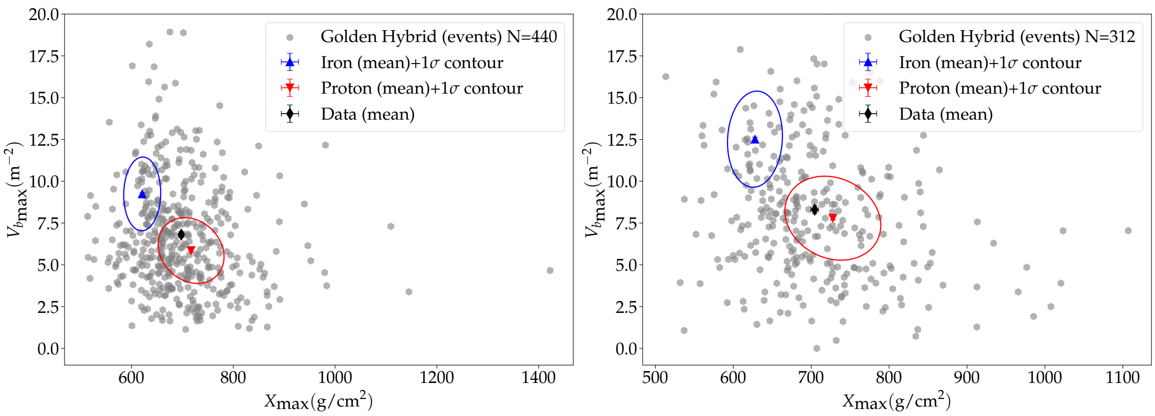

Figure 6.11: (Left) Event-by-event distribution of $V_{b_{\max}}$ as a function of $X_{\max}$ in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ (left) and $\log_{10}(E/\mathrm{eV}) \in [17.6, 17.7)$ (right), showing real Golden hybrid data together with proton and iron primaries simulated with EPOS-LHC high-energy hadronic interaction model.

the values obtained from simulations for pure proton, helium, nitrogen, and iron primaries. Different colours correspond to the three high-energy hadronic interaction models used in the simulations: EPOS-LHC, QGSJet-II.04, and Sibyll-2.3c. The horizontal black line indicates the value of $\tau(V_{b_{\max}}, X_{\max})$ measured from data, while the grey band represents the statistical uncertainty obtained from bootstrap resampling.

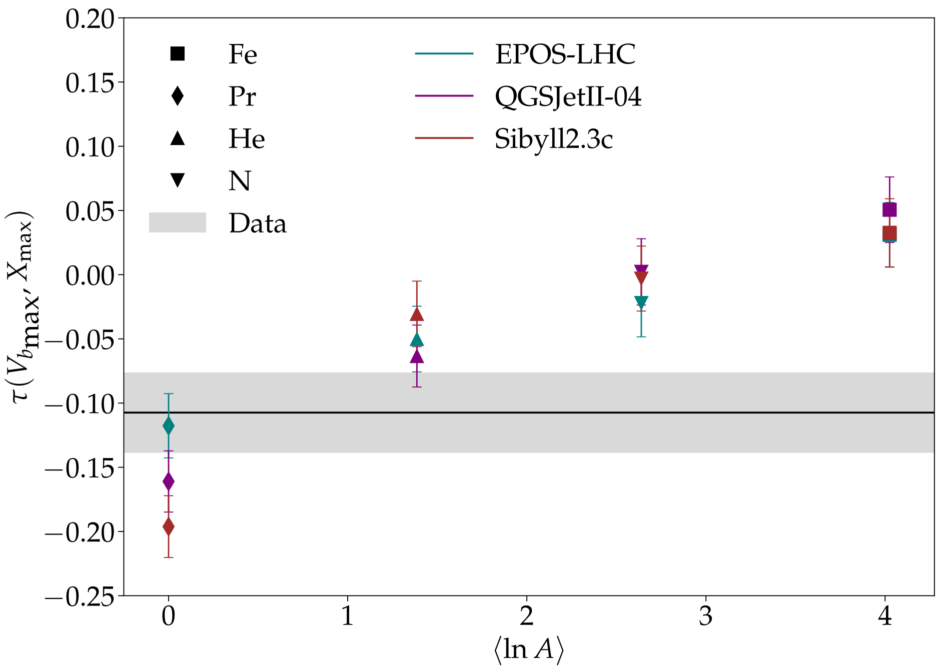

Figure 6.12: $\tau(V_{b_{\max}}, X_{\max})$ as a function of the mean logarithmic mass $\langle \ln A \rangle$ for the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$. The simulated primaries (p, He, N, and Fe) are shown in different markers, where the colours correspond to different high-energy hadronic interaction models. The $\tau(V_{b_{\max}}, X_{\max})$ corresponding to the data is shown in a black line with the errors in a grey band.

In the present case, the statistical uncertainty of the data measurement remains sufficiently large that none of the models or composition scenarios can be excluded within the considered energy interval. Moreover, the number of Golden Hybrid events available for this energy bin is $N = 440$, as shown in the left panel of Fig. 6.11. This sample size is insufficient to reject the hypothesis of a pure composition according to the minimum number of events required for discrimination, calculated in Eq. (6.6) and shown in Fig. 6.10.

6.6 Summary and conclusions

This chapter investigated the use of the Kendall rank correlation coefficient $\tau(V_{b_{\max}}, X_{\max})$ between the depth of shower maximum $X_{\max}$ and the muon-related observable $V_{b_{\max}}$ as a composition-sensitive observable. Among the correlation estimators considered, τ was identified as the most suitable metric due to its robustness to outliers, weak dependence on the underlying distribution of the variables considered, and comparatively small statistical uncertainty. Its statistical fluctuations were shown to scale as $1/\sqrt{N}$, and bootstrap studies demonstrated that its sampling distribution is well approximated by a Gaussian for fixed energy. The estimator was also found to exhibit minimal sensitivity to the choice of high-energy hadronic interaction model, making it a stable observable for composition studies.

The dependence of $\tau(V_{b_{\max}}, X_{\max})$ on primary mass was examined using simulated samples of proton, helium, nitrogen, and iron primaries, as well as an Auger-inspired mixed composition. A monotonic relationship between $\tau(V_{b_{\max}}, X_{\max})$ and the mean logarithmic mass $\langle \ln A \rangle$ was observed, confirming that the correlation between $X_{\max}$ and $V_{b_{\max}}$ encodes information about the mass scale of the primary particle. This enables statistical discrimination between pure and mixed-composition scenarios.

The sensitivity of the method was quantified by modelling the difference in correlation coefficients between pure and mixed compositions, accounting for statistical fluctuations using bootstrap-derived uncertainties. The minimum number of events required to reject the pure-composition hypothesis at a 5σ level was derived, and the uncertainty on this requirement was obtained through the propagation of the Gaussian uncertainty on the correlation coefficient. Because the required event count depends on the reciprocal of the separation in correlation coefficients, the resulting distribution of the minimum event requirement is asymmetric and was treated using a reciprocal-normal model.

Finally, the statistical requirement was compared with the number of events expected to be recorded by Pierre Auger Observatory as a function of energy. This comparison shows that the available statistics exceed the required minimum below $\log_{10}(E/\mathrm{eV}) \lesssim 17.9$ for an exposure of 10 years. In this energy range, the correlation-based method provides sufficient statistical power to distinguish between a well-mixed composition containing both light and heavy primaries and a relatively pure composition dominated by nuclei of similar mass for simulations generated with EPOS-LHC high-energy hadronic interaction model.

An exploratory application of the method to real FD–SD Golden hybrid events was also performed by computing Kendall's correlation coefficient between $X_{\max}$ and $V_{b_{\max}}$ in fixed energy bins with expectations from simulated pure primaries. Within current statistics, the statistical uncertainty of $\tau(V_{b_{\max}}, X_{\max})$ computed for Golden hybrid data remains sufficiently large to provide a statistically significant exclusion of pure-composition scenarios. This result is consistent with the minimum event requirement derived earlier and confirms that the present hybrid dataset does not yet provide sufficient statistical power for a definitive composition determination using the $X_{\max}$–$V_{b_{\max}}$ correlation alone. Nevertheless, the observed behaviour demonstrates the feasibility of the method and provides a direct empirical consistency check of the simulation-based expectations.

Table 6.1: Comparison of correlation coefficients considered in this work. The table summarises statistical assumptions, robustness properties, and practical interpretation for experimental data.

Property	r_p	r_s	r_{GH}	τ
Data used	Values	Ranks	Ranks	Pairwise ranks
Association measured	Linear	Monotonic	Rank-order deviation	Concordance
Distribution assumption	Normal (approx.)	None	None	None
Sensitivity to outliers	High	Moderate	Low	Low
Sensitivity to nonlinearity	High	Low	Low	Low
Uses metric information	Yes	No	No	No
Robust to skewed data	No	Yes	Yes	Yes
Range	$[-1, 1]$	$[-1, 1]$	$[-1, 1]$	$[-1, 1]$
Efficiency (Gaussian data)	Highest	High	Lower	High
Robustness (non-normal data)	Poor	Good	Very good	Very good

CHAPTER 7

Correlation based composition analysis using X_{max} measured by surface detector

In the previous chapter, the mass composition of UHECR was investigated using the correlation between the muon-related observable $V_{b_{\mathrm{max}}}$ and the maximum depth of the shower X_{max} measured by the FD. Although FD measurements provide precise information on the longitudinal development of air showers, their operation is limited to clear nights without the moon, resulting in a duty cycle of approximately 15%. In contrast, the SD operates continuously, providing wide-sky coverage and significantly larger event statistics. Therefore, extension of correlation-based composition methods to an SD-based framework provides an opportunity to substantially improve statistical precision in composition studies.

Estimating X_{max} using only SD observables is intrinsically challenging. Unlike FD measurements, which directly probe the longitudinal development of the electromagnetic component, SD-based estimators rely on indirect correlations between ground-level signals and shower evolution. As a result, such estimators are affected not only by finite resolution but also by reconstruction biases that depend on shower properties, including the primary mass. The early approaches developed by Pierre Auger Observatory exploit the risetime of SD signals as a proxy for the development of showers [68]. The risetime is sensitive to the temporal spread of particles at ground, which in turn depends on the relative contributions of the electromagnetic and muonic components. As explained in Section 5.3.6, the mapping between risetime and X_{max} introduces a composition-dependent bias. Consequently, the reconstructed X_{max} cannot be interpreted as a simple smeared version of the true value, but rather as a biased estimator whose response varies with primary mass. This has important implications for correlation studies, since in this case, the measured correlation coefficient does not purely reflect the underlying physics of shower development but is contaminated by the reconstruction procedure. This argument extends naturally to rank-based correlation estimators such as Kendall's τ, which is used throughout this work. Kendall's τ is defined through the relative ordering of pairs of events and is therefore invariant under strictly monotonic transformations of the observables. However, this invariance does not protect against reconstruction biases that alter the event ordering itself. In the presence of composition-dependent bias, the transformation from $X_{\mathrm{max}}^{\mathrm{MC}}$ to $X_{\mathrm{max}}^{\mathrm{rec}}$ is not globally monotonic throughout the full data set. Instead, it depends on the primary mass and shower properties, thereby introducing mixing of orderings across different subsets of events. As a result, pairs of events that are concordant (or discordant) in terms of $(V_{b_{\mathrm{max}}}, X_{\mathrm{max}}^{\mathrm{MC}})$ may change their ordering when expressed in terms of $(V_{b_{\mathrm{max}}}, X_{\mathrm{max}}^{\mathrm{rec}})$. This leads to a modification of Kendall's τ that cannot be interpreted solely as an effect of finite resolution.

By contrast, if the reconstruction is unbiased, then the impact on Kendall's τ is well defined. In this case, the smearing introduces random rank inversions whose probability increases with resolution $\sigma[X_{\max}^{\rm SD}]$, leading to a gradual dilution of the correlation strength. Crucially, this effect is symmetric and does not introduce artificial correlations, but only reduces the magnitude of τ. Therefore, the degradation of Kendall's τ can be directly interpreted in terms of detector resolution.

More recent Pierre Auger Observatory developments aim to overcome the limitations of risetime-based estimators. Methods based on air-shower universality provide a physically motivated mapping between SD observables and $X_{\max}$ [79], while deep neural network (DNN) approaches exploit the full time traces recorded by the WCD [67]. These techniques significantly improve resolution and extend $X_{\max}$ measurements to higher energies. However, even with these advanced approaches, the reconstructed $X_{\max}$ still shows a residual dependence on composition, and dedicated selections are required to control biases. Furthermore, existing SD-based $X_{\max}$ reconstructions have been developed for the 1500 m array. Their direct application to the 750 m array, which probes lower energies and different shower geometries, would require dedicated training and validation. In the absence of such calibration, this work adopts a phenomenological approach in which the true MC value of $X_{\max}$ is smeared with a Gaussian resolution;

$$X_{\max}^{\rm SD} = X_{\max}^{\rm MC} + \epsilon_{X_{\max}} \quad \text{where} \quad \epsilon_{X_{\max}} \sim \mathcal{N}\left(0, \sigma^2_{X_{\max}^{\rm SD}}\right). \tag{7.1}$$

This ensures that the transformation from $X_{\max}^{\rm MC}$ to $X_{\max}^{\rm smeared}$ is unbiased and preserves the statistical independence between the noise term and $V_{b_{\max}}$. As a consequence, any observed reduction in Kendall's τ can be unambiguously attributed to finite resolution effects, rather than to estimator-induced biases. Three representative resolutions are considered to bracket the performance expected from realistic SD-based reconstructions at infill energies. This approach provides a controlled framework to quantify how detector resolution degrades the sensitivity of correlation-based composition observables.

The primary objective of this chapter is to investigate whether Kendall's rank correlation coefficient between $V_{b_{\max}}$ and $X_{\max}$ remains a robust composition-sensitive observable when $X_{\max}$ is measured with a realistic SD-like resolution. In particular, the analysis examines how detector resolution affects the strength of the correlation, its statistical properties, and its sensitivity to the primary mass distribution. Both scenarios of pure and mixed composition are considered, allowing the dependence of the estimator on the mean logarithmic mass $\langle \ln A \rangle$ and the mass dispersion $\sigma[\ln A]$ to be explored. All the analyses in this chapter are completely simulation-based.

Beyond assessing sensitivity to composition, the statistical behaviour of the estimator is also examined, including its uncertainty scaling and its discriminating power as a function of event statistics. By comparing the correlation coefficient of different simulated composition mixtures obtained with a realistic SD-like resolution that spans a range of mean masses and dispersions, the analysis evaluates the extent to which SD-based measurements can constrain the primary mass distribution.

Therefore, this chapter provides a systematic investigation of correlation-based composition analysis in a SD context. The results establish the performance of the Kendall correlation observable under realistic resolution conditions and assess its potential as a high-statistics tool for CR mass composition studies. The structure of this chapter is as follows. Section 7.1 introduces the construction of an SD-like observable for the depth of shower maximum by applying controlled Gaussian smearing to the MC $X_{\max}$ and examines the resulting joint distributions of $X_{\max}$ and $V_{b_{\max}}$ under different resolution scenarios. Section 7.1 also investigates the dependence of τ obtained with different resolutions on $X_{\max}$ in relation to the composition mixture and detector resolution, as well as its statistical behaviour. Section 7.2 explores the relationship between the correlation coefficient and the mean logarithmic mass $\langle \ln A \rangle$, demonstrating the sensitivity of the estimator to the total mass scale of the primary composition. Finally, Section 7.3 summarises the main findings and discusses the implications for studies of high-statistics composition using SD measurements.

7.1 Kendall's correlation coefficient between $X_{\max}^{\mathrm{SD}}$ and $V_{b_{\max}}$

As discussed in Chapter 5, the determination of $X_{\max}$ from SD data using the $\langle \Delta \rangle$ method exhibits a bias towards heavier primary compositions and is therefore not directly suitable for composition inference. To emulate a measurement of $X_{\max}$ similar to SD while avoiding model-dependent reconstruction biases, a Gaussian smearing is applied directly to the MC shower value of $X_{\max}$. The smearing is implemented with a zero mean and fixed standard deviations of 20, 40, and 60 $\mathrm{g\,cm^{-2}}$. The value of 20 $\mathrm{g\,cm^{-2}}$ corresponds to the typical resolution achieved by the measurements of FD in Pierre Auger Observatory [59], providing an optimistic reference scenario. The values of 60 $\mathrm{g\,cm^{-2}}$ are motivated by the resolution obtained in reconstruction studies based on $X_{\max}$ obtained with SD using $\langle \Delta \rangle$ method [68], and represent more realistic and conservative performance scenarios for measurements similar to SD. The value of 40 $\mathrm{g\,cm^{-2}}$ corresponds to an intermediate value. The resulting fluctuating observable is denoted $X_{\max}^{\mathrm{SD}}$.

The joint distribution of the observables $X_{\max}^{\mathrm{SD}}$ and $V_{b_{\max}}$ for proton and iron primaries in a single energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$ is shown in the right panel of Fig. 7.1. The energy considered in this analysis corresponds to the corrected reconstructed energy described in Section 4.2. Restricting the analysis to a narrow energy interval minimises residual energy-dependent effects, ensuring that the observed correlation structure reflects intrinsic shower-to-shower fluctuations rather than variations driven by energy. The simulations used to generate these data are described in Section 4.1.2.

A clear separation between the proton and iron populations is visible in the $X_{\max}^{\mathrm{SD}}$–$V_{b_{\max}}$ plane. As described in Chapter 2, showers initiated by heavier primaries develop earlier in the atmosphere, leading to smaller values of $X_{\max}^{\mathrm{SD}}$ and producing a larger number of muons at ground level, resulting in higher values of $V_{b_{\max}}$. The earlier development arises both from the larger interaction cross section of heavy nuclei and, more importantly, from the smaller energy per nucleon in the superposition picture, which causes the electromagnetic cascade to reach its maximum at smaller atmospheric depths. Conversely, lighter primaries penetrate deeper before reaching the shower maximum and generate fewer muons. This opposite mass dependence produces a characteristic anti-correlation between $X_{\max}^{\mathrm{SD}}$ and $V_{b_{\max}}$, which forms the basis for using their correlation as a composition-sensitive observable.

Using the simulated data sets described in Section 4.1, the correlation between $X_{\max}^{\mathrm{SD}}$ and $V_{b_{\max}}$ was evaluated using Kendall's τ, following the methodology of Section 6.1. The analysis was performed for different $X_{\max}$ resolutions to quantify the impact of $X_{\max}^{\mathrm{SD}}$ resolution on the estimator. All calculations were carried out within a single reconstructed energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ and repeated across additional bins to verify that the observed behaviour is not driven by residual energy dependence.

The left panel of Fig. 7.2 shows Kendall's rank correlation coefficient $\tau(V_{b_{\max}}, X_{\max}^{\mathrm{SD}})$ as a function of the proton fraction C_p defined in Eq. (6.1) for a two-component composition mixture. The results are presented for three assumed resolutions of the maximum depth of the shower, $X_{\max}^{\mathrm{SD}}$, indicated by purple (20 $\mathrm{g\,cm^{-2}}$), red (40 $\mathrm{g\,cm^{-2}}$) and green (60 $\mathrm{g\,cm^{-2}}$) markers. The error bars represent bootstrap uncertainties.

For all resolutions, the correlation is negative across the full composition range in the considered energy bin, indicating a robust anti-correlation between the surface observable $V_{b_{\max}}$ and the longitudinal development parameter $X_{\max}^{\mathrm{SD}}$. The magnitude of the anti-correlation

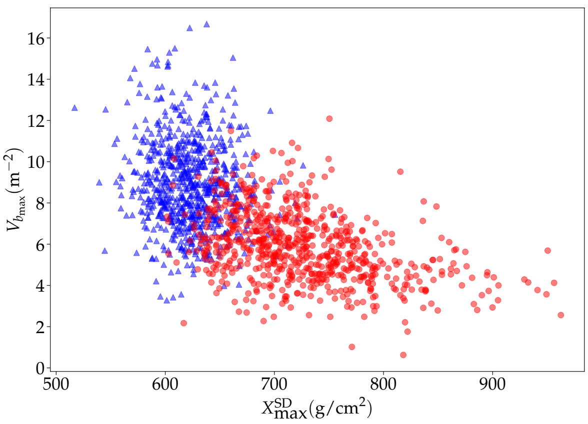

X_{max} resolution = 20 g cm^{-2}

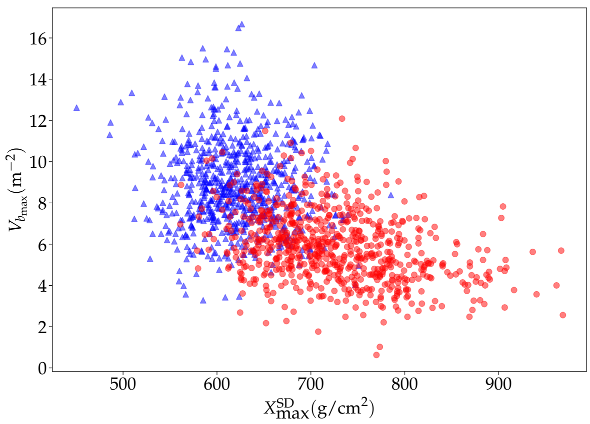

X_{max} resolution = 40 g cm^{-2}

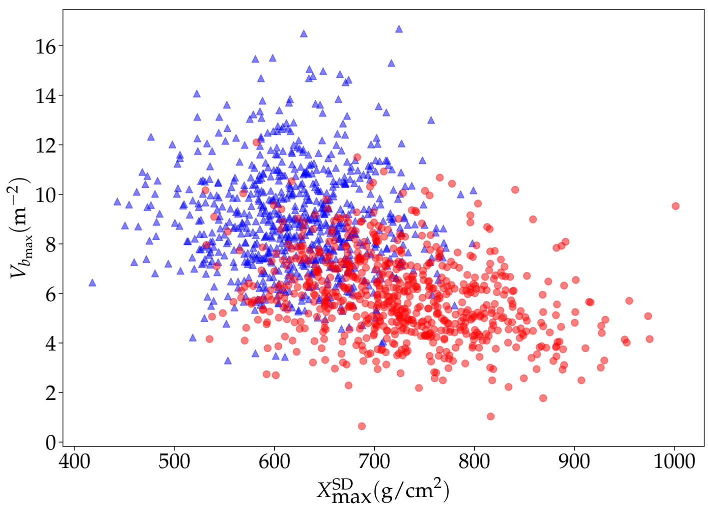

X_{max} resolution = 60 g cm^{-2}

Figure 7.1: $V_{b_{\mathrm{max}}}$ as a function of $X_{\mathrm{max}}^{\mathrm{SD}}$ for proton and iron primaries in a single energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$, for different assumed X_{max} resolutions. The simulations correspond to EPOS-LHC high-energy hadronic interaction model.

increases as the composition evolves from pure iron ($C_p = 0$) to mixed compositions, reaches a minimum around $C_p \sim 0.6$, and then weakens slightly toward the pure proton composition. A clear systematic dependence on $X_{\mathrm{max}}^{\mathrm{SD}}$ resolution is observed. Finer resolution yields stronger anti-correlations. This behaviour reflects the expected degradation of rank sensitivity as shower-maximum information is increasingly smeared.

