

**Measurement of
the Composition of Cosmic Rays and
the Proton-Proton Inelastic Cross Section
at Ultrahigh Energies
with the Pierre Auger Observatory**

Zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften

von der KIT-Fakultät für Physik des
Karlsruher Instituts für Technologie (KIT)

genehmigte

Dissertation

von

M.Sc. Olena Tkachenko

aus Charkiw, Ukraine

Tag der mündlichen Prüfung: 26/01/2024

Referent: Prof. Dr. Ralph Engel

Korreferent: Prof. Dr. Ulrich Husemann

Betreuer: Dr. Michael Unger

*To the memory of those
who made the ultimate sacrifice
for the survival of Ukraine
and its people.*

Abstract

The existence of cosmic rays with energies beyond 10^{20} eV has been known for over half a century, but understanding their astrophysical origin remains one of the major unsolved problems in modern astrophysics. Studies of the mass composition and hadronic interaction properties of cosmic rays can improve our knowledge of the nature of ultrahigh-energy cosmic rays. The mass composition is one of the key observables to constrain possible astrophysical sources of cosmic rays, and the study of hadronic interactions at energies far beyond man-made accelerators provides fundamental insights into the physics of soft QCD interactions. However, neither the mass composition nor the interaction cross sections are well measured at ultrahigh energies, and their estimation typically requires assumptions on one of these quantities to determine the other.

In this work, the cross section and mass composition analyses are updated with the latest data of the Pierre Auger Observatory. The mass composition estimation employs a new approach, utilizing the Bayesian Markov Chain Monte Carlo sampling algorithm.

Avoiding the shortcomings of separate analyses, this work presents a combined approach for the independent and simultaneous self-consistent estimate of the cosmic-ray composition and hadronic interactions for the first time. A standard mass composition measurement from the depth of the atmospheric shower maximum, X_{\max} , as observed by the Fluorescence Detector of the Pierre Auger Observatory, is performed under the assumption of the modified proton-proton interaction cross section. Our results suggest that the inferred proton-proton inelastic cross section at ultrahigh energies agrees well with the standard extrapolations of accelerator data. The cosmic-ray composition determined in this work confirms earlier estimates. Due to the high statistics of the full Phase I data set of the Pierre Auger Observatory and the marginalization over the cross section and X_{\max} scale, inherent in our novel analysis method, the complex evolution of mass groups with energy is now known with higher confidence.

Zusammenfassung

Seit mehr als einem halben Jahrhundert ist die Existenz kosmischer Strahlung mit Energien jenseits von 10^{20} eV bekannt. Jedoch bleibt das Verständnis ihres astrophysikalischen Ursprungs eines der größten ungelösten Probleme der modernen Astrophysik. Untersuchungen zur Massenzusammensetzung und hadronischen Wechselwirkungen der kosmischen Strahlung können unser Wissen über die Natur der ultrahochenergetischen kosmischen Strahlung verbessern. Die Massenzusammensetzung ist von großer Bedeutung, um mögliche astrophysikalische Quellen von kosmischer Strahlung einzugrenzen. Die Untersuchung hadronischer Wechselwirkungen bei Energien weit über denen, die von Menschen gebauten Beschleunigern erreicht werden können, liefert grundlegende Erkenntnisse über die Physik der Wechselwirkungen bei der weichen Quantenchromodynamik.

In dieser Arbeit werden die Analysen des Wirkungsquerschnitts und der Massenzusammensetzung mit den neuesten Daten des Pierre-Auger-Observatoriums aktualisiert. Zur Schätzung der Massenzusammensetzung wird ein neuer Ansatz verwendet, der den Bayes'schen Markov-Chain-Monte-Carlo-Sampling-Algorithmus nutzt.

In dieser Arbeit wird zum ersten Mal ein kombinierter Ansatz für die unabhängige und gleichzeitige selbstkonsistente Abschätzung der Zusammensetzung der kosmischen Strahlung und der hadronischen Wechselwirkungen vorgestellt, der die Unzulänglichkeiten getrennter Analysen vermeidet. Eine Standardmessung der Massenzusammensetzung aus der Tiefe des atmosphärischen Schauersmaximums, X_{\max} , wie es vom Fluoreszenzdetektor des Pierre-Auger-Observatoriums beobachtet wird, wird unter der Annahme eines modifizierten Proton-Proton-Wechselwirkungsquerschnitts durchgeführt. Unsere Ergebnisse deuten darauf hin, dass der abgeleitete inelastische Proton-Proton-Wechselwirkungsquerschnitt bei ultrahohen Energien gut mit Standard-Extrapolationen von Beschleunigerdaten übereinstimmt. Die in dieser Arbeit ermittelte Zusammensetzung der kosmischen Strahlung bestätigt frühere Abschätzungen. Aufgrund der hohen statistischen Aussagekraft des vollständigen Phase-I-Datensatzes des Pierre-Auger-Observatoriums und der Marginalisierung über den Wirkungsquerschnitt und die X_{\max} -Skala, die unserer neuen Analyseverfahren inhärent ist, ist die komplexe Entwicklung der Massengruppen mit der Energie nun mit größerer Sicherheit bekannt.

Contents

Introduction	9
1 Cosmic rays and extensive air showers	11
1.1 Galactic and Extragalactic cosmic rays	11
1.2 Development of extensive air showers in atmosphere	15
1.2.1 Heitler model for electromagnetic cascades	16
1.2.2 Heitler-Matthews model for hadronic cascades	18
1.3 Hadronic interaction models	22
1.3.1 High-energy models	23
1.3.2 Modified hadronic interactions	29
1.3.3 Hadronic cross sections and Glauber theory	32
2 The Pierre Auger Observatory	35
2.1 The surface detector	36
2.2 The fluorescence detector	39
2.2.1 Event reconstruction	44
2.2.2 Resolution and Systematics	48
2.3 Atmospheric monitoring	48
2.4 The Pierre Auger Observatory Upgrade: AugerPrime	50
2.5 Data sets and data selection	51
3 Measurement of Mass Composition	55
3.1 MCMC Sampling of the X_{\max} distribution	56
3.1.1 Fitting algorithm	57
3.1.2 Test of fraction fitting with MCMC algorithm	64
3.1.3 Confidence intervals and coverage probability	65
3.2 Application to data	71
3.2.1 Comparison to previous results	75

CONTENTS

3.2.2	Composition fractions for different data sets	77
3.2.3	Effect of including additional particle groups	78
3.2.4	Fits with a free X_{\max} scale	80
3.3	Results	84
3.3.1	Composition fractions	84
3.3.2	Further applications of the MCMC mass composition fit	87
3.4	X_{\max} -dependent X_{\max} bias	91
4	Measurement of the proton-proton cross section from the tail of the X_{\max} distribution	99
4.1	Method	99
4.2	Bias and uncertainties	102
4.3	Results	103
5	Combined estimation of the cosmic-ray mass composition and proton-proton inelastic cross section	109
5.1	Rescaling of the proton-proton interaction cross section	110
5.2	Fitting algorithm	113
5.3	Effect of the cross section modification on the mass composition fit	118
5.4	Performance of the method	120
5.5	Results	136
	Summary and Outlook	147
	Appendix	169
	Acknowledgment	191

Introduction

One of the main open questions of astrophysics and astroparticle physics is the origin and nature of ultrahigh-energy cosmic rays. These particles are the most energetic particles known in the universe, and neither the mechanism nor the site of their acceleration is known. Low-energy cosmic rays are detectable using various techniques with balloon- or space-borne detectors. However, observing cosmic-ray particles at the highest energies is only feasible with ground-based detector arrays spanning large areas due to the low flux of ultrahigh-energy cosmic rays reaching the ground. As primary cosmic-ray nuclei traverse the atmosphere, they initiate a cascade of interactions, forming extensive air showers composed of secondary particles. These showers are then detected by arrays of particle detectors, such as the Pierre Auger Observatory and the Telescope Array, which measure the lateral particle densities at the ground. In addition to particle detectors, these arrays are equipped with fluorescence telescopes to detect the light emission resulting from the collisions of cosmic ray particles with atmospheric nitrogen molecules and thus to directly observe the longitudinal development of the air shower.

The atmospheric depth at which the number of particles in a shower reaches its maximum, X_{\max} , is an observable which is very sensitive to the cosmic-ray mass composition and the characteristics of hadronic interactions during the development of the extensive air shower in the atmosphere. The shower maximum carries information about the type of nuclei that initiated the shower, as it is proportional to the logarithm of the mass of the primary particle. At the same time, X_{\max} is also determined by the depth of the first interaction in the atmosphere, which is directly related to the interaction cross section of the particle initiating the shower with the air. However, neither the mass composition nor the interaction cross sections are well measured at ultrahigh energies, and their estimation typically requires some assumptions on one of these quantities to determine the other. Currently, the standard analysis, estimating the mass composition of cosmic rays

at ultrahigh energies, relies on the accuracy of the existing hadronic interaction models. On the other hand, to estimate the particle interaction cross sections, the current analyses assume a specific composition to interpret the data.

The main focus of this thesis is to measure the cosmic-ray mass composition and the proton-proton interaction cross section. Rather than performing separate analyses that depend on each other's results, a new approach for estimating both the cosmic ray mass composition and proton-proton interaction cross section simultaneously and independently is presented. The measurement of the mass composition of cosmic rays is of great importance for understanding their astrophysical nature and the mechanisms responsible for the acceleration to ultrahigh energies, and the estimation of the interaction cross sections provides us with a window to probe the interaction properties at energies beyond the reach of existing accelerator experiments.

After a brief introduction to cosmic rays, hadronic interactions, and the Pierre Auger Observatory, a new approach to fitting the mass composition using Bayesian inference with the Markov Chain Monte Carlo algorithm is discussed. Extensive tests were conducted to assess the performance of the method across various composition scenarios and to investigate its sensitivity to the different fit ingredients. The fractions obtained from the fit to the full Phase I data of the Pierre Auger Observatory are compared under the assumption of the different hadronic interaction models and for varying numbers of particle species in the mixture. Additionally, the limits on the fraction of protons and photons, and the cosmic-ray rigidity are calculated.

Next, a method for rescaling the cross section of proton-proton interactions is discussed, and the effect of the corresponding modifications on the X_{\max} distribution is analyzed. The latter is crucial for a mass composition fit under the assumption of modified interaction properties. In such a way, it is possible to perform a standard mass composition estimation for the range of modified cross sections and estimate the best-fit combination of the studied variables. Additionally, the X_{\max} scale is treated as another fit parameter to account for both systematic uncertainties in the data as well as the theoretical uncertainties associated with particle production in air showers. Afterwards, the performance of the combined approach is assessed, and the level of bias in the fit results is calculated.

Finally, the mass composition and cross section fit is applied to recent data of the Pierre Auger Observatory. The cosmic-ray mass composition, proton-proton interaction cross section, and a shift in the X_{\max} scale are estimated, and the physics implications of the obtained values are discussed.

Chapter 1

Cosmic rays and extensive air showers

The observed characteristics of cosmic rays play a crucial role in describing the cosmic environment. Studies of the properties of cosmic rays are essential for describing the interstellar medium as cosmic rays affect the abundances of the isotopes of the light elements and ionize the interstellar atoms and molecules through their numerous interactions with interstellar matter [1]. Furthermore, cosmic rays may also play an important role in galaxy formation and evolution, for example, by affecting star formation, through the feedback mechanism [2]. Data from cosmic ray experiments provide insights into the origin and composition of particles at the highest observed energies, as well as fundamental information about the physics behind their acceleration and propagation through astrophysical plasmas.

1.1 Galactic and Extragalactic cosmic rays

Cosmic rays cover a wide range of energies [3], from 10^7 eV over 10^{20} eV. Since the discovery of cosmic rays, a number of different techniques have been used to measure the flux of particles at each energy, which, displayed as a function of energy, gives a cosmic ray energy spectrum.

The energy spectrum of cosmic rays is well described as a steeply declining power law over more than ten decades in energy:

$$\frac{dN}{dE} \sim E^\gamma, \quad (1.1)$$

where γ denotes the spectral index. The spectrum is shown in Fig. 1.1, where the top points are all-particle cosmic-ray spectrum. The fluxes of neutrinos, electrons, positrons, and γ -emission are also displayed.

The spectrum is relatively featureless, with several breaks resulting in changes in the power-law slope. Below ~ 20 GeV, the flux of Galactic cosmic rays is modulated by the magnetic field in the solar-wind plasma [4]. The three transition points in the spectrum define three energy intervals. The first hardening of the spectrum, resulting in the change of the spectral index from $\gamma \approx -2.7$ to $\gamma \approx -3.0$ is observed around 3 PeV. It was discovered in 1958 by Kulikov and Khristiansen [5] and is usually referred to in the literature as the knee. As shown by the KASCADE experiment, the measured steepening of the total particle spectrum may reflect the decrease in the flux of the light elements[6]. The origin of this hardening remains an open question, and there is a wide range of possible theories explaining the observed feature (see [7] for an overview). One of the first explanations for the PeV hardening was the cut-off in the rigidity¹ of the sources, producing the bulk of cosmic rays below 10^{15} eV [8]. In general, theories explaining the origin of the knee in cosmic-ray spectra can be divided into several groups [9, 10]. The so-called astrophysical models relate the break in the spectrum to changes in the source acceleration and/or propagation in the interstellar medium. Another class of models attempts to explain the steepening of the spectrum due to changes in particle interactions with increasing energy [11] but is disfavoured since \sqrt{s} is less than LHC center-of-mass-energy in Galactic cosmic rays. The origin of the knee could also be attributed to the assumption of the spectrum and composition of primary cosmic rays. One of the most promising and well-studied scenarios for the origin of the knee is based on the maximum energy achievable by diffusive shock acceleration in supernova remnants, which is believed to be the main acceleration process for Galactic cosmic rays [12]. Other common explanations are, for example, an increase in the escape of cosmic rays from the galaxy due to changes in the diffusion properties of charged particles [13] or the collective reacceleration of cosmic-ray particles originating from the Galactic disk by the Galactic wind [14]. In addition, some studies suggest that interactions with the interstellar medium may be responsible for the spectral break.

Around 20 PeV, some hardening in the spectrum has recently been observed, with a change in the spectral slope to $\gamma \approx -2.9$ [16, 17], followed by a softening in a spectrum at $\approx 10^{17}$ eV [18, 16, 19, 20, 21, 22], with a spectral slope above

¹The rigidity of a particle with charge Ze and momentum p (energy E) is $R = pc/(Ze) \simeq E/(Ze)$

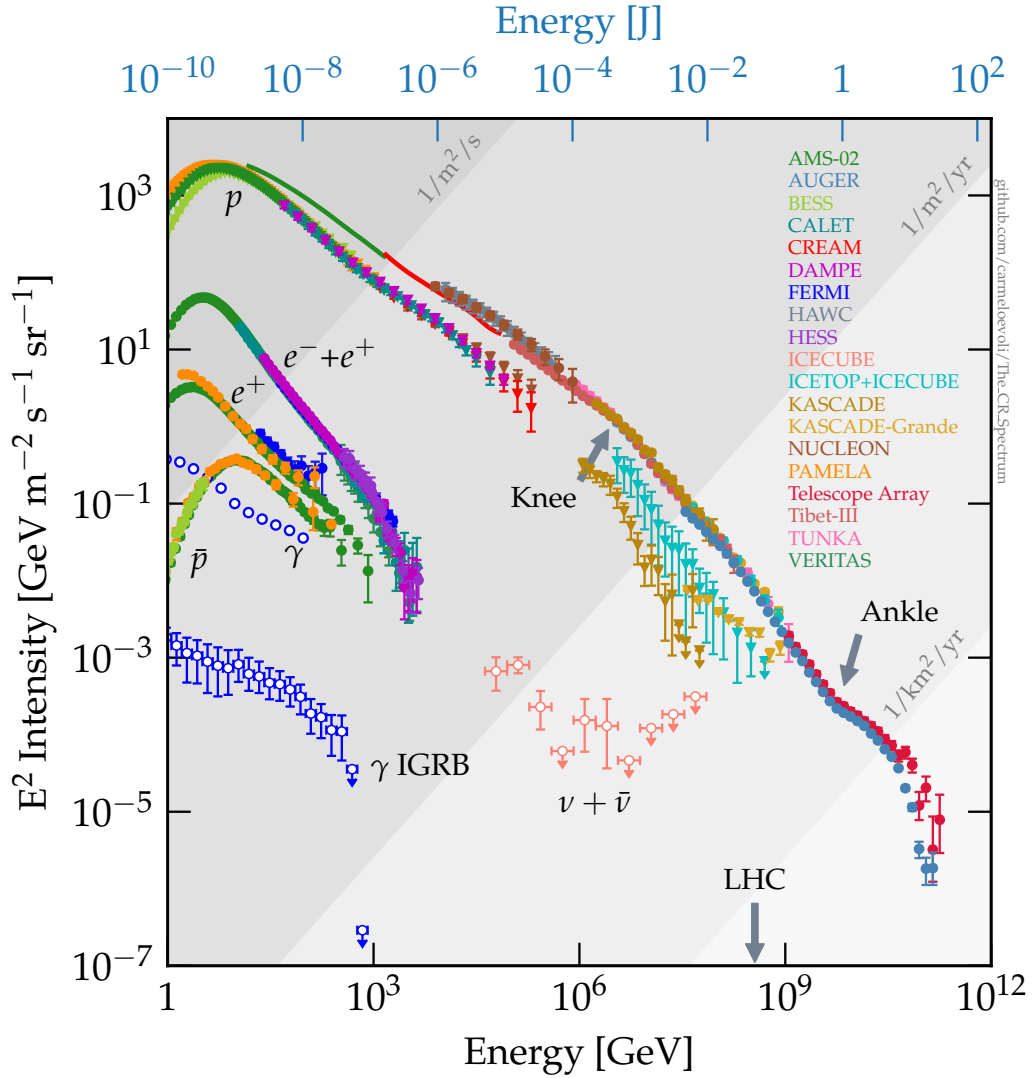



Figure 1.1: The cosmic-ray energy spectrum. The results of the several experiments operating at different energies are shown. For low energies the spectra of different primary nuclei and for high energy cosmic rays the all-particle energy spectra are plotted. The spectral features seen on the plot are discussed in the text. Figure from [15].

the break of $\gamma \approx -3.3$. It is referred to in the literature as the second knee, or the iron knee, marking the rigidity cut-off for the Galactic iron [23, 24, 25].

At around $10^{18.5}$ eV, there is a hardening in the cosmic ray spectrum known as the ankle [26, 20, 27] with a change in the spectral index to ≈ -2.6 [28, 29]. One of the standard interpretations for the break at EeV energies is known as the *ankle model* for the transition from the soft Galactic to the harder extragalactic component of the cosmic ray flux (see, e.g., [30, 31]). Within this model, the fluxes of both Galactic and extragalactic cosmic rays become equal at the onset of the cosmic ray ankle, suggesting the need for the existence of an additional Galactic component [32] that is accelerated to energies beyond the maximum energy predicted by the Standard Model. The *ankle model* assumes that extragalactic cosmic rays dominate the flux above the ankle. At the same time, the Galactic cosmic rays below the break are dominated by the heavier nuclei, i.e. iron. The latter one, however, contradicts the measurements of the average depth of the shower maximum observed by HiRes [33] and the Pierre Auger Observatory [34], which suggest that cosmic rays are primarily composed of lighter nuclei not only above but also below the ankle within the 1-5 EeV range.

Alternatively, the transition between the Galactic and extragalactic cosmic-ray components could be characterized by the *dip model* [35, 36]. In the *dip model*, the transition starts at the second knee and lasts until the e^+e^- pair production dip above 1 EeV, caused by the energy losses due to the interactions of the extragalactic protons with cosmic microwave background photons. In the context of the *dip model*, there is a transition from heavy to light between 0.1 and 1 EeV, whereas above the dip energy, an extragalactic component dominated by lighter nuclei is anticipated. This agrees with composition measurements in the dip region but disagrees with the Pierre Auger Observatory results above 4 EeV [36]. The transition from Galactic to extragalactic cosmic rays can also be explained within the *mixed composition model* [37, 38]. In the *mixed model*, the composition above the ankle is assumed to be a mixture of different nuclei. At the same time, the Galactic cosmic rays at the end of their spectrum are dominated by iron. Depending on the chosen model parameters, the transition may occur either at the ankle or at the energy below it. In the latter case, there might not be a need for an additional Galactic component beyond what is predicted by the Standard Model of Galactic cosmic rays. Nonetheless, in the majority of *mixed composition models*, the composition remains predominantly proton-dominated up to ≈ 50 EeV, thus failing to explain the mixed composition reported by the Pierre Auger Observatory at $E > 10$ EeV. A potential way around this could be an assumption of the enhanced abundances of heavier nuclei, exceeding the current expectation for the

source composition [39]. At the end of the all-particle spectrum, the suppression of the cosmic ray flux has been seen in the data of different experiments (see, e.g., [27, 40, 41, 42]). This steepening observed at $E \gtrsim 5 \cdot 10^{19}$ eV is consistent with the so-called GZK-cutoff predicted by Greisen, and Zatsepin and Kuzmin [43, 44]. They discovered that cosmic-ray particles lose energy when interacting with the cosmic microwave background (CMB). The interaction is happening through the Δ -resonance, resulting in the production of pions:



or



The pions and neutrons produced by the Δ resonance then decay into very high-energy photons and neutrinos. Because of this energy loss, which amounts on average to 20%, the distance a proton with an energy above 10^{20} eV can travel is limited by about 100 Mpc. Therefore, the detection of protons with energies above 10^{20} eV would imply either a propagation distance smaller than the GZK volume or the existence of particles with much higher energies [45]. If heavier nuclei dominate at ultrahigh energies, the interaction of cosmic rays with CMB photons would occur via the giant dipole resonance leading to an energy loss length that is similar to the one of protons for iron nuclei, and even smaller for intermediate-mass nuclei. However, it is not entirely clear whether the observed suppression is a definitive confirmation of the GZK cutoff or is due to the proximity to the maximum limit to which cosmic sources can accelerate particles [46, 47].

1.2 Development of extensive air showers in atmosphere

Due to the low flux of ultrahigh-energy cosmic rays, they cannot be observed directly in space, but are detected indirectly by studying the particle cascades they produce in the atmosphere. As cosmic rays pass through the atmosphere, they undergo a series of interactions with atmospheric molecules. These interactions lead to the formation of cascades of secondary particles that evolve both longitudinally and laterally, see Fig. 1.2. The majority of the secondary particles produced are either charged or neutral pions, the latter of which rapidly decay into photons. The photons can then convert into electron-positron pairs, and more photons can be produced by bremsstrahlung of high energy electrons and positrons. The charged

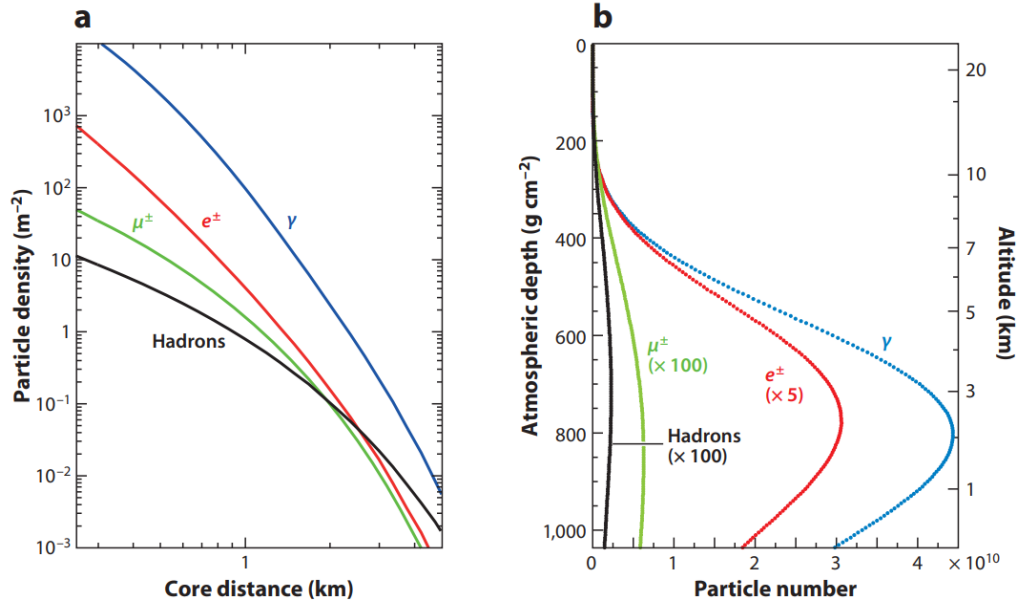


Figure 1.2: Example of the average lateral (left) and longitudinal (right) profiles of the vertical air showers simulated with CORSIKA for a proton primary at 10^{19} eV. Figure from [48].

pions have a longer lifetime, so they undergo numerous interactions before finally decaying into neutrinos and muons, the latter of which are detected at ground level.

1.2.1 Heitler model for electromagnetic cascades

The development of electromagnetic air shower cascades can be qualitatively well described by a simple splitting developed by Heitler [49] and Carlson and Oppenheimer [50].

The electromagnetic particles entering the atmosphere are subjected to pair-production and bremsstrahlung. After traversing the splitting length of $d = \lambda_r \ln 2$, where λ_r is the radiation length of about 37 g/cm^2 , the electron undergoes one-photon bremsstrahlung. Half of the electron's initial energy is then lost to the photon produced. The splitting of the photon after it has traveled the same mean free path d occurs via the pair production process, with the energy split equally between the outgoing electron and positron. Therefore, the number of secondary

particles doubles for each splitting length. Fig. 1.3 (left) shows a schematic representation of the evolution of the electromagnetic air shower for the first $n = 4$ splittings. The cascade process stops when the energy per particle reaches critical energy ξ_e^c , below which the ionization energy losses exceed the radiation losses, and no splitting is possible. In air, the critical energy is 85 MeV. Thus, the Heitler profile can be written as

$$N(X) = \begin{cases} 2^{X/d} & \text{if } X \leq n_e^c d; \\ 0 & \text{if } X > n_e^c d, \end{cases} \quad (1.4)$$

where n_e^c is the maximum number of iterations at which the critical energy is reached:

$$n_e^c \ln 2 = \ln(E_0/\xi_e^c). \quad (1.5)$$

The depth $X_{\max}^{(\text{EM})} = n_e^c d$ corresponds to the point at which the maximum number of secondary particles in the cascade, $N_{\max} = E_0/\xi_e^c$, is reached. Thus, the maximum number of particles in the shower is directly proportional to the initial energy, and the depth at which the shower reaches its maximum is logarithmically dependent on the initial energy of the particle initiating the shower:

$$\begin{aligned} N_{\max} &\sim E_0 \\ X_{\max} &\sim \ln E_0. \end{aligned} \quad (1.6)$$

Although the Heitler model provides a very simplified description of the evolution of electromagnetic cascades, it gives good qualitative agreement with the predictions from air shower simulations, and the estimated X_{\max} is not contradicted by the calculations from the cascade equations describing the propagation of the cosmic-ray showers in the atmosphere [52]. However, the number of electrons predicted by the Heitler model, $N_e = \frac{2}{3}N_{\max}$, exceeds the actual number [7], leading to an overestimation of the electron/photon ratio. This occurs because the model is unable to account for all the intricacies involved in the development of an air shower, such as the simultaneous emission of the multiple bremsstrahlung photons or the ranging out of electrons and positrons. A straightforward correction for this overestimation is a correction factor in the denominator, the specific value of which depends on the characteristics of the detector being considered.

1. Cosmic rays and extensive air showers




Figure 1.3: Schematic view of the development of the air shower cascade. Left: an example of a photon-induced electromagnetic shower. Right: an example of a proton-induced hadronic shower. In the electromagnetic shower, wavy lines represent photons. In the hadronic shower, dashed lines correspond to neutral particles, and solid lines to charged particles. Figure from [51].

1.2.2 Heitler-Matthews model for hadronic cascades

The development of hadron-induced cascades (see Fig. 1.3 (right)) is far more complicated than electromagnetic showers due to the numerous particle decays and the complexity of the hadronic multiparticle productions. Although it is possible to write a set of cascade equations for hadron-induced showers, it is not feasible to find their analytical solution. However, just like for the electromagnetic showers, the development of the hadron-initiated cascades can also be explained using a framework similar to the Heitler model. An approximation for the hadronic cascades was established by Matthews [51] and is referred to as the Heitler-Matthews model.

After the proton enters the atmosphere and travels the distance $\lambda_i \ln 2$, where λ_i is the interaction length of a particle, it interacts with an air nucleus. For pions with energies of 10 - 1000 GeV, the interaction length in air is $\lambda_i \approx 120 \text{ g/cm}^2$. After each interaction n_{tot} new particles are created, each with an energy E/n_{tot} , of which two-thirds are charged particles (charged pions, π^\pm) and one-third are neutral particles (neutral pions π^0) [48]. Due to the isospin symmetry, neutral pions immediately decay into two photons, triggering electromagnetic showers. The charged pions continue to interact with air nuclei, each time passing through another layer of the atmosphere with a thickness of $\lambda_i \ln 2$. The multiplications of particles stop when the energy of a pion falls below a critical energy ξ_c^π , which

happens when the interaction length becomes greater than the decay length of pions in air. At this point, the pion undergoes decay, producing a muon and a neutrino. After the n interactions the total number of charged pions will be $N_\pi = n_{\text{ch}}^n = (2/3n_{\text{tot}})^n$ with a total energy of

$$E_\pi^{\text{tot}} = \left(\frac{2}{3}\right)^n E_0. \quad (1.7)$$

The remaining energy would go into the electromagnetic component:

$$E_{\text{EM}}^{\text{tot}} = \left[1 - \left(\frac{2}{3}\right)^n\right]. \quad (1.8)$$

The energy per single pion is then

$$E_\pi = \frac{E_0}{\left(\frac{3}{2}n_{\text{ch}}\right)^n}. \quad (1.9)$$

Note that the pion multiplicity is assumed to be constant throughout the cascade evolution, whereas it has been shown that, similarly to the accelerator measurements for pp and $p\bar{p}$ collisions with multiplicity evolving as $\sim E^{0.2}$, the multiplicity in pion-air interactions also slowly increases with energy [53]. The leading hadron fragments are not taken into account in the calculations as well.

From Eq. 1.9, one can calculate the critical energy and the number of collisions needed to reach it (see [51]). Assuming that the shower was initiated by a particle of 10^{15} eV and that the multiplicity of charged particles in the pion-air interaction is $n_{\text{ch}} = 10$ (so there are five neutral pions and a total of 15 new particles created after each interaction), after four interactions the energy per pion would be $E_\pi = 20$ GeV. Since at this energy the decay length of the charged pion is around 1 km, and the distance between the beginning and end of the fourth atmospheric layer, obtained from the exponential atmospheric profile, is about 1.8 km, the probability of decay would exceed the probability of the next interaction. Therefore, the corresponding critical energy would be the energy after four interactions $\xi_c^\pi = 20$ GeV as discussed above. Generally, the critical energy decreases with an increase in the initial energy. The number of interactions the pion will undergo before reaching the critical energy can be calculated as

$$n_c = \frac{\ln[E_0/\xi_c^\pi]}{\ln\left[\frac{3}{2}n_{\text{ch}}\right]} = 0.85 \lg \left[\frac{E_0}{\xi_c^\pi} \right], \quad (1.10)$$

1. Cosmic rays and extensive air showers

For showers initiated by primary particles with energies between 10^{14} and 10^{17} eV, the critical energy would be in the 10-30 GeV range, and n_c would take values between 3 and 6.

The number of muons produced after the critical energy of pions is reached:

$$N_\mu = N_\pi = n_{\text{ch}}^{n_c} = \left(\frac{E_0}{\xi_\pi \zeta_c} \right)^\beta, \quad (1.11)$$

where the β parameter is

$$\beta = \frac{\ln n_{\text{ch}}}{\ln \frac{3}{2} n_{\text{ch}}}, \quad (1.12)$$

so it depends only on n_{ch} and varies between 0.85 and 0.92 [53], resulting in a little less than linear increase in the number of muons with energy. For the case $n_{\text{ch}} = 10$, as assumed above, β is approximately 0.85. At the end of the cascade, all the primary energy is transferred to the electromagnetic component and the muons:

$$E_0 = \xi_c^\pi N_{\text{max}} + \xi_\pi^\pi N_\mu. \quad (1.13)$$

This expression is a reflection of the conservation of energy. Considering that the number of electrons is $N_e = N_{\text{max}}/g$, where g is a correction factor taking into account the detector properties when comparing the estimated electron number to the measured value, it can be re-written as

$$E_0 = 0.85 \text{ GeV} (N_e + 24 N_\mu). \quad (1.14)$$

Based on the primary type, the energy distribution between the electromagnetic and muonic components would be different. Since Eq. 1.14 shows only a linear dependence on the number of muons and electrons, the initial energy can be easily determined if both these quantities are measured.

The depth of the shower maximum in a hadronic cascade can be expressed by considering the first generation $n = 1$ of electromagnetic subshowers [48]:

$$X_{\text{max}} \approx \lambda_i + \lambda_r \ln \frac{E_0}{2 n_{\text{tot}} \xi_c^\pi} = \lambda_i + \lambda_r \ln \frac{E_0}{3 n_{\text{ch}} \xi_c^\pi} = \lambda_i + X_{\text{max}}^\gamma - \lambda_r \ln[2 n_{\text{ch}}]. \quad (1.15)$$

Adding the higher hadronic generations would not change the structure of Eq. 1.15, but for an accurate estimate of X_{max} , it is essential to account for all




Figure 1.4: Decomposition of the observed X_{\max} distribution. The X_{\max} distribution can be represented as a convolution of the distribution of the depths of the first interaction point, X_1 , and of the difference $\Delta X_1 = X_{\max} - X_1$. Figure from [54].

generations of the cascade from their origin points, as well as the influence of the leading particle. Utilizing Eq.1.15 would typically result in an underestimation of X_{\max} by approximately $2\lambda_i$ [51]. The elongation rate for the hadronic shower is then

$$\lambda = \lambda^\gamma + \frac{d}{d \lg E_0} [\lambda_i \ln 2 - \lambda_r \ln(3n_{\text{ch}})]. \quad (1.16)$$

Therefore, the elongation rate for the hadron-initiated shower will differ from the electromagnetic elongation rate by the factor depending on the multiplicity n_{ch} and hadronic interaction length λ_i : the hadronic elongation rate will be smaller for larger n_{ch} and will increase with increase in λ_i . For protons, the calculations result in $\lambda^p = 58 \text{ g/cm}^2$, which is smaller by 27 g/cm^2 than the electromagnetic elongation rate.

The depth X_{\max} can be decomposed into the distribution of the depths of the first interaction, X_1 and the $\Delta X_1 = X_{\max} - X_1$, see Fig. 1.4. From this, the original X_1 -distribution can be derived from the measured X_{\max} and $\Delta X_1 = X_{\max} - X_1$ obtained from the air shower simulations. The average depth of the first interaction, $\langle X_1 \rangle$ is equal to the interaction length λ_i .

By knowing the depth of the first interaction point in the air, it is possible to estimate the corresponding particle interaction cross section. For the interaction of protons in the atmosphere, X_1 is distributed as [55]

$$\frac{dp}{dX_1} = \frac{1}{\lambda_{p\text{-air}}} e^{-X_1/\lambda_{p\text{-air}}}, \quad (1.17)$$

1. Cosmic rays and extensive air showers

So, if the distribution of the depth of the first interaction point in air is measurable, it is possible to compute $\lambda_{p\text{-air}}$. The proton-air cross section is then

$$\sigma_{p\text{-Air}} = \frac{\langle m_{\text{Air}} \rangle}{\lambda_i}, \quad (1.18)$$

where the mean mass of air is $\langle m_{\text{Air}} \rangle \approx 14.45m_p$ and λ_i then has a meaning of the interaction length of protons in air.

The above derivations are specific to the basic scenario where the primary particle is a proton. In the case of the nucleus with atomic mass $A > 1$, the so-called superposition model can be applied, where each nucleus with energy E_0 and is treated as A single nucleus with energy E_0/A initiating A independent showers. Accordingly, the predictions for the number of electrons and muons and for the depth of the shower maximum can be generalised by replacing E with a reduced energy E_0/A :

$$N_\mu^A = N_\mu^p A^{1-\beta}, \quad (1.19)$$

$$X_{\text{max}}^A = X_{\text{max}}^p - \lambda_r \ln A. \quad (1.20)$$

where the superscript (p) refers to the expression derived for the proton-initiated hadronic shower. Therefore, both the number of muons and the depth at which the air shower reaches its maximum will depend on the primary nucleus. For the same primary energy, heavier nuclei will produce more muons. The depth of the shower maximum will be shallower. It is, therefore, possible to use the measured N_μ and X_{max} to distinguish between the showers initiated by the different nuclei primaries and by measuring X_{max} and E , $\ln A$ can be obtained.

1.3 Hadronic interaction models

The energies of cosmic rays cover a wide range, well beyond the maximum energy achievable by existing collider facilities such as the Large Hadron Collider (LHC) at CERN. In addition, most of the interactions in the air shower take place in the forward region, e.g. with pseudorapidity of $\eta > 5$, while most accelerators are designed to be sensitive to the intermediate range, and it is technically challenging to study forward particle production. Both factors pose significant difficulties for studying particle interactions in EAS with existing collider experiments. Thus,

the interpretation of cosmic ray measurements at ultrahigh energies must rely on extrapolations from data obtained at lower energies.

Several Monte Carlo generators have been developed and used to model the hadronic interactions in EAS cascades and predict air shower observables. Since it is currently impossible to characterize the soft hadronic interactions in air showers from the first principles, the existing hadronic interaction models are phenomenological and utilize the effective theory approach. The most frequently used high-energy models are EPOS-LHC [56], QGSJETII [57], and Sibyll [58, 59]. The recent versions of these models belong to the post-LHC generation of hadronic interaction models, which means they have been fine-tuned and re-parameterized based on data obtained from experiments at the LHC. In air shower simulations with Monte Carlo codes such as CORSIKA [60] or CONEX [61], in addition to high-energy interaction models, low-energy interaction models are also included to account for interactions below 200 GeV. Some commonly used low energy models are for example FLUKA [62] and UrQMD [63]. Since the main interest of this work is the cosmic ray data at the highest energies, and the studied observables are not affected by the choice of the low-energy model, only the main features of the high-energy models and the implications of the differences between them for the interpretation of the cosmic ray measurements are discussed.