The right panel of Fig. 7.2 shows $\sigma\sqrt{N}$ as a function of the number of events N for three composition scenarios: pure iron, pure proton, and a mixture of 50/50. The colours again denote the assumed $X_{\mathrm{max}}^{\mathrm{SD}}$ resolution. $\sigma\sqrt{N}$ remains approximately constant over the explored range of N, indicating that the per-sample uncertainty is largely independent of sample size, as expected for a fixed underlying distribution. Pure iron exhibits the highest $\sigma\sqrt{N}$, followed by proton, while the mixed composition yields the lowest values, reflecting reduced contrast between composition classes. To remain conservative, the statistical uncertainty associated with $\tau\left(V_{b_{\mathrm{max}}}, X_{\mathrm{max}}^{\mathrm{SD}}\right)$ is taken to be $0.7/\sqrt{N}$ for the remainder of this work. This value corresponds to the worst-case scenario: pure iron primaries. The bootstrap results show no significant dependence of this scaling on energy within the range considered.

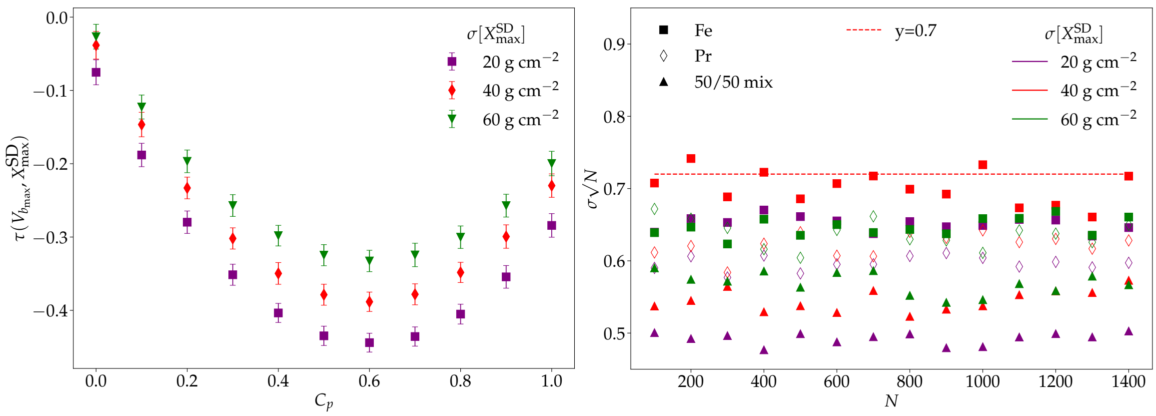

Figure 7.2: (Left) Kendall's τ between $X_{\max}^{\mathrm{SD}}$ and $V_{b_{\max}}$ as a function of proton fraction C_p. Colours indicate different $X_{\max}^{\mathrm{SD}}$ resolutions. (Right) Statistical uncertainty expressed as $\sigma\sqrt{N}$ as a function of sample size N. Both plots correspond to a fixed energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ and simulations generated with EPOS-LHC high-energy hadronic interaction model.

7.2 Mass composition analysis

Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max}^{\mathrm{SD}})$ was evaluated for four pure primary compositions (proton, helium, nitrogen, and iron), as well as for the simulated AugerMix composition, within the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ for the EPOS-LHC high-energy hadronic interaction model. The resulting values are presented as a function of the mean logarithmic mass $\langle \ln A \rangle$ in Fig. 7.3.

Figure 7.3 exhibits a clear and systematic dependence of $\tau(V_{b_{\max}}, X_{\max}^{\mathrm{SD}})$ on $\langle \ln A \rangle$. As the composition transitions from light to heavy primaries, the correlation coefficient varies monotonically, reflecting the mass dependence of the development of the shower. This behaviour arises from the opposite scaling of $X_{\max}$ and $V_{b_{\max}}$ with primary mass: heavier primaries interact earlier in the atmosphere, leading to lower values of $X_{\max}$ and enhanced muon production at ground level, whereas lighter primaries penetrate deeper and produce fewer muons. These opposing trends induce a well-defined ordering in the event sample, which is captured by Kendall's τ through the relative balance of concordant and discordant pairs. The monotonic behaviour observed for pure compositions demonstrates that $\tau(V_{b_{\max}}, X_{\max}^{\mathrm{SD}})$ tracks the underlying mass scale of primary particles. The AugerMix composition yields the most negative value of $\tau(V_{b_{\max}}, X_{\max}^{\mathrm{SD}})$, consistent with the broader distribution of the shower characteristics expected from a heterogeneous mixture of primaries. The enhanced anti-correlation reflects the increased overlap of shower populations with distinct development patterns, which amplifies rank discordance between $V_{b_{\max}}$ and $X_{\max}^{\mathrm{SD}}$ with different resolutions. This behaviour confirms that the correlation coefficient provides a robust composition-sensitive observable capable of discriminating between different mass scenarios.

Filled markers in Fig. 7.3 correspond to results obtained using reconstructed energy, while open markers indicate calculations performed with the true MC energy. A close agreement between the two is observed only for the AugerMix composition. This is expected since the SD energy calibration is derived by assuming an AugerMix-like composition. For pure primary samples, a clear deviation between reconstructed and true energy results is observed, indicating the presence of a composition-dependent energy bias. This bias affects both $V_{b_{\max}}$ and $X_{\max}^{\mathrm{SD}}$, altering their joint ordering and consequently modifying the measured correlation coefficient. This behaviour demonstrates that the correlation observable is not intrinsically independent of energy reconstruction but instead inherits a dependence on

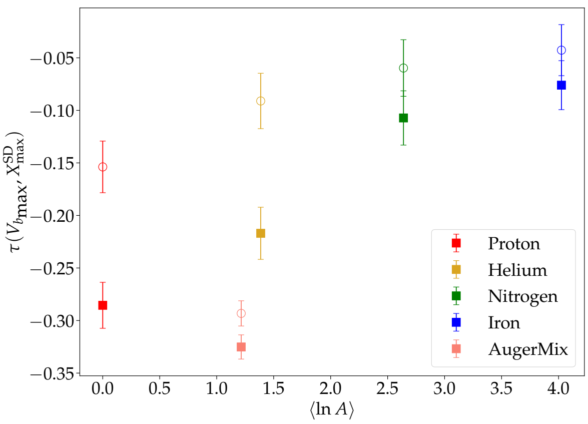

X_{max} resolution = 20 g cm^{-2}

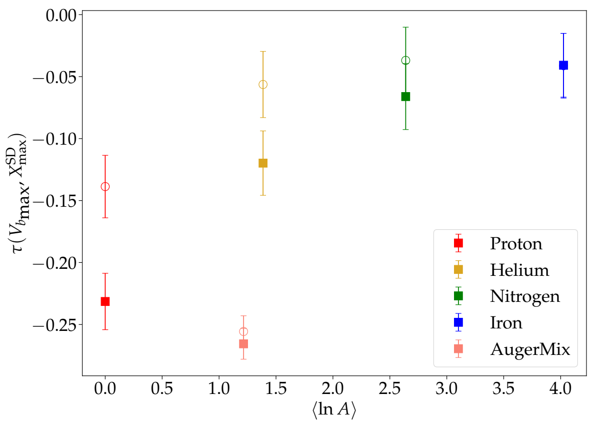

X_{max} resolution = 40 g cm^{-2}

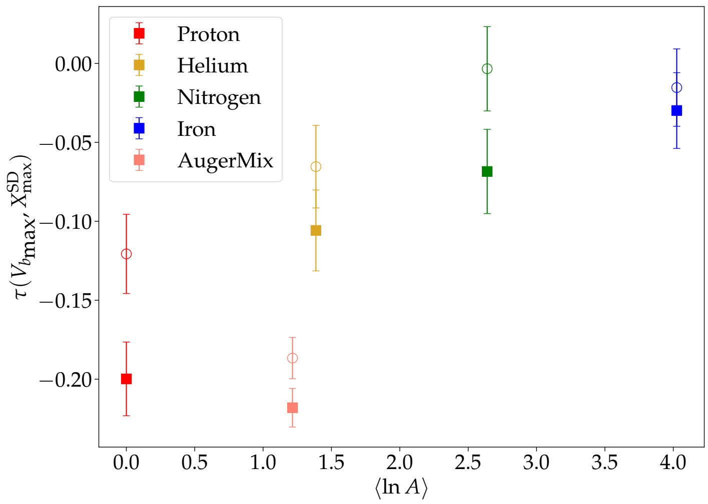

X_{max} resolution = 60 g cm^{-2}

Figure 7.3: Kendall's correlation coefficient $\tau(V_{b_{\mathrm{max}}}, X_{\mathrm{max}}^{\mathrm{SD}})$ as a function of the mean logarithmic mass $\langle \ln A \rangle$ for pure primary compositions and the AugerMix model in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ for EPOS-LHC high-energy hadronic interaction model. Filled markers correspond to reconstructed energy, while open markers show results obtained using true MC energy. Different X_{max} resolutions are shown in different panels and captioned accordingly.

the assumed composition through the calibration procedure. The apparent stability of the correlation for the AugerMix case reflects the consistency between the calibration model and the underlying composition, rather than a general insensitivity to reconstruction effects.

The impact of the X_{max} resolution on the correlation measurement is further illustrated in Fig. 7.4. While the monotonic dependence on $\langle \ln A \rangle$ is preserved across all resolutions, the magnitude of the anti-correlation decreases as the resolution degrades from 20 to 60 g cm^{-2}. This reflects the progressive loss of ordering information in $X_{\mathrm{max}}^{\mathrm{SD}}$, which reduces the contrast between concordant and discordant event pairs. In particular, the distinction between pure and mixed compositions is reduced at larger resolutions, indicating an increasing degeneracy as detector smearing washes out intrinsic shower-to-shower fluctuations.

Although an X_{max} resolution of 20 g cm^{-2} is comparable to that achieved by FD measurements, the applicability of this method in a purely SD framework is limited by the energy calibration procedure. Since the SD energy scale is derived by assuming an AugerMix composition, the reconstructed energy exhibits a composition-dependent bias when applied

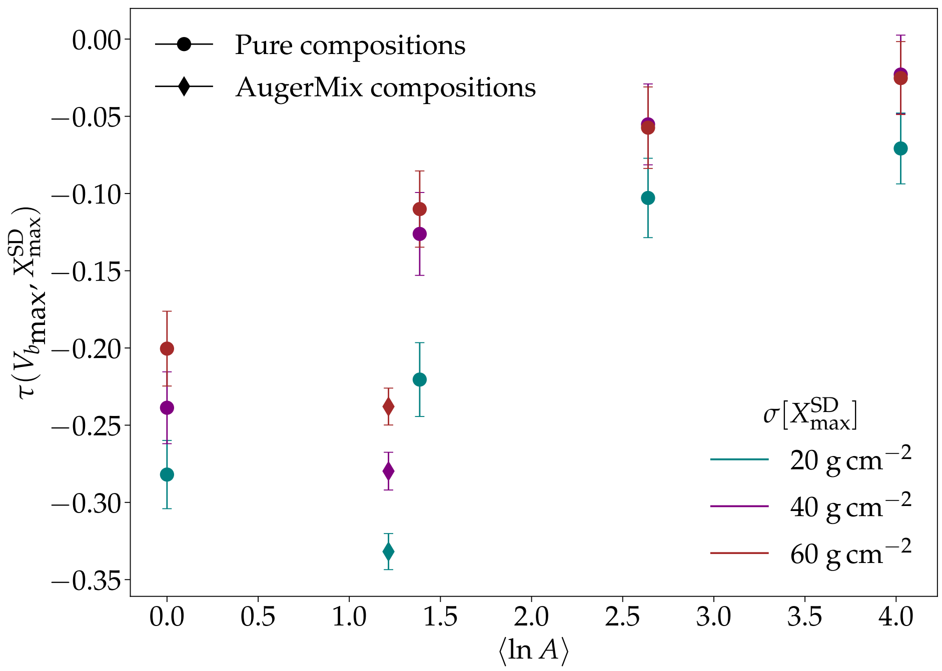

Figure 7.4: Kendall's correlation coefficient $\tau(V_{b_{\max}}, X_{\max}^{\rm SD})$ as a function of the mean logarithmic mass $\langle \ln A \rangle$ for pure primary compositions and the AugerMix model in the energy bin $\log_{10}(E/\mathrm{eV}) \in [17.5, 17.6)$ for EPOS-LHC high-energy hadronic interaction model. Different $X_{\max}$ resolutions are shown in different colours.

to pure primaries. This bias propagates into both observables entering the correlation, leading to systematic distortions in their joint distribution. As a consequence, the extracted correlation coefficient is not solely determined by intrinsic shower physics but also by the assumed calibration model.

This behaviour was further verified by repeating the analysis using the alternative muon density observable V_b^*, as well as across different high-energy hadronic interaction models. In all cases, the use of SD reconstructed energy for pure primary compositions resulted in a persistent composition-dependent bias, indicating that this effect is robust and not specific to a particular observable or interaction model. Although V_b^* includes an explicit energy normalisation term designed to remove the leading energy dependence of the observable, it does not eliminate the composition-dependent bias introduced by the SD energy reconstruction. Since the reconstructed energy itself is systematically biased in a mass-dependent way, the normalisation procedure effectively rescales V_b^* by a composition-correlated factor. This induces distortions in the event-by-event ordering of V_b^*, which directly impact the Kendall correlation. Consequently, the use of V_b^* does not mitigate the observed bias, indicating that the effect originates from the underlying energy calibration rather than from a simple energy scaling of the observable.

This introduces an important limitation for composition studies based on SD observables. Even in the presence of an $X_{\max}$-like resolution comparable to that of the FD, the mass-dependent energy bias prevents a straightforward and unbiased interpretation of the correlation in terms of primary composition. Therefore, any application of this method to data requires a careful treatment of energy reconstruction effects and their dependence on the underlying mass composition. A detailed investigation of these effects lies beyond the scope of this thesis and is not pursued further here, but constitutes an important aspect to be addressed in future work.

7.3 Summary and conclusions

This chapter investigated the use of Kendall's rank correlation coefficient between the surface muon observable $V_{b_{\max}}$ and the depth of shower maximum $X_{\max}^{\text{SD}}$ as a composition-sensitive probe within a SD-based framework. Given the limited duty cycle of the fluorescence detector, such an approach offers the potential for significantly increased event statistics. To emulate SD-like measurements of $X_{\max}$ without introducing reconstruction artefacts, Gaussian smearing was applied to the true MC shower maximum with resolutions of 20, 40, and 60 g cm^{-2}.

The joint distributions of $X_{\max}^{\text{SD}}$ and $V_{b_{\max}}$ exhibit a clear anti-correlation driven by their opposite dependence on primary mass: heavier primaries interact earlier in the atmosphere and produce larger muon densities, while lighter primaries penetrate deeper and yield fewer muons. This intrinsic physical relationship is reflected in the rank ordering of events and forms the basis for using Kendall's τ as a composition-sensitive observable. The anti-correlation persists under realistic detector resolutions, although its magnitude decreases as the $X_{\max}$ resolution worsens.

The statistical behaviour of the estimator follows the expected $1/\sqrt{N}$ scaling, with a normalisation determined by the most conservative resolution and composition scenario. The discrimination power per sample remains approximately constant, implying that improvements in sensitivity are primarily driven by increased statistics. This behaviour is consistent with the results obtained in the previous chapter using $X_{\max}$ measurements from the FD, indicating that the statistical properties of the correlation estimator are largely independent of the specific implementation of the $X_{\max}$ observable.

A systematic and monotonic dependence of $\tau\left(V_{b_{\max}}, X_{\max}^{\text{SD}}\right)$ on the mean logarithmic mass $\langle \ln A \rangle$ was observed for pure compositions, demonstrating that the correlation tracks the underlying mass scale of the primary particles. Mixed compositions yield more negative correlation values due to the increased overlap of showers with distinct development characteristics, which enhances rank discordance between $X_{\max}^{\text{SD}}$ and $V_{b_{\max}}$.

An important limitation arises from the energy calibration procedure. Since the SD energy scale is derived assuming an AugerMix composition, the reconstructed energy exhibits a composition-dependent bias for pure primaries. This bias propagates into both $V_{b_{\max}}$ and $X_{\max}^{\text{SD}}$, modifying their joint ordering and affecting the measured correlation. Consequently, the apparent stability of the correlation with respect to energy reconstruction is only valid when the assumed composition matches the calibration model. This introduces an intrinsic model dependence, which complicates the use of this observable for unbiased composition measurements in a purely SD-based analysis.

In summary, the Kendall correlation between $V_{b_{\max}}$ and $X_{\max}^{\text{SD}}$ provides a statistically robust and physically well-motivated, composition-sensitive observable. It is sensitive to the mean of the primary mass distribution and remains effective under realistic detector resolutions. However, a detailed treatment of energy reconstruction effects and their composition dependence lies beyond the scope of this thesis and is not explored further, but it represents an important direction for future work.

CHAPTER 8 Reconstruction of muon lateral distribution functions

The muon content of the shower is reconstructed by fitting a parametrised MLDF to the signals recorded by the detectors. From the fitted distribution, the total number of muons in the shower, N_μ, is obtained. The fit is performed by maximising a likelihood function that models the detector response to an expected muon density. Once the best-fit parameters are obtained, the MLDF is evaluated at a fixed reference distance from the shower axis, providing a single stable estimator of the shower muon content [45]. For the UMD at Pierre Auger Observatory, this estimator is defined as the reconstructed muon density at a reference distance of 450 m from the shower core. This distance was found to be optimal for an array with a hexagonal spacing of 750 m; it minimises the variance of the reconstructed muon content and therefore produces the most stable observable. The optimality was established by analysing reconstruction fluctuations across multiple detector-level realisations generated from the same underlying simulated MLDF; i.e, repeated samplings of identical showers with independent Poisson and detector response noise [47].

This chapter presents two novel methods for reconstructing particle densities using muon counts from UMDs at various distances from the shower axis. These two methods were developed using MC air-shower simulations. Simulated muon content measurements obtained from the UMD were used to demonstrate these techniques. In addition, a comparative analysis of these approaches was performed with previously established methodologies documented in the literature. This chapter presents a new reconstruction method that uses only the ADC information. In addition, it proposes a hybrid reconstruction method suitable for detectors with binary and ADC modes. This chapter serves as an outlook chapter, and the reconstruction methods developed here can provide improved estimates of muon density at 450 m, which may serve as an alternative normalisation observable to be combined with X_{max} in future correlation analyses.

Section 8.1 describes the characterisation of the ADC output. Section 8.2 discusses the MC shower library and the simulation of ADC mode. Section 8.3 outlines the likelihoods used in the reconstruction process and describes the reconstruction processes developed previously. The performance of the two proposed methods is discussed in Section 8.4, along with a comparison with previous reconstruction methods. The chapter is concluded in Section 8.5.

This chapter is strongly based on the following publication:

- *Reconstruction of air shower muon lateral distribution functions using integrator and binary modes of underground muon detectors*, V. V. Kizakke Covilakam, A. D. Supanitsky, D. Ravignani, Eur. Phys. J. C (2023) 83:1157 [42]

8.1 Characterization of the ADC output

The output of ADC is linear and can be expressed as the arithmetic sum of individual signals [44]. The total charge associated with n muons can be expressed as

$$Q = \sum_{i=1}^{n} q_i, \tag{8.1}$$

where q_i represents the charge of a single muon. The number of muons that hit the detector (n) follows a Poisson distribution ($P(n|\mu)$) with parameter μ,

$$P(n|\mu) = \exp(-\mu)\frac{\mu^n}{n!}. \tag{8.2}$$

The average number of muons in Eq. (8.2) is given by $\mu = \rho_\mu A \cos\theta$, where ρ_μ[1] is the mean muon density of an air shower on the shower plane, A is the active area of the UMD module, and θ is the zenith angle of the air shower. The probability density function that describes the total detector output can be expressed as

$$f(Q|\mu) = \sum_{n=0}^{\infty} f(Q|n) \times P(n|\mu), \tag{8.3}$$

where $f(Q|n)$ is the total charge distribution of n muons. The mean and standard deviation of Q are, respectively, given by

$$\langle Q \rangle = \mu \langle q \rangle, \tag{8.4}$$

$$\sigma^2[Q] = \mu \left(\varepsilon^2[q] + 1\right) \langle q \rangle^2. \tag{8.5}$$

Here, $\langle q \rangle$ is the mean charge corresponding to a single incident muon and $\varepsilon[q]$ corresponds to the relative error of the charge deposited by a muon, which is the ratio of the standard deviation to the mean ($\varepsilon[q] = \sigma[q]/\langle q \rangle$). The estimator of the number of muons $\widehat{\mu} = Q/\langle q \rangle$ is an unbiased variable with variance $\sigma^2[\widehat{\mu}] = \sigma^2[Q]/\langle q \rangle^2$.

As stated in Ref. [37], the uncertainties in estimating the muon density of air showers are mainly governed by Poisson fluctuations in the muon content in the low-occupancy regime, where the number of muons recorded by an individual detector is small. For higher muon multiplicities, detector effects, such as signal pile-up and overlap, introduce additional uncertainties that can become comparable to or larger than the Poisson fluctuations. In particular, Eq. (8.5) highlights that ADC contributes to the uncertainty of an ideal Poissonian detector because the fluctuations in $\widehat{\mu}$ are greater than those expected from the Poisson distribution. In binary mode, the resolution plateaus as the number of particles increases, while in ADC mode the resolution improves at higher particle densities. To clarify the detector resolution, the resolution of the ADC mode ($\sigma[\widehat{\mu}]/\mu$) is compared with that of the binary mode and the ideal Poisson detector in Fig. 8.1. To calculate the resolution of the binary mode, the single window (also known as the infinite window strategy) likelihood developed in Ref. [46] was used. The plots show that the binary mode performs similarly to an ideal Poisson detector at low numbers of muons. However, as the number of muons increases, the ADC mode outperforms the binary mode.

For the ADC mode, the charge of a single muon is treated as a random variable whose logarithm follows a normal distribution. This choice is motivated by the fact that the detector response is governed by several multiplicative processes - energy deposition, scintillation

[1] The muon density is obtained by averaging the muon density over a circular ring at a fixed radial distance from the shower axis in the shower plane.