1.3.1 High-energy models

All high-energy interaction models are generally based on the Gribov-Regge field theory (GRFT) [64]. In GRFT, multiple hard and soft particle interactions occur simultaneously, involving the exchange of microscopic parton cascades, each treated as an emission of the effective color-neutral object called Pomeron [65]. Within the GRFT framework, parton cascades are considered to have both perturbative and non-perturbative parts. The amplitude of the Pomeron can be analyzed in two different ways: one approach, referred to as the soft+hard approach, considers it as a combination of pure hard and soft components, while the other approach treats it as a convolution of these two components (known as the semi-hard approach) [66]. Within the semi-hard approach, a general Pomeron (Parton ladder) is assumed, consisting of the purely phenomenological soft component and a hard or semi-hard component with the perturbative part parameterized following the Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) formalism [67]. This approach is implemented in both the EPOS-LHC and QGSJETII models. In the context of the EPOS-LHC model, an additional mechanism for particle production exists, which involves the decay of off-shell remnants [68]. In the central

collisions involving hadrons and nuclei at very high energies, high parton densities are achieved, leading to substantial nonlinear corrections to the interaction dynamics, such as screening and saturation. In the EPOS-LHC model, the nonlinear effects due to Pomeron-Pomeron interactions are treated via the modified amplitude tuned to reproduce the observed cross section and multiplicity data [69]. On the other hand, in the QGSJETII model, the non-linearities are explicitly described in individual hadronic and nuclear collisions using an enhanced Pomeron diagram [70].

The Sibyll hadronic interaction model and an older DPMJET model [71] take a different approach toward the treatment of the particle interactions. They combine the minijet approach [72] commonly found in Monte Carlo generators used in collider physics, together with some aspects of GRFT. In the Sibyll model, an energy-dependent transverse momentum cutoff separates the hard and soft interaction regions, allowing the high-density phase space to be avoided. The cutoff also takes into account the nonlinear effects. Above the cutoff, the cross sections for the production of parton jets can be calculated using QCD. In the initial versions of the Sibyll model, it allowed for the possibility of multiple hard interactions with only one soft interaction. However, starting with the Sibyll 2.1 version [73], the range of phase space for soft interactions has been extended, permitting multiple soft interactions as well.

While all these models share a common conceptual framework, their different approaches to the underlying physics of hadronic interactions result in different predictions for cosmic-ray observables and varying degrees of agreement with the LHC data. Significant differences are present not only between different models but also between different versions of a single model. The choice of either the semi-hard or minijet approach has a direct impact on the predictions for parameters such as elasticity (the largest energy fraction carried by the leading particle) and pseudorapidity (the measure of a particle angle relative to the beam), with Sibyll 2.3d and QGSJETII-04 having a narrower pseudorapidity, resulting in some disagreement with LHC data in the TeV region [74]. For cosmic-ray energies up to about 100 TeV, the variation in multiplicity predictions between all models is minimal. However, as energies exceed the TeV range, there is a difference in multiplicity of up to 20-30% between the different models for p-air and π -air interactions.

Since all the recent models are tuned to the latest LHC measurements, they describe the proton-proton cross sections measured at the accelerator facilities well. Moreover, beyond the LHC limit, the differences in the predictions of different models are minor, as can be seen in Fig. 1.5 (bottom), where the proton-proton

cross sections from the air shower simulations with the discussed MC generators are compared with the results from the collider experiments (see for reference [75, 76, 77, 78, 79, 80]) and with the measurement from the Pierre Auger Observatory.

A tendency similar to the change in the multiplicity with energy can be seen in the inelastic proton air and pion air interactions cross sections. For instance, up to the energies reached at the LHC, all the models accurately reproduce accelerator data and show no substantial difference. However, when extended to extrapolated energy ranges, the predictions of different models begin to diverge. Fig. 1.5 (top) shows the predictions from air shower simulations with different models in comparison to cosmic-ray data. The most significant difference is between the two versions of the Sibyll generator, Sibyll 2.1 and Sibyll 2.3d models, where the former produces cross sections that are approximately 10-15% larger than the predicted by the earlier version. Since in the Sibyll MC generator the proton-air cross sections are calculated based on the proton-proton cross sections via the Glauber formalism [81], the decrease in the proton-air cross sections in the Sibyll 2.3d model is partly due to the smaller proton-proton cross sections. Due to the smaller proton-proton cross sections, there is some additional reduction in the proton-air cross section. This reduction is due to the inclusion of inelastic shielding, where a particle's electric charge is shielded by surrounding particles, which was omitted in the earlier Sibyll version. The proton-air cross sections calculated with the EPOS-LHC model fall between the Sibyll 2.3d and Sibyll 2.3c extrapolations, and a minimal difference is present between the predictions obtained with Sibyll 2.3d and QGSJETII-04 models for proton air production cross sections.

The differences between the models in the description of the cascade evolution and the predictions of the interaction properties are also reflected in the estimated depth of the maximum air shower. The depth of the shower maximum X_{\max} is primarily determined by the position of the first interaction, and is also influenced by other factors such as diffraction, multiplicity, and elasticity. Consequently, one can expect the corresponding changes in X_{\max} due to the variations in these factors. Fig. 1.6 (left) shows the difference in the average X_{\max} for pure proton and iron showers between the different models. Overall, the shift between all models is nearly constant with energy. The smallest and the largest $\langle X_{\max} \rangle$ values are obtained from the simulations with the QGSJETII-04 and Sibyll 2.3d models, respectively. Sibyll 2.3d is about 20-25 g/cm² deeper than QGSJETII-04 and also 10 g/cm² deeper than EPOS-LHC, leading to an interpretation of the data towards a heavier composition. Additionally, Sibyll 2.3d provides the largest discrimination between proton and iron showers. The increase in $\langle X_{\max} \rangle$ predicted by the

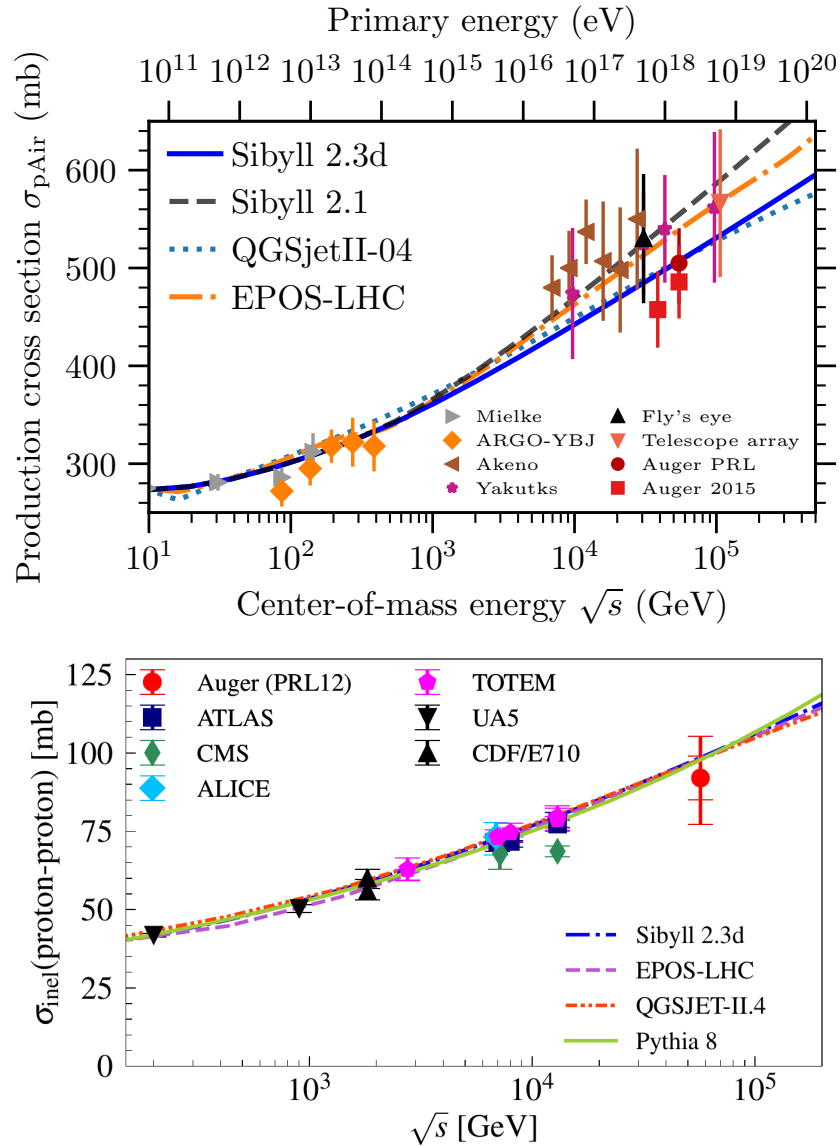


Figure 1.5: The evolution of the proton-air and proton-proton production cross section with energy. Top: predictions of the proton-air interaction cross section compared to measurements with EAS experiments. Figure from [59]. Bottom: predictions of the proton-proton interaction cross section compared to measurements with EAS and particle accelerator experiments.

latest version of Sibyll is attributed to the larger interaction length, smaller multiplicity, and larger elasticity for the proton interactions in air [59]. The differences in the mean fluctuations of X_{\max} , $\sigma(X_{\max})$, are shown in Fig. 1.6 (bottom). The QGSJETII-04 and Sibyll 2.3d predictions are once more two extremes, with the difference between them of $\approx 20 \text{ g/cm}^2$. Apart from Sibyll 2.3d, other models discussed in this chapter predict roughly the same fluctuations, with $\Delta\sigma(X_{\max})$ on the order of a few g/cm^2 . In the case of the iron shower, the $\sigma(X_{\max})$ shows relatively little sensitivity to the choice of model and does not exceed 10 g/cm^2 , with EPOS-LHC having the smallest $\sigma(X_{\max})$.

The number of muons at the ground level N_μ is susceptible to the properties of the hadronic interactions and, therefore, differs between models. Muons are produced by decaying hadrons at the very end of the cascade development. Therefore, the number is mainly affected by the secondary particle multiplicity and the baryon production rate, with about 30% more muons produced in subshowers initiated by baryons [48]. A higher multiplicity and a higher baryon production rate would have the effect of increasing the number of muons at ground level. As mentioned above, the number of muons is also significantly affected by the type of primary particle that initiates the air shower. With an increase in the particle's mass initiating the shower, more muons are expected to be produced.

In general, the prediction of the number of muons is consistent across all post-LHC models. However, none of them aligns with the observed number of muons recorded by various cosmic-ray experiments, such as the Pierre Auger Observatory [85, 86], Telescope Array [87], and IceCube Observatory [88], resulting in too few muons predicted in these models. This discrepancy between the predictions from the air shower simulations and the cosmic ray data starts above 10^{16} eV and increases linearly with the logarithm of primary energy, reaching 30% - 60% at 10^{19} eV . It can be noted that some other experiments, such as EAS-MSU [89] and KASCADE-Grande, do not report any muon deficit. There is also no discrepancy between the simulations and measurements of the Yakutsk EAS array [90], except for some decrease in the muon density observed for the inclined showers at $\approx 10^{19} \text{ eV}$ in the latest data [90]. These differences might be related to different muon energy thresholds and zenith angle ranges for which these results were reported.

The existence of the muon deficit, also known as the Muon Puzzle, suggests that there are some shortcomings in the existing models of hadronic interactions. Numerous studies have attempted to explain and solve the discrepancy, but have yet to be successful so far, see [91, 66] for discussion on the various attempts to resolve the Muon Puzzle. The observed effect is assumed to originate from soft-

1. Cosmic rays and extensive air showers




Figure 1.6: Comparison of model predictions for the moments of the X_{\max} distribution with data from the Pierre Auger Observatory [82, 83]. Top: mean X_{\max} , $\langle X_{\max} \rangle$. Bottom: X_{\max} fluctuations, $\sigma(X_{\max})$. Figures from [84].

QCD interactions, given that the infrequent occurrence of hard scatterings leading to the production of new heavy particles is insufficient to account for the differences. Given the strong sensitivity of the muon count to the fraction of energy carried away by electromagnetic component, a reduction in this fraction would release more energy available for muon production and thus lead to an increase in the number of muons [66] at the end of the cascade. While preserving consistency with data for other air shower observables, the desired effect may be achieved by reducing the π^0 fraction. The decrease in the secondary pion fraction may be attributed to collective hadronization, a consequence of the formation of quark-gluon plasma blobs [92, 93]. Given the early onset of the deficit, there is potential to investigate the problem in the high-energy hadron collisions at the LHC. In the recent ALICE data, an increase in the strangeness production was measured at medium rapidities [94], providing a potential lead in resolving the Muon Puzzle. It has been demonstrated in [95] that an empirical toy model constructed based on the ALICE measurements can address the muon deficit issue. It is important to note that, generally, none of the existing interaction models provides sufficient agreement with the observed trends in hadronic interactions according to the ALICE data. Further constraints on multi-hadron production can be gained from forward neutrino flux measurements at FASER ν [96], and new experiments at the Forward Physics Facility (FPF) [97].

1.3.2 Modified hadronic interactions

As outlined in the previous section, the development of extensive air showers is highly sensitive to the properties of hadronic interactions at ultrahigh energies. Specifically, key factors include the inelastic cross section for hadrons in air, the hadron multiplicity, elasticity, and the ratio of electromagnetic to hadronic energy flow [98]. The change in these properties could have a significant impact on the prediction of the observables of the air shower and on the interpretation of the cosmic ray data. The effects of modifications in hadronic interactions have been extensively studied and addressed in various works, such as the analyses conducted by R. Ulrich et al. [99] and J. Blazek et al. [100, 101]. Such works also address the issue of whether the discrepancies among existing interaction models can be used as uncertainties in our understanding of the hadronic interaction models, particularly assessing whether they exceed the actual systematic uncertainties or are too limited in scope to cover all possibilities. The key findings of the studies on the modified properties of hadronic interactions are briefly summarized below, given their high relevance to the subject of this thesis.

1. Cosmic rays and extensive air showers

In the work by Ulrich et al. [102], the individual properties of hadronic interactions are modified by an energy-dependent factor:

$$f(E, f_{19}) = 1 + (f_{19} - 1)F(E), \quad (1.21)$$

where

$$F(E) = \begin{cases} 0 & \text{if } E \leq E_{\text{tr}}; \\ \frac{\lg(E/E_{\text{tr}})}{\lg(10^{19}\text{eV}/E_{\text{tr}})} & \text{if } E > E_{\text{tr}}, \end{cases} \quad (1.22)$$

The factor is defined in such a way that is equal to 1 below the threshold energy E_{tr} , where the model predictions are tuned to the LHC data. Beyond the threshold energy, deviations from the properties defined by the model increase logarithmically, reaching a value of f_{19} at 10^{19} eV. These modifications were implemented into the CONEX air shower program to investigate the impact of deviations from the original model on the observables of extensive air showers (EAS). The factor $f(E, f_{19})$ is evaluated after each hadronic interaction within the cascade. If its value deviates from unity, the properties of individual secondary particles are altered using the resampling algorithm. These modifications have been implemented in the Sibyll generator, and since it is a semi-superposition model, the same resampling algorithms can be applied to showers initiated by any nuclei [103], taking into account the energy scaling with the atomic mass number $E = E_0/A$, where E_0 is the total energy of the projectile nucleus. The cross sections for nucleus-air interactions, calculated within the framework of the modified hadronic interaction model, are linked to the corresponding nucleon-air (proton-air) interaction cross sections through Glauber theory [81].

Figs. 1.7 and 1.8 show the effect of the rescaled interaction properties, such as production cross section, multiplicity, and elasticity, on the predictions of air shower observables for proton and iron primaries, respectively. As can be seen from the plot, the modification in the inelastic cross section has the largest effect on the depth of the maximum air shower development, e.g. on the mean $\langle X_{\text{max}} \rangle$ and standard deviation $\sigma(X_{\text{max}})$. The change in the cross section by 50% may reduce or increase (depending on the direction of the change) the predictions for the $\langle X_{\text{max}} \rangle$ by more than 100 g/cm² for proton-initiated shower and by more than 40 g/cm² for iron-initiated showers. For the fluctuations of X_{max} , significant changes are only observed for primary protons when the cross section is reduced, leading to an increase in the standard deviation of more than 100 g/cm² for a cross section twice as small. When the cross section is 50% larger than predicted by the unmodified model, there is only a small decrease in the fluctuations of about 10

g/cm², with no sensitivity to further increases in the cross section value. Regarding iron, the changes in the standard deviation are nearly symmetrical for both the decrease and increase in the cross section within the rescaling parameter range, resulting in a change of 3-4 g/cm² compared to the predictions from simulations with standard Sibyll cross sections. The elasticity and multiplicity also impact the predictions for X_{\max} , mainly affecting the $\langle X_{\max} \rangle$ with less dependence in the X_{\max} fluctuations than for the changes in the cross sections. The decrease in multiplicity leads to more deep showers with a rate close to the effect of the change in cross sections within a narrow range of f_{19} values close to the original model. As the evolution in X_{\max} with the multiplicity becomes more complicated than a negative logarithmic function for a larger significant deviation from the standard cross section values, the dependence starts to behave more like an inverse function with the further increase/decrease in f_{19} . Elasticity has an effect opposite to the multiplicity on $\langle X_{\max} \rangle$, resulting in shallower showers for larger values. For a rescaling factor twice as large, the increase is around 30-40 g/cm² for both primary protons and iron. While the standard deviation shows no dependence on varying the multiplicity, increasing the elasticity would also result in larger fluctuations. However, as the elasticity becomes less than the original mode, the sensitivity of X_{\max} to the changes in the elasticity gradually decreases, with almost no effect on the X_{\max} fluctuations for f_{19} below unity.

In addition to the effect on the depth of the maximum air shower development, the changes in the interaction properties will also affect the number of electrons and muons observed at ground level. There is both a direct change from the interaction properties themselves and a change due to the effect on X_{\max} : for larger X_{\max} values, more electrons will be observed than for shallower showers. Note that in the case of muons, X_{\max} has a much smaller impact on the predicted N_{μ} since muons are heavily attenuated in the atmosphere compared to the electrons. While the number of electrons is mainly affected by X_{\max} , the direct modifications in the interaction properties are more relevant for muons. Thus, it can be seen that the mean $\log_{10} N_e$ as well have a change correlated with that of X_{\max} for the modifications in cross section, multiplicity, and elasticity. However, the effect of the interaction properties themselves also can not be neglected for the number of electrons, since the fluctuations in $\log_{10} N_e$ do not strictly follow the evolution of $\langle X_{\max} \rangle$ with modifications in the original model, except for the multiplicity, which, indeed, has the largest influence on the predictions for the number of electrons at the ground level. The cross sections also affect the fluctuations in the $\log_{10} N_e$, but only if they are smaller than the extrapolation in the original model. Thus, either the multiplicity or the interaction cross section can be reduced

1. Cosmic rays and extensive air showers




Figure 1.7: The effect of the modified hadronic interaction properties on air shower observables for the proton-initiated cascade: shower maximum X_{\max} (left), number of electrons N_e (middle), and number of muons N_μ (right). The averages of 1000 showers simulated at $10^{19.5}$ eV are shown. Figure from [102].

to effectively make the fluctuations in the predicted number of electrons.

As was mentioned above, the interaction properties directly affect the number of muons, and there are fewer correlations with X_{\max} . In contrast to both $\langle X_{\max} \rangle$ and mean $\log_{10} N_e$, which show an anti-correlation with the multiplicity, the mean $\log_{10} N_\mu$ increases for the large multiplicity values. However, the effect of cross sections and elasticity is much less pronounced for the mean and the fluctuations of $\log_{10} N_\mu$ seem to be solely determined by the elasticity.

The same effects as seen for the number of muons and electrons in the proton-initiated showers are also present for the iron primaries but to a much lesser extent.

1.3.3 Hadronic cross sections and Glauber theory

The Glauber multiple scattering theory, introduced in [103], provides a link between hadron-hadron interactions and interactions involving nuclei, such as nucleus-air interactions in air shower simulations. One of the main assumptions in Glauber theory is that the eikonal functions of the individual scattering processes could be added up linearly [104]. For example, for a scattering process with two target nucleons, the elastic amplitude will be

1. Cosmic rays and extensive air showers




Figure 1.8: The effect of the modified hadronic interaction properties on air shower observables for the iron-initiated cascade: shower maximum X_{max} (left), number of electrons N_e (middle), and number of muons N_μ (right). The averages of 1000 showers simulated at $10^{19.5}$ eV are shown. Figure from [102].

$$G_{\text{Glauber}}(\vec{b}) = 1 - \exp^{-\chi_1(\vec{b}) - \chi_2(\vec{b})} = 1 - (1 - \Gamma_1)(1 - \Gamma_2), \quad (1.23)$$

where $\Gamma_{1,2} = 1 - \exp^{-\chi_{1,2}}$ is an impact parameter written in terms of the eikonal function $\chi_{1,2}$.

The scattering amplitude of a hadron h interacting with a nucleus of mass number A is then

$$f_{fi}^{hA}(s, q^2) = \frac{ik}{2\pi} \int \exp^{i\vec{q}\vec{b}} \psi_f^*(\vec{r}_1 \dots \vec{r}_A) \Gamma_{hA}(\vec{b}, \vec{s}_1 \dots \vec{s}_A) \psi_i(\vec{r}_1 \dots \vec{r}_A) d^2b \prod_{j=1}^A d^3r_j, \quad (1.24)$$

where ψ_i and ψ_f are wave functions of the nucleus in the initial and final states. \vec{r}_j and \vec{s}_j represent the positions of the nucleons in the nucleus and their projections on the plane perpendicular to the momentum vector. From the optical theorem, the total and elastic cross sections are then

$$\sigma_{hA}^{\text{tot}} = 2\Re \int \tilde{\Gamma}_{hA}(\vec{b}) d^2b \quad (1.25)$$

$$\sigma_{hA}^{\text{ela}} = \int |\tilde{\Gamma}_{hA}(\vec{b})|^2 d^2b, \quad (1.26)$$

1. Cosmic rays and extensive air showers

with an impact amplitude of

$$\tilde{\Gamma}_{hA}(\vec{b}) = 1 - \prod_{j=1}^A [1 - \int \Gamma_{hN}(\vec{b} - \vec{s}_j) \rho_j(\vec{r}_j) d^3 r_j]. \quad (1.27)$$

In the last expression $\tilde{\Gamma}_{hN}$ is the amplitude describing the interaction of the hadron h with nucleon N , and ρ_j is the single nucleon density.

The production cross section, also referred to as the inelastic cross section, can be obtained by subtracting the elastic σ_{hA}^{ela} and quasi-elastic $\sigma_{hA}^{\text{inela}}$ scattering from the total cross section. The sum of the latter two can be expressed as

$$\sigma_{hA}^{\text{ela}} + \sigma_{hA}^{\text{inela}} = \int \left| 1 - \prod_{j=1}^A [1 - \Gamma_{hN}(\vec{b} - \vec{s}_j)] \right|^2 \left(\prod_{j=1}^A \rho_j(\vec{r}_j) d^3 r_j \right) d^2 b. \quad (1.28)$$

With the superposition approach the expression above could be generalized for any nucleus-nucleus interaction cross sections:

$$\sigma_{AB}^{\text{ela}} + \sigma_{AB}^{\text{inela}} = \int \left| \prod_{j=1}^A \prod_{k=1}^B [1 - \Gamma_{NN}(\vec{b} - \vec{s}_j \vec{\tau}_k)] \right|^2 \left(\prod_{j=1}^A \rho_j(\vec{r}_j) d^3 r_j \right) \left(\prod_{k=1}^B \rho_k(\vec{r}_k) d^3 r_k \right) d^2 b, \quad (1.29)$$

with the corresponding total and elastic cross sections:

$$\sigma_{AB}^{\text{tot}} = 2\Re \int \tilde{\Gamma}_{AB}(\vec{b}) d^2 b \quad (1.30)$$

$$\sigma_{AB}^{\text{ela}} = \int |\tilde{\Gamma}_{AB}(\vec{b})| d^2 b. \quad (1.31)$$

The positions of nucleons in the second nucleus with mass number B are denoted by \vec{t}_k and $\vec{\tau}_k$.

In air shower simulations, Glauber theory provides the simplest and fastest way to link the proton-proton and proton-nucleus interaction properties. Its implementation in the Sibyll MC generator allows easy conversion of any modifications in the model extrapolations for proton-proton to any nucleus-nucleus interactions.

Chapter 2

The Pierre Auger Observatory

The Pierre Auger Observatory, located near the city of Malargue in the province Mendoza in Argentina, at the height of 1400 m above sea level, is the world's largest observatory for studying ultrahigh-energy cosmic rays [105]. The concept of building an observatory of such a scale and scientific potential was initially proposed by Jim Cronin and Alan Watson at the International Cosmic Ray Conference in Dublin in 1991. The observatory's design, including the scientific discussion and cost estimate, was completed in 1995. The construction of the observatory started in the Southern Hemisphere in 2002 after extensive six months of studies on concept validation and the detector performance with the Engineering Array consisting of 32 surface detector stations and a single fluorescence telescope. Construction of the Observatory was finalized in 2008, and since January 2004, it has been continuously detecting the most energetic particles reaching the Earth.

The observatory has a hybrid design consisting of the surface detector [106] measuring the cosmic-ray particles at the ground level (SD), and the fluorescence detector (FD) [107] observing the development of the longitudinal air shower profile in the atmosphere. The layout of the Pierre Auger Observatory is shown in Fig. 2.1. The SD array consists of 1600 water Cherenkov detectors (WCD). It is overlooked by 24 fluorescence telescopes located at four sites - Los Leones (LL), Los Morados (LM), Loma Amarilla (LA), and Coihueco (CO), with six telescopes at each location. In addition, three high-elevation telescopes (HEAT) were built at the Coihueco site. Near it, the AMIGA (Auger Muon and Infill for the Ground Array) detector has been installed, consisting of the muon counters buried underground and the more densely packed array of 60 water Cherenkov stations at a distance of 750 m.

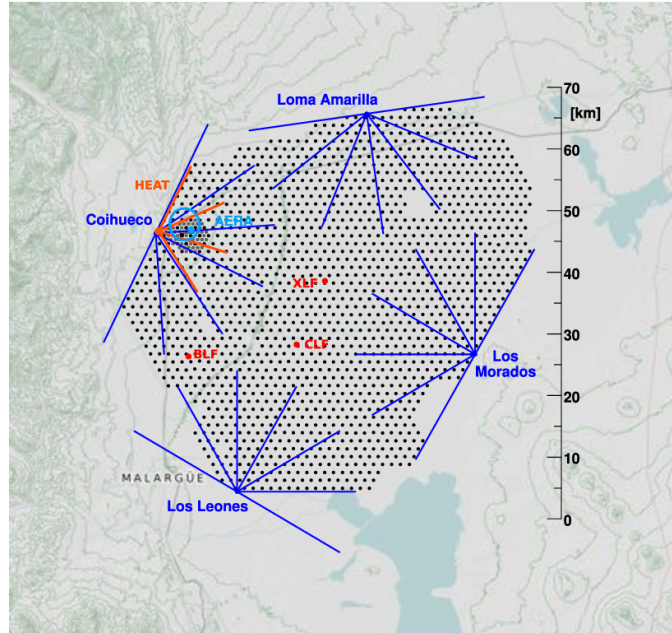


Figure 2.1: Layout of the Pierre Auger Observatory. Each black dot corresponds to one water Cherenkov detector. Four fluorescence detector sites (Los Morados, Los Leones, Coihueco, and Loma Amarilla) are shown in blue, with additional extra high elevation (HEAT) telescopes at the Coihueco site shown in red. The lines of the corresponding colors indicate the field of view of each of the 24+3 fluorescence telescopes. Near Coihueco, the denser Infill array and Engineering Radio Array (shown in blue) are located. Two laser (XLF and CLF) and one balloon (BLF) facilities are marked with red dots. Figure from [108].

The standard range of energies measured by Observatory is $\gtrsim 10^{18}$ eV. Combining the HEAT telescopes and the 750 m array allows one to extend the observatory's sensitivity down to lower energies of $\gtrsim 10^{17}$ eV.

2.1 The surface detector

The 1660 surface detector stations of the regular array (SD-1500) are spaced 1.5 km apart on a hexagonal grid and cover an area of about 3000 km². Located 6 km away from the Coihueco site, there is an Infill array consisting of 61 stations spread over 23.5 km² with additional stations nested within the standard

grid, reducing the spacing between the WCDs to 750 m. The detection efficiency at the trigger level reaches 100% for energies above $3 \cdot 10^{18}$ eV for the SD-1500 array, and the full efficiency range can be extended down to above $\approx 3 \cdot 10^{17}$ eV with the Infill array. The duty cycle of the SD is close to 100%.

Each WCD station consists of a cylindrical tank, 3.6m in diameter and 1.5 m high, filled with 12000 liters of purified water. The water level in the tank reaches up to around 1.2 m, which allows the detection of inclined air showers with zenith angles of up to 80° . The water is enclosed by a sealed liner (Tyvek bag) with an opaque, low-density polyethylene reflective inner surface. When an energetic charged particle passes through the water volume, Cherenkov light is emitted. The light is then diffusely reflected from the liner surface and detected by one of three photomultiplier tubes (PMTs), with properties suitable for long-term measurements of Cherenkov radiation with a peak in the 350-450 nm range (see [109] for the discussion on the required PMTs characteristics), mounted on top of the liners. The liner has three windows through which the PMTs can see the water volume. Two signals are taken from the PMTs, one from the PMT anode and one from the last dynode, the latter being inverted and amplified to 32 times the anode charge gain. The signals are digitized at a 40 MHz sampling frequency by 10-bit flash analog-to-digital converters (FADCs). The combination of these two signals provides sufficient dynamic range to cover the full range of fluxes detected close to the shower core (~ 1000 particles per μs) and far from the shower core (\sim one particle per μs). The emitted Cherenkov light is measured in the units of a vertical-equivalent muon (VEM). The signals obtained from the cosmic-ray-induced EAS particles are selected by a complex hierarchical trigger system [110]. The trigger chain at each surface station includes two levels of local triggers, T1 and T2, which operate independently, and a third level, T3, which is formed at the Central Data Acquisition System (CDAS) and initiates central data acquisition for further selection of physical (T4) and accurate (T5) events from the stored data [111], see below for detailed discussion.

At level 1, the data acquisition is triggered in the station and temporarily stored in case a T3 will be formed. The T1 has two different trigger modes for detecting electromagnetic and muonic components of the EAS. In the first mode, which is designed to select large signals, both close and spread in time, a threshold trigger (TH) is used. It is formed when each of the three PMTs has a signal above $1.75I_{VEM}^{peak}$ in coincidence, and the trigger rate is around 100 Hz. In the second mode with the 'Time-over-Threshold' (ToT) trigger, the sequences of small and spread-out signals are detected. The ToT trigger has a rate of less than 2 Hz,

2. The Pierre Auger Observatory

mainly attributed to the coincidence of the signals from two muons within the duration of the sliding window. While the TH mode is the most sensitive for detecting highly inclined muonic showers, the ToT mode is efficient for detecting close low-energy showers of electromagnetic origin or high-energy distant EAS. The ToT mode requires a coincidence of at least two PMTs with a minimum of 13 bins out of 120 FADC bins above a threshold of $0.2I_{\text{VEM}}^{\text{peak}}$. The corresponding time spread of 325 ns allows us to reject the muon background since it is larger than the average signal duration of a single muon of 150 ns.

At the second trigger level, the trigger rate is decreased to approximately 20 Hz per detector due to the bandwidth limitations of the communication systems. While all ToT triggers of the first level are reaching the T2 level, the additional selection is applied to the TH triggers, requiring a higher threshold of $3.2I_{\text{VEM}}^{\text{peak}}$ in coincidence. For the initiation of the T3 trigger for the data acquisition at CDAS, spatial and temporal coincidence of the T2 triggers is required. The first condition is met by requiring at least three neighboring detectors to be ToT-triggered, e.g., a detector should have its closest and second closest neighbors also triggered. There should be a particular spread between the first and subsequent triggers to fulfill the timing condition.

At the T4 level, physical events, e.g., real showers, are selected from the T3 set. In addition to the ToT trigger criteria at the T3 level, which already selects predominantly actual events due to the low muon background, additional requirements are placed on the geometry of the three nearby triggered stations. Furthermore, the timing of the trigger should be compatible with the speed of light. These criteria ensure a selection efficiency above 98% for events with zenith angles below 60° . The efficiency could increase to $\approx 100\%$ with a compact configuration (a detector station has at least three triggered tanks out of the closest six neighbors) and timing requirements imposed on the four nearby stations. In the end, only the events that could be properly reconstructed are to be selected. This is done with a fiducial (quality) trigger, T5, which rejects the events with the wrong core position and, therefore, with the wrong energies. This is particularly relevant for events recorded at the edge of the array, where a shower part could be missing, with the real core outside the array and the reconstructed one inside. To ensure the correct energy reconstruction, only the events well contained in the array are selected by the T5 trigger. This is achieved by requiring all six of the closest neighbors of the detector with the highest signal to be in operation. The resulting effective area is reduced by 10%.

The calibration of the surface detector station is done with the measurements




Figure 2.2: The water Cherenkov detector station of the surface detector. Left: Photograph of the WCD station. Figure from [112]. Right: Schematic view of one WCD station. Figure from [113]. A detailed description of the design and operation of the WCD can be found in the text.

of atmospheric muons to define the conversion between 1 VEM and the measured signal [114]. The peak of the light produced by the muons is measured from the charge (the voltage amplitude peak) histograms. The position of the peak for each PMT and the sum of three PMTs is determined, 1.03 and 1.09 VEM, respectively, and can be used to convert the PMT electronics signal to one VEM.

The station is powered by a solar photovoltaic system consisting of two 55 Wp solar panels and two 105 Ampere-hour (Ah), 12 V batteries, producing on average 10 W power. For the transmission of the data to the communication towers, which in turn provide the link with CDAS and the correct timing, the communication radio and GPS antennas are installed at the top of the WCD.

2.2 The fluorescence detector

As the charged cosmic rays pass through the atmosphere, they interact with atmospheric molecules. These interactions excite the rotational and vibrational modes of the nitrogen molecules, resulting in ultraviolet emission in the $\approx 300\text{-}430$ nm range, which can then be collected with fluorescence telescopes. The atmosphere acts as a calorimeter for determining the energy of EAS particles, and the number of emitted photons is proportional to the energy deposited in the atmosphere. Measured as a function of the atmospheric slant depth X , defined as the amount of material traversed by the shower, the emission rate gives us the development of

2. The Pierre Auger Observatory




Figure 2.3: The FD site of the Pierre Auger Observatory. Left: Photograph of the FD station | by Tobias Schulz. Right: Schematic layout of the FD station. Figure from [107].

the longitudinal profile of the energy deposit in the atmosphere dE/dX .

As mentioned above, there are 27 fluorescence telescopes - six at each of the four FD sites, plus three additional high-elevation telescopes at the Coihueco site, extending the measured energy range to about 10^{17} eV. An example of an FD station and a schematic view of its interior is shown in Fig 2.3. The field of view of each telescope covers an angular range of 30° by 30° in azimuth and elevation. Due to their relative orientation towards the field, a single FD station's total coverage is 180° in azimuth. The HEAT telescopes can see elevation angles from 30° up to 58° due to the increased inclination. Due to the safety limitations on the PMT illumination, the FD can operate only during clear nights with a low moon fraction and night sky background, which limits the FD duty cycle to around 19%. Due to the weather conditions, this number lowers to 15% [108].

The telescope's optical system consists of the aperture system, including the aperture itself, a filter and a corrector ring, a segmented mirror, and a fast-timing camera [115]. A circular diaphragm with a UV-transparent filter is installed at the bay entrance. The filter selects the part of the spectrum with a wavelength below 410 nm, thus reducing the background noise signal from visible photons while having high efficiency in the range of nitrogen fluorescence. The corrector ring with an aspherical lens minimizes the coma aberration and corrects the spherical aberration caused by the mirror without increasing the spot size. At the entrance of the telescope bay, there is a shutter, which is opened during the telescope operating hours on suitable nights. There are also curtains to prevent the telescope from being exposed to daylight in the event of an electronic or shutter failure. Normally, shutters close automatically with rain/snow or strong wind. After passing through

the aperture, the light is focused on the segmented mirror. The area of the mirror is around 13 m^2 . The mirror segments are either made of aluminum alloy sheets (LL and LM sites) and have a rectangular shape or glass with a reflective aluminum layer coated with SiO_2 for mechanical protection and have a hexagonal shape (CO and LA sites). The average reflectivity of the mirror segments is more than 90% in the part of the spectrum between 300 and 400 nm. The light collected by the mirror is then focused onto the photomultiplier camera. The camera comprises 440 hexagonal pixels of almost 45 mm in size arranged on a 22 by 20 spherical focal surface [116].

Each pixel contains a photomultiplier tube mounted in the center of the pixel and surrounded by light collectors. The field of view of a single pixel is 1.5° . The light collectors have a simplified classic Winston cone structure with a so-called "Mercedes" stars, ensuring a smooth transition between adjacent pixels and good light collection efficiency of around 93%. The photomultiplier unit is an 8-stage PMT with $\approx 25\%$ quantum efficiency for the wavelength range between 350 and 400 nm. The high voltage for the PMT is provided by an HV-divider chain in the head electronics (HE) unit. The signal driver circuitry, another unit element, maintains the gain stability by keeping the bias from the divider- introduced below 1%. The HE unit is connected to the PMTs via the distribution boards, which provide power to the PMT electronics and receive the output signal. The distribution boards are installed at the back of the camera without affecting the telescope's field of view, and each board is connected to 44 PMTs.