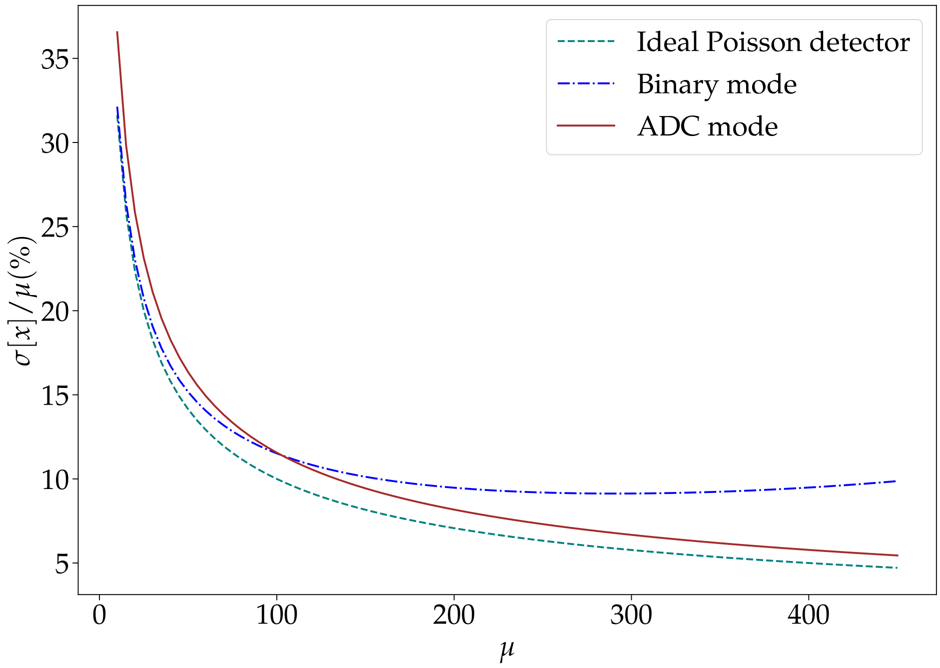

Figure 8.1: Comparison of the detector resolution for the ADC mode, binary mode, and an ideal Poisson detector. The variable x in $\sigma[x]$ refers to the estimator of the average number of muons corresponding to each mode.

light production, photon transport, and SiPM amplification, whose combined fluctuations naturally lead to a log-normal charge distribution [40]. The corresponding probability density function is therefore well described by a two-parameter log-normal model of the form

$$f(q) = \mathrm{LN}(m, \theta^2) \equiv \frac{1}{\sqrt{2\pi}\,\theta\, q} \exp\left[-\frac{(\ln q - m)^2}{2\theta^2}\right], \tag{8.6}$$

where m is the scale parameter and θ is the shape parameter, and can be calculated in terms of $\langle q \rangle$ and $\varepsilon[q]$ as,

$$m = \ln\left[\frac{\langle q \rangle}{\sqrt{1+\varepsilon^2[q]}}\right], \tag{8.7}$$

$$\theta = \sqrt{\ln(1+\varepsilon^2[q])}. \tag{8.8}$$

If the shape parameter remains small ($\theta^2 \lesssim 1$), which is the case for UMDs in Pierre Auger Observatory, the sum of n log-normal terms is expected to behave like a log-normal variable, regardless of the value of n, as cited in Ref. [49]. The same is verified in Appendix D.1.

When at least one muon is detected ($n \geq 1$), the distribution $f(Q|n)$ of the total charge Q corresponding to n muons is,

$$f(Q|n) \cong \mathrm{LN}(m_n, \theta_n^2). \tag{8.9}$$

The parameters m_n and θ_n are estimated from the parameters corresponding to the distribution function of an incident muon (m and θ),

$$m_n = m + \frac{\theta^2}{2} + \ln\left[\frac{n}{\sqrt{1+\frac{\exp(\theta^2)-1}{n}}}\right], \tag{8.10}$$

$$\theta_n = \sqrt{\ln\left[1+\frac{\exp(\theta^2)-1}{n}\right]}. \tag{8.11}$$

In the linear region of the ADC channel, the probability density function of Q for a given value of μ can be derived from Eqs. (8.3) and (8.9) and expressed as a compound distribution of the form

$$\begin{aligned} f_{\mathrm{C}}(Q|\mu) \cong \delta(Q)\exp(-\mu) + \sum_{n=1}^{\infty} \frac{1}{\sqrt{2\pi}\,\theta_n\, Q}\exp\left(-\mu\right)\frac{\mu^n}{n!} \\ \times \exp\left[-\frac{(\ln Q - m_n)^2}{2\theta_n^2}\right]\Theta(Q), \end{aligned} \tag{8.12}$$

where $\delta(Q)$ represents the Dirac delta function, which accounts for the possibility of zero signals due to the absence of muons. In the detector response, this term represents signals compatible with the constant electronic baseline of the ADC channel. The function $\Theta(Q)$ is the Heaviside step function. In the worst-case scenario where all muons arrive in the same time bin, a single 10 m^2 module causes the LG channel to saturate at approximately 362 muons. In contrast, the HG channel saturates at a much lower value of around 85 muons [39]. The actual detection range is wider because each UMD comprises three such modules.

In regions of higher particle density, i.e., for large values of μ (or $\widehat{\mu}$), it is expected that the probability density function of the total charge for a given value of μ can be approximated by a Gaussian distribution (see Appendix D.2),

$$f_{\mathrm{G}}(Q|\mu) \cong \frac{1}{\sqrt{2\pi}\,\sigma[Q]}\exp\left[-\frac{(Q-\langle Q\rangle)^2}{2\,\sigma^2[Q]}\right], \tag{8.13}$$

where $\langle Q\rangle$ and $\sigma[Q]$ can be obtained from Eqs. (8.4) and (8.5), respectively.

8.2 Simulations

Reconstruction methods were developed using a MC shower library generated with CORSIKA v7.7100 [52]. The high-energy hadronic interaction model EPOS-LHC [12] and the low-energy hadronic interaction model FLUKA [53] were used. Showers of iron and proton primaries were simulated with energies in the range $\log_{10}(E/\mathrm{eV}) = [17.5, 19]$ in steps of $\Delta\log_{10}(E/\mathrm{eV}) = 0.25$ for zenith angles of 30° and 45°. The azimuth angles were randomly generated between $-180°$ and 180°. During simulation, a statistical thinning of 10^{-6} was applied to reduce the number of tracked particles [52]. 35 protons and 30 iron showers were produced for each energy and zenith angle. Although experimental evidence suggests a muon deficit in simulated air showers [83], this does not affect the comparison of different strategies to reconstruct the MLDF. Since all strategies are applied to the same set of simulations (with the same muon deficit), the comparison is made on a relative basis, and the muon deficit is a common bias across all methods. This deficit in the simulations seems to be related to the lack of knowledge of hadronic interactions at higher energies, since the high-energy hadronic interaction models used in the simulations extrapolate the accelerator data to CR energies.

The average MLDF and the average arrival time distribution of the muons as a function of their distance to the shower axis of the simulated showers were obtained for each primary type, energy and zenith angle of the simulated showers. Muons with energies greater than 1 GeV/cos θ_μ were selected to account for soil-shielding effects. The average MLDFs were fitted with a KASCADE-Grande MLDF [36] to obtain the input MLDF, which is the a priori muon lateral distribution function that is fed into the likelihood procedure. In Fig. 8.2, the arrival time histogram for iron showers is presented at a zenith angle of 30° and energy $10^{18.5}$ eV at three different distances. These histograms represent the normalised arrival-time distribution, where each bin corresponds to the fraction μ_i/μ of muons arriving within a 10-ns time interval. The muons were observed to arrive more evenly spaced in time, farther away from the shower axis.

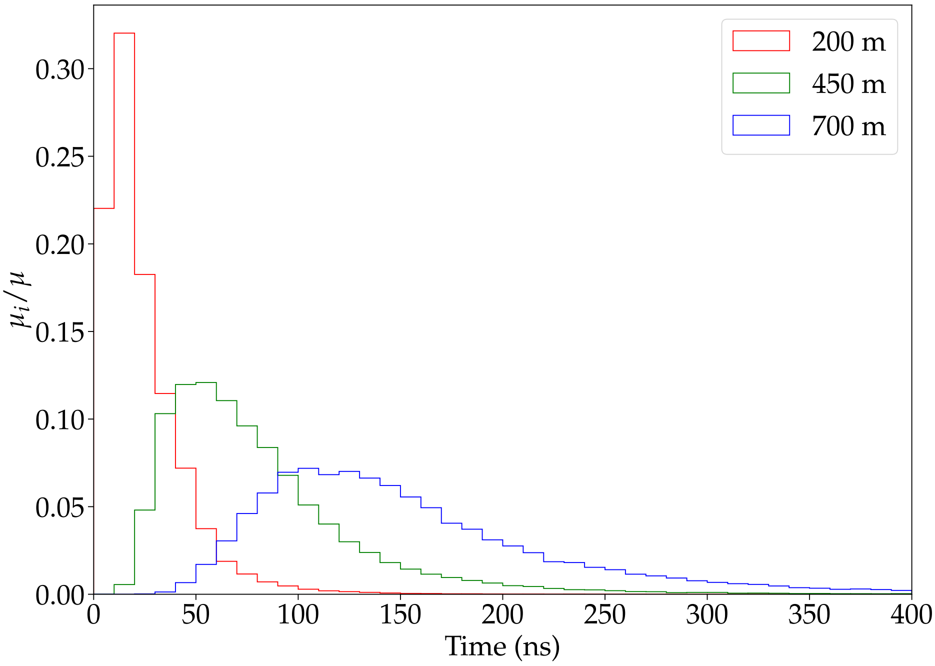

Figure 8.2: Average muon arrival time histogram showing the muon fraction in each time bin at three different distances from the shower axis for iron showers at a zenith angle of 30° and energy of $10^{18.5}$ eV.

The MLDF fit and the average muon time distribution are obtained as a function of the distance from the shower axis for each primary particle; the energy level and the zenith angle served as input for the detector simulations. The binary mode, including the pile-up effect, was simulated as described in Ref. [46]. The muons were counted within a 40-ns time window.

A simulation programme originally developed in Ref. [46] is used to simulate the response of the UMD in binary mode. In this simulation, the number of muons crossing a 30 m^2 detector located 2.5 m underground–corresponding to the UMD configuration of Pierre Auger Observatory–is recorded for each simulated air shower. The detectors are placed on a regular grid with a spacing of 750 m, consistent with the simulated array layout implemented in the reconstruction code. Each muon counter is segmented into 192 scintillator bars and reads out in time windows of 40 ns (inhibition windows), enabling both binary and time-resolved measurements of the muon signal. This framework has been extended to include the characteristics of the ADC output, as described in Section 8.1. The normalised signal as a function of time generated by a single muon can be characterised as a function resembling a logarithmic normal distribution, denoted as $s(t, t_0)$,

$$s(t, t_0) = \frac{\Theta(t - t_0)}{\sum_{i=1}^{N} \exp\left[-\frac{\ln^2((t_i - t_0)/\tau)}{2\theta_{\mathrm{T}}^2}\right] \frac{\Delta t}{t_i - t_0} (t - t_0)} \times \exp\left[-\frac{\ln^2((t - t_0)/\tau)}{2\theta_{\mathrm{T}}^2}\right], \tag{8.14}$$

where $\Theta(t)$ is the Heaviside function, Δt is the sampling time (6.25 ns for this work), t_0 is the signal start time, $m_{\mathrm{T}} = \ln(\tau/\mathrm{ns})$ and θ_{T} are the parameters of a log-normal-like function, and N is the number of signals. Note that $s(t, t_0)$ has the dimension of $[\mathrm{time}]^{-1}$.

Once the ADC reaches saturation, the linear proportionality between the signal charge and the number of injected muons is lost. To determine the saturation level of the ADC signal, one can calculate the maximum of the function $s(t, t_0)$ for a specific number of muons required to saturate the ADC channel (N_μ),

$$S_{\mathrm{L}} = \frac{N_\mu \langle q \rangle}{\sum_{i=1}^{N} \exp\left[-\frac{(\ln^2((t_i - t_0)/\tau)}{2\theta_{\mathrm{T}}^2}\right] \frac{\Delta t}{t_i - t_0}} \frac{\exp\left(\frac{\theta_{\mathrm{T}}^2}{2}\right)}{\tau}. \tag{8.15}$$

This estimate implicitly assumes that the muons arrive simultaneously, such that their individual signals overlap maximally and produce the largest possible peak amplitude. If the muons are distributed in time, the overlap between pulses is reduced, and a larger number of muons would be required to reach the saturation level.

In this study, the LG channel of the UMD at Pierre Auger Observatory was considered. The saturation level is determined by $N_\mu = 1086$, which represents three times the number of muons required by a single module to reach saturation [39]. The parameters of the log-normal-like function were selected as $\ln(\tau/\mathrm{ns}) = 4$ and $\theta_{\mathrm{T}} = 0.4$. These values correspond to a particular set of laboratory measurements performed using four 4 m-long scintillator strips placed inside an aluminium container. The optical fibres were glued to each strip and coupled through an optical connector to a calibrated SiPM array, which consisted of 64 SiPM connected to the front end of the electronics kit and sealed into a PVC box. A muon telescope triggered the electronics when the muon crossed the four bars. The muon pulses measured by ADC were recorded at different positions on the muon telescope along the bars, and the average signal as a function of time for one muon was obtained (see Ref. [40] for more details). Other suitable parameter values for the log-normal-like function would produce similar outcomes. Detector saturation occurs when the signal produced in the detector exceeds the dynamic range of the readout electronics, causing the ADC output to reach its maximum value and the proportionality between the recorded signal and the number of muons to be lost. The detector saturation establishes the upper limit of the ADC detection range.

The ADC signal (in units of ADC counts/time corresponding to n muons) is obtained by considering the unsaturated signal,

$$S_{\mathrm{US}}(t) = \sum_{i=1}^{n} q_i \, s(t, t_i), \tag{8.16}$$

where n, t_i, and q_i are obtained by sampling the Poisson distribution, the corresponding average muon arrival time distribution, and the charge distribution of one muon (see Eq. 8.6), respectively. Finally, the signal is obtained from the following expression,

$$S(t) = \begin{cases} S_{\mathrm{US}}(t) & \text{if} \quad S_{\mathrm{US}}(t) < S_{\mathrm{L}}, \\ S_{\mathrm{L}} & \text{if} \quad S_{\mathrm{US}}(t) \geq S_{\mathrm{L}}. \end{cases} \tag{8.17}$$

The values of the parameters corresponding to the single-muon charge distribution of Eq. (8.6) considered are $m = 5$ and $\theta = 0.5$ [40]. In this case, the values correspond to the same laboratory measurements of the muon detector. In this work, the ADC signal is simulated by generating the response of individual muons and summing their contributions in time. The number of muons is sampled from a Poisson distribution, while their arrival times and charges are drawn from the corresponding arrival-time and single-muon charge distributions. The resulting unsaturated signal is obtained by summing the individual muon pulses. Finally, the detector response is modelled by imposing the ADC saturation level, such that signals exceeding the saturation threshold are clipped at S_{L} in Eq. (8.15). Fig. 8.3 shows an unsaturated and a saturated pulse for the above scenario corresponding to a simulated event.

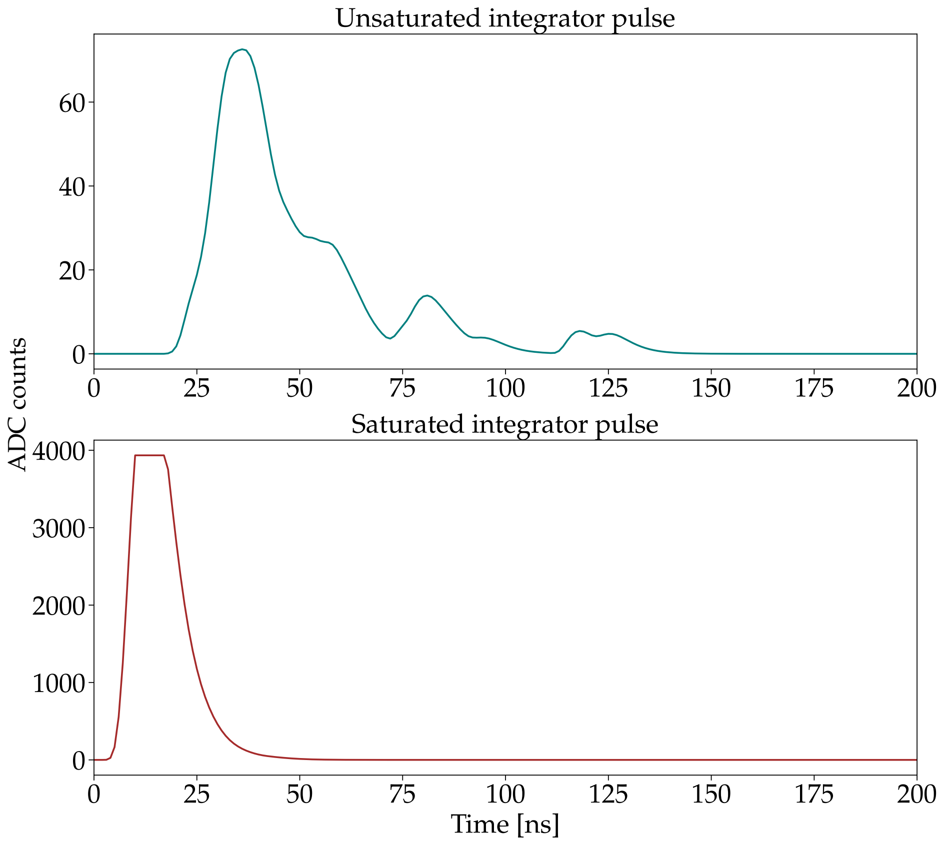

Figure 8.3: Simulated unsaturated and saturated pulses corresponding to a simulated event in the LG channel of the ADC for an iron shower of 10^{19} eV and $\theta = 30°$. The distances to the shower axis of the unsaturated and saturated pulses are ~ 876 m and ~ 211 m, respectively.

The total signal or charge in a given detector (Q) in Eq. (8.1) is obtained by adding the signal evaluated at each discrete time value, obtained according to the sampling time of the ADC, multiplied by Δt.

8.3 Likelihoods and reconstruction methods

In this section, two reconstruction methods are developed: the ADC reconstruction method, which is based only on the ADC information, and the ADCProfile reconstruction method, which combines the ADC and binary outputs. The proposed ADCProfile reconstruction method extends the profile reconstruction method that was developed in Ref. [47].

The MLDF denoted by $\mu(r;\vec{p})$ depends on two free parameters in $\vec{p} = (\mu_0, \beta)$ and the distance between the detector and the shower axis (r). Here, μ_0 represents the normalisation parameter, while β is the slope. The MLDF is factorised by μ_0 and a form function $g(r;\beta)$ that rapidly decreases with increasing distance from the shower axis,

$$\mu(r;\vec{p}) = \mu_0 \frac{g(r;\beta)}{g(r_0;\beta)}, \tag{8.18}$$

where $g(r;\beta)$ is,

$$g(r;\beta) = \left(\frac{r}{r_1}\right)^{-\alpha} \left(1 + \frac{r}{r_1}\right)^{-\beta} \left(1 + \left(\frac{r}{10r_1}\right)^2\right)^{-\gamma}. \tag{8.19}$$

Fixed values of $\alpha = 0.75$, $r_1 = 320$ m, and $r_0 = 450$ m were chosen. The parameter γ is a fixed value obtained from the input MLDF and was found to be almost constant for the primary energy, the zenith angle, and the primary type. The parameter γ describes the behaviour of the MLDF at large distances from the shower axis, where the statistics are low, and the detector either has few muons or is not triggered. In this study, γ is fixed at -4.18. This value is determined by fitting the γ values obtained from the input MLDFs for both proton and iron primaries over the energy range under study at a 30° zenith angle using a zero-degree polynomial, as shown in Fig. 8.4.

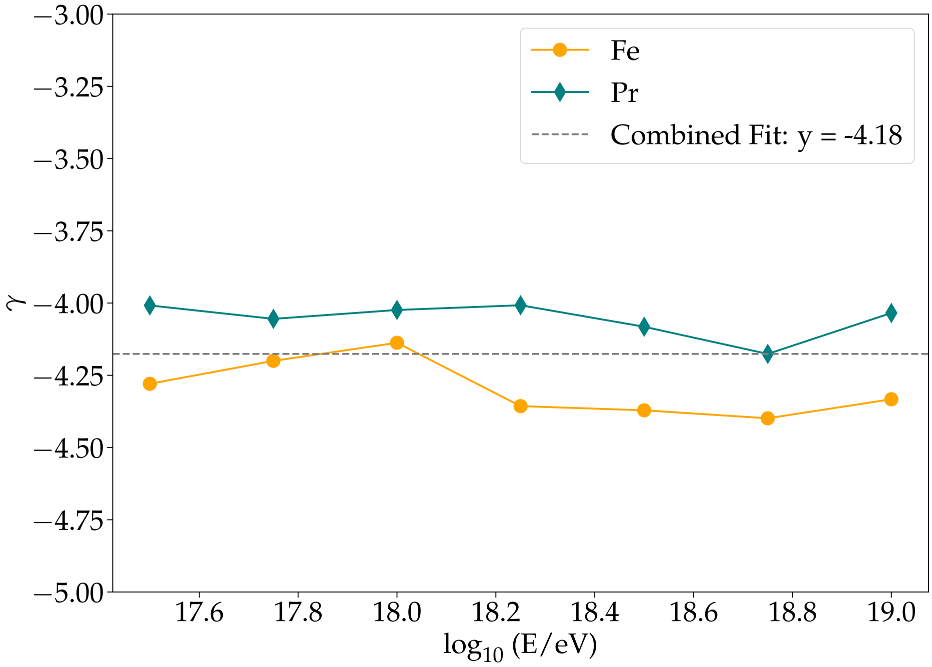

Figure 8.4: The γ values obtained from input MLDFs for proton and iron primaries at different energies at 30° zenith angle. The fit was performed using a zero-degree polynomial for the combined primaries.

The free-fit parameters in $\vec{p}$ are obtained by minimising the following function,

$$-2\ln\mathcal{L}_{\mathrm{T}} = -2\sum_i \ln\left[\mathcal{L}(\mu(r_i,\vec{p}))\right]. \tag{8.20}$$

The sum runs over the detectors in this context. To choose $\mathcal{L}$ for ADC reconstruction, the estimated number of muons was considered in the LG channel of the ADC. In contrast, for

ADCProfile reconstruction, $\mathcal{L}$ is selected depending on the estimated number of muons in the various time windows of the binary traces.

Current reconstruction methods for reconstructing the MLDF from UMD use binary information. These methods are briefly discussed in the following subsections, and new reconstruction methods using the ADC are proposed.

8.3.1 Time-dependent reconstruction method

The time-dependent likelihood method is the most straightforward method for reconstructing the MLDF with binary mode data [45]. The stations are divided into three classes based on the number of activated segments[2]. A segment (or bar) is *on* if it has a signal compatible with one or more muons, and in binary mode, a segment is considered activated when the time series of logical signals generated by the discriminator electronics matches a predefined pattern associated with a muon. The stations are marked silent if they have fewer than 3 segments *on*[3] and saturated if more than 124 segments are *on*[4]. Good stations are neither silent nor saturated. Here, an approximate likelihood is used for each class. The combined likelihood is given by

$$\begin{aligned}\mathcal{L}(\vec{p}) = &\prod_{i=1}^{N_{\text{sat}}} \frac{1}{2}\left[1 - \text{erf}\left(\frac{n_i^C - \mu(r_i;\vec{p})}{\sqrt{2\mu(r_i;\vec{p})}}\right)\right] \times \prod_{i=1}^{N_{\text{good}}} e^{-\mu(r_i;\vec{p})}\frac{\mu(r_i;\vec{p})^{n_i^{\text{corr}}}}{n_i^{\text{corr}}!} \\ &\times \prod_{i=1}^{N_{\text{sil}}} e^{-\mu(r_i;\vec{p})}\left(1 + \mu(r_i;\vec{p}) + \frac{1}{2}\mu(r_i;\vec{p})^2\right),\end{aligned} \tag{8.21}$$

where the first, second, and third terms correspond to the saturated, good, and silent stations, respectively. r_i is the distance of the i-th station to the shower axis, and N_{sat}, N_{good}, and N_{sil} are the numbers of saturated, good and silent stations, respectively. n_i^C is the number of segments *on*, and n_i^{corr} is the number of muons after applying the correction and is given by,

$$n_i^{\text{corr}} = \sum_j \frac{\ln(1 - \frac{k_j}{n})}{\ln(1 - \frac{1}{n})}, \tag{8.22}$$

where the sum runs over all time bins, k_j is the number of segments *on* in the j-th time bin, and n is the total number of segments.