The signals received by the PMT is processed with a three-level trigger system. In the first stage, the signal is digitized at 10 Hz and stored for further reading. The rate of each pixel (hit rate) is measured and adjusted to 100 Hz by changing the pixel trigger threshold depending on the background. At the second level, a trigger is formed if the obtained after the first level geometrical pattern resembles the cosmic ray pattern. The track segment should follow a straight pattern and have at least 5 pixels in length. Fig 2.4 shows the patterns used to identify a cosmic-ray signal. The observed track is compared with the predicted geometries, including their rotations and mirror reflections. Since not all tracks pass through the center of each pixel and provide sufficient illumination of the pixel, and there is a chance of a defective PMT, only four out of five triggered pixels are required. The events passing the second level are stored in the pixel memory and analyzed in software at the third trigger level by PCs connected to the mirrors (MirrorPC). If the event comes from the calibration or test systems or is produced by an artificial light source, it is sent to the PCs connected to the eyes (EyePC) and stored in the corresponding raw data files. On the contrary, before being sent to the Eye-




Figure 2.4: The patterns for the cosmic-ray tracks identification in FD. Figure from [107].

PCs, the actual air shower events are passed through the third trigger level at the MirrorPCs. At the third level, the data is cleaned of the noise events that survived the hardware triggers and the randomly triggered pixels. The software also rejects triggers from lightning or muons hitting the camera. The efficiency of the algorithm is relatively high, with about 94% of the background events correctly identified and removed and less than 1% of real air shower events erroneously removed. The events surviving the software trigger are sent to the eyePCs. The coincident events from the telescopes at the same site are merged. After that, a hybrid trigger for the SD is generated and sent by the eyePCs.

A calibration of the FD is required to convert the recorded ADC counts into the light flux at the aperture. For this, a pixel response to an influx of incident photons from the solid angle covered by that pixel needs to be measured [117]. Such an end-to-end calibration includes various factors of the telescope optics, such as aperture projection, UV filter transmittance, light collection efficiency, quantum efficiency, PMT gain, and the reflectivities of optical surfaces and telescope mirrors etc., should be considered. Three different calibration methods are used at the Pierre Auger Observatory: absolute calibration [118], relative calibration [119], and wavelength (spectral) calibration [120].

The absolute end-to-end calibration of the FD is performed with a calibrated portable drum-shaped light source 2.5 m in diameter. It consists of two UV LEDs that emit pulsed light of known intensity at a wavelength of 375 ± 12 nm. The surface of the drum is a good Lambertian source, evenly illuminating all pixels in the camera over the entire angular range of the camera relative to the drum. The light source is mounted inside a Teflon cylinder on the front of the drum. The light output for each LED is monitored by the silicon detector at the top of the cylinder.

The drum itself is calibrated with a silicon photodiode. The charge from the reference PMTs, placed on the drum axis, is recorded for each pulse, providing a distribution of the observed integrated flux. Another PMT is installed at the optical

bench. It records the diffuse emission from the pulsed LEDs with a neutral density filter and an intensity adjusted to match the measured intensity of the drum. At each intensity setting, the second measurement is done with the PMT replaced by a NIST-calibrated photodiode and the neutral density filter removed. This gives a linear relationship between the PMT response and the photodiode current, which can then be used to calculate the photon flux, knowing the technical characteristics of the LEDs and photodiode.

An absolute calibration of the fluorescence detector is also checked with vertical laser shots at 337 and 355 nm. Assuming perfect Rayleigh scattering with a known cross section of the atmosphere, it is possible to calculate the flux of photons at the telescope aperture. It is, therefore, possible to estimate the response of each pixel to the number of photons detected from a segment of the laser beam. The track produced by the laser beam is similar to that of cosmic rays. However, this method can be less practical since only a few pixels are illuminated at a time, and the calibration of the entire camera would require frequent changes in laser position.

The relative calibration is performed every night before and after the data taking. It allows for monitoring the changes from night to night and the seasonal variations in the response of the PMTs [121]. The light from three xenon flash light sources (referred to as A, B, and C) located at each FD site is distributed via optical fibers to each mirror. Each light source in the system monitors the different detector components, so the properties and characteristics of the three LEDs differ depending on the purpose. The light from source A is transmitted to the light diffusers installed at the center of each mirror to illuminate all the pixels in the camera simultaneously. It monitors the stability and linearity of the camera system. The light from the second source is transmitted to the diffusers installed at the center of two sides of the camera. It is then reflected from the mirror to the camera for checking the camera PMTs and the stability of mirror reflectivity. The light from the third source is flashed at the diffusers installed outside the entrance aperture on the telescope doors. The light from the diffusers illuminates a reflective Tyvek screen located outside of the UV filter. This screen reflects the light towards a mirror, and the camera records the reflected light. At this calibration step, the overall stability of the telescope is checked, including mirror reflectivity and PMT response.

Multi-wavelength calibration employs a similar to end-to-end absolute calibration technique and is used to estimate the wavelength dependence of the spectral response of the fluorescence detector. Drum intensity measurements are made at several wavelengths in the nitrogen region (320, 337, 355, 380 and 405 nm) us-

ing the same reference PMTs as for drum calibration. The combination of the PMT response and the quantum efficiency, measured in the lab, gives a quantity proportional to the number of photons emitted. The relative uncertainty at each wavelength is approximately 5%. The final calibration curve is normalized to 375 nm.

2.2.1 Event reconstruction

The reconstruction of the events detected by the hybrid detector includes two stages: geometrical reconstruction and shower profile and energy reconstruction [107]. The first step of the geometrical event reconstruction is to determine the shower detector plane (SDP) consisting of the shower axis and the FD. The orientation of the shower axis is fixed by two parameters - the impact parameter R_p and the angle between the axis and the horizontal χ_0 , and its position is defined by the timing information of the camera pixels. If the angular velocity remains unchanged over the range of R_p and the range of observed elevation angles χ_0 , it is not easy to constrain these parameters from FD alone, since then there is no single pair of values that fits the angular velocity. Therefore, the information of the time of arrival of the shower front at ground level is used to improve the accuracy of the geometry reconstruction [122]. Such a so-called hybrid geometry reconstruction is achieved by using the measured time of a single SD station close to the shower axis. The resulting accuracy is 50 m in the core location and around 0.6° in the angular resolution. The uncertainty of the geometry reconstruction is evaluated with the CLF facility with a well-defined beam geometry.

With a known geometry and fluorescence efficiency, the light detected by FD can be converted into an energy deposit as a function of the slant depth. In the standard approach, the contributions from the different light sources are carefully estimated using an analytical approach [123], the contributions of fluorescence and Cherenkov light are treated equally as a signal, which is possible due to the observed universality of the energy spectra of electrons and positrons, with which the number of charged particles (responsible for Cherenkov light) can be related to the energy deposit (responsible for Fluorescence light). An example of the measured light at aperture and the reconstructed light contributions, and energy deposit profile is shown in Fig. 2.5. Here the detected light at the telescope aperture is shown as points with error bars and the reconstructed contributions from different sources of light: Fluorescence light, direct and scattered Cherenkov light. Light scattered at aerosols and air molecules is labeled as "Mie" and "Rayleigh", respectively. A small fraction of multiple-scattered fluorescence light is indicated

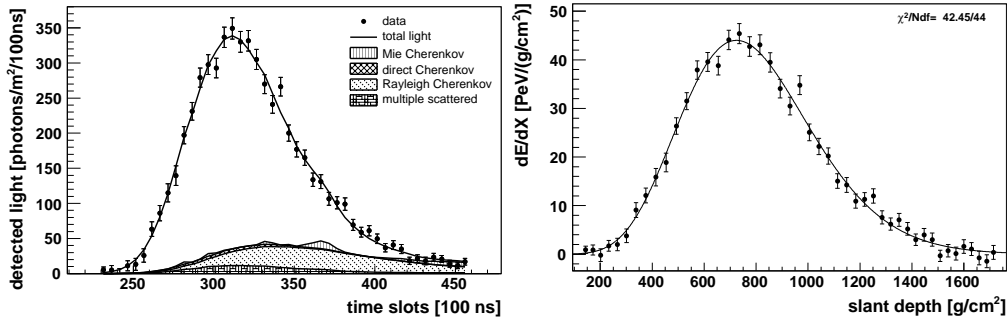


Figure 2.5: Example of a reconstructed shower profile (event 931431). Figure from [105].

as well. The reconstructed longitudinal profile is shown on the right panel of Fig. 2.5. The data points denote the reconstructed energy deposit in the atmosphere and the line is the fit with the Gaisser-Hillas function, Eq. 2.2. The energy deposited by an air shower, which is measured as a function of the traversed depth, is commonly referred to as the longitudinal air shower development profile. Fig. 2.6 shows an example of the longitudinal air profile for iron and proton primaries simulated at 10^{19} eV. The slant depth X is obtained by integrating the density of air along the direction of arrival of the air shower through the curved atmosphere,

$$X(l) = \int_l^\infty \rho(\mathbf{r}(l')) dl, \quad (2.1)$$

where $\rho(\mathbf{r}(l))$ is the density of air at a point with longitudinal coordinate l along the shower axis.

By integrating the longitudinal profile, one can derive a calorimetric measurement of the shower's energy. The atmospheric depth at which the longitudinal profile reaches its maximum, known as X_{\max} , is the parameter most sensitive to the mass of the primary particle and the properties of the hadronic interactions.

To obtain X_{\max} , the measured longitudinal air shower profile is reconstructed with a profile function, such as e.g. Gaisser-Hillas (GH) [124], Greisen [125], and 'Gaussian in Age' [22] functions. The Gaisser-Hillas profile, which is commonly used in the reconstruction of the Pierre Auger Observatory data, can be parameterized as [124]:




Figure 2.6: An example of the longitudinal air shower profile for iron (red dotted lines) and proton (black solid lines) primaries. Figure from [126].

$$N(X)_{\text{GH}} = N_{\text{max}} \left(\frac{X - X_0}{X_{\text{max}} - X_0} \right)^{\frac{(X_{\text{max}} - X_0)}{\lambda}} \exp \left(\frac{X_{\text{max}} - X}{\lambda} \right), \quad (2.2)$$

where N_{max} is the number of particles at shower maximum, X_0 is the depth of the start of the profile and λ is a parameter determining the shape of the profile.

The function can be rewritten in a Gaussian-like form by introducing new variables $N' = N/N_{\text{max}}$ and $X' = X - X_{\text{max}}$, and redefining the fit parameters as $R = \sqrt{\lambda/|X_0 - X_{\text{max}}|}$ and $L = \sqrt{|\lambda(X_0 - X_{\text{max}})|}$ [127], which would then correspond to the characteristic width and an asymmetry parameter respectively:

$$N' = \left(1 - \frac{RX'}{L} \right)^{R-2} \exp \left(\frac{X'}{LR} \right). \quad (2.3)$$

Based on the concept of the "universal air shower profile", the average longitudinal profile can be used instead of the individual profiles. This reduces the statistical uncertainties, e.g. in the light collection or X_{max} determination, affecting the estimation of the profile shape parameters, making it possible to fit L and

R parameters simultaneously [128]. Using the average profile also helps to reduce the systematic uncertainties, allowing an event-by-event analysis.

The elongation rate $D_e = \frac{d\langle X_{\max} \rangle}{d \ln E}$ [129, 130], defined as an evolution in the position of the mean depth of shower maximum, $\langle X_{\max} \rangle$, as a function of energy E , shows a change in the primary composition with an increase in energy. It can be written as [131]:

$$D_e = (1 - B)\lambda_r \left(1 - \frac{d\langle \ln A \rangle}{d \ln E} \right), \quad (2.4)$$

where B shows a dependence on the hadron-air interactions, λ_r is the radiation length for electromagnetic cascades in air, and A is a mass number. For convenience, the elongation rate is commonly defined as a change in average X_{\max} per decade of energy, $d\langle X_{\max} \rangle / d \lg E$. For a photon-initiated shower in the electromagnetic cascade the elongation rate is equal to $D = 85 \text{ g/cm}^2$.




Figure 2.7: The evolution of the first moments of the X_{\max} distribution, the mean $\langle X_{\max} \rangle$ (left) and standard deviation $\sigma(X_{\max})$ (right) as a function of energy from measurements with the fluorescence and surface detectors of the Pierre Auger Observatory and Telescope Array. The results of the Telescope Array were shifted in energy to match the energy scale of the Pierre Auger Observatory. Figure from [47].

The average X_{\max} , $\langle X_{\max} \rangle$, can also be expressed as the sum of the contributions from each mass group [132]:

$$\langle X_{\max} \rangle = \sum_{s=1}^N f_s \mu_s \quad (2.5)$$

2. The Pierre Auger Observatory

where f_s is the fraction of nuclear species s and μ_s is the corresponding mean of the X_{\max} distribution.

Similarly, for the standard deviation of the X_{\max} distribution:

$$\sigma(X_{\max}) = \sqrt{\sum_{s=1}^N f_s V_s + \sum_{s=1}^N f_s \mu_s^2 - \left(\sum_{s=1}^N f_s \mu_s \right)^2}, \quad (2.6)$$

where V_s is the variance of nuclear species s .

Fig. 2.7 shows an example of the measured from the cosmic-ray data mean, $\langle X_{\max} \rangle$, and standard deviation, $\sigma(X_{\max})$. As can be seen, the observed trend indicates an increasing presence of the heavier component at higher energies.

2.2.2 Resolution and Systematics

The performance of the FD profile reconstruction has been extensively studied since both the energy and X_{\max} scales of the observatory are based on reconstructed FD events. The uncertainty of the energy scale is 14% [133], and is dominated by the absolute optical calibration of the telescopes. The energy resolution is about 8%, almost independent of the energy [134].

Of particular importance for this thesis is the X_{\max} reconstruction. A detailed discussion on the detector resolution and systematics of X_{\max} can be found in [34]. The corresponding energy dependence is shown in Fig. 2.8. As can be seen, the resolution is dominated by the photo-electron statistics of the profile (labelled "detector" in Fig. 2.8 (left)), and at high energies, contributions from the uncertainty of measurement of aerosols in the atmosphere become non-negligible.

Reconstruction uncertainties at low energies and atmospheric uncertainties at high energies dominate the systematics of the X_{\max} scale. While the different contributions shown in Fig. 2.8 (right) add up quadratically to a near-constant value, their energy dependence could be an overall X_{\max} scale uncertainty and an energy-dependent X_{\max} shift. To consider this, the covariance matrix of the systematic X_{\max} uncertainty was determined in [34] and will be used in some of the studies discussed below.

2.3 Atmospheric monitoring

Atmospheric conditions play an essential role in the reconstruction of extensive air showers, as they affect both the development of the air shower in the atmosphere

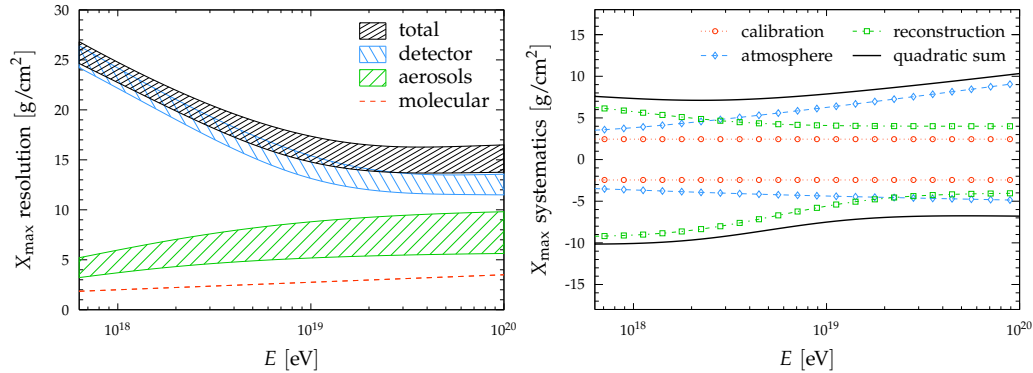


Figure 2.8: Resolution (left) and systematic uncertainty (right) of the X_{\max} measurement with Fluorescence Detector of the Pierre Auger Observatory. Figure from [34].

and the measurement of the the fluorescence yield, which depends on the local pressure, temperature and humidity [135]. For example, atmospheric conditions could affect SD measurements, such as the energy estimate of the SD - the signal measured 1000 m from the shower core due to their effect on the particle density at the ground, as well as the SD trigger probability could be altered by the varying atmospheric pressure [136]. Fluorescence telescopes are particularly sensitive to atmospheric conditions, as their measurements are affected by optical attenuation due to light absorption and scattering due to the molecular atmospheric density (Rayleigh scattering) and the aerosol density (Mie scattering) [137]. Therefore, a number of instruments for monitoring the atmospheric conditions were deployed in the field, allowing the accurate estimation of the variable properties of the atmosphere.

The installed instruments can be roughly divided into the following three groups:

- Monitoring the state of the atmosphere: weather stations, and the Global Data Assimilation System (GDAS);
- Estimation of the aerosols content: laser facilities - Central Laser Facility (CLF) and eXtreme Laser Facility (XLF), two Aerosol Phase Function Monitors (APFs) and the F(/ph)otometric Robotic Telescope for Atmospheric Monitoring (FRAM);
- Monitoring the presence of clouds: four cloud cameras, four LIDAR sta-

2. *The Pierre Auger Observatory*

tions, FRAM, CLF lidar, and XLF.

The information on the local ground-level weather conditions is provided every 5 minutes by weather stations installed at all four FD sites. The stations monitor the temperature, pressure, humidity, and wind speed. Since they could only do measurements at the ground level, the information on the vertical profiles of the temperature, pressure, and humidity is obtained from the global atmospheric model GDAS, which was developed by NOAA's National Centers for Environmental Prediction. The data from GDAS [138] are published once a week for every three hours, and it covers a vertical distance of up to about 26 km.

The aerosol content of the atmosphere is estimated based on the scattering of CLF and XLF laser beams shot horizontally during the FD operation. Measurements are taken every 15 minutes. They are compared with simulations and clear conditions to estimate the vertical aerosol optical depth. Additionally, the backscatterings in the atmosphere are measured by elastic LIDARS, installed at each FD site [139]. Shots are fired both vertically and horizontally toward the FD stations. The latter measures the light attenuation at the ground level. LIDARs are also used to identify clouds. As the beam passes through the cloud layer, it is scattered back and recorded by a signal receiver. The time measurements allow us to identify the height of the clouds. As clouds have higher temperatures and produce more infrared light, they can also be seen with infrared cameras. Another technique used by the Observatory for aerosol attenuation measurement and cloud tracking is based on stellar photometry. This is done with the F(otometric) R(obot)ic Telescope for Atmospheric Monitoring (FRAM), which has been operating at the Los Leones site since 2005 [140].

2.4 The Pierre Auger Observatory Upgrade: Auger-Prime

After 15 years of operation, an upgrade of the observatory was planned to improve the performance of the observatory and enhance the sensitivity of the detectors. The original upgrade plan included upgrading the SD electronics, installing the Surface Scintillator Detector (SSD) and Radio Detector (RD) on each WCD, deploying the Underground Muon Detector (UMD) and extending the FD duty cycle [108]. The upgrade is expected to clarify the origin of the suppression of the cosmic ray flux above $4 \cdot 10^{19}$ eV, improve the estimation of the primary composition, and access the properties of the hadronic interaction at the highest energies.

The SSDs are the main part of the upgrade and will provide a complementary and independent measurement of particle densities at ground to the WCDs. While muons dominate the WCD signal, the SSD is more sensitive to electromagnetic particles. Thus, the combination of these two detectors will allow better differentiation between the air shower's different components, muonic and electromagnetic, at ground level, providing a better estimation of the energy scale and mass composition [141]. The dynamic range of the SD was improved by adding a fourth small PMT at each WCD station. The old unified boards of the SD electronic were replaced with the new ones, characterized by faster sampling and signal processing, better timing accuracy, and an increased gain and dynamic range. With the new electronics, it will be possible to process signals from all the upgrade detectors simultaneously.

The Auger Engineering Radio Array (AERA) [142], consisting of 150 radio detector stations and designed to test the capabilities of detecting radio emission from cosmic rays, has been in operation for several years before the start of the upgrade. Based on the AERA performance and the results obtained, extending the Radio Detector over the entire SD array was proposed by placing radio antennas on top of the upgraded WCDs. The radio upgrade will better detect inclined showers and complement the mass sensitivity studies with the WCD+SSD combination. Another component of the AugerPrime project is completing the full SD-750 array with UMDs [140], which serves as a tool for directly measuring muons. The upgraded AMIGA detector will improve the accuracy of muon detection and provide a means of cross-checking the results obtained with the SD.

At the time of writing this thesis, the commissioning of the AugerPrime upgrade was nearing completion, with all the SSDs installed and already taking data.

2.5 Data sets and data selection

To measure the cosmic-ray mass composition, we use the hybrid events collected by the Pierre Auger Observatory, i.e., the events that were seen with FD and triggered at least one SD station. The data selection procedure is based on the extensive description provided in [34]. Only data recorded during stable operation and optimal atmospheric conditions are considered. This includes the requirements of the vertical aerosol optical below 0.1 at 3 km above the ground level, and the reconstruction of the longitudinal air shower profile is not affected by clouds. Another requirement is that the depth of the shower maximum must fall within the geometrical field of view of the telescope, and its reconstruction uncertainty

2. The Pierre Auger Observatory

should be below 40 g/cm^2 . Only the showers with the difference between SD trigger probabilities for proton and iron of less than 5% are selected to avoid mass composition bias. The contribution from Cherenkov light should be at most 20%. Lastly, applying the fiducial field of view cuts (FoV) ensures the high quality of events without X_{max} selection bias, regardless of the shower geometry and X_{max} values.

The slant depth range is defined by the lower and upper limits, X_{low} and X_{up} respectively, where X_{max} of each event is reconstructed with a resolution better than 40 g/cm^2 . The acceptance of the shower is contingent on this interval being sufficient to include most portion of the X_{max} distribution. Cuts (boundaries), $X_{\text{low}}^{\text{cut}}$ and $X_{\text{up}}^{\text{cut}}$, are further imposed on the limits of the slant depth range. The cuts for the field of view limits are estimated from the behavior of the mean X_{max} as a function of the values of the sections, since the true X_{max} is unknown at this stage of the reconstruction, and are set at values where the deviation of $\langle X_{\text{max}} \rangle$ from the asymptotic value begins to exceed 5 g/cm^2 .

The results shown in this thesis will exploit the use of three different datasets:

- the data used in the X_{max} analysis by the Pierre Auger Collaboration and published in [34, 82], covering the data-taking period between December 2004 and December 2012 with a total number of events surviving all the cuts of 19759 for the energies above $10^{17.8} \text{ eV}$. We will refer to this data set as to PRD14 in the following;
- the data used for the updated analysis presented at the International Cosmic Ray Conference (ICRC) in 2019 [143], covering the data-taking period during 2004-2017 with a total number of events surviving all the cuts of 31085 for the energies above $10^{17.8} \text{ eV}$. Thus, the statistics in this data set is about 1.5 higher the one available in the PRD14. This data set will be referred to as ICRC19;
- the most recent data, covering the entire Phase I of the Pierre Auger Observatory operations between December 2004 and December 2021, presented at the ICRC conference in 2023 [144]. Phase I has a total of 74760 events for the energies above $10^{17.8} \text{ eV}$. This is the main data set used in this thesis and it is referred to as ICRC23.

In the listed datasets, energies below $10^{17.8} \text{ eV}$, measured by the Pierre Auger Observatory with HEAT/Coihueco (HeCo) telescopes, are excluded. For the discussion on the mass composition analysis with the low energy extension of

the Pierre Auger Observatory, one can refer to the measurements presented at ICRC 2017 [145], which cover the energy range $10^{17.2} - 10^{17.8}$ eV for the period between June 2010 and December 2015.

In addition to the increased statistics due to more years of observations, other differences between the datasets can affect the interpretations. One such factor is the change in the event reconstruction. With an increase in the precision of the fluorescence yield measurements over the years, the reconstruction of the longitudinal air shower profile is being updated, affecting the X_{\max} and energy scale determination (see, e.g [133] for the reconstruction used in ICRC19 data and [146] for the reconstruction used in ICRC23 data). In addition, measurements of the vertical aerosol optical depth profile at the Pierre Auger Observatory have recently been reviewed and updated [147].

2. *The Pierre Auger Observatory*

Chapter 3

Measurement of Mass Composition

One of the most robust techniques for estimation of the mass composition is based on the depth of atmospheric shower maximum, X_{\max} , as measured by the fluorescence detector of the Pierre Auger Observatory. By comparing the X_{\max} distribution to the predictions from air shower simulations with varying fractions of nuclear mass groups, the best-fit composition can be estimated with, for example, a binned maximum-likelihood estimator (MLE), which is a standard approach used so far in Auger, see e.g. [82, 148].

In this chapter, we discuss a slightly different approach for fitting the composition fractions, namely instead of using a minimization algorithm we apply the Markov Chain Monte Carlo (MCMC) method to investigate the composition fractions with Bayesian techniques. The MCMC approach has several advantages for mass composition measurements. Firstly, MCMC can be applied to global optimization problems, and it will not get stuck in a local minimum (at least theoretically if the number of samples is large enough and/or the sampling steps are set appropriately). Furthermore, MCMC can deal with a large number of highly correlated parameters which is numerically impossible with standard gradient minimizers such as MINUIT [149]. This can be very useful for composition studies if e.g. one aims at marginalizing over all 26 charges from a proton to iron instead of using mass groups or if in addition to the nuclear fractions one also wishes to fit properties of hadronic interactions. Most importantly, MCMC allows to sample of the posterior probability density function of the estimated fractions. Thus, it is easy to marginalize over the mass composition for derived quantities. Examples of such derived quantities are the average rigidity (energy per charge) of cosmic rays and the first two moments of the X_{\max} distribution, as will be discussed later. A future application could be the fit of the proton-proton cross section, with un-

certainties marginalized over the mass composition and other nuisance parameters such as the systematic uncertainties of the energy- and X_{\max} -scale or assumptions made in the calculation of nucleus-air cross sections.

The chapter covers the following content. Firstly, the MCMC sampling algorithm and the mass composition fitting procedure are introduced. Next, we assess the accuracy of the fitting method and investigate any potential biases in the results. Then, we evaluate the impact of X_{\max} distribution binning and the template generation method on the results. We also compare the fractions predicted from the different datasets (see Sec. 2.5 for the description of the considered data). Finally, we present the mass composition derived from the Phase I data of the Pierre Auger Observatory. Additionally, we explore potential applications of the MCMC fitting algorithm and the results obtained.

3.1 MCMC Sampling of the X_{\max} distribution

The name of the Markov chain Monte Carlo combines two properties: Monte Carlo and Markov chain. The basis of the Monte Carlo method lies in sampling from the posterior probability density function to infer the quantity of interest, i.e., the probability distribution of the fitted model given the data. In Monte Carlo methods, the sampling process does not start at the mean but involves generating a significant number of samples to average. Indeed, sampling directly from the posterior distribution can be challenging, especially in high-dimensional problems. The use of Markov chains simplifies the sampling process in such cases. In Markov chains, each subsequent step is influenced solely by the preceding step and is not dependent on any values prior to it. This property, known as the Markov property, imparts a memoryless characteristic to the process, meaning that future sampling would depend only on the current value and not the preceding steps. Combining Monte Carlo sampling and the Markov chain algorithm is a powerful tool for approximating complex probability distributions.

The two commonly used algorithms in MCMC are the Metropolis-Hastings (MH) algorithm and the Gibbs sampling algorithm, the latter being a special case of the MH algorithm. In the MH method, the step (state) is accepted or rejected based on the so-called acceptance ratio, which is the ratio of the probability of the proposed sample under the target distribution to the probability of the current state. If the probability of the proposed sample is equal to the acceptance probability, the sample is accepted and becomes a new state. Otherwise, the state remains unchanged. The Gibbs algorithm, on the other hand, updates only one variable

at a time in the new state, considering the current values of the other variables. In this study, we adopt the Metropolis-Hastings algorithm. Here, we only briefly introduced the Bayesian fitting with MCMC algorithm discussed and outlined its main components. A more extensive description of the MCMC technique can be found in [150]

There exist several libraries in Python for performing MCMC sampling, the most commonly used of which are the EMCEE [151] and PYMC [152] modules. Due to its widespread application in astrophysics (more than 8165 citations on NASA/ADS as of the writing of this work), we decided to use EMCEE library, but verified that both modules give identical results.

3.1.1 Fitting algorithm

The required input to any MCMC sampler consists of two parts, and calculates and returns the sum of the logarithm of the prior, $p(\theta)$, of the fit parameters θ , and the logarithm of the likelihood function $p(X|\theta)$, where X is the data. The main output of the MCMC algorithm are samples of the posterior distribution:

$$p(\theta|X) \propto p(X|\theta)p(\theta). \quad (3.1)$$

The prior $p(\theta)$ is set to have a form of the uniform distribution within the allowed parameter range, which corresponds to the so-called uniform uninformative prior. In the case of infinite MCMC statistics, the best sample with such a prior is equivalent to the maximum likelihood estimator.

To measure the mass composition we use the depth of maximum of extensive air showers, X_{\max} , which is one of the observables most sensitive to the ultrahigh-energy cosmic-ray composition. As there are fluctuations in the properties of first few hadronic interactions in the in the cascade, it is not possible to accurately estimate the primary on the event-by-event basis, so instead the distribution of the measured X_{\max} values is used. Assuming that nuclei of mass A produces the distribution $F_i(X_{\max})$ the overall distribution is then a sum of individual [34]:

$$F(X_{\max}) = \sum_i f_s F_i(X_{\max}), \quad (3.2)$$

where f_s is a fraction of primary particle of type s .

The relation between the true and observed X_{\max} distribution is

$$F_{\text{obs}}(X_{\max}) = \int_0^{\infty} F(X_{\max}) \mathcal{E}(X_{\max}) R(X_{\max}^{\text{rec}} - X_{\max}) dX_{\max}, \quad (3.3)$$

3. Measurement of Mass Composition

where ε is a detector efficiency and R is the detector resolution that relates the true X_{\max} to the reconstructed, X_{\max}^{rec} . In the case of the ideal detector, ε would be a constant, and a δ -function can approximate R .

To describe the measured number of entries n_i in the X_{\max} histogram at depth bin i , we use the same Poisson likelihood as in the previous Auger mass composition analyses [82, 153]:

$$-\ln L \equiv -\ln p(X|\theta) = \sum_i (C_i - n_i + n_i \ln \frac{n_i}{C_i}), \quad (3.4)$$

where C_i is the model prediction defined via

$$C_i = \frac{N_{\text{data}}}{\sum_s \sum_i f_s X_{s,i}} \sum f_s X_{s,i}. \quad (3.5)$$

Here f_s are the fractions of nuclear species included in the fit, N_{data} is the total number of events in the data and $X_{s,i}$ is the binned X_{\max} distribution template. The template is a simulated distribution of X_{\max} values including the effects of detector acceptance and resolution according to Eq. 3.3. It can be parameterized as [82]

$$X_{s,j}^m = \sum_n^{N_{\text{MC}}} a(X_{s,n}^t) p_j(X_{s,n}^t) / N_{\text{MC}}, \quad (3.6)$$

where $a(X_{s,n}^t)$ is an acceptance weight for the n th event of s particle species and $p_j(X_{s,n}^t)$ is the probability for each X_{\max} bin j .

By definition, the composition fractions f_s can have values between 0 and 1 and should sum up to 1. To make sure that these conditions are fulfilled, we do not fit the fraction f_s , but the auxiliary variables ζ_i that encode these constraints by the following relation to the fractions (see e.g. [154]):

$$\begin{aligned} f_1 &= \zeta_1 \\ f_2 &= (1 - \zeta_1)\zeta_2 \\ &\dots \\ f_{i-1} &= (1 - \zeta_1)(1 - \zeta_2)\dots(1 - \zeta_{i-2})\zeta_{i-1} \\ f_i &= (1 - \zeta_1)(1 - \zeta_2)\dots(1 - \zeta_{i-2})(1 - \zeta_{i-1}) \end{aligned} \quad (3.7)$$

and set the prior on the values of ζ_i to be uniform between 0 and 1. The geometrical interpretation of this transformation is to identify a composition mix with a point \mathbf{u} on an n -dimensional unit sphere. The fractions can then be identified with

the squared length of the Cartesian coordinates u_i and ζ_i as the squared directional cosines

$$\begin{aligned} f_1 &= u_1^2 = \cos^2 \varphi_1 \equiv \zeta_1 \\ f_2 &= u_2^2 = \sin^2 \varphi_1 \cos^2 \varphi_2 \equiv (1 - \zeta_1)\zeta_2 \\ &\dots \end{aligned} \tag{3.8}$$

for which $\mathbf{u}^2 = \sum u_i^2 = 1$ and $u_i \in [0, 1]$.

The X_{\max} distribution templates $X_{s,i}$ are generated with the CONEX air shower program [155] and then modified by the X_{\max} acceptance and resolution [82]. To account for the detector acceptance each bin of the template ϕ_i is multiplied by the detector acceptance $\varepsilon_{\text{rel}}(X_{\max_i})$:

$$\phi'_i = \phi_i \cdot \varepsilon_{\text{rel}}(X_{\max_i}), \tag{3.9}$$

calculated for the bin centers with the following standard parameterization [34]:

$$\varepsilon_{\text{rel}}(X_{\max}) = \begin{cases} e^{-\frac{X_{\max}-x_1}{\lambda_1}} & ; \quad X_{\max} \leq x_1, \\ 1 & ; \quad x_1 < X_{\max} \leq x_2, \\ e^{-\frac{X_{\max}-x_2}{\lambda_2}} & ; \quad X_{\max} > x_2. \end{cases} \tag{3.10}$$

ε_{rel} gives the probability to measure a shower as a function of X_{\max} . Due to the fiducial event selection, it has a flat central part between x_1 and x_2 . Outside of this range, the efficiency decreases exponentially, and the measured X_{\max} distribution is, therefore, expected to be biased in the tail of the distribution.

The detector resolution can be parameterized as the sum of two Gaussian distributions $G(\sigma)$:

$$R(X_{\max}^{\text{rec}} - X_{\max}) = fG(\sigma_1) + (1 - f)G(\sigma_2), \tag{3.11}$$

where f , σ_1 , and σ_2 are parameters of the X_{\max} resolution, and the bin center corresponds to the expectation of the normal distribution. Thus, if we consider the X_{\max} distribution in the range between 0 and 2000 g/cm² with a bin width of 2 g/cm², the distribution is approximated with 1000 double Gaussians. For each bin, the resolution function is evaluated for the X_{\max} values in the range 0-2000 g/cm² with steps of 1 g/cm².

3. Measurement of Mass Composition




Figure 3.1: The detector acceptance as a function of logarithmic energy for the following species: photons (orange pentagons), H (red diamonds), He (cyan triangles), N (grey squares) and Fe (blue circles). The lighter the species, the more it is affected by the detector acceptance.

To have the correct number of counts in the bins each of the double Gaussians R_i is divided by the sum over the distribution range $\sum_i R_i$ to normalize it and then multiplied by the corresponding probability density function values of the corrected for acceptance effects X_{\max} template ϕ'_i :

$$R'_i = \frac{R_i}{\sum_i R_i} \cdot \phi'_i. \quad (3.12)$$

Then, the renormalized double Gaussian distributions R'_i are summed up. We rebin the obtained distribution accordingly the selected binning in the data in order to have the same number of entries per X_{\max} interval in the templates. For example, if the selected binning is 20 g/cm^2 , bins are grouped by 10, as one bin with a width of 20 grams can be composed of 10 bins with a width of 2 grams, and the bin counts are summed within each group. The corrected for acceptance

and resolution binned template $X_{s,i}$ for a certain energy bin and a single species s is weighted according to its species fraction f_s to form model prediction C_i for the comparison with data in each X_{\max} bin.

The resolution and acceptance slightly alter the shape of the X_{\max} templates. Moreover, the acceptance changes the normalization of the individual components. But since the fiducial field of view cuts assure a large range of constant acceptance, the corrections are small for nuclear primaries. As can be seen in Fig. 3.1, the total selection efficiency, i.e. the integral over $\epsilon_{\text{rel}}(X_{\max})$ weighted by the normalized X_{\max} distribution is close to 100% for elements $A \geq 2$, and $\geq 90\%$ for protons. Only photons are severely affected by the acceptance with a selection efficiency dropping to as low as 50% at ultrahigh energies for the standard X_{\max} selection, which was of course not optimized for photon searches.

The fit runs over the X_{\max} range between 0 and 2000 g/cm² and the default binning of the distribution, unless noted otherwise, is 10 g/cm².

The assumed uninformative prior is quite common in the Bayes statistics, however here we set a uniform flat prior for the parameters ζ_i , via which the fractions are defined, and this does not necessarily result in the same uniform prior for the fractions.¹

To make sure that the appearance of the prior will not affect the fit we consider two other options of defining the prior and/or the model prediction in the likelihood without introducing the auxiliary variables ζ_i in the fitting procedure, which allows us to have non-informative prior for the fractions as well. As a first alternative implementation, we studied the fit of $N - 1$ fractions with the uniform prior on $[0,1]$ and defined the N th fraction as $f_N = 1 - \sum_{i=1}^{N-1} f_i$ while requiring the f_N to be within the $[0, 1]$ interval. Another option we investigated was to fit all fractions but modify the model prediction in such a way that the X_{\max} templates for species are individually normalized and the double sum in the Eq. (3.5) can be omitted, so the likelihood itself will guarantee that the resulting fractions will add up to 1 without imposing any additional constraints. We checked all three discussed options and found that the estimated fractions for the best-fit point estimate discussed in the next section are very similar and it can be concluded that the particular choice of the prior is not relevant for the result.