8.3.2 Single window reconstruction method

The single-window likelihood method reconstructs the MLDF with an array of segmented detectors without considering detector timing. This method uses the exact likelihood of the number of muons in a detector given the number of segments with a signal [46]. Each segment can take two distinct states depending on the number of particles: *on* if hit by one or more muons, and *off* otherwise. According to Poisson, the probability that a segment is *off* is $q = e^{-\mu/n}$, and the probability that a segment is *on* is $p = 1 - q$, where n is the number of segments. This method is based on the exact likelihood of the number of muons in a detector given the number of segments with a signal.

[2]In the case of UMD, the scintillator strip is the generic segment of a segmented detector; the three counters at each array position are equivalent to a single detector divided into 192 segments that cover 30 m^2.

[3]The silent limit of 3 was set to reject accidental triggers caused by background muons crossing two strips.

[4]This limit has to be set because the adopted likelihood underestimates data errors. At 124 segments, the statistical uncertainty is 35% larger than the one used in the reconstruction. The underestimation grows with the number of segments *on*.

Considering that the states of the segments are independent of each other, the probability of k segments out of n, *on*, follows a binomial distribution,

$$L(\mu;k) = \binom{n}{k} e^{-\mu}(e^{\mu/n} - 1)^k \tag{8.23}$$

The maximum likelihood estimator of the number of expected muons ($\hat{\mu}$) is given by,

$$\hat{\mu} = -n\ln(1 - \frac{k}{n}). \tag{8.24}$$

If $k = n$, the likelihood tends to unity when μ increases, and the maximum likelihood estimator of μ tends to infinity. In this case, the likelihood sets a lower bound to the number of muons allowed in the LDF fit, and these detectors are labelled saturated. Hence, a detector is considered saturated if all segments have a signal. A detector is considered silent when fewer than 3 segments have a signal. The likelihood of silent detectors is

$$\mathcal{L}(\mu;k<3) = e^{-\mu}\left(1 + n\left(e^{\mu/n} - 1\right) + \frac{n(n-1)}{2}\left(e^{\mu/n} - 1\right)^2\right). \tag{8.25}$$

8.3.3 The profile reconstruction method

The profile reconstruction method uses detector timing and a profiling technique to reconstruct the MLDF from binary mode data [47]. Considering that the number of segments *on* in each time bin (k_i) is independent of each other, the likelihood of the number of muons (μ_i) in the i-th bin is given by a Binomial distribution (Eq. (8.23)) and the likelihood of all time bins ($L(\boldsymbol{\mu})$) is the product of single-window likelihoods,

$$L(\boldsymbol{\mu}) = \prod_{i=1} L_i(\mu_i) \tag{8.26}$$

where i runs over the time bins and $\boldsymbol{\mu}$ = (μ_1, μ_2, etc.). To overcome the lack of knowledge about the signal time distribution $d\mu/dt$ and derive the μ_i's from μ, the μ_i's are considered nuisance parameters. In this method, a profiling technique is applied to search for nuisance parameters that maximise the likelihood under the restriction $\sum \mu_i = \mu$,

$$L_p(\mu) = \max_{\sum \mu_i = \mu} L(\boldsymbol{\mu}) \tag{8.27}$$

and the profile likelihood ratio is

$$\lambda(\mu) = \frac{L_p(\mu)}{L_{max}} \tag{8.28}$$

where L_{max} denotes the global maximum of the likelihood calculated without any restriction on $\boldsymbol{\mu}$. This maximum is achieved when $\hat{\boldsymbol{\mu}}$ = ($\hat{\mu}_1$, $\hat{\mu}_2$, ...), all obtained from Eq. (8.24).

For silent stations ($\mu < 3$), the profile likelihood reduces to the Poisson probability of observing at most two muons, yielding a time-independent contribution identical to that used in the standard reconstruction.

$$\mathcal{L}(\mu<3) = \prod_{i=1}^{N_{\mathrm{sil}}} e^{-\mu(r_i;\vec{p})}\left(1 + \mu(r_i;\vec{p}) + \frac{1}{2}\mu(r_i;\vec{p})^2\right) \tag{8.29}$$

8.3.4 ADC reconstruction method

If the estimated number of muons in the ADC channel (denoted as $\widehat{\mu}$) is less than or equal to 3, the station is labelled as non-triggered. This is done to eliminate any accidental triggers that random coincidences may cause. In this case, the probability is calculated by integrating Eq. (8.12) from $-\infty$ to $Q_{\min}$,

$$\mathcal{L}(\mu; Q) = \left(1 + \frac{1}{2}\sum_{n=1}^{\infty} \operatorname{erfc}\left(\frac{m_n - \ln(Q_{\min})}{\sqrt{2}\theta_n}\right)\frac{\mu^n}{n!}\right) \times \exp(-\mu) \tag{8.30}$$

Here $Q_{\min} = 3\langle q\rangle$ and $\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$, where $\operatorname{erf}(x)$ is the error function.

If a station triggers and the estimated number of muons in the LG ADC channel is less than 200, a compound likelihood function is taken into account,

$$\mathcal{L}(\mu; Q) = \sum_{n=1}^{\infty} \frac{1}{\sqrt{2\pi}\,\theta_n\, Q} \exp\left[\frac{-(\ln Q - m_n)^2}{2\theta_n^2}\right] \times \frac{\mu^n}{n!}\exp(-\mu). \tag{8.31}$$

When the estimated number of muons in the ADC channel is equal to or greater than 200, and the station is not saturated, a Gaussian likelihood is taken into consideration (refer to Appendix D.2). This is done to improve the reconstruction performance,

$$\mathcal{L}(\mu; Q) = \frac{1}{\sqrt{2\pi\sigma^2[Q]}} \exp\left[-\frac{(Q - \langle Q\rangle)^2}{2\sigma^2[Q]}\right], \tag{8.32}$$

where Eqs. (8.4) and (8.5) give the mean and variance, respectively.

If a particular station is saturated, only a lower bound on the true signal is available. Therefore, the likelihood is constructed as the upper tail of the Gaussian distribution, obtained by integrating Eq. (8.32) from Q to $+\infty$. Here, Q refers to the integral of a pulse, where the upper portion is clipped (refer to the bottom panel of Fig. 8.3),

$$\mathcal{L}(\mu; Q) = \frac{1}{2}\left(1 - \operatorname{erf}\left(\frac{Q - \langle Q\rangle}{\sqrt{2\sigma^2[Q]}}\right)\right), \tag{8.33}$$

where Eqs. (8.4) and (8.5) give the mean and variance, respectively. This corresponds to the probability of observing a signal equal to or larger than Q. Consequently, saturated stations provide a one-sided constraint, effectively acting as a lower limit on μ.

8.3.5 ADCProfile reconstruction method

In this approach, when the binary mode of a muon detector exhibits a high occupancy–i.e., when the number of bars with signal k reaches the saturation threshold ($k < 124$) in at least one inhibition window–the information from the binary channel becomes unreliable. In this regime, the likelihood used for both triggered and saturated stations in the ADC likelihood method (Eqs. (8.32) and (8.33)) is adopted instead. If $k < 124$ in all time windows of the trace, the profile likelihood of Ref. [47] is used. This choice of k was motivated by the relative uncertainty in estimating the number of incident muons in the same inhibition window being less than $\sim 16\%$ for $k \leq 124$ [48]. The non-triggered stations were defined and treated as described in Ref. [47] and use the likelihood in Eq. (8.29).

Table 8.1 summarises several likelihood functions reported in the literature and provides a comparison with the likelihoods proposed in this work.

The compound likelihood in Eq. (8.31) is compared to the profile likelihood from Ref. [47], used in the ADCProfile reconstruction approach and shown in Fig. 8.5. The plots correspond

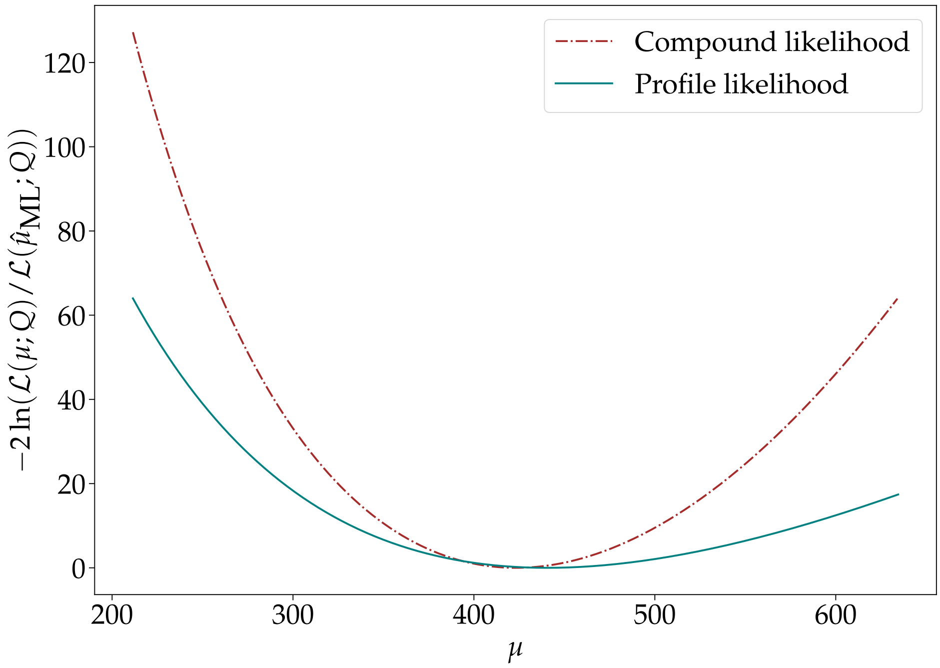

Figure 8.5: Comparison of the likelihoods used in this work. The plots correspond to iron primary at 10^{18} eV at 30° zenith angle. The maximum likelihood estimator of μ is obtained from profile likelihood and has the value $\widehat{\mu} = 423$

to a value of $\widehat{\mu}_{\rm ML} = 423$, which is the maximum likelihood estimator of μ obtained from profile likelihood.

As seen in Fig. 8.5, both likelihoods exhibit a well-defined minimum near the maximum likelihood estimator, indicating consistent reconstruction of the signal. However, the compound likelihood displays a noticeably sharper curvature around the minimum compared to the profile likelihood, reflecting a stronger constraint on μ. This difference arises from the fact that the compound likelihood incorporates the full ADC information, including the charge fluctuations of individual muons, whereas the profile likelihood relies solely on binary segmentation data. For muon counters with 192 segments and a time resolution of 40 ns per inhibition window, the binary response becomes increasingly saturated at high occupancies, leading to a loss of information. In contrast, the ADC-based compound likelihood retains sensitivity in this regime, resulting in improved resolution. Consequently, while both methods provide consistent estimates of μ, the compound likelihood offers enhanced statistical power, particularly at large muon densities where the binary detector response approaches saturation.

8.4 Results

8.4.1 Performance of the reconstruction methods

In this section, the performance of the proposed reconstructions is evaluated. To simulate the number of muons arriving at the detector in each shower, a Poisson distribution was sampled with the parameter μ given by the average MLDF. The time distributions were sampled to simulate the arrival times of the muons. Each realisation is referred to as an event. For each primary type, zenith angle, and primary energy, 10,000 events were simulated, including a simulation of binary and ADC acquisition modes. The number of muons as a function of the distance to the shower axis for simulated events was adjusted using the profile, ADC, and ADCProfile likelihoods and compared against an ideal counter [46]. An ideal counter measures the true number of muons crossing its sensitive area without any pile-up effects[5]. In this case, each muon is detected independently with 100% efficiency, providing a direct measurement of the underlying Poisson-distributed muon number and defining the best-case scenario for the achievable resolution of a muon detector. A simple Poisson likelihood is used to describe an ideal detector,

$$\mathcal{L}(\mu; k) = e^{-\mu}\frac{\mu^k}{k!}, \tag{8.34}$$

where k is the number of counted particles.

The numerical minimisation software library MINUIT [81], implemented in the ROOT data analysis framework [82], was used to minimise log-likelihoods.

Fig. 8.6 displays an example of an MLDF fit achieved using the ADC reconstruction method (left panel) and the ADCProfile reconstruction method (right panel). The presented data show the number of muons in the UMD for iron primaries at 10^{19} eV and 30° zenith angle, with each station labelled according to the likelihood used in the reconstruction process. The left panel plot indicates the Gaussian likelihood (see Eq. (8.32)) used for stations near the shower axis. In contrast, the Compound likelihood (see Eq. (8.31)) was used for the remaining stations. The plot also shows a few non-triggered stations and a saturated station. On the other hand, the plot on the right panel shows that the ADC reconstruction method was used for stations near the shower axis, while the profile reconstruction method was applied to the remaining stations.

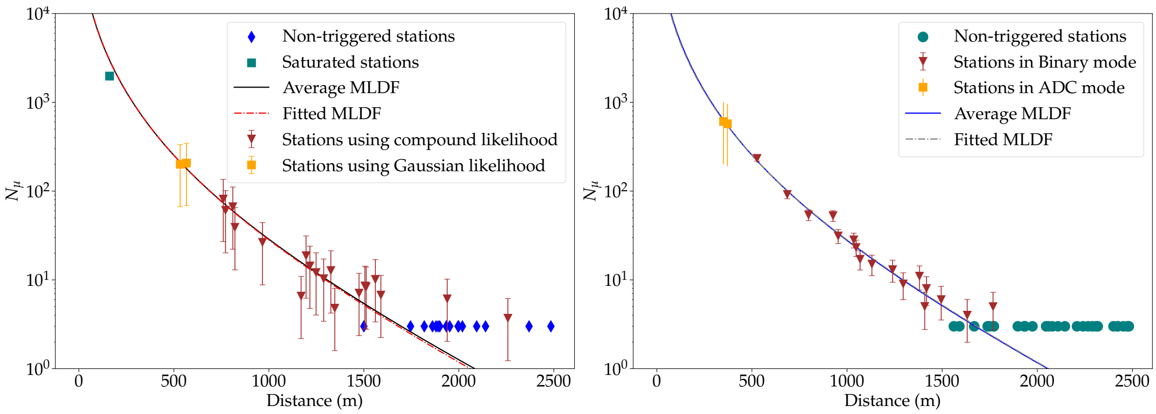

Figure 8.6: The muon lateral distribution function was fitted to the simulated data using the ADC reconstruction method (left panel) and the ADCProfile reconstruction method (right panel). The data correspond to the estimated number of muons in the UMD calculated by a simulation of an iron primary at 10^{19} eV and 30° zenith angle. The solid lines correspond to the MLDF input used in the simulation. Note that the plots correspond to two different events.

8.4.2 Saturated events as functions of logarithmic energy

Fig. 8.7 displays the percentage of saturated events as a function of logarithmic energy. An event is categorised as saturated when it involves at least one saturated muon detector. In both the ADCProfile and ADC approaches, detectors are viewed as saturated when the estimated count of muons in the ADC channel exceeds the saturation limit of the LG channel, as described by Eq. (8.15). In the profile reconstruction technique, a detector is saturated if the station contains 192 bars of signals in a single time window. The plot shows a noticeable

[5] Under-counting from nearly simultaneous hits in the same strip.

increase in saturated events with primary energy. Moreover, the fraction of saturated events appeared similar for all considered reconstruction methods, primarily due to the use of 40-ns-wide inhibition windows in binary mode. However, the fraction of saturated events for profile reconstruction is smaller when the inhibition window is shorter, as explained in Appendix D.3.

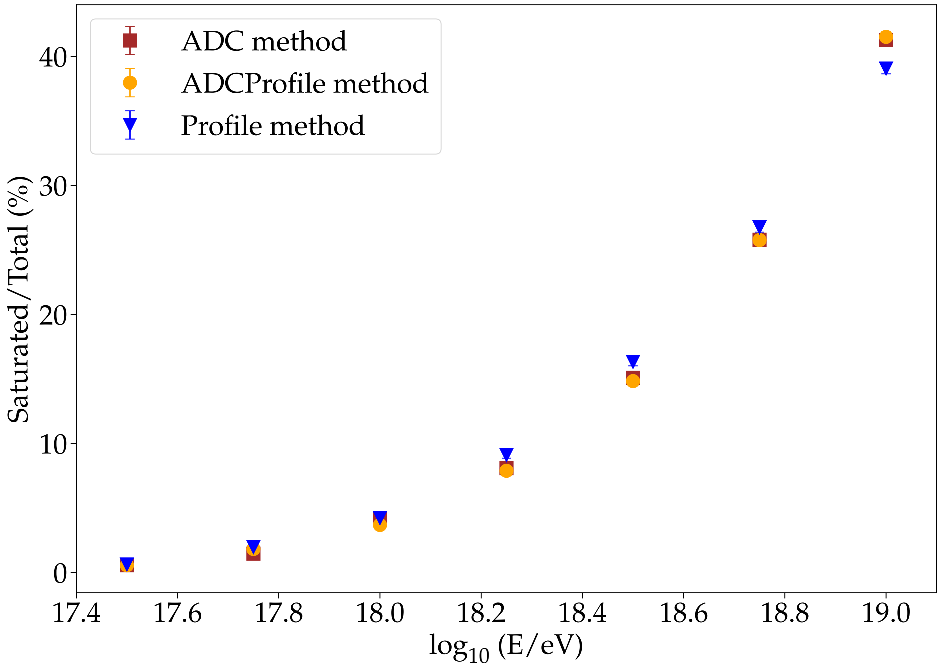

Figure 8.7: Fraction of saturated events for iron primaries at 30° zenith angle as a function of the logarithm of primary energy for the ADC, ADCProfile, and profile reconstruction methods, including error bars. The error bars are smaller than the marker size.

8.4.3 Relative bias and standard deviation

The reconstructed $\widehat{\mu}_{\mathrm{ML}}(450)$ (the maximum likelihood estimator of the average number of muons at 450 m from the shower axis) varies with different reconstructions of the same shower due to fluctuations in the detector response, including Poisson fluctuations in the number of detected muons, segmentation effects (pile-up in the 192 scintillator bars), time binning in 40 ns inhibition windows, and, where applicable, fluctuations in the ADC signal. Fig. 8.8 shows the distributions of the reconstructed $\widehat{\mu}_{\mathrm{ML}}(450)$ for iron primaries at $10^{18.5}$ eV and 30° zenith angle, considering various reconstruction methods and including saturated events. The data corresponding to the ideal counter were fitted with a Gaussian distribution, where the dotted line represents the input value of $\mu(450)$.

To assess the reconstruction performance, it is necessary to compare the input value $\mu(450)$ to the corresponding value fitted to the simulated data. The relative bias is defined as the difference between the input $\mu(450)$ and the average of $\widehat{\mu}_{\mathrm{ML}}(450)$ calculated over all reconstructions, i.e., $B = \mu(450) - \langle\widehat{\mu}_{\mathrm{ML}}(450)\rangle$. The bias is computed as an ensemble average over simulated showers and therefore includes the effect of shower-to-shower fluctuations in the true muon content at fixed energy and zenith angle and represents the systematic uncertainty in the estimation of $\mu(450)$. Fig. 8.9 displays the relative bias (left panel) and relative standard deviation (right panel), $\varepsilon(450) = \sigma[\widehat{\mu}_{\mathrm{ML}}](450)/\mu(450)$, of $\widehat{\mu}_{\mathrm{ML}}(450)$ as functions of the logarithmic energy for different reconstructions, including saturated events. Here, $\langle\widehat{\mu}_{\mathrm{ML}}\rangle$ and $\sigma[\widehat{\mu}_{\mathrm{ML}}](450)$ are estimated from the Gaussian distribution fit of $\widehat{\mu}_{\mathrm{ML}}(450)$

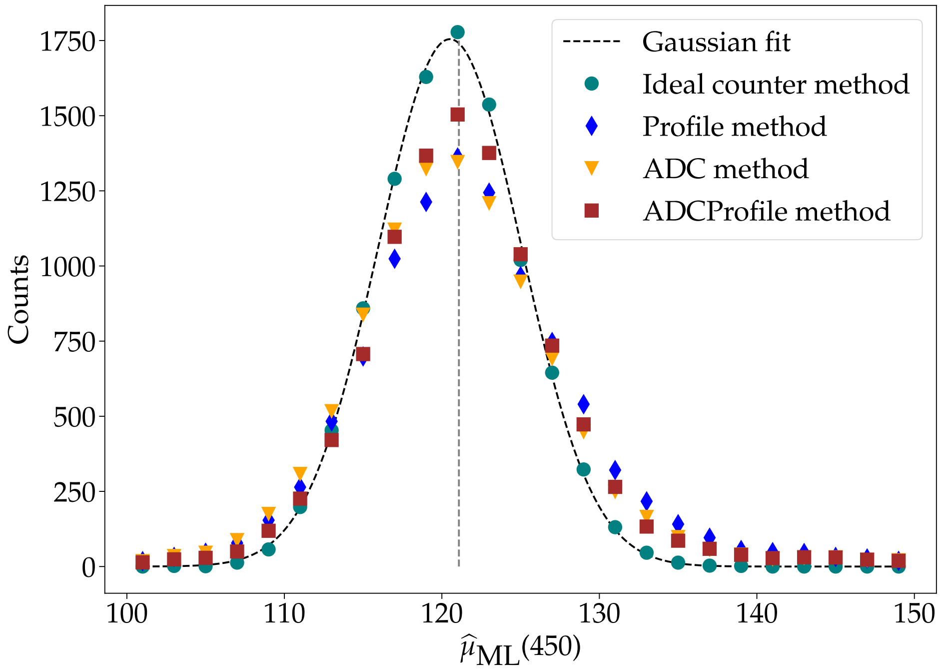

Figure 8.8: Histogram of the reconstructed $\widehat{\mu}_{\rm ML}(450)$ for iron primaries at $10^{18.5}$ eV and 30° zenith angle for the different reconstruction methods considered. The Gaussian distribution is fitted to the ideal counter data.

histograms. The figure shows that the relative bias of the ADC and ADCProfile reconstruction methods is less than 2% at lower simulated energies and lies within 1% for energies greater than $10^{17.5}$ eV. The value of $\varepsilon(450)$ decreases with primary energy due to the increase in muons triggering more detectors and the smaller relative fluctuations at different detectors. The values of $\varepsilon(450)$ for ADC reconstruction are similar to those for the profile reconstruction. Specifically, $\varepsilon(450)$ for the ADC reconstruction is larger than that for the profile reconstruction at low energies but is smaller at high energies. The $\varepsilon(450)$ for the ADCProfile method takes smaller values or is equal to those in the ADC and profile reconstructions. Note that the $\varepsilon(450)$ for the ADCProfile differs by no more than $\sim 2\%$ from that of the ideal counter in the considered energy range. This indicates that the ADCProfile reconstruction operates close to the fundamental resolution limit set by an ideal detector, with detector and reconstruction effects contributing at most a $\sim 2\%$ degradation in performance.