¹Uniform priors on ζ_i in Eq. (3.7) correspond to a uniform distribution of f_1 , but to skewed (but identical) distribution peaked at zero for all the other fractions. We also investigated the fraction prior using the method of normalized Gaussians [156] to pick the directional cosines. In this case, the pdfs of all fractions are equal, but again skewed and peaking at zero. The fact that this method results in equal prior for all fractions makes it conceptually more attractive and could be investigated further in the future.

3. Measurement of Mass Composition

The testing of the fitting algorithm will be performed with the latter option of defining the likelihood, while for the fit with the actual data we use the parameterization with the auxiliary variables ζ_i . The reason for this is that the latter option with the normalized templates does not allow us to account for the acceptance effects in the templates, which we do not necessarily need for the tests with the simulated data, but which we can not ignore for the fit to the observable.

As mentioned above, the main output of the MCMC fit is the posterior distribution, from which we can derive the point estimates of the fractions and the associated statistical uncertainties. Examples of the projections of MCMC samples as the corner plot and as the posterior distribution are shown in Fig. 3.2 for the four-component fit to the PRD14 data, at energies of $10^{17.8} - 10^{17.9}$ eV. It is also possible to access the correlations between the different nuclei from the corner plot. In the example given, there is clearly a strong anti-correlation between H and He, as well as between He and N.

There are several ways to report numerical values summarizing the posterior distribution, and one can use, for example, mean, mode, or median, all of which are calculated directly from the distribution or the maximum a posteriori (MAP) estimator, which, as its name suggests, is based on the maximum of the posterior distribution. In case of the uninformative prior used here, MAP is equivalent to the maximum likelihood value. Depending on the shape of the posterior distribution, i.e. whether we have a symmetric or a skewed distribution, this or that point estimate of the parameter of the interest may be more adequate to summarize the distribution, but the most common and recommended is the use of the median (50th percentile) of the distribution when reporting the point value of the fit parameters. Regarding the calculation of the statistical uncertainties, one of the most common ways are to quote the uncertainties based on the 16th and 84th percentiles of the distribution or to use the 68% highest posterior density interval, which constitutes the 1-sigma confidence range. The latter one is more applicable in a case of a skewed posterior distribution, as it allows to obtain the one-sided confidence limits. For the cross-checks, the mass composition fit is also performed with the minimization algorithm MINUIT, in particular with the Python-based *iminuit* [157] interface for the MINUIT2 C++ library, for which the parameter estimates are free from statistical fluctuations associated with the finiteness of the MCMC chain. The walkers, i.e. the various chains that the algorithm uses to explore the parameter space, for the auxiliary ζ parameters in the MCMC fit, which are the members of the ensemble moving around a given initial position, are initialized with a uniform random distribution on [0,1]. The two main indicators of the performance of the MCMC sampler are the convergence of the chain and the independence

3. Measurement of Mass Composition

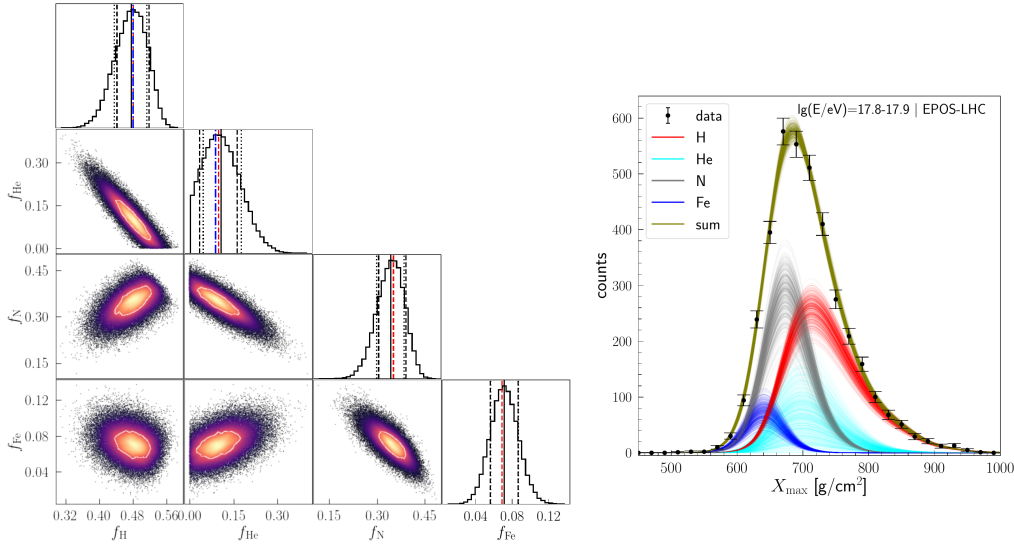


Figure 3.2: Example of the MCMC posterior samples for the 4-component fit to the PRD14 data at energies of $10^{17.8}$ - $10^{17.9}$ eV. Left: The corner plot for the MCMC draws from the posterior distribution showing the 2-D correlations between the parameters and the distributions of fitted parameters (diagonal). The contours in the density plots denote the 1-sigma level containing 39.3% of the volume, which corresponds to 68% for the 1-D distribution. The following point estimates are shown for the distributions: median (solid black line), mode (dash-dotted blue line), and MAP (dashed red line). The Equal-tailed and the Highest Density 1-sigma credible intervals are indicated with the dotted and dashed black lines respectively. Right: The posterior predictive for four species: H (in red), He (in cyan), N (in gray), Fe (in blue). The combination of four species is shown in olive and compared with the data.

3. Measurement of Mass Composition

of the samples. For this, we check the acceptance fraction, autocorrelation, and Gelman-Rubin statistic \hat{R} [158] to see if the chain has converged. The acceptance fraction α_f is the ratio between the accepted number of steps and the total number of steps. In general, there are no particular limits for the acceptance fraction. Still, it is advisable to have it between 0.2 and 0.5 since at very small α_f , there will be very few independent samples. In contrast, almost all samples will be accepted at $\alpha_f \sim 1$, preventing the converging to the target density. The autocorrelation time allows us to evaluate the number of steps needed to produce enough independent samples. The smaller the number, the better, meaning that fewer iterations of the probability distribution function are needed. In addition to the thinning of the chain, the autocorrelation time could also be reduced by adjusting the parameters of the stretch-move algorithm implemented in the *emcee* module. The \hat{R} , evaluated from the square root of the ratio between the within-chain and between-chain variance, provides another means to evaluate the chain's convergence; as a rule of thumb, it should be less than 1.05. A larger \hat{R} value would indicate a suboptimal mixing of the chain. Thus, based on the performance of the MCMC sampling, the number of steps in each MCMC chain was set to 45000 for the 4-component fit. To stabilize the chain, we discharge the first 5000 steps as burn-in and to reduce the auto-correlation, the MCMC chain is thinned by storing only every 15th step in the output. The length of the chain and the burn-in period should be adjusted based on the number of fitted parameters. More parameters, especially if they are strongly correlated, would necessitate a more extended sampling.

Since the fractions are defined via ζ variables we have only three fit parameters in the MCMC when a fit for 4 composition fractions is performed. We then convert the samples from the posterior distributions obtained for ζ variables into the posterior distributions for fractions using the Eq. (3.7), so that we can get the point estimates of fractions and the corresponding statistical uncertainties. Such conversion and an error propagation from ζ to fractions is actually an example of an application of posterior distributions.

3.1.2 Test of fraction fitting with MCMC algorithm

To test the MCMC fit performance we applied the fit to ensembles of the data simulated with X_{\max} distributions drawn from the Gumbel function parameterization [159, 160] for different mixtures of four particle species at an energy of $10^{17.8}$ eV. The number of events in the templates is $2 \cdot 10^6$ per species, and the number of events in the simulated data sets is varied from 100 to $2 \cdot 10^5$. For each composition, we repeat the generation of the data set and apply the fitting pro-

cedure N times. In total, for each composition scenario, we generated $N = 100$ data sets and then calculated the average of the fractions over the N realizations. In Fig. 3.3 the results for the following benchmark composition scenarios are shown:

- a) equal composition: H:He:N:Fe: 0.25:0.25:0.25:0.25;
- b) highly unequal composition - low He and N: H:He:N:Fe: 0.6:0.05:0.1:0.25;
- c) highly unequal composition - low He and Fe: H:He:N:Fe: 0.45:0.05:0.45:0.05.

On the x -axis the total number of events N_{tot} in the data is shown and the y -axis is the average of the fraction values over 100 fit realizations. Four different point estimates are shown: the median, mode, and the MAP from the MCMC fit and also the MLE found by MINUIT. As can be seen, all point estimates are biased for the unequal composition scenarios if the number of events is low, and the more unequal the composition is, the more significant the bias. In the limit of large statistics, the fraction estimators converge to the generated fractions. Asymptotically, the fitted fractions are unbiased within a couple of percent.


The median has the most significant bias, while the MAP is the least biased point estimator. As expected, the MAP is similar to the MLE obtained by the gradient search with MINUIT. Two examples of posterior distributions with a significant bias in the point estimates of the fractions are shown in Fig. 3.4. The posterior distribution displayed in Fig. 3.4 (left) looks like a normal symmetric distribution, so all three discussed point estimates are almost identical. However, their values are significantly less than the actual simulated fractions. On the other hand, the posterior distribution shown in Fig. 3.4 (right) is strongly skewed with point estimates that are, in this case, all biased except for the median, which is in a reasonable agreement with the simulated fraction

The overall performance of the MCMC fitting algorithm is good and compatible with gradient minimizers. Some bias is possible only in case of the very low statistics in data and the extreme composition scenarios with a nearly pure composition. The latter one could be due to the physical non-negativity constraint on the estimated composition fractions [161].


3.1.3 Confidence intervals and coverage probability

As an additional performance check, we also looked at the coverage probability of the 68%, 90%, and 95% credible intervals. The coverage probability is the


3. Measurement of Mass Composition



(a) $H:He:N:Fe: 0.25:0.25:0.25:0.25$.



(b) $H:He:N:Fe: 0.6:0.05:0.1:0.25$.



(c) $H:He:N:Fe: 0.45:0.05:0.45:0.05$.

Figure 3.3: Average fitted fractions for three composition scenarios as a function of the number of detected events, N_{tot} . The following point estimates of fractions are shown: the median (gray triangles), mode (purple circles), MAP (blue squares), and the MLE fractions from MINUIT (red diamonds). The true fraction is shown as the green dashed line.

probability that a certain confidence level of a statistical estimator contains the true value. Ideally, the coverage probability should be equal to the chosen confidence level. The credible interval is a Bayesian equivalent to the frequentist confidence interval. In the case of symmetric posterior density and flat prior, the credible interval will be numerically close to the confidence interval. Unlike the confidence

3. Measurement of Mass Composition




Figure 3.4: Examples of posterior distributions resulting in biased point estimates for the tests with the simulated data. The following point estimates are shown: the median (solid black line), mode (dash-dotted blue line), and MAP (dashed red line). The true fractions are shown as green dotted lines. The 68 % ETI and HDI intervals are shown in black dotted and dashed lines, respectively.

Right: H posterior, H:He:N:Fe = 0.45:0.05:0.45:0.05.

Left: He posterior, H:He:N:Fe = 0.6:0.05:0.1:0.25.

interval, which reflects how often the actual value lies within the interval limits, the credible interval tells us what is the posterior probability that the true value can be found within this range, which corresponds to the significance of 1-, 2-, and 3-sigma. Regarding the calculation of statistical uncertainties, there are two options for defining the uncertainty: one can quote the uncertainties based on the so-called equal-tailed interval (ETI) or the highest density interval (HDI). In the case of the ETI, the probability of having a value below the interval is the same as above it and necessary includes the median so that the 1-sigma credible interval (68%) will have 16% of the distribution on both sides of the limits. The second option, the Highest Density Interval, is defined so that all points within the interval have a higher probability density than points outside. By construction, the HDI always includes the mode of the distribution. The ETI and HDI methods for computing credible result in identical values if the distribution is symmetric. If the distribution is skewed, the intervals will be different. Since the HDI is generally recommended when the posterior distribution is skewed, which happens occasionally with the obtained posterior distributions for fractions, we use the HDI to report the 1-sigma uncertainty on the estimated parameters. However, a point estimate of fractions may lie outside the boundaries of the credible interval. This

3. Measurement of Mass Composition

is especially true if the MAP is used to report the value since it is not related to a property of the distribution, unlike the other estimates, but the MCMC is a draw with the largest posterior. The corner plot with the example of the MAP lying outside the 68% HDI credible interval and both mode and MAP lying outside the central (ETI) credible interval is shown in Fig. 3.5.

As can be seen, in this example the posterior distributions are highly skewed towards lower values (to the left) for two species, He and Fe, so it looks like we have only half of the usual distribution due to the physical limit at 0. Thus, the 16th percentile of the distribution (lower border of the central 68% credible interval) falls around the maximum of the posterior distribution, where, by definition, we also expect to have the MAP and mode, so for such a distribution it is better to report the HDI as a 1-sigma uncertainty. Nevertheless, though using the interval constructed with the HDI method undoubtedly has advantages for a skewed posterior distribution, we still can have the MAP point estimate outside the interval for a more or less symmetric distribution. A possible alternative to the issue with the point estimate not being inside the credible interval can be switching to the one-sided limit if the two-sided credible interval does not contain a point estimate, e.g., if the fraction value lies outside the limit on the lower side of the distribution we only report the upper limit, and vice versa - only the lower limit is considered if the point estimate has a higher value than the upper edge of the two-sided credible interval. However, it may result in unrealistically large uncertainties and, correspondingly, an over-coverage of the confidence intervals.

The coverage probabilities for three standard credible intervals for the simulated data with two different compositions are shown in Fig. 3.6. Coverage probability is a frequentist concept and quantifies the fraction of realizations that the interval contains the true value of interest,

$$p_{\text{coverage}} = \frac{N(\text{true value within interval})}{\text{number of realizations}} \quad (3.13)$$

Ideally, the coverage probability equals the confidence level, i.e., for an interval at 68% C.L., the coverage probability should also be 68%. While the confidence level quantifies the *desired* frequency of "type-I errors", the coverage probability measures the *actual* occurrence of these errors.

A closer inspection of the resulting coverage probabilities leads to the qualitative statement that the coverage probability behavior varies depending on the composition and available data statistics. While for some species it tends to increase or decrease as the number of events in the simulated data changes, for other species it seems that the coverage probability fluctuates around a constant value.






Figure 3.5: An example of the corner plot showing the 2-D correlations between the parameters and the distributions of fitted parameters (diagonal) with some point estimates outside the 1σ credible interval. The fit 4-component fit was performed with the PRD14 data at the energies of $10^{18.6} - 10^{18.7}$ eV. The contours in the density plots denote the 1-sigma level containing 39.3%, i.e. $1 - \exp(-0.5)$, of the volume, which corresponds to 68% for the 1-D distribution. The following point estimates are shown: the median (solid black line), mode (dash-dotted blue line), and MAP (dashed red line). The dotted and dashed black lines indicate the 1σ ETI and HDI uncertainty ranges.

We can not expect good coverage at low N_{tot} where point estimates are biased. In particular, there is over-coverage for N and Fe in the low-He/low-N scenario shown in Fig. 3.6b and a significant under-coverage for He and N in the low-He/low-Fe scenario displayed in Fig. 3.6c. Moreover, in both scenarios displayed, we have an over-coverage at large N_{tot} at the 90% and 95% levels.


3. Measurement of Mass Composition



(a) $H:He:N:Fe: 0.25:0.25:0.25:0.25$.



(b) $H:He:N:Fe: 0.6:0.05:0.1:0.25$.



(c) $H:He:N:Fe: 0.45:0.05:0.45:0.05$

Figure 3.6: Coverage probabilities for three different composition scenarios. Left: 68% credible interval. Center: 90% credible interval. Right: 95% credible interval.

It can be noted that the data for the benchmark compositions were generated for the exact number of events without any fluctuations. To see whether this can affect the coverage probability in some way, we simulated the data following the multinomial distribution for one of the composition scenarios, thus adding fluctuations to it. We then compared the calculated coverage probability with the one for the fit to the data simulated with the exact number of events. For this check, we increased the number of realizations up to 1000 for more accuracy. The obtained coverage probabilities differ only slightly, with the values for the data simulated with the multinomial distribution being slightly smaller than for the data with the fixed number of events. The coverage probability almost did not change with the increase in the number of realizations from 100 to 1000 as well.

Summarizing the above, the bias in the fit is on average small and within the statistical uncertainty for typical event statistics and composition fractions as observed in the data, but in general depends on the composition and the number of events in the data set. Nearly unbiased results are observed when the MAP is used to report the outcome of the MCMC fitting procedure or when the fit is done with the MLE, which is equivalent to the prior used here. This bias also affects the considered coverage probability of the credible intervals making a specific contribution to the observed over-coverage/under-coverage for different composition scenarios with a varying size of the data set.

3.2 Application to data

In this section, we apply the fitting procedure to the most recent Auger data and discuss the results obtained. The parameters of the MCMC fit and the way of constructing the model predictions, including the log-likelihood, remain the same as in Sec. 3.2.1.

The mass composition fits are performed using a default mix of H, He, N, and Fe, but also with different combinations of particle species. Several hadronic interaction models are considered in the analysis, i.e., EPOS-LHC, QGSJETII-04, Sibyll2.3c, and the latest Sibyll 2.3d. In light of the release of the newer Sibyll version during the course of this study, the analysis of the most recent data does not use the previous version of the Sibyll MC generator. For a discussion of the differences between the predictions of air shower observables from the different models, see the section on hadronic interaction models. In addition, given that QGSJETII-04 does not describe the Auger data well, it was decided not to use it for the composition estimates from the full Phase I data.

3. Measurement of Mass Composition




Figure 3.7: Comparison of fractions and the quality of the fit for the different binning of the X_{\max} distribution as a function of logarithmic energy. The data set is ICRC19 and the fit is done with the Sibyll 2.3d hadronic interaction model. The bottom panel shows the quality of the fit $P_{\chi^2}(D, \nu)$ estimated from the χ^2 statistics.

Since we found the MAP point estimate to be the least biased among all discussed options, we decided to use it to report the 'data point' of the fitted fractions. Two distinct approaches can be employed to evaluate the quality of the fit. According to [82], it can be defined as the probability of obtaining a worse fit, e.g., a larger negative likelihood, than that obtained with the data, assuming that the distribution predicted by the fit results is correct. However, estimating the p-values for the fit in this way, along with the corresponding systematic uncertainties, can be a time-consuming procedure, particularly when the p-values are small, as it would require generating a substantial amount of simulated data. Alternatively, we can use the fact that the deviance D defined by the logarithm of the Poissonian likelihood is approximately χ^2 -distributed [153]. Then, considering the number

of degrees of freedom ν , defined as the difference between the number of bins and the number of the fitted parameters, we can evaluate the fit quality, denoted in a text as $P_{\chi^2}(D, \nu)$, from the χ^2 statistics. Defining the degrees of freedom in the likelihood, Eq.3.2, is not straightforward since also bins with template entries < 1 are included in the sum. When determining the number of degrees of freedom, we consider only the bins in which the entry values in the model prediction exceed 0.1 or in which more than zero events were observed in the data. Thus, the resulting quality of the fit depends on the number of degrees of freedom, so one should be cautious with this approach when comparing the output of the fits to the X_{\max} distributions with different binning. The difference in the quality of the fit depending on the binning of the X_{\max} distribution is shown in the bottom panel of Fig. 3.7.