Fig. 8.10 displays the relative bias (left panel) and relative standard deviation (right panel) as a function of the distance to the shower axis for iron primaries at an energy of $10^{18.25}$ eV and a 30° zenith angle. The dotted line indicates the UMD reference distance at 450 m. The ADC and ADCProfile reconstruction methods exhibit a bias similar to that of an ideal counter at distances closer to the shower axis. This is because the ADC, which dominates this region, is more suitable for measuring high muon density than the binary mode. The ADC and ADCProfile methods yielded smaller standard deviations than the profile method at distances closer to the shower axis. This is because the ADC dominates in the ADCProfile method at distances less than 450 m from the shower axis, whereas the binary mode dominates at larger distances. The ADCProfile reconstruction method had a relative standard deviation comparable to that of the profile reconstruction method at distances further from the shower axis. Based on the relative bias and relative standard deviation, the ADCProfile method provides the best overall reconstruction performance, although the improvement over the other methods is marginal, and all approaches yield comparable results.

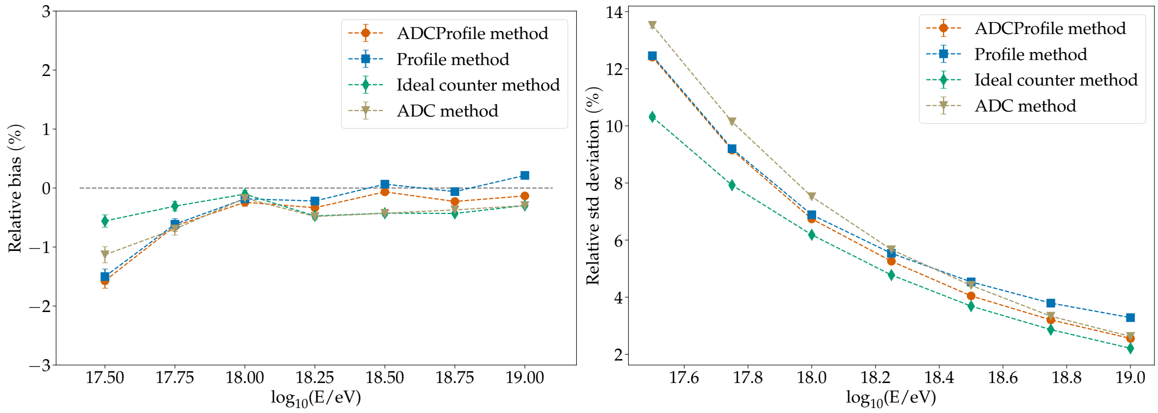

Figure 8.9: Relative biases (left) and standard deviations (right) of the reconstructed $\widehat{\mu}_{\rm ML}(450)$ as a function of the logarithm of the primary energy for ADC, ADCProfile, and profile reconstructions compared to an ideal counter. The plot corresponds to the iron primaries at 30° zenith angle. The error bars are included for all reconstruction methods.

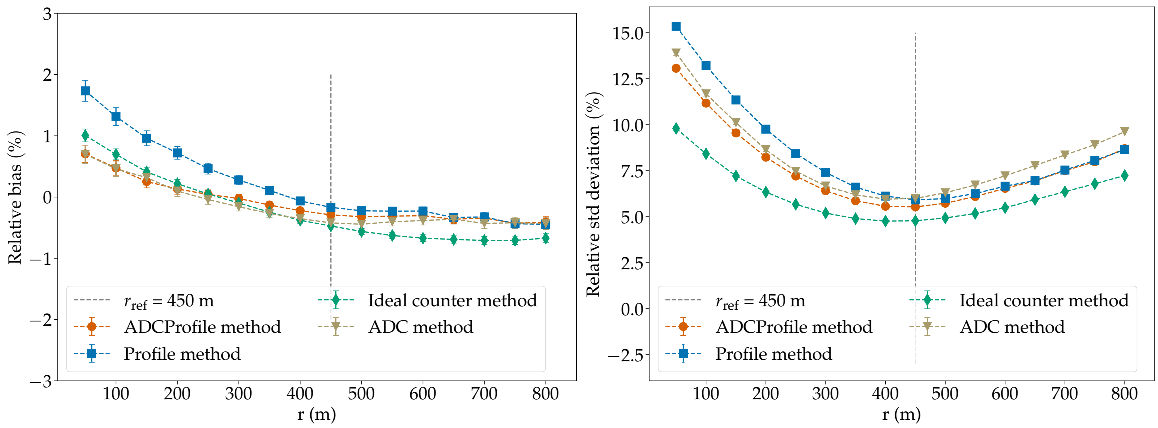

Figure 8.10: Relative bias (left) and relative standard deviation (right) of the reconstructed $\widehat{\mu}_{\rm ML}(\rm r)$ as functions of distance to the shower axis for iron primaries at 30° zenith angle and $10^{18.25}$ eV energy. The dotted line represents the reference distance of the UMD. The ADC, ADCProfile, and profile reconstruction methods were compared against an ideal counter. The error bars are included for all reconstruction methods.

8.4.4 Coverage probability

The coverage probability quantifies the quality of the reconstruction method in terms of the confidence interval of $\hat{\mu}_{\rm ML}(450\ \rm m)$. This is the fraction of events in which the confidence interval includes the true value from the simulated MLDF. In the reconstruction, the 1σ errors of the MLDF normalisation μ_0 and slope parameter β are calculated for each event. The parabolic parameter error is obtained directly from the fit procedure [81]. Fig. 8.11 displays the coverage of $\widehat{\mu}_{\rm ML}(450)$ for the reconstructions of iron primaries at a 30° zenith angle and different energies using different reconstruction methods. The dotted line at 68% corresponds to the 1σ Gaussian confidence interval coverage. As depicted in the figure, the coverage probabilities of all reconstructions are close to one another and to the Gaussian reference.

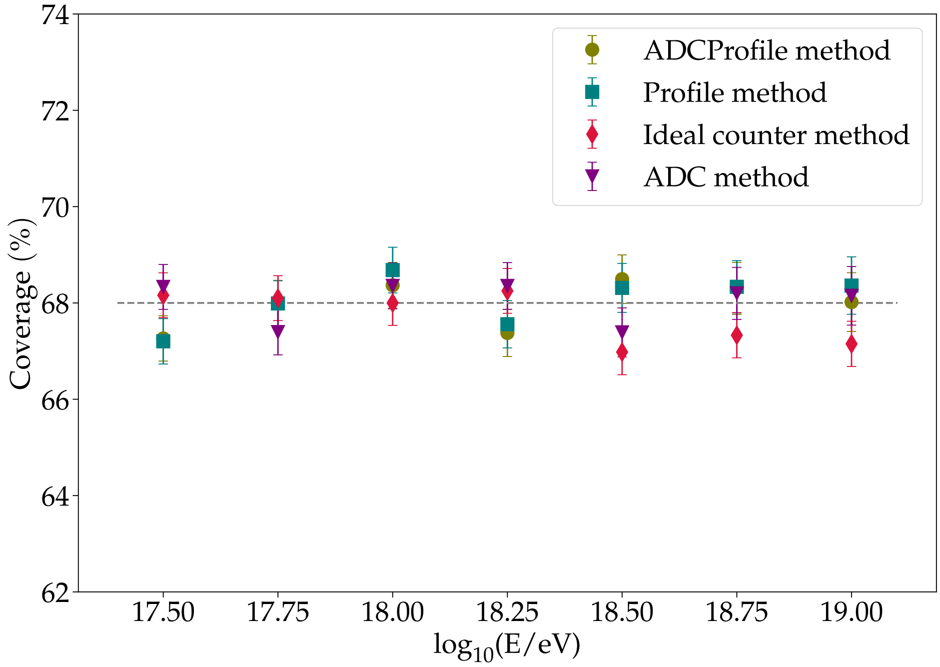

Figure 8.11: Coverage probability of the 1σ confidence interval of $\widehat{\mu}_{\mathrm{ML}}(450)$ as a function of the logarithm of energy for iron primaries at 30° zenith angle. The dotted line at 68% corresponds to the coverage of the 1σ confidence interval of the Gaussian likelihood. The ADC, ADCProfile, and profile reconstruction methods were compared against an ideal counter.

8.4.5 Effect of saturated events

To thoroughly examine the impact of saturated events on the reconstruction process, Fig. 8.12 shows the reconstructed value of $\widehat{\mu}_{\mathrm{ML}}(450)$ for iron primaries with energy of $10^{18.5}$ eV and a 30° zenith angle. The ADC method was used for reconstruction, and the blue and red histograms indicate the reconstructed non-saturated and saturated events, respectively. The distribution of saturated events exhibits a pronounced tail towards larger values of $\widehat{\mu}_{\mathrm{ML}}(450)$, reflecting the one-sided nature of the likelihood used for saturated signals. Since saturation provides only a lower bound on the signal, fluctuations in the reconstructed charge can lead to overestimation of μ, resulting in an extended high-μ tail.

The relative bias (left panel) and relative standard deviation (right panel) against the logarithm of primary energy after excluding saturated events for all reconstruction methods are plotted in Fig. 8.13. The figure shows that the ADC, ADCProfile, and profile reconstruction methods perform similarly. The relative standard deviation obtained with ADC and ADCProfile is closer to an ideal counter, particularly at higher energies.

In this analysis, the saturation level in the LG channel of the ADC was set to 1086 muons. To evaluate the performance of the ADC reconstruction method at various saturation levels, 10,000 events were reconstructed at lower ($\sim$ 300 muons) and higher ($\sim$ 2500 muons) saturation levels. Fig. 8.14 compares the relative bias (left panel) and relative standard deviation (right panel) of $\widehat{\mu}_{\mathrm{ML}}(450)$ at different ADC saturation levels. The plot indicates that although higher ADC saturation levels yield better relative standard deviations, especially at higher energies, this improvement is minimal and insignificant relative to the unrealistic saturation level chosen.

It is worth noting that the results obtained for iron nuclei at 45°, as well as for protons at 30° and 45°, are similar to those obtained for iron nuclei at 30°. The ADCProfile reconstruction outperformed the ADC and profile reconstructions across the analysed energy range.

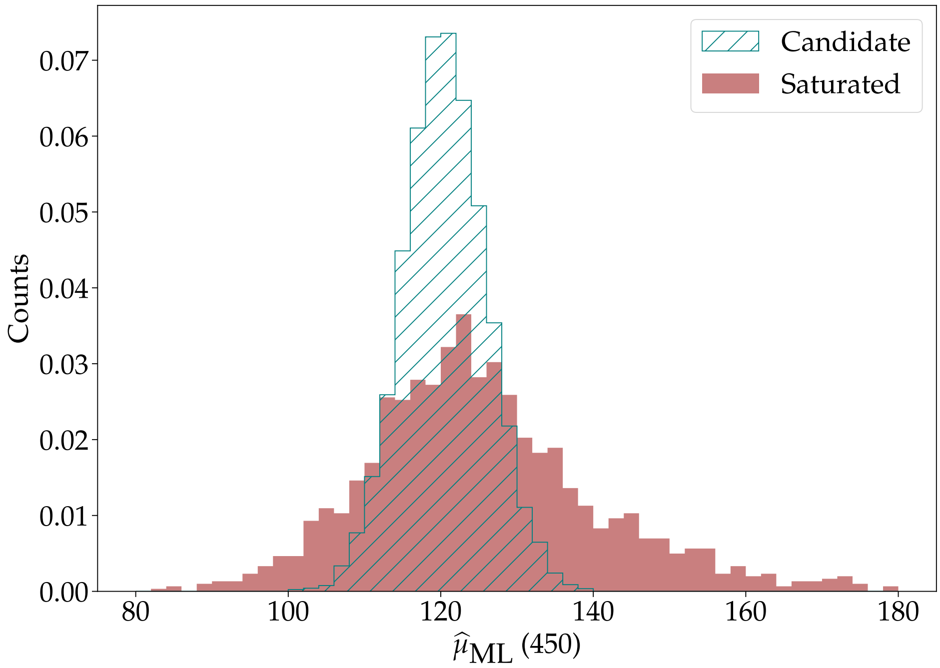

Figure 8.12: Distributions of the reconstructed $\widehat{\mu}_{\mathrm{ML}}(450)$ for iron primaries at $10^{18.5}$ eV energy and 30° zenith angle for the ADC reconstruction method. The blue and red histograms correspond respectively to the reconstructed non-saturated and saturated events.

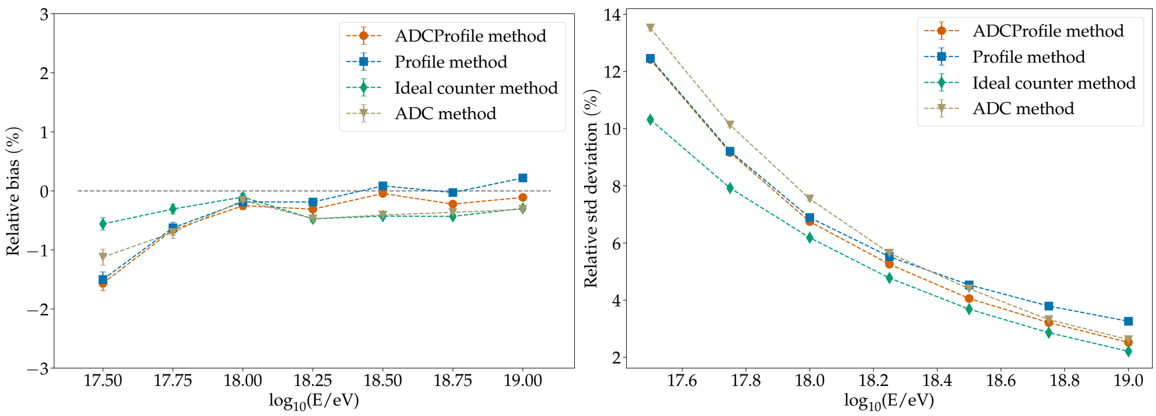

Figure 8.13: Relative bias (left panel) and the relative standard deviation (right panel) of the reconstructed $\widehat{\mu}_{\mathrm{ML}}(450)$ against the logarithmic energy for iron primaries at 30° zenith angle for different reconstruction methods, excluding saturated events.

The relative standard deviation of $\widehat{\mu}_{\mathrm{ML}}(450)$ for the ADCProfile reconstruction is closest to that of an ideal counter.

8.5 Summary and conclusions

The electronics of segmented muon detectors can operate in two independent acquisition modes: binary and ADC modes. Both acquisition modes are currently used by the UMD at Pierre Auger Observatory. In the past, muon detectors at AGASA also had both acquisition modes, which could be used independently or in combination to estimate the density of muons hitting the detector.

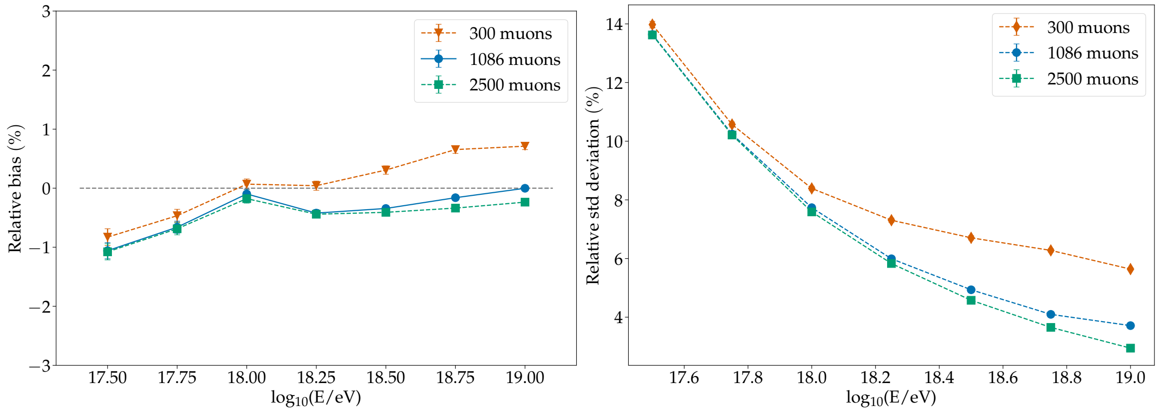

Figure 8.14: A comparison of the relative bias (left panel) and relative standard deviation (right panel) of the reconstructed $\widehat{\mu}_{\mathrm{ML}}(450)$ with the ADC reconstruction method at three different saturation levels of the LG channel as a function of the logarithm of energy.

In this chapter, two new methods are developed to reconstruct the muon lateral distribution function based on the ADC mode. It was found that the performance of the ADC method was similar to that of the profile method developed for the binary mode. Specifically, at lower energies, the profile method slightly outperforms the ADC-based method, whereas at high energies, the ADC-based method performs marginally better than the profile method.

A second method that combines both acquisition modes was also developed. The profile likelihood is used for muon detectors with relatively low muon densities, while the ADC likelihood is applied at high muon densities, where the binary response becomes saturated. The performance of this combined method is found to be similar to or better than that of the individual profile and ADC methods, depending on the primary energy. In particular, the detector resolution–defined as the standard deviation of the reconstructed number of muons at 450 m from the shower axis–is very close to that of an ideal counter across the considered energy range.

Overall, these results indicate that the combined method provides the most accurate reconstruction of the MLDF, effectively exploiting the complementary strengths of both acquisition modes. This is the first study to systematically investigate such a hybrid reconstruction approach. The performances of the different reconstruction methods were determined from simulations, and the combined method is well-suited for application to experimental data, where its consistency with other reconstruction approaches can be further validated.

Table 8.1: Summary of MLDF reconstruction methods.

Method	UMD mode	Likelihood formulation
Original (time-dependent)	Binary	• Combined likelihood (good, silent, saturated): Eq. (8.21) • Good: Poisson term • Silent: Poisson low-count probability ($n < 3$) (included in Eq. (8.21)) • Saturated (lower limit): one-sided Gaussian (included in Eq. (8.21))
New (single window)	Binary	• Binomial likelihood: Eq. (8.23) • Estimator: Eq. (8.24) • Silent: low-count likelihood (Eq. (8.25))
Profile	Binary	• Time-bin likelihood product: Eq. (8.26) • Profile likelihood: Eq. (8.27) • Likelihood ratio: Eq. (8.28) • Silent: Poisson low-count probability ($n < 3$) Eq. (8.29)
ADC (this thesis)	ADC	• Silent ($n_{\mathrm{LG}} \leq 3$): Eq. (8.30) • Intermediate ($3 < n_{\mathrm{LG}} < 200$): Eq. (8.31) • High ($n_{\mathrm{LG}} \geq 200$): Gaussian Eq. (8.32) • Saturated (lower limit): Gaussian tail Eq. (8.33)

ADCProfile (this thesis)	Binary + ADC	• Non-saturated (binary): Profile likelihood Eq. (8.27)
		• Silent: Eq. (8.29)
		• Binary saturation ($k \geq 124$): switch to ADC likelihood
		• ADC regimes:
		  • $3 < n_{\mathrm{LG}} < 200$: Eq. (8.31)
		  • $n_{\mathrm{LG}} \geq 200$: Eq. (8.32)
		  • ADC saturation (lower limit): Eq. (8.33)

CHAPTER 9

Summary and outlook

This thesis presented a comprehensive investigation of the muonic component of EAS measured with the UMD of Pierre Auger Observatory [10, 14]. The primary objectives were to improve the sensitivity to mass composition and to test the consistency between air-shower simulations and experimental data in the energy range $E/\text{eV} \in (10^{17.3}, 10^{18.7})$ eV. The work addressed three interconnected aspects of UHECR physics: (i) a quantitative assessment of the muon deficit through a direct comparison of measured and simulated radial muon densities, (ii) the construction of statistically robust composition-sensitive observables based on muon measurements and their correlation with the depth of shower maximum, X_{max}, and (iii) the development of improved muon reconstruction techniques. Together, these studies establish a coherent framework that links detector-level muon measurements, statistical inference of the primary mass composition, and reconstruction methodology.

The first major component of this work consisted of a detailed characterisation of the MLD measured with the UMD. Radial muon density profiles were constructed in bins of reconstructed energy and zenith angle and compared directly with detector-level simulations generated under mixed-composition scenarios constrained by fluorescence-detector measurements of mass fractions [5]. A key methodological feature of the analysis was the use of bin-by-bin comparisons of profiled muon densities without imposing an explicit parametric functional form for the MLDF. This approach preserved sensitivity to genuine radial shape differences between the data and the simulation. The comparison between measured and simulated muon density profiles was formulated as a single-parameter multiplicative scaling test. The fitted scale factor quantifies the normalisation mismatch between predicted and observed muon densities while preserving sensitivity to shape differences. The results show a systematic increase in the required scale factor with energy, indicating that the discrepancy between simulations and data grows toward higher energies. This behaviour is consistent across different scale estimators.

Near the reference distance of 450 m, where different MLD predictions tend to intersect, the discrepancy is well described by a predominantly multiplicative normalisation shift. Systematic uncertainties were evaluated at the radial-profile level and propagated to inferred scale factors. The dominant contributions arise from energy calibration and detector response modelling. Statistical uncertainties are negligible in comparison. Data and simulations are compatible with systematic uncertainties. Within systematic uncertainties, data and simulations remain formally compatible, although all contemporary high-energy hadronic interaction models considered in this work require a rescaling of the predicted muon densities. Therefore, this analysis confirms that current simulations systematically underestimate the muon content of air showers and that the discrepancy depends both on energy and on the high-energy hadronic interaction model.

The second component of this work focused on constructing observables sensitive to primary mass using the spatial structure of the muon distribution. A new observable, V_b, defined as a weighted sum of muon densities referenced to the optimal UMD distance scale, was introduced and systematically optimised. The optimised observable $V_{b_{\max}}$ captures both the total muon content and the radial structure of the shower without requiring an explicit MLDF parametrisation. The discrimination power between light and heavy primaries increases with energy, exhibits weak dependence on zenith angle, and remains robust across hadronic interaction models. Importantly, $V_{b_{\max}}$ shows stable behaviour in different hadronic interaction models, indicating limited model sensitivity at the observable level. Energy-normalised variants do not improve separation, implying that the intrinsic energy scaling of the muon signal carries physically relevant composition information.

In parallel, the SD risetime asymmetry method was evaluated as an indirect estimator of $X_{\max}$. Although the method reproduces expected qualitative trends, the reconstructed values exhibit composition-dependent bias and reduced separation power relative to true fluorescence measurements. This limits its direct applicability for precision composition inference without additional model-dependent corrections.

Building on these observables, the thesis investigated the statistical correlation between muon content and shower maximum depth as a composition-sensitive probe. Rank-based estimators were studied in detail, and Kendall's rank correlation coefficient [21, 74] was identified as the most suitable estimator due to its robustness against global normalisation shifts and monotonic transformations. Simulations demonstrate a monotonic relationship between the preferred correlation coefficient and the mean logarithmic mass of the primary composition. The method, therefore, enables discrimination between pure- and mixed-composition scenarios. The statistical power of the approach was quantified by deriving the minimum number of events required to reject the pure-composition hypotheses at a fixed significance. Using the measured CR flux and the expected exposure of Pierre Auger Observatory [4], sufficient statistics are available at lower energies $\log_{10}(E/\mathrm{eV}) \lesssim 17.8$ to reject the hypothesis of pure composition. However, current hybrid datasets remain statistically limited for definitive constraints at higher energies. An exploratory comparison with experimental data confirms consistency with simulated composition mixtures but does not yet allow statistically significant discrimination. The framework nonetheless establishes a statistically well-defined and model-robust method for future high-statistics analyses.

To exploit the nearly continuous duty cycle of the SD, the correlation method was extended to a SD framework by introducing controlled resolution smearing of $X_{\max}$ to emulate realistic SD-based reconstruction. The intrinsic anti-correlation between muon content and the depth of the shower maximum is preserved under realistic detector resolution, although its magnitude decreases with increasing uncertainty. The correlation coefficient remains sensitive not only to the mean mass but also to the dispersion of the mass distribution. Its statistical uncertainty scales as expected with event count, and the estimator shows stability under energy reconstruction uncertainties and detector-resolution effects. These properties make correlation-based approaches well-suited for future high-statistics composition studies using SD measurements, including potential $X_{\max}$ estimators derived from advanced reconstruction techniques such as machine-learning approaches.

Accurate characterisation of the MLDF requires reliable reconstruction of muon densities from detector signals. The final component of this thesis developed new reconstruction methods that exploit the dual acquisition capabilities of the UMD electronics, which operate in binary and analogue-to-digital converter (ADC) modes. Two reconstruction approaches based on ADC measurements were developed and validated using simulated detector responses. Their performance was compared with the established profile-likelihood reconstruction method used for binary-mode data. The method based on ADC exhibits a comparable overall performance, with slightly improved resolution at high energies and slightly reduced performance at low energies. A hybrid reconstruction scheme that combined both acquisition modes was also introduced. In this approach, binary-mode likelihoods are applied in the low-density regime, while ADC-based likelihoods are used at high muon densities. The combined method achieves performance comparable to or better than either individual approach and approaches the resolution of an idealised detector across the explored energy range. These results demonstrate that ADC information provides a viable and competitive basis for MLDF reconstruction and that combining acquisition modes improves stability under varying detector occupancy conditions. The method can be applied directly to experimental data and provides a foundation for improved muon-based observables.

9.1 Overall Conclusions

The results of this thesis provide a comprehensive investigation of muon-based observables in UHECR air showers. Several key conclusions emerge:

- Air-shower simulations systematically underestimate the muon content measured by the UMD, and the discrepancy increases with energy. All contemporary high-energy hadronic interaction models considered require rescaling relative to the data.
- The optimised observable $V_{b_{\max}}$ provides a robust, model-independent measure of muon content and radial structure with strong discrimination power between light and heavy primaries and is a useful tool for composition analysis. The $X_{\max}$ obtained from the $\langle\Delta\rangle$ method has a bias dependent on composition, and hence this data-driven method is not useful for composition studies.
- The statistical correlation between muon observables and shower maximum is a powerful composition-sensitive quantity that is robust to calibration uncertainties. Comparison with expected detector exposure shows that sufficient statistics are available at lower energies $\log_{10}(E/\mathrm{eV}) \lesssim 17.8$ to reject the hypothesis of pure composition.
- Correlation-based methods remain effective under realistic SD resolution and provide sensitivity to both mean mass and mass dispersion.
- New reconstruction methods based on the ADC and hybrid acquisition modes provide reliable and high-resolution estimates of muon densities. This can be incorporated into the future to improve composition analysis.

9.2 Outlook

Future improvements in detector exposure, reconstruction accuracy, and hadronic interaction modelling will enhance the statistical power of correlation-based composition studies. In particular, high-statistics SD measurements combined with improved modelling of hadronic interactions will allow increasingly precise constraints on the primary mass distribution. Continued comparison of radial muon profiles with simulations will also provide critical input for refining hadronic interaction models at energies beyond the reach of accelerators.

The combined ADC-binary MLDF reconstruction method developed in this thesis can provide improved estimates of $\rho_\mu(450)$, which may serve as an alternative normalisation observable to be combined with $X_{\max}$ in future correlation analyses. The systematic integration of muon measurements, reconstruction methodology and statistical inference developed here contributes to the larger effort to understand the origin and nature of UHECR.

APPENDIX A

Additional plots of the comparison of the muon content between data and simulations

A.1 Additional energy bins

The density profile of muons and residuals is shown as a function of $\log_{10}(r/\mathrm{m})$ for different energy bins for various high-energy hadronic interaction models in the zenith angle range $[0°, 45°]$ in the following subsections. (Top) Mean muon density as a function of radial distance compared with the scaled AugerMix reference model. Different curves correspond to scale factors derived from various fit windows, including the full radial range and restricted intervals. The grey band indicates the systematic uncertainty of the detector. The multiple fitted curves correspond to scale factors obtained from different radial fit ranges or definitions, including f_{full}, f_1, f_2, and f_{450}. (Bottom) Residuals between data and the scaled model, with dashed lines marking statistical confidence limits, and the grey region showing the systematic band.

EPOS-LHC

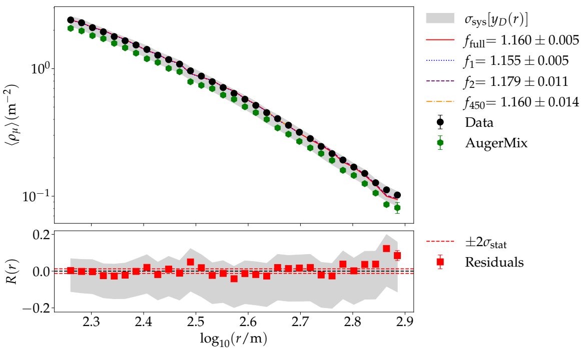

$\log_{10}(E/\mathrm{eV}) \in [17.3, 17.5)$

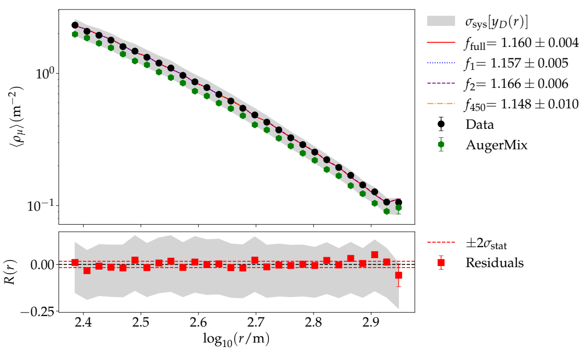

$\log_{10}(E/\mathrm{eV}) \in [17.5, 17.7)$

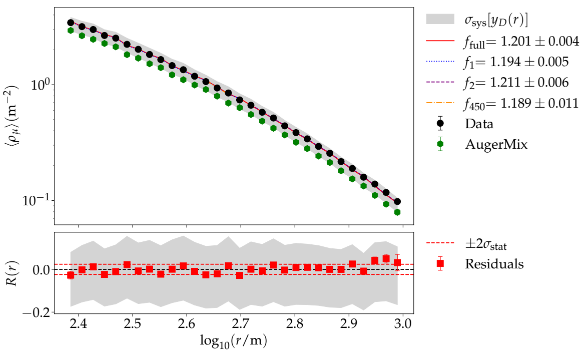

$\log_{10}(E/\mathrm{eV}) \in [17.7, 17.9)$

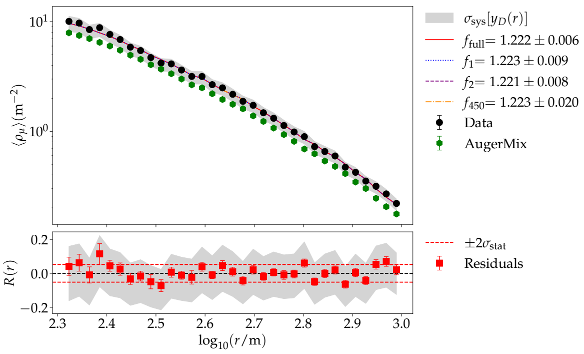

$\log_{10}(E/\mathrm{eV}) \in [18.1, 18.3)$

Figure A.1: Muon density profile and residuals are shown as a function of $\log_{10}(r/\mathrm{m})$ for different energy bins for EPOS-LHC high-energy hadronic interaction model in the zenith angle range $[0°, 45°]$

Sibyll-2.3c

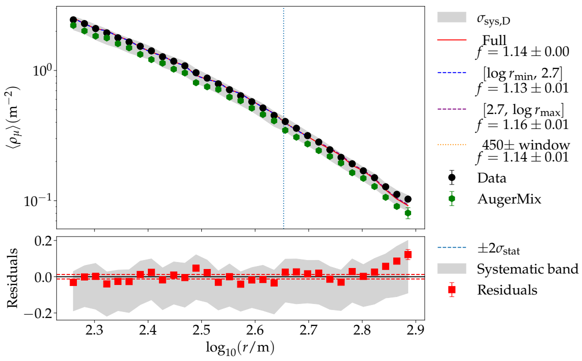

$\log_{10}(E/\mathrm{eV}) \in [17.3, 17.5)$

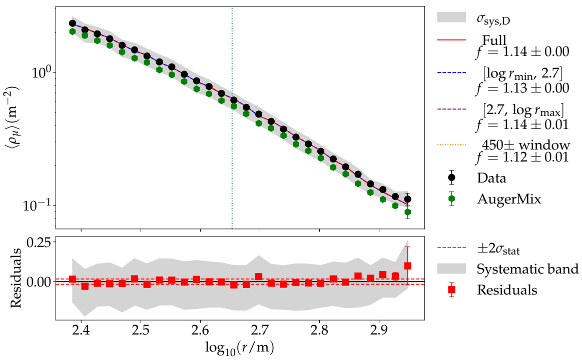

$\log_{10}(E/\mathrm{eV}) \in [17.5, 17.7)$

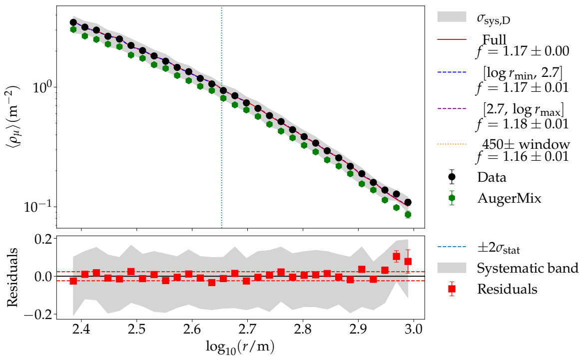

$\log_{10}(E/\mathrm{eV}) \in [17.7, 17.9)$

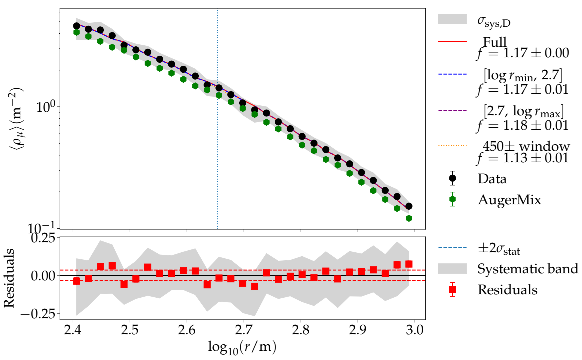

$\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$

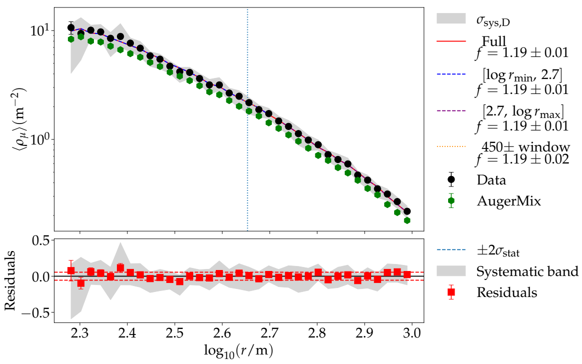

$\log_{10}(E/\mathrm{eV}) \in [18.1, 18.3)$

Figure A.2: Muon density profile and residuals are shown as a function of $\log_{10}(r/\mathrm{m})$ for different energy bins for Sibyll-2.3c high-energy hadronic interaction model in the zenith angle range $[0^\circ, 45^\circ]$

QGSJet-II.04

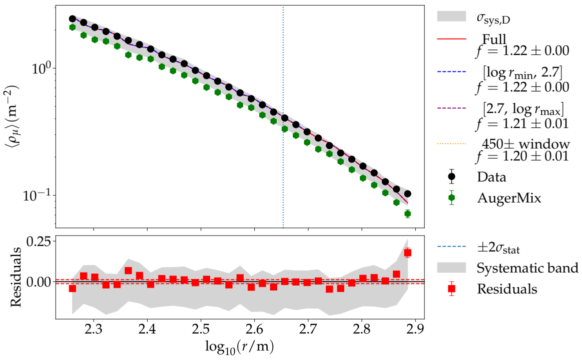

$\log_{10}(E/\mathrm{eV}) \in [17.3, 17.5)$

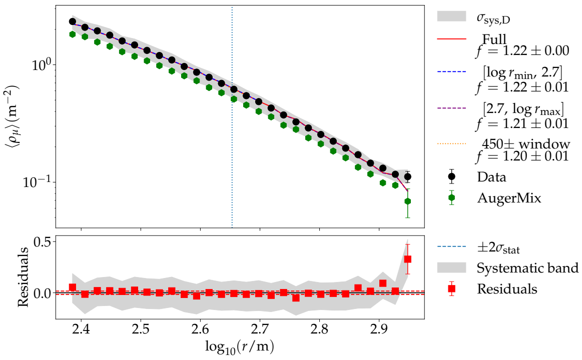

$\log_{10}(E/\mathrm{eV}) \in [17.5, 17.7)$

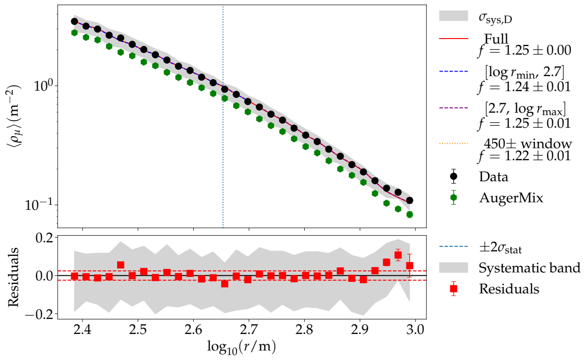

$\log_{10}(E/\mathrm{eV}) \in [17.7, 17.9)$

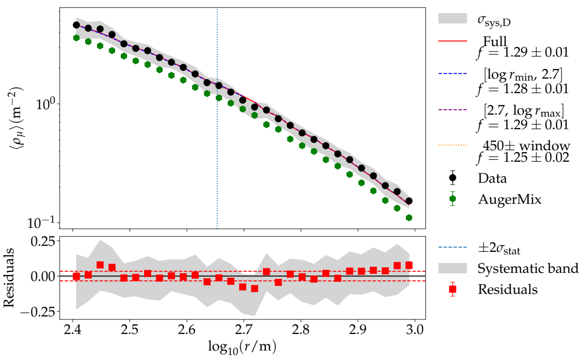

$\log_{10}(E/\mathrm{eV}) \in [17.9, 18.1)$

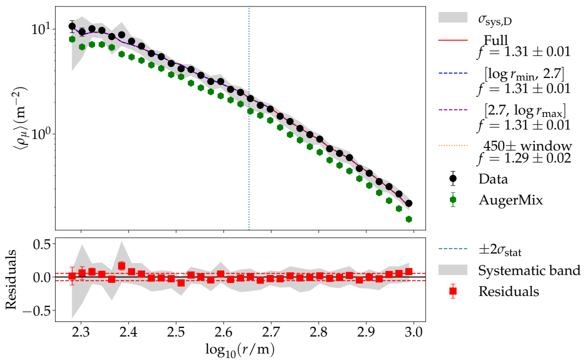

$\log_{10}(E/\mathrm{eV}) \in [18.1, 18.3)$

Figure A.3: Muon density profile and residuals are shown as a function of $\log_{10}(r/\mathrm{m})$ for different energy bins for QGSJet-II.04 high-energy hadronic interaction model in the zenith angle range $[0°, 45°]$

APPENDIX B

Muon density based observables

B.1 Discrimination power of observables: Fixed energies

To study the maximum degree of separation between proton and iron distributions without an energy bias, a MC shower library of proton and iron showers generated with CORSIKA v7.5 with EPOS-LHC as the high-energy hadronic interaction model and FLUKA [53] as the low-energy hadronic interaction model with fixed energies of $\log_{10}(E/\mathrm{eV}) \in [17.5, 18.5]$ in steps of $\log_{10}(E/\mathrm{eV}) = 0.25$ and two zenith angles $\theta = 32°, 38°$ was used as input for the simulation and reconstruction of the UMD and SD using $\overline{\text{Off}}\underline{\text{line}}$. The merit factor defined in Eq. (5.4) for V_b was evaluated for different b values ranging from 0 to 6 in steps of 0.1 for different bins in $\log_{10}(E/\mathrm{eV})$, as shown in Fig. B.1.

As seen in Fig. B.1, the merit factor increases with increasing energy, which implies that the ability of the parameter V_b to separate the iron and proton primaries increases with increasing energy. The value of the index b, which gives the maximum merit factor, also varies with energy. To select an average index value b that maximises the merit factor over a wide range of energies, it is necessary to redo the analysis with a continuous energy simulation library.

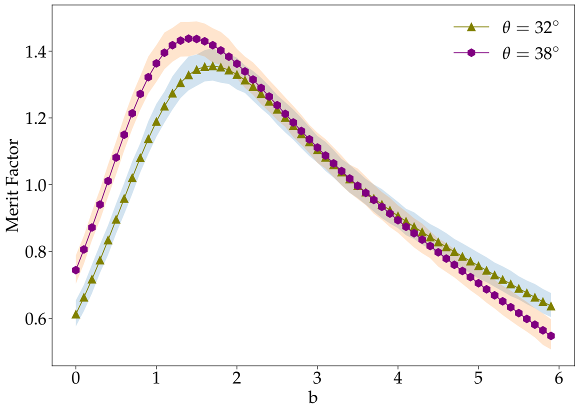

$\log_{10}(E/\mathrm{eV}) = 17.5$

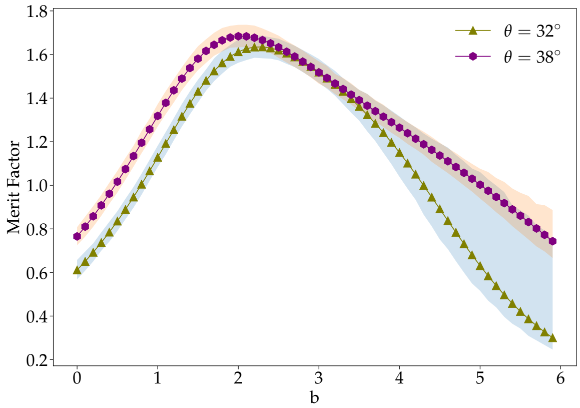

$\log_{10}(E/\mathrm{eV}) = 18.0$

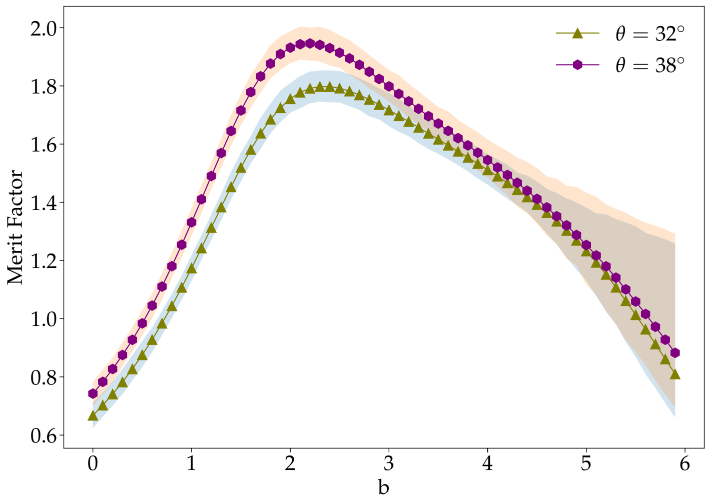

$\log_{10}(E/\mathrm{eV}) = 18.5$

Figure B.1: The merit factor defined in Eq. (5.4) for V_b evaluated for different b values ranging from 0 to 6 in steps of 0.1 for different bins in $\log_{10}(E/\mathrm{eV})$.

APPENDIX C

Additional tests for the correlation analysis

C.1 Statistical uncertainty of correlation coefficients

The statistical uncertainty of the correlation estimators for EPOS-LHC high-energy hadronic interaction model was discussed in Section. 6.2 with Fig. 6.2. The same was evaluated for high-energy hadronic interaction models QGSJet-II.04 (left panel) and Sibyll-2.3c (right panel) and is shown in Fig. C.1. The behaviour of $\sigma\sqrt{N}$ is shown as a function of the sample size N. The results obtained are in agreement with the results obtained for EPOS-LHC high-energy hadronic interaction model and further cement the use of a constant statistical uncertainty of $0.67/\sqrt{N}$ for Kendall's correlation coefficient, irrespective of the high-energy hadronic interaction model used.

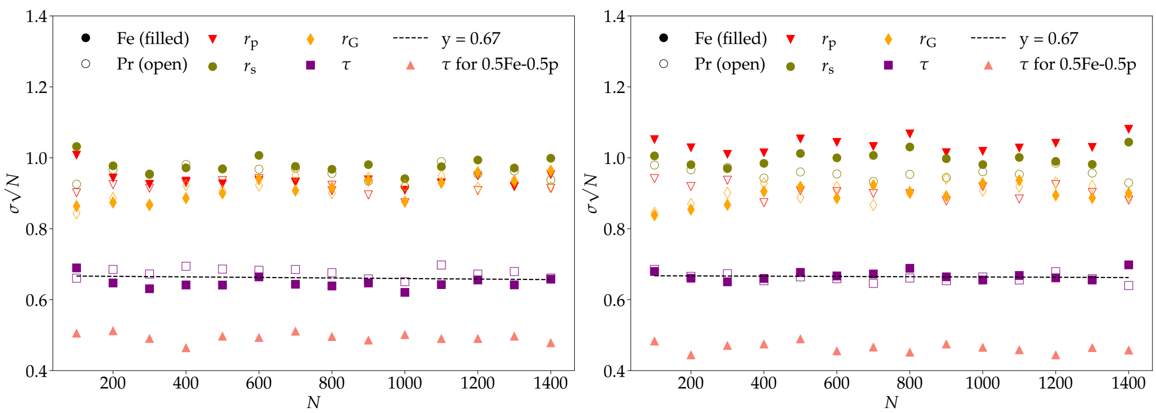

Figure C.1: Statistical uncertainty of the estimators, shown as a function of sample size N for high-energy hadronic interaction models QGSJet-II.04 (left panel) and Sibyll-2.3c (right panel).

C.2 Verifying the Gaussian behaviour of Kendall's correlation coefficient

To characterise the statistical behaviour of the correlation estimator $\tau(V_{b_{\max}}, X_{\max})$, bootstrap resampling was performed within a fixed energy bin $\log_{10}(E/\mathrm{eV}) \in [18.0, 18.1)$. For each primary species, the event sample was repeatedly resampled with replacement and $\tau(V_{b_{\max}}, X_{\max})$ was calculated for each resample. This procedure generates an empirical sampling distribution for the estimator, representing the statistical fluctuations expected for a finite number of events. The resulting bootstrap distributions for the proton and iron primaries are shown in Fig. C.2, together with Gaussian probability density functions constructed from the sample mean and standard deviation of the bootstrap values.

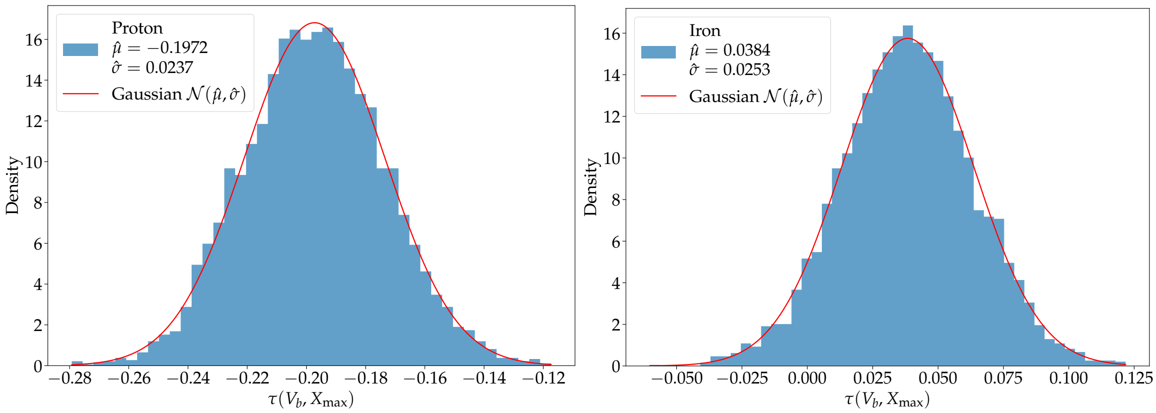

Figure C.2: Bootstrap sampling distributions of $\tau(V_{b_{\max}}, X_{\max})$ for proton and iron primaries in a fixed energy bin. Red curves show Gaussian probability density functions constructed from the sample mean and standard deviation of the bootstrap values.

Statistical analysis assumes that the bootstrap sampling distribution of the correlation coefficient is approximately Gaussian. This assumption is tested using normal Q–Q plots as shown in Fig. C.3, which compare the ordered bootstrap values with the theoretical quantiles of a normal distribution with the same mean and variance. If the data are normally distributed, the points lie along a straight line. For both proton and iron primaries, the Q–Q plots show excellent linear agreement over the full central range, with only minor deviations in the extreme tails. These small tail deviations are expected for bounded statistics such as τ and do not significantly affect central uncertainty estimates. The Q–Q analysis, therefore, supports modelling the correlation coefficient as a Gaussian distribution within each energy bin. This validation is essential for the subsequent propagation of uncertainties and for deriving the reciprocal-normal distribution of the minimum event requirement.

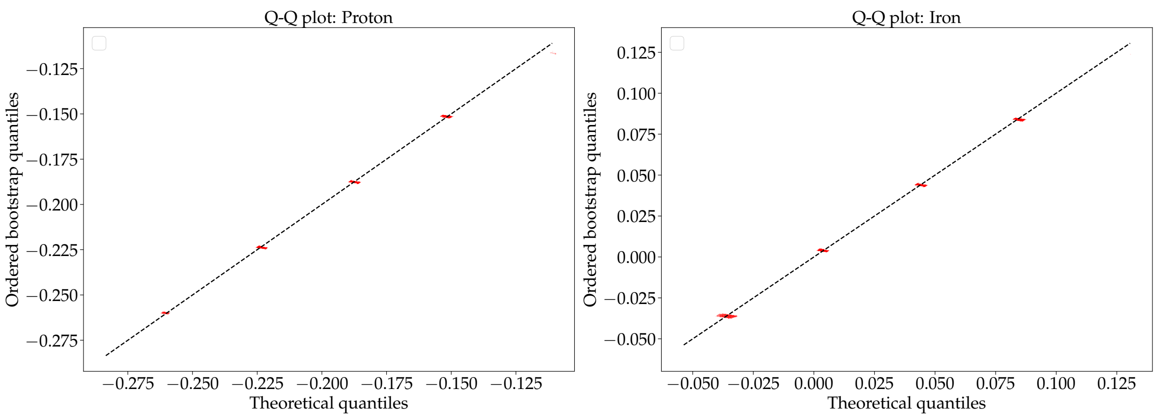

Figure C.3: Normal Q-Q plots of bootstrap estimates of $\tau(V_{b_{\max}}, X_{\max})$ for proton and iron primaries. The dashed line indicates the expected relation for a Gaussian distribution with the same mean and variance. The near-linear behaviour confirms the approximate normality of the sampling distribution.

The combined evidence from histograms and Q–Q plots supports the approximation in $\tau \sim \mathcal{N}(\mu_\tau, \sigma_\tau)$ within a fixed energy bin. In particular, the difference between the two correlation estimates can be modelled as a Gaussian, allowing analytic uncertainty propagation for derived quantities such as $\Delta\tau$ and $\sqrt{N_{\text{min}}}$. Thus, the bootstrap study provides the statistical foundation for the Gaussian modelling assumptions used in the main analysis.

C.3 Reciprocal transformation of a Gaussian-distributed correlation

Figure C.4 shows the probability density of the transformed variable $x = \sqrt{N_{\text{min}}} = A/\Delta\tau$, obtained from the reciprocal transformation of a Gaussian-distributed correlation difference. The histogram represents the distribution generated by sampling MC, while the solid curve shows the reciprocal analytical normal probability density derived from the assumed Gaussian behaviour of $\Delta\tau$. The excellent agreement between the numerical and analytical curves confirms that the reciprocal transformation of a normally distributed variable produces the expected asymmetric, right-skewed distribution. This skewness reflects the non-linear propagation of uncertainty under the reciprocal mapping. Small values of $\Delta\tau$ correspond to large values of x, generating an extended high-value tail. Physically, this behaviour implies that rare statistical fluctuations leading to small composition separation require disproportionately large event samples to achieve the same level of significance. The agreement between MC and the analytical results validates the reciprocal-normal model used to describe the statistical distribution of the minimum event requirement.

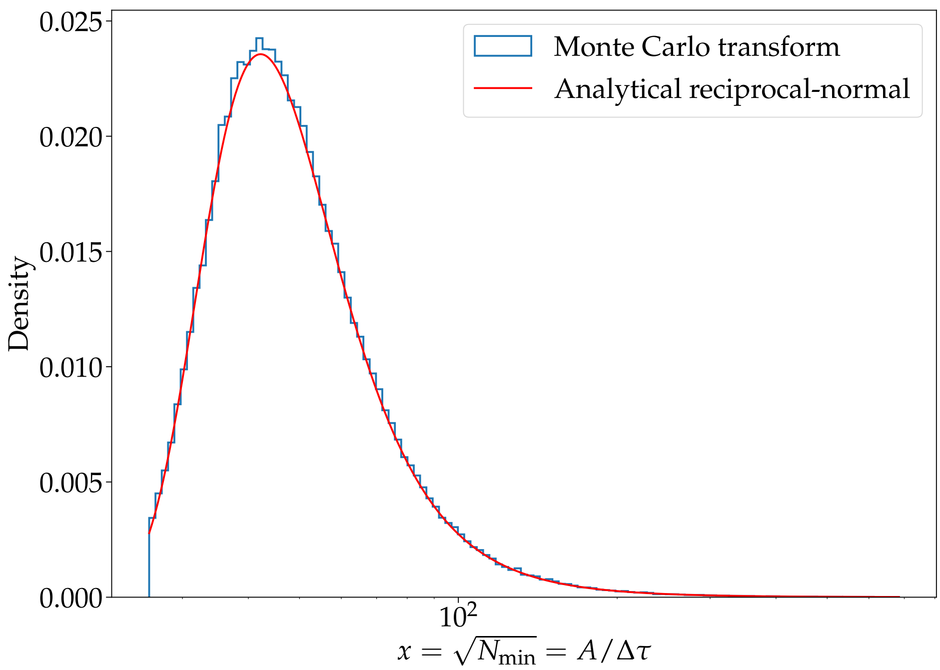

Figure C.4: Probability density of the transformed variable $x = \sqrt{N_{\text{min}}} = A/\Delta\tau$. The histogram shows the distribution obtained from MC sampling of a Gaussian-distributed $\Delta\tau$, while the solid curve represents the analytical reciprocal-normal probability density. The close agreement confirms the validity of the reciprocal-normal model for the minimum event requirement.

C.4 Cumulative distribution function of the mean estimator of Kendall's correlation coefficient

The cumulative distribution function of $x = \sqrt{N_{\text{min}}} = A/\Delta\tau$ where $A = 5 \times 0.67$ is given by

$$F(x) = \begin{cases} \frac{1}{2}\left(\mathrm{Erf}\left(\frac{\langle\Delta\hat{\tau}\rangle - \frac{A}{x}}{\sqrt{2}\,\sigma_{\Delta\hat{\tau}}}\right) - \mathrm{Erf}\left(\frac{\langle\Delta\hat{\tau}\rangle}{\sqrt{2}\,\sigma_{\Delta\hat{\tau}}}\right)\right), & x < 0, \\ \frac{1}{2}\left(1 - \mathrm{Erf}\left(\frac{\langle\Delta\hat{\tau}\rangle}{\sqrt{2}\,\sigma_{\Delta\hat{\tau}}}\right)\right), & x = 0, \\ 1 + \frac{1}{2}\left(\mathrm{Erf}\left(\frac{\langle\Delta\hat{\tau}\rangle - \frac{A}{x}}{\sqrt{2}\,\sigma_{\Delta\hat{\tau}}}\right) - \mathrm{Erf}\left(\frac{\langle\Delta\hat{\tau}\rangle}{\sqrt{2}\,\sigma_{\Delta\hat{\tau}}}\right)\right), & x > 0. \end{cases} \tag{C.1}$$

where $\mathrm{Erf}(x)$ corresponds to the error function.

APPENDIX D

Additional tests for the reconstruction methods

D.1 Verification of log-normal character of the sum of n log-normal variables

As shown in Eq. (8.9), when $\theta^2 \lesssim 1$, the sum of n log-normal terms is expected to exhibit log-normal behaviour. The simulations confirmed that Q has log-normal characteristics.

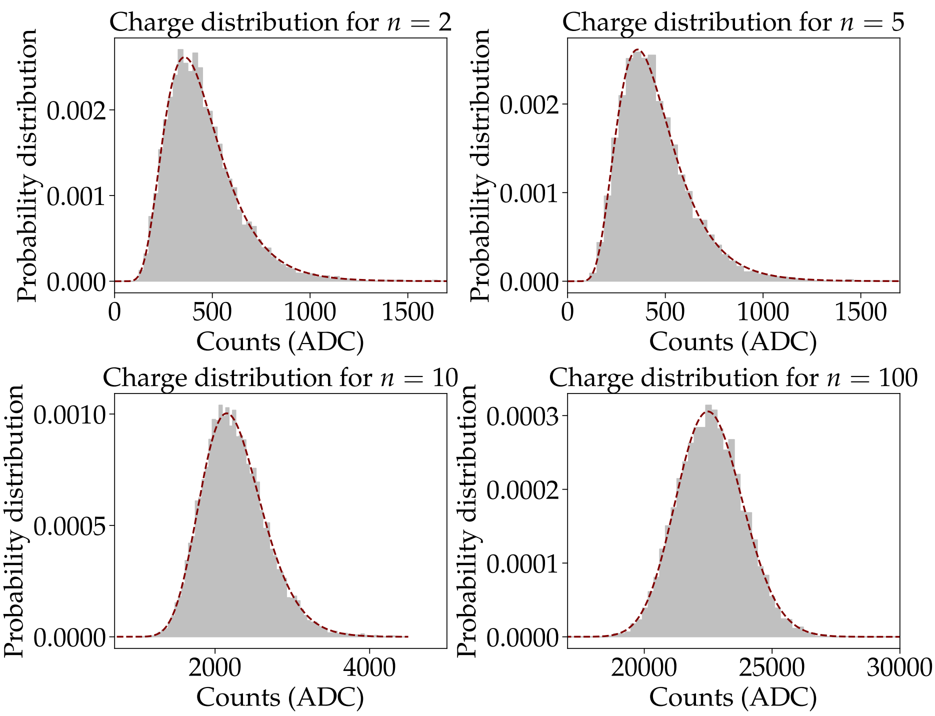

Figure D.1: Comparison of the simulated charge distribution with the analytical expression in Eq. (8.9) for the LG channel of the UMDs at Pierre Auger Observatory.

Fig. D.1 compares histograms of the sum of n random log-normal variables with the corresponding log-normal functions (dotted lines) with parameter values $m = 5$ and $\theta = 0.5$ in Eqs. 8.10 and 8.11 [40]. The figure shows that Eq. (8.9) is an excellent approximation for the LG channel. Therefore, when the variance of the distribution is small ($\theta^2 \lesssim 1$), the sum of n log-normal terms behaves similarly to the sum of narrowly distributed random variables, regardless of the value of n. The figures reveal that for smaller values of n, the distribution has a prominent tail towards positive total-charge values, indicating a larger shape parameter. However, for $n = 100$, the tail of the distribution was almost negligible, and the shape parameter was insignificant.

D.2 Comparison of Gaussian and Compound likelihoods

To compare $f_C(Q|\mu)$ in Eq. (8.12) with a Gaussian distribution $f_G(Q|\mu)$ in Eq. (8.13) more efficiently, a parameter of the following form is considered,

$$\tilde{\varepsilon} = 100 \times \left(\frac{f_G}{f_C} - 1\right) \text{ in \%}. \tag{D.1}$$

Fig. D.2 shows the plot of $\tilde{\varepsilon}$ as a function of the z-score for various values of μ. The z-score is a statistical measurement that describes the relationship of a value to the mean of a group of values. The z-score is measured in terms of standard deviations from the mean and is defined as follows,

$$\text{z-score} = \frac{Q - \langle Q \rangle}{\sigma[Q]}. \tag{D.2}$$

From Fig. D.2, the absolute value of $\tilde{\varepsilon}$ is less than 5% for the LG channel of the ADC in the $2\,\sigma$ region of the distribution for $\mu = 200$. Hence, it is safe to conclude that the $f_C(Q|\mu)$ distribution can be approximated by a Gaussian distribution for values of μ of the order of or greater than 200.

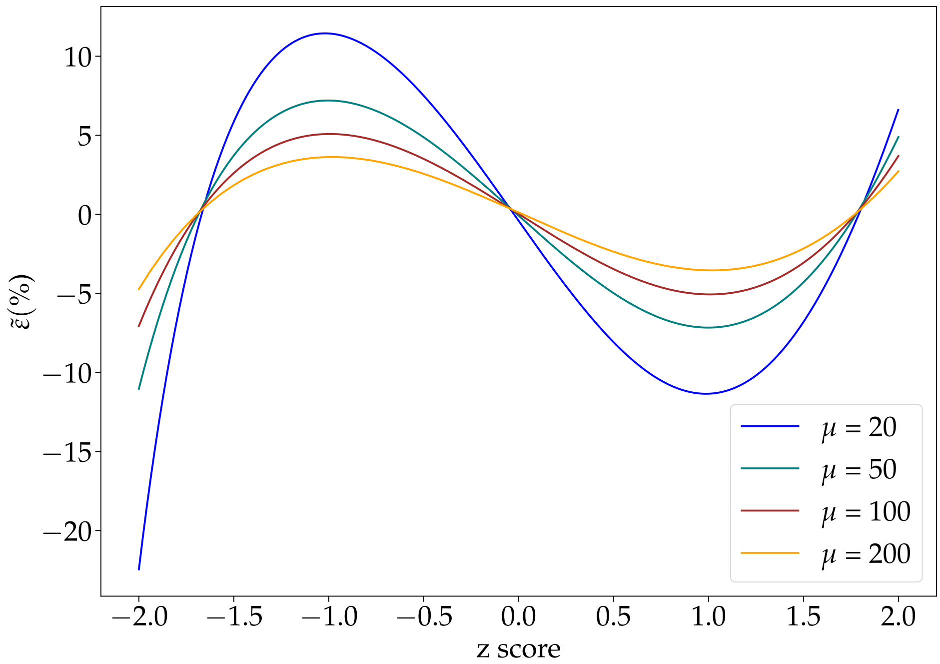

Figure D.2: $\tilde{\varepsilon}$ as a function of the z-score of the LG channel.

D.3 Effect of inhibition window on binary mode saturation

The selection of the inhibition window in binary mode affects the number of saturated events. To study this effect in detail, 10,000 events were reconstructed using the Profile reconstruction method with an inhibition window of 25 ns and compared with the 40-ns window used in this analysis. Fig. D.3 displays the fraction of saturated events plotted against the logarithmic energy. As seen in the figure, the choice of the inhibition window affects the fraction of saturated events in the Profile reconstruction method. The fraction of saturated events was higher for the Profile method with a 40 ns window and increased rapidly with energy compared to the 25 ns window.

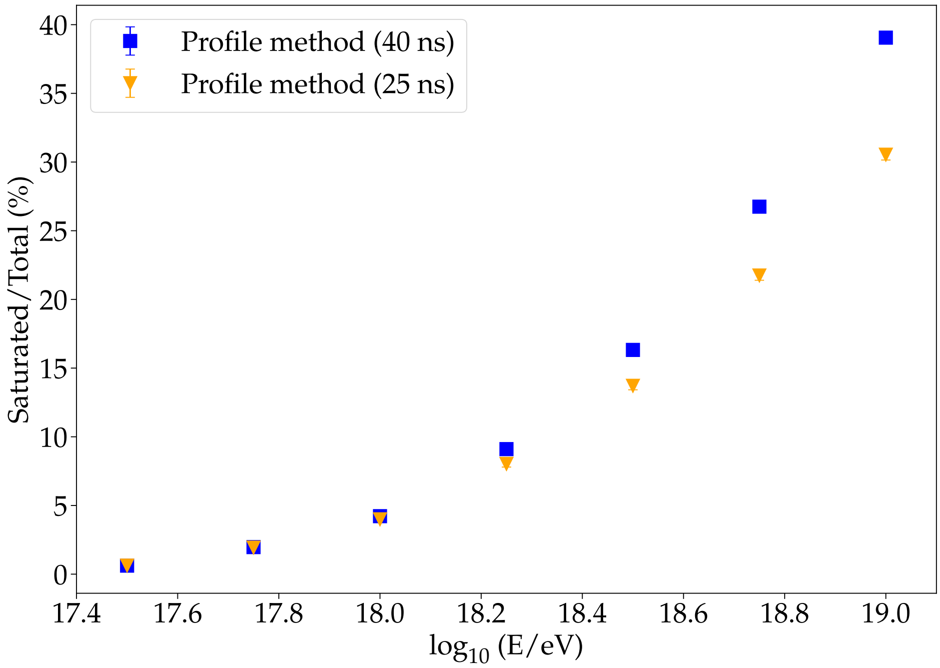

Figure D.3: Fraction of saturated events for iron primaries at 30° zenith angle as a function of the logarithm of primary energy for the Profile reconstruction method at 25 and 40 ns inhibition windows with error bars. The error bars are smaller than the marker size.

List of Tables

List of Figures

Bibliography

[1] T. K. Gaisser, R. Engel, E. Resconi, *Cosmic Rays and Particle Physics*, Cambridge University Press (2016)doi:10.1017/CBO9781139192194.

[2] S. Navas *et al.* (Particle Data Group), *Review of Particle Physics*, Phys. Rev. D **110** (2024) 030001doi:10.1103/PhysRevD.110.030001.

[3] K. Kotera, A. V. Olinto, *The Astrophysics of Ultrahigh-Energy Cosmic Rays*, Annu. Rev. Astron. Astrophys. **49** (2011) 119–153doi:10.1146/annurev-astro-081710-102620.

[4] Pierre Auger Collaboration, *Measurement of the cosmic-ray energy spectrum above* 2.5×10^{18} *eV using the Pierre Auger Observatory*, Phys. Rev. D **102** (2020) 062005doi:10.1103/PhysRevD.102.062005.

[5] Pierre Auger Collaboration, *Features of the Energy Spectrum of Cosmic Rays above* 2.5×10^{18} *eV Using the Pierre Auger Observatory*, Phys. Rev. Lett. **125** (2020) 121106doi:10.1103/PhysRevLett.125.121106.

[6] Pierre Auger Collaboration, *Depth of Maximum of Air-Shower Profiles at the Pierre Auger Observatory*, Phys. Rev. D **90** (2014) 122005doi:10.1103/PhysRevD.90.122005.

[7] Pierre Auger Collaboration, *The Fluorescence Detector of the Pierre Auger Observatory*, Nucl. Instrum. Meth. A **620** (2010) 227–251doi:10.1016/j.nima.2010.04.023.

[8] Pierre Auger Collaboration et al., *Observation of a large-scale anisotropy in the arrival directions of cosmic rays above* 8×10^{18} *eV*, Science **357** (2017) 1266-1270DOI:10.1126/science.aan4338.

[9] T. K. Gaisser, A. M. Hillas, *Reliability of the Method of Constant Intensity Cuts for Reconstructing the Average Development of Vertical Showers*, In: Proceedings of the 15th International Cosmic Ray Conference (ICRC 1977), August 13–26, Plovdiv, Bulgaria. 1977.

[10] Pierre Auger Collaboration, *Calibration of the underground muon detector of the Pierre Auger Observatory*, JINST **16** (2021) P04003doi:10.1088/1748-0221/16/04/P04003.

[11] F. Salamida, *Highlights from the Pierre Auger Observatory*, PoS **ICRC2023** (2023), 016 doi:10.22323/1.444.0016.

[12] T. Pierog *et al.*, *EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider*, Phys. Rev. C **92** (2015) 034906doi:10.1103/PhysRevC.92.034906.

[13] C. Meurer and N. Scharf, *HEAT - A low energy enhancement of the Pierre Auger Observatory*, Astrophysics and Space Sciences Transactions **7** (2011) 183-186 10.5194/astra-7-183-2011.

[14] Pierre Auger Collaboration, *The Pierre Auger Cosmic Ray Observatory*, Nucl. Instrum. Meth. A **798** (2015) 172–213doi:10.1016/j.nima.2015.06.058.

[15] A. Coleman, J. Eser, E. Mayotte, F. Sarazin, F. G. Schröder, D. Soldin, T. M. Venters, R. Aloisio, J. Alvarez-Muñiz and R. Alves Batista, *et al.*, *Ultra high energy cosmic rays: The intersection of the Cosmic and Energy Frontiers*, Astropart. Phys. **149** (2023) 102819 doi:10.1016/j.astropartphys.2023.102819.

[16] S. Ostapchenko, *Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: QGSJET-II model*, Phys. Rev. D **83** (2011) 014018doi:10.1103/PhysRevD.83.014018.

[17] F. Riehn *et al.*, *Hadronic interaction model Sibyll 2.3d and extensive air showers*, Phys. Rev. D **102** (2020) 063002doi:10.1103/PhysRevD.102.063002.

[18] S. T. Hahn, *Methods for Estimating Mass-Sensitive Observables of Ultra-High Energy Cosmic Rays using Artificial Neural Networks*, PhD thesis (2023) Karlsruher Institut für Technologie (KIT) 10.5445/IR/1000154770.

[19] Pierre Auger Collaboration, *Testing Hadronic Interactions at Ultrahigh Energies with Air Showers Measured by the Pierre Auger Observatory*, Phys. Rev. Lett. **117** (2016) 192001doi:10.1103/PhysRevLett.117.192001.

[20] Pierre Auger Collaboration, *Muons in Air Showers at the Pierre Auger Observatory*, Phys. Rev. D **91** (2015) 032003doi:10.1103/PhysRevD.91.032003.

[21] W. J. Conover, *Practical Nonparametric Statistics*, Wiley (3rd ed., 1999).

[22] S. Mollerach, E. Roulet, *Progress in high-energy cosmic ray physics*, Prog. Part. Nucl. Phys. **98** (2018) 85–118 doi:10.1016/j.ppnp.2017.10.002.

[23] T. K. Gaisser, T. Stanev, S. Tilav, *Cosmic Ray Energy Spectrum from Measurements of Air Showers*, Front. Phys. **8** (2013) 748–758 doi:10.1007/s11467-013-0319-7.

[24] V. Berezinsky, A. Z. Gazizov, S. I. Grigorieva, *Dip in UHECR spectrum as signature of proton interaction with CMB*, Phys. Lett. B **612** (2005) 147–153 doi:10.1016/j.physletb.2005.02.058.

[25] L. A. Anchordoqui, *Ultra-high-energy cosmic rays*, Phys. Rep. **801** (2019) 1–93 doi:10.1016/j.physrep.2019.01.002.

[26] R. Ruffini, G. V. Vereshchagin, S. S. Xue, *Cosmic absorption of ultra high energy particles*, Astrophys. Space Sci. **361** (2016) 82 doi:10.1007/s10509-016-2668-5.

[27] R. Aloisio, V. Berezinsky, A. Gazizov, *Transition from galactic to extragalactic cosmic rays*, Astropart. Phys. **39–40** (2012) 129–143 doi:10.1016/j.astropartphys.2012.09.007.

[28] K. Greisen, *End to the Cosmic-Ray Spectrum?*, Phys. Rev. Lett. **16** (1966) 748–750 10.1103/PhysRevLett.16.748.

[29] G. T. Zatsepin, V. A. Kuzmin, *Upper limit of the spectrum of cosmic rays*, JETP Lett. 4 (1966) 78-80, Pisma Zh.Eksp.Teor.Fiz. 4 (1966) 114-117.

[30] A. Haungs *et al.*, *KCDC: the Kascade Cosmic Ray Data Centre*, Journal of Physics: Conference Series, 632 01 2011 DOI 10.1088/1742-6596/632/1/012011.

[31] A. M. Hillas, *The Origin of Ultra-High-Energy Cosmic Rays*, Annu. Rev. Astron. Astrophys. **22** (1984) 425–444 doi:10.1146/annurev.aa.22.090184.002233.

[32] M. Teshima, *Expanded array for giant air shower observation at Akeno*, Nucl. Instrum. Methods Phys. Res. A **247** (1986) doi:10.1016/0168-9002(86)91324-0.

[33] K. Shinozaki, M. Teshima, *AGASA results*, Nucl. Phys. B Proc. Suppl. **136** (2004) doi:10.1016/j.nuclphysbps.2004.10.045.

[34] N. Hayashida, *Muons (≥ 1 GeV) in large extensive air showers of energies between $10^{16.5}$ eV and $10^{19.5}$ eV observed at Akeno*, J. Phys. G **21** (1995) 1101doi:10.1088/0954-3899/21/8/008.

[35] T. Abu-Zayyad *et al.*, *The surface detector array of the Telescope Array experiment*, Nucl. Instrum. Meth. A **689** (2012) 87–97 doi:10.1016/j.nima.2012.05.079.

[36] W. D. Apel *et al.*, *The KASCADE-Grande experiment*, Nucl. Instrum. Methods Phys. Res. A **620** (2010) 202–216 doi:10.1016/j.nima.2010.03.147.

[37] B. Wundheiler, *The AMIGA Muon Counters of the Pierre Auger Observatory: Performance and Studies of the Lateral Distribution Function*, PoS (ICRC2015) **236** (2016) doi:10.22323/1.236.0324.

[38] F. Gesualdi, A. D. Supanitsky, *Estimation of the number of counts on a particle counter detector with full time resolution*, Eur. Phys. J. C **82** (2022) 925 doi:10.1140/epjc/s10052-022-10895-9.

[39] Pierre Auger Collaboration, *Design and implementation of the AMIGA embedded system for data acquisition*, JINST **16** (2021) T07008 doi:10.1088/1748-0221/16/07/T07008.

[40] A. M. Botti, *Determination of the chemical composition of cosmic rays in the energy region of 5 EeV with the AMIGA upgrade of the Pierre Auger Observatory*, PhD thesis, Karlsruher Institut für Technologie (KIT) (2019) doi:10.5445/IR/1000100543.

[41] A. M. Botti, *The AMIGA underground muon detector of the Pierre Auger Observatory – performance and event reconstruction*, PoS (ICRC2019) **358** (2019) 0411 doi:10.22323/1.358.0411.

[42] V. V. Kizakke Covilakam, A. D. Supanitsky, D. Ravignani, *Reconstruction of air shower muon lateral distribution functions using integrator and binary modes of underground muon detectors*, Eur. Phys. J. C **83** (2023) 1157 doi:10.1140/epjc/s10052-023-12344-7.

[43] M. Figueira, *An improved reconstruction method for the AMIGA detectors*, PoS (ICRC2017) **301** (2017) 0396 doi:10.22323/1.301.0396.

[44] Pierre Auger Collaboration, *Design, upgrade and characterization of the silicon photomultiplier front-end for the AMIGA detector at the Pierre Auger Observatory*, JINST **16** (2021) P01026 doi:10.1088/1748-0221/16/01/P01026.

[45] A. D. Supanitsky *et al.*, *Underground muon counters as a tool for composition analyses*, Astropart. Phys. **29** (2008) 461–470 doi:10.1016/j.astropartphys.2008.05.003.

[46] D. Ravignani, A. D. Supanitsky, *A new method for reconstructing the muon lateral distribution with an array of segmented counters*, Astropart. Phys. **65** (2015) 1–10 doi:10.1016/j.astropartphys.2014.11.007.

[47] D. Ravignani, A. D. Supanitsky, D. Melo, *Reconstruction of air shower muon densities using segmented counters with time resolution*, Astropart. Phys. **82** (2016) 108–116 doi:10.1016/j.astropartphys.2016.06.001.

[48] A. D. Supanitsky, *Estimation of the number of muons with muon counters*, Astropart. Phys. **127** (2021) 102535 doi:10.1016/j.astropartphys.2020.102535.

[49] M. Romeo, V. Da Costa, F. Bardou, *Broad distribution effects in sums of lognormal random variables*, Eur. Phys. J. B **32** (2003) 513–525 doi:10.1140/epjb/e2003-00131-6.

[50] K. Kamata and J. Nishimura, *The Lateral and the Angular Structure Functions of Electron Showers*, Prog. Theor. Phys. Suppl. **6** (1958) 93-155 doi:10.1143/PTPS.6.93.

[51] K. Greisen, *Cosmic Ray Showers*, Annual Review Nuclear and Particle Science. **10** (1960) 63-108 10.1146/annurev.ns.10.120160.000431.

[52] D. Heck *et al.*, *CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers*, FZKA Report **6019** (1998) url.

[53] A. Ferrari *et al.*, *FLUKA: a multi-particle transport code*, CERN Yellow Report **2005** (2005) doi:10.2172/877507.

[54] M. Bleicher *et al.*, *Relativistic Hadron-Hadron Collisions in the Ultra-Relativistic Quantum Molecular Dynamics Model*, J. Phys. G **25** (1999) 1859–1896 doi:10.1088/0954-3899/25/9/308.

[55] S. Ostapchenko, *QGSJET-II: physics, recent improvements, and results for air showers*, EPJ Web Conf. **52** (2013) 02001 doi:10.1051/epjconf/20125202001.

[56] F. Riehn *et al.*, *The hadronic interaction model Sibyll 2.3c and Feynman scaling*, PoS (ICRC2017) **301** (2017) arXiv:1709.07227.

[57] T. Pierog, R. Engel, D. Heck, S. Ostapchenko and K. Werner, *Latest Results from the Air Shower Simulation Programs CORSIKA and CONEX*, Proceedings of the 30th International Cosmic Ray Conference, ICRC 2007 **4** (2007) arXiv:0802.1262.

[58] Pierre Auger Collaboration, *The energy spectrum of cosmic rays beyond the turn-down around* 10^{17} *eV as measured with the surface detector of the Pierre Auger Observatory*, Eur. Phys. J. C **81** (2021) 966 10.1140/epjc/s10052-021-09700-w.

[59] J. Bellido, *Depth of maximum of air-shower profiles at the Pierre Auger Observatory: Measurements above* $10^{17.2}$ *eV and Composition Implications*, PoS (ICRC2017) **2017** (506) doi:10.22323/1.301.0506.

[60] Justin M. Albury, José A. Bellido, Bruce R. Dawson, *Exploring the energy threshold for full trigger efficiency of the Surface Detector with Hybrid events*, Auger internal note GAP-2018-038 (2018).

[61] J. de Jesús, *Study of the mass composition of cosmic rays with the Underground Muon Detector of AMIGA*, Karlsruher Institut für Technologie (KIT) doi:10.5445/IR/1000182143.

[62] F. Gesualdi, *The muon content of atmospheric air showers and the mass composition of cosmic rays*, Karlsruher Institut für Technologie (KIT) doi:10.5445/IR/1000155822.

[63] E. Rodriguez, N. González, and B. Wundheiler, *Digging up pinocchios with the Muon Lateral Distribution Function*, Auger internal note GAP-2026-010 (2026).

[64] B. Dawson, *The Energy Scale of the Pierre Auger Observatory*, PoS 36th International Cosmic Ray Conference **2019** (231) doi:10.22323/1.358.0231.

[65] G. Ros *et al.*, *A new composition-sensitive parameter for Ultra-High Energy Cosmic Rays*, Astropart. Phys. **35** (2011) 140–151 doi:10.1016/j.astropartphys.2011.06.011.

[66] N. González *et al.*, *A muon-based observable for a photon search at 30–300 PeV*, Astropart. Phys. **114** (2019) 48–59 doi:10.1016/j.astropartphys.2019.06.005.

[67] Pierre Auger Collaboration, *Measurement of the Depth of Maximum of Air-Shower Profiles with energies between* $10^{18.5}$ *and* 10^{20} *eV using the Surface Detector of the Pierre Auger Observatory and Deep Learning*, Phys. Rev. D **111** (2025) 022003 10.1103/PhysRevD.111.022003.

[68] Pierre Auger Collaboration, *Inferences on mass composition and tests of hadronic interactions from 0.3 to 100 EeV*, Phys. Rev. D **96** (2017) 122003 doi:10.1103/PhysRevD.96.122003.

[69] P. Sánchez Lucas, *The* $\langle\Delta\rangle$ *Method: An estimator for the mass composition of ultra-high-energy cosmic rays*, Universidad de Granada PhD Thesis (2016) 10481/44643.

[70] Pierre Auger Collaboration, *Direct measurement of the muonic content of extensive air showers between* 2×10^{17} *and* 2×10^{18} *ev at the Pierre Auger Observatory*, The European Physical Journal C **80** (2020) 1–19 doi:10.1140/epjc/s10052-020-8055-y.

[71] Pierre Auger Collaboration, *Evidence for a mixed mass composition at the 'ankle' in the cosmic-ray spectrum*, Physics Letters B7622016288-295 10.1016/j.physletb.2016.09.039.

[72] T. Gnambs, *A Brief Note on the Standard Error of the Pearson Correlation*, Collabra: Psychology (2023) doi:10.1525/collabra.87615.

[73] P. Sprent, *Applied Nonparametric Statistical Methods*, Springer Dordrecht (2011) 978-94-009-1223-6.

[74] M. G. Kendall, *A New Measure of Rank Correlation*, Biometrika **30** (1938) 81–93 doi:10.1093/biomet/30.1-2.81.

[75] M. G. Kendall, *The Treatment of Ties in Ranking Problems*, Biometrika **33** (1945) 239–251 doi:10.2307/2332303.

[76] R. A. Gideon, R. A. Hollister, *A Rank Correlation Coefficient Resistant to Outliers*, J. Am. Stat. Assoc. **82** (1987) 656–666 doi:10.1080/01621459.1987.10478480.

[77] Pierre Auger Collaboration, *Depth of maximum of air-shower profiles at the Pierre Auger Observatory. I. Measurements at energies above* $10^{17.8}$ *eV*, Physical Review D, **90** (2014) doi:10.1103/PhysRevD.90.122005 .

[78] Pierre Auger Collaboration, *Depth of maximum of air-shower profiles at the Pierre Auger Observatory. II. Composition implications*, Phys. Rev. D **90** (2014) 122006 doi:10.1103/PhysRevD.90.122006.

[79] M. Stadelmaier, *Reconstructing Air-Shower Observables using a Universality-Based Model*, PoS UHECR2024 (2025), 117 doi:10.22323/1.484.0117.

[80] A. Yushkov for the Pierre Auger Collaboration, *Mass composition of cosmic rays above* $10^{17.2}$ *eV from hybrid data*, PoS (ICRC2019) (2019) 482.

[81] F. James, M. Roos, *Minuit - a system for function minimization and analysis of the parameter errors and correlations*, Comput. Phys. Commun. **10** (1975) 343–367 doi:10.1016/0010-4655(75)90039-9.

[82] M. Antcheva *et al.*, *ROOT - A C++ framework for petabyte data storage, statistical analysis and visualization*, Comput. Phys. Commun. **180** (2009) 2499–2512 doi:10.1016/j.cpc.2009.08.005.

[83] J. Albrecht *et al.*, *The Muon Puzzle in cosmic-ray induced air showers and its connection to the Large Hadron Collider*, Astrophys. Space Sci. **367** (2022) 27 doi:10.1007/s10509-022-04054-5.

[84] P. Picozza *et al.*, *PAMELA – A payload for antimatter matter exploration and light-nuclei astrophysics*, Astropart. Phys. **27** (2007) 296–315 doi:10.1016/j.astropartphys.2006.12.002.

[85] J. Chang *et al.*, *The DArk Matter Particle Explorer mission*, Astropart. Phys. **95** (2017) 6–24 doi:10.1016/j.astropartphys.2017.08.005.

[86] S. Ting, *The Alpha Magnetic Spectrometer on the International Space Station*, Nucl. Phys. B Proc. Suppl. **243–244** (2013) 12–24 doi:10.1016/j.nuclphysbps.2013.09.001.

[87] Pierre Auger Collaboration, *The Pierre Auger Observatory and its Upgrade*, Science Reviews - From the End of the World **1** (2020) doi:10.52712/sciencereviews.v1i4.31.

[88] V. F. Hess, *Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten*, Physikalische Zeitschrift **13** (1912) 1084–1091.

[89] V. F. Hess (translated and commented), *On the Observations of the Penetrating Radiation during Seven Balloon Flights*, arXiv:1808.02927.

[90] P. Carlson, *Discovery of Cosmic Rays*, AIP Conf. Proc. **1516** (2013) 9–15doi:10.1063/1.4791733.

[91] W. Heitler, *The Quantum Theory of Radiation*, Clarendon Press, Oxford (1954).

[92] J. Matthews, *A Heitler model of extensive air showers*, Astroparticle Physics **22** (2005) 387–397doi:10.1016/j.astropartphys.2004.09.003.

[93] Aartsen, M.G. and Ackermann, M. and Adams, J. and Aguilar Sánchez et al., *Astrophysical neutrinos and cosmic rays observed by IceCube*, Advances in Space Research **62** (2017) doi:10.1016/j.asr.2017.05.030.

[94] K.-H. Kampert, M. Unger, *Measurements of the cosmic ray composition with air shower experiments*, Astropart. Phys. **35** (2012) 660–678 doi:10.1016/j.astropartphys.2012.02.004.

[95] J. Knapp, D. Heck, S. J. Sciutto, M. T. Dova, M. Rameika, *Extensive air shower simulations at the highest energies*, Astropart. Phys. **19** (2003) 77–99 doi:10.1016/S0927-6505(02)00214-6.

[96] S. J. Sciutto, *AIRES: A system for air shower simulations*, arXiv:astro-ph/9911331.

[97] D. Soldin, *Update on the Combined Analysis of Muon Measurements from Nine Air Shower Experiment*, PoS ICRC2021 (2021) 349 10.22323/1.395.0349.

[98] I. Allekotte *et al.* (Pierre Auger Collaboration), *The Surface Detector System of the Pierre Auger Observatory*, Nucl. Instrum. Meth. A **586** (2008) 409–420 doi:10.1016/j.nima.2007.12.010.

[99] J. Abraham *et al.* (Pierre Auger Collaboration), *Trigger and Aperture of the Surface Detector Array of the Pierre Auger Observatory*, Nucl. Instrum. Meth. A **613** (2010) 29–39 doi:10.1016/j.nima.2009.11.018.

[100] A. Aab *et al.* (Pierre Auger Collaboration), *Reconstruction of events recorded with the surface detector of the Pierre Auger Observatory*, JINST **15** (2020) P10021 doi:10.1088/1748-0221/15/10/P10021.

[101] S. Argirò *et al.*, *The Offline Software Framework of the Pierre Auger Observatory*, Nucl. Instrum. Meth. A **580** (2007) 1485–1496 doi:10.1016/j.nima.2007.07.010.

[102] J. Abraham *et al.* (Pierre Auger Collaboration), *Upper limit on the cosmic-ray photon flux above 10^{19} eV using the surface detector of the Pierre Auger Observatory*, Astropart. Phys. **29** (2008) 243–256 doi:10.1016/j.astropartphys.2008.01.003.

[103] Pierre Auger Collaboration, *The scintillator surface detector of the Pierre Auger Observatory: performance and first results*, JINST **20** (2025) P08002 doi:10.1088/1748-0221/20/08/P08002.

[104] F. G. Schröder, *Radio detection of high-energy cosmic rays with the Auger Engineering Radio Array*, Nucl. Instrum. Meth. A **824** (2016) 648–653 doi:10.1016/j.nima.2015.11.050.

[105] Pierre Auger Collaboration, *The Auger Engineering Radio Array: detector performance and air-shower reconstruction*, JINST **20** (2025) P12017 doi:10.1088/1748-0221/20/12/P12017.

[106] Pierre Auger Collaboration, *Extraction of the muon signals recorded with the surface detector of the Pierre Auger Observatory using recurrent neural networks*, JINST **16** (2021) 07016 10.1088/1748-0221/16/07/P07016.

[107] Pierre Auger Collaboration, *The Pierre Auger Observatory Upgrade - Preliminary Design Report*, 2016 eprint 1604.03637.

[108] M. Platino, M. Hampel et al., *AMIGA at the Auger Observatory: The scintillator module testing system*, Journal of Instrumentation **6** (2011) P06006 10.1088/1748-0221/6/06/P06006.

[109] Pierre Auger Collaboration, *Reconstruction of events recorded with the surface detector of the Pierre Auger Observatory*, Journal of Instrumentation **15** (2020) P10021 doi:10.1088/1748-0221/15/10/P10021 .

[110] D. Newton, J. Knapp, A. Watson, *The optimum distance at which to determine the size of a giant air shower*, Astroparticle Physics **26** (2007) 414–419 doi:10.1016/j.astropartphys.2006.08.003.

[111] H. P. Dembinski, et al., *A likelihood method to cross-calibrate air-shower detectors*, Astroparticle Physics **73** (2016) 44–5110.1016/j.astropartphys.2015.08.001.

[112] A. Etchegoyen and Pierre Auger Collaboration, *AMIGA, Auger Muons and Infill for the Ground Array*, arXiv preprint arXiv:0710.1646 (2007).

[113] A. Pla-Dalmau, A. Bross and V. V. Rykalin, *Extruding plastic scintillator at Fermilab*, IEEE Nuclear Science Symposium Conference Record **1** (2003) 102 - 104 Vol.1 10.1109/NSSMIC.2003.1352007.

[114] Pierre Auger Collaboration, *Muon counting using silicon photomultipliers in the AMIGA detector of the Pierre Auger Observatory*, Journal of Instrumentation **12** (2017) P03002 10.1088/1748-0221/12/03/P03002.

[115] A.D.Supanitsky et al., *Effect of multiple reusing of simulated air showers in detector simulations*, Astroparticle Physics **30** (2008) 264-269 10.1016/j.astropartphys.2008.10.001.

[116] H. P. Dembinski, *Computing mean logarithmic mass from muon counts in air shower experiments*, Astroparticle Physics **102** (2018) 89–94 10.1016/j.astropartphys.2018.05.008.

[117] H. Dembinski, T. Pierog, *Future Proton-Oxygen Beam Collisions at the LHC for Air Shower Physics*, PoS(ICRC2019)235 **2019** (235) doi:10.22323/1.358.0235.

[118] Pierre Auger Collaboration, *Prototype muon detectors for the AMIGA component of the Pierre Auger Observatory*, Journal of Instrumentation **11** (2016) P02012 10.1088/1748-0221/11/02/P02012.

Acknowledgments

The process of writing this thesis has been quite arduous: it involved a pandemic, relocations, an unreasonable dependence on caffeine, and what can only be described as a deeply questionable sleep schedule. Nevertheless, I got through it—somehow. If this thesis exists in its current form, it is largely because of the people who supported me along the way.

I would like to thank my supervisor, Daniel Supanitsky. This thesis would not have been possible without your encouragement, trust, and patience. Thank you for guiding me toward better opportunities and for knowing when to push and when to let things evolve. I also thank Ralph Engel for his guidance and valuable input on this work. A special thanks to Diego Ravignani, David Schmidt, and Steffen Hahn for their support and discussions.

I am grateful to the funding agencies—DAAD, and HIRSAP—for recognising the value of this programme and supporting it. I also thank CNEA and CONICET for granting me this PhD position and providing a place where I could learn, work and occasionally panic productively. Many thanks to Adrián Sedoski, Sabine Bucher, Anna Kendrick, Marie-Christine Mundt, Consuelo Vitón, Sabrina Coller, and Tadeo Lopez for always being there to help—whether the issue was administrative or technical.

I am also grateful to Alberto Etchegoyen and Aníbal Gattone for giving me this opportunity. Thanks to Flavia Gesualdi for sharing her knowledge (and patience). Thanks to Carmina, Marina, Ezequiel, Gaby, Joaquín, and Fer—working alongside you made everyday institute life not only more pleasant, but genuinely enjoyable.

Although my time in Germany was limited, it left a lasting impression and broadened my perspective in ways I did not expect. I would like to thank Emily, Tobias, Paul, and Fiona—thank you for your warmth, kindness, and for making me feel at home. Special thanks to Steffan and Karla for the home-cooked meals. Thank you to Maike Etzold for making sure I always had a home to return to in Karlsruhe—it meant more than I can properly express.

I am grateful to everyone with whom I shared this journey, including the students, researchers, and staff of both ITeDA and KIT, as well as the members of the Pierre Auger Collaboration. Being part of this community has been a privilege.

I would like to thank my parents, Girisa and Nandini, for their patience and support. I thank my brother Vishnu for always having my back. I would also like to thank my grandmother, Susheela, for her constant love and support. I would like to thank my other grandmother, Leelambika—although she is no longer here, her love and support have remained with me throughout this journey. To all my cousins, uncles, and aunts—thank you for being there in all the ways that matter.

I would also like to thank my friends outside the academy who supported me throughout this period: Edwin, Anirudh, Ramya, Brishel, Francisco and my Junín family—thank you for keeping me grounded and reminding me that there is, in fact, a world outside of research.

Finally, I want to express my deepest gratitude to Marcos for his patience, understanding, and unwavering support. You have been my constant through all the uncertainty, and I could not have done this without you. I also thank his parents, grandparents, and sister for always making me feel welcomed and at home.