For most of the analysis in this chapter, we will use the CONEX air shower simulation programme to generate the X_{\max} distribution templates. Alternatively, as air shower simulations with sufficient statistics can be time-consuming, the Gumbel function parameterization used in the previous section to assess the method performance is also commonly utilized. As shown in Fig. 3.8, the estimated mass composition obtained using the CONEX-simulated X_{\max} distribution templates and the Gumbel distribution parameterization overall shows nearly no difference with an exception for one energy bin. The figure also demonstrates that the MAP and MLE give the same results, as expected.

In addition to the statistical uncertainty, we also consider the systematic uncertainty, in particular the X_{\max} -related one, when reporting the final results. In the PRD14 paper [82], the uncertainties on X_{\max} , the energy scale, and the parameterization of the detector resolution and acceptance (see Eq. 3.11 and Eq. 3.10, respectively) are taken into account. For the fit discussed in the current Chapter, we considered only the effect of the uncertainty on the X_{\max} measurements as a first approximation, which gives the most significant contribution to the total systematic uncertainty. To calculate the X_{\max} -related uncertainty, we consistently varied the shift in the X_{\max} scale δX_{\max} in the range from $-\sigma$ and $+\sigma$ with a step of $0.1(X_{\max}^{\text{up}} - X_{\max}^{\text{low}})$. We then fitted the shifted and rebinned data with the *iminuit* python minimizer. We then considered the largest and the smallest obtained fractions as the lower and upper systematic uncertainty limits. Similarly, the uncertainty limits were evaluated for the p-values by taking the smallest and largest p-values obtained for the fits with different X_{\max} -scales within the systematic uncertainty. Alternatively, we can also include the X_{\max} -scale as a free parameter constrained by a Gaussian prior. This approach fits nicely into the spirit of the MCMC approach to marginalize over all parameters, leading to similar systematic

3. Measurement of Mass Composition

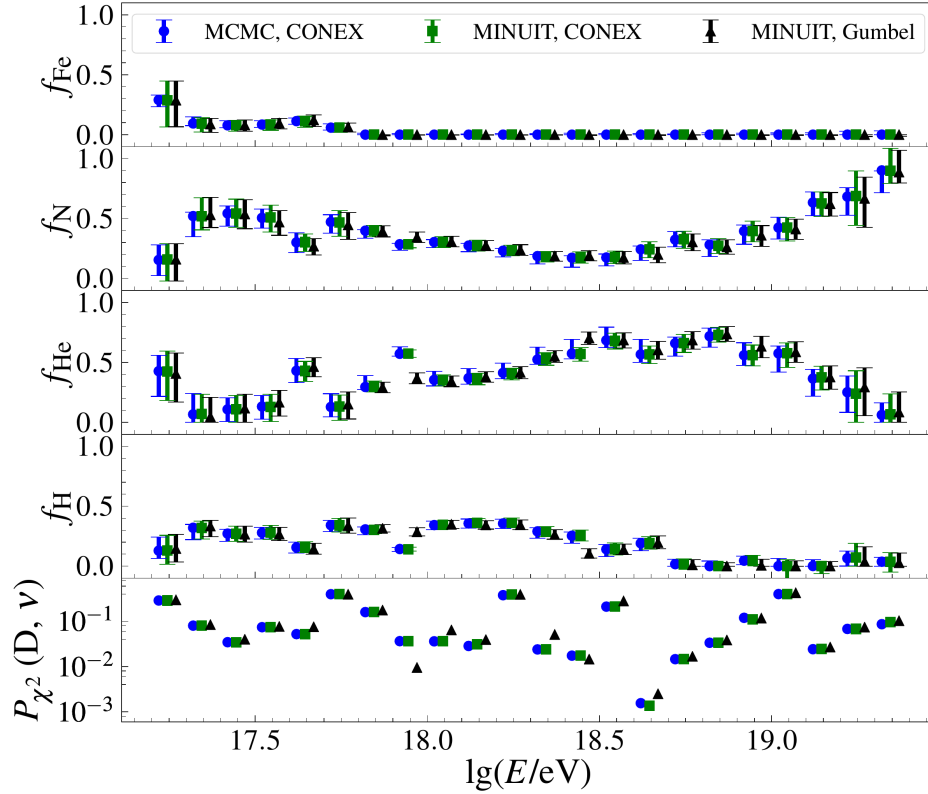


Figure 3.8: Comparison of the composition fit with CONEX X_{max} and Gumbel X_{max} distribution templates applied to the data from [145] for the low-energy part of the fit and from [143] for the high-energy part. In addition, the fit results obtained with the MCMC algorithm are compared to the prediction with the gradient minimizer.

uncertainties for the fractions, but it turns out that the best-fit scale-shifts are non-zero, i.e., that the data *prefers* to be shifted with respect to the model predictions. While this is interesting in itself, it goes beyond the inference of the fractions for the nominal X_{max} and model scales, and will be discussed in Sec.3.2.4.

Regarding the spectrum, we assume a spectral index of 1.1 for energies below 10^{18} eV and a spectral index of 2.2 for energies above 10^{18} eV, except for the last integral energy bin, which has a spectral index of 4.7. These values are obtained from the fit to the energy spectrum dN/dE of the selected FD event, which differs from the SD energy spectrum reported in [162] for the events arriving at Earth.

3.2.1 Comparison to previous results

As a first step of the analysis, the MCMC algorithm was applied to the PRD14 data. In Fig. 3.9 the comparison between the fractions obtained with the MCMC algorithm and the composition fractions derived in [82] for the PRD14 data set with the standard MLE approach is shown. As can be seen, there is good overall agreement. Choosing any of the other point estimates discussed above, we obtain very similar results except for the highest energies, where using the mode and median as point estimates gives fractions that differ more from the PRD14 values. Note that at some energies the MAP estimate is outside the 68% credible interval, and we then report the upper or lower physical limit of the fraction as the corresponding statistical uncertainty. On average, the uncertainties obtained with the MCMC error bars exceed the statistical uncertainties from the PRD14 mass composition fit. These small differences are within the precision of the coverage discussed in 3.1.3.

With the estimated fractions, we can then calculate the first two moments of the X_{\max} distribution, namely the mean $\langle X_{\max} \rangle$ and the standard deviation $\sigma(X_{\max})$ to compare with the data using Eqs. 2.5 and 2.6, respectively.

Since the MCMC fit gives us the posterior distributions for fractions, for calculating the X_{\max} moments, we may not use the point estimates of fractions. Instead, we can obtain the posterior distributions for the moments by plugging the samples from the distributions for fractions into the Eqs. 2.5 and 2.6, for which we can calculate the X_{\max} moments with any point estimate. The statistical uncertainties on the moments can be determined exactly as for the fractions by quoting the limits of the highest density interval of the distribution. This way of estimating the moments of the X_{\max} distribution illustrates one of the MCMC advantages - the construction of arbitrary posterior distributions from the samples and, therefore, a straightforward way to propagate uncertainties to arbitrary quantities derived from the fit values.

The comparison of the calculated from the MCMC fit X_{\max} moments and the moments from the PRD14 data set is shown in Fig. 3.10. Here, as a point estimate of the moment, we report the value corresponding to the MAP fractions. Alternatively, the mean or mode of the computed distribution for moments can be used for reporting the point estimates. As the one-sigma statistical uncertainty, we consider the HDI interval of the posterior distributions of the X_{\max} moments. Similar to the fractions, the point estimate is slightly above the credible interval for some energies. Since there are no physical limits for moments, we cannot redefine the statistical uncertainty as we did for fractions by quoting the upper/lower physical

3. Measurement of Mass Composition

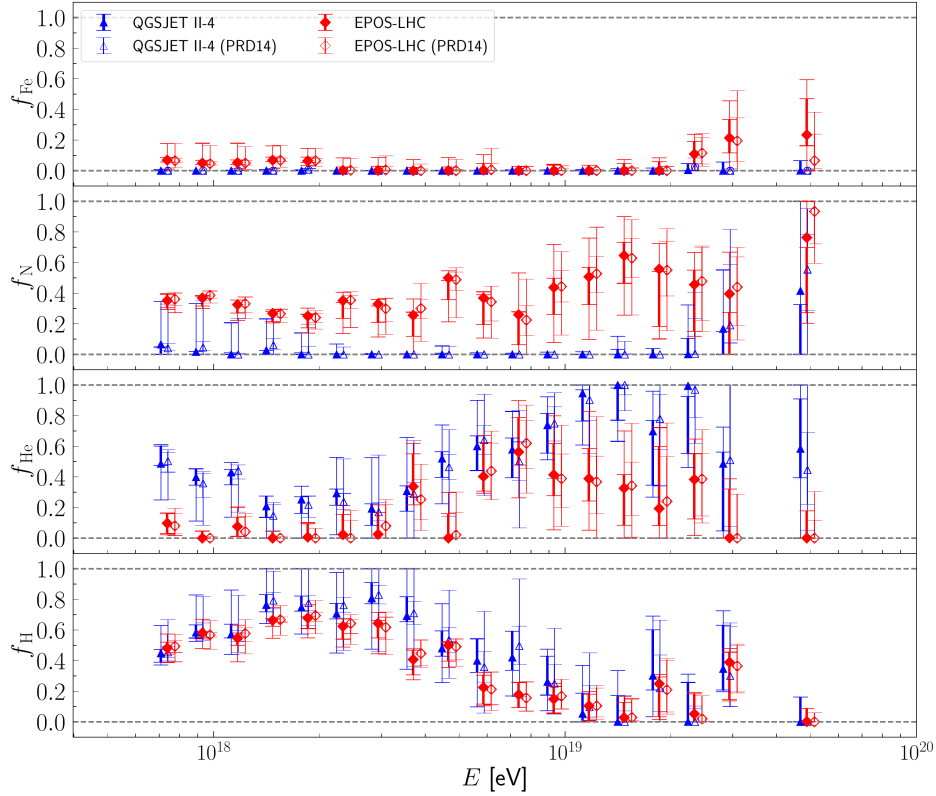


Figure 3.9: Comparison of the fractions derived with MCMC to the the previous Auger measurements [82] for the QGSJETII-04 (blue triangles) and EPOS-LHC (red diamonds) hadronic interaction models. The MCMC fractions are shown as filled symbols, and the PRD14 fractions are shown as open symbols of the corresponding colors. Inner error bars denote statistical, and outer error bars correspond to the total uncertainties.

limit in such a case. As can be seen, there is good agreement between the calculated X_{\max} moments and the data. However, systematic differences exist between data and calculations below 10^{18} eV, where the fitted composition model consistently underpredicts the $\langle X_{\max} \rangle$ by ≤ 5 g/cm². The comparison of the moments of the data and the fitted distribution is a powerful test of the agreement of data and the interaction model, in addition to the goodness-of-fit estimate discussed at the beginning of Sec. 3.2.




Figure 3.10: The first two moments of the X_{\max} distribution calculated from the MCMC fit (blue squares) to the PRD14 data (green circles): the mean $\langle X_{\max} \rangle$ (left) and the standard deviation $\sigma(X_{\max})$ (right). The calculated moments are shown as green circles and the data as blue squares.

3.2.2 Composition fractions for different data sets

In Fig. 3.11, the comparison of the estimated fractions for all three datasets discussed in Sec.2.5 for the EPOS-LHC and Sibyll 2.3d hadronic interaction models is shown. As can be seen from the figure, the overall trend in the evolution of the fractions with energy remains the same. In general, the fractions derived from the different datasets do not differ much. Given the broad systematic uncertainty range (see, e.g., Fig 3.16 for estimating the total uncertainty, which includes statistical and systematic uncertainties), we can say that the values agree within the uncertainty range. The results for the ICRC23 and ICRC19 datasets agree within the statistical uncertainty. However, it should be noted that there are some differences in the point values of the fractions, especially for He and N nuclei. The differences are also more significant for the fractions fit performed with the Sibyll 2.3d interaction model. In particular, the fraction of helium estimated from the PRD14 data is smaller than that in the ICRC19 and ICRC23 data sets. Simultaneously, as the He and N fractions are anti-correlated, the fraction of nitrogen nuclei is higher in the PRD14 data. The most substantial amount of nitrogen is seen in the most recent data, whereas the predictions for nitrogen in the ICRC19 data are between PRD14 and ICRC23. The heavier composition in the ICRC23 data is also in line with the changes in data $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$, see [144] for

3. Measurement of Mass Composition

the comparison of the X_{\max} distribution moments for PRD14 and ICRC23 data. For the predictions with the Sibyll 2.3d model, the fitted He fraction is nearly 0 at most energies below $10^{18.2}$ eV for the PRD14 data, while in two other datasets, it is estimated to be between 30 - 40%. The difference of around 30% between the fractions from ICRC23 and PRD14 data estimated with the Sibyll 2.3d model persists throughout the entire energy range, except for several energy bins in the middle. There are also more protons at the lowest energies in the PRD14 data than in the later data sets. The difference between the ICRC19 and ICRC23 data in the prediction for composition fractions reaches at most no more than 10-15% at the highest nitrogen contribution. A similar trend in the discrepancies between the different data sets is also observed in the fit with the EPOS-LHC model, but to a lesser extent, with ICRC19 and ICRC23 data showing good agreement within the statistical uncertainties for the energy range considered. The contribution of iron is compatible with zero for all datasets. It is worth noting that the ICRC19 X_{\max} data indicates a Fe component at ultrahigh energies, which was less pronounced in previous fits with smaller statistics and correspondingly coarser energy bins at these energies. No iron was visible in the PRD14 fraction fits, and only Sibyll2.3 showed a small iron contribution at UHE in the fits presented at ICRC17. In our fit to the PRD14 data, there is also some iron at the highest energies, but this could be due to the change in energy binning compared to the analysis in the corresponding paper, as here we use the same binning at the highest energies for all three datasets in question, which results in a smaller number of events in the two last energy bins for the PRD14 data than in the corresponding paper.

3.2.3 Effect of including additional particle groups

As mentioned in the introduction to this chapter, the MCMC algorithm allows marginalization over a large parameter space. Thus, we can include other nuclei in the fit to see what effect they might have without worrying that many parameters might complicate the fitting procedure. The conventional fit usually consists of four species corresponding to four particle groups, but adding the Si nucleus or photons into the mix is of great interest because the Si mass group is dominating at ultrahigh energies in some astrophysical fits of spectrum and composition, see, e.g. [163]. Figs. 3.12 and 3.13 display several different fits with a varied composition mix, incorporating additional nuclei such as photons, oxygen, and silicon alongside the four standard nuclei.

The addition of different nuclei to the mix has a minor effect on the predicted composition, which is also reflected in the quality of the fit. If nitrogen is present

3. Measurement of Mass Composition

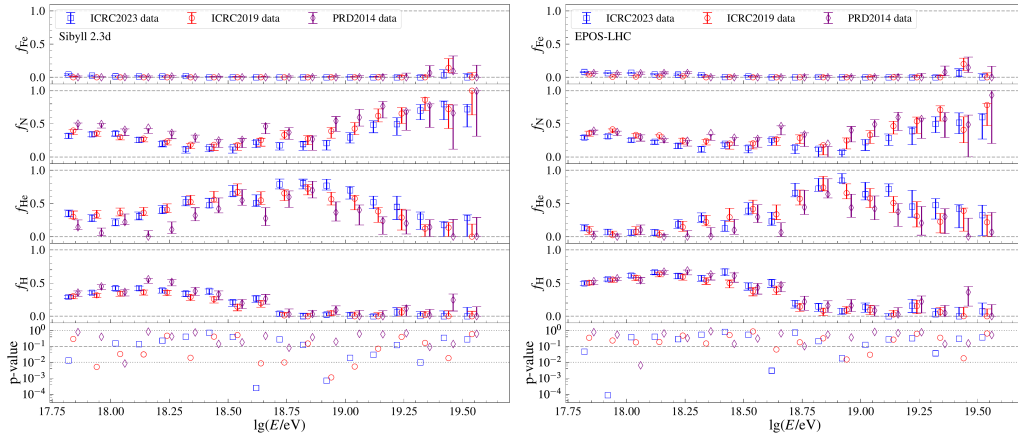


Figure 3.11: Comparison of the estimated composition fractions for the PRD14, ICRC19 and ICRC23 datasets with the Sibyll 2.3d (left) and EPOS-LHC (right) hadronic interaction models. Only statistical uncertainties are shown.

in the mix, oxygen can replace nitrogen in the fit, as these two elements are highly correlated. Adding Si to the fit does not alter the overall results, and the quality of the fit remains unchanged. The fitted Si fraction is consistent with zero throughout the energy range, except for the last integral energy bin, where there is some contribution from Si instead of Fe. This is in line with astrophysical fits that prefer Si at these energies. Including photons in the fit results in the estimated non-zero photon fractions contributing up to 5-6% to the mix. In addition, with photons in the fit, we see some improvement in fit quality over the nearly all energy range, particularly around $10^{18.6}$ eV, where the four-component conventional mix fits the data poorly. The cosmogenic origin of such a large photon fraction conflicts with the inferred (heavy) composition at ultrahigh energies. Furthermore, it disagrees with the photon limits derived from hybrid data, which, due to the additional use of surface detector data, have better sensitivity than the X_{\max} -only study presented here. Instead of a photonic origin, it is thus more likely that the models used to fit the X_{\max} distribution underpredict the number of deep showers, i.e., showers with large X_{\max} values, and this deficit is alleviated here by adding a photon component. A combined composition and cross section fit, discussed in Chap. 5, could show if a lower cross section can lead to a consistent photon fraction between hybrid and X_{\max} -only analyses.

3. Measurement of Mass Composition




Figure 3.12: The comparison of the fit results for the different particle species in the mix to the conventional four-component fit with H, He, N, and Fe nuclei (black stars). The additional components shown are O, Si and photons (γ). Only statistical uncertainties are displayed. The bottom panels shows the quality of the fit $P_{\chi^2}(D, v)$ estimated from the χ^2 statistics.

3.2.4 Fits with a free X_{\max} scale

Typically, the estimation of the mass composition by fitting the data to the templates constructed with either of the hadronic interaction models relies on the assumption that the prediction of the X_{\max} scale from the hadronic interaction models is correct. However, it is well known that existing models face challenges in accurately predicting observations of air showers, and moreover, the uncertainties on the predicted X_{\max} scale are larger than about a third of the difference between the $\langle X_{\max} \rangle$ of proton and iron nuclei [164, 165]. Therefore we investigated the possibility to include a shift in the X_{\max} scale δX_{\max} as an additional parameter in the composition fit. We apply a shift in the data rather than the templates, so no rebinning is required in the templates. The MCMC prior is redefined accordingly to include δX_{\max} , for which we set the upper and lower bounds to -40 and

3. Measurement of Mass Composition




Figure 3.13: The comparison of the fit results for the different particle species in the mix to the conventional 4-component fit with H, He, N, and Fe nuclei (black stars). The additional components shown are O, Si and photons (γ). Only statistical uncertainties are displayed. The bottom panels shows the quality of the fit $P_{\chi^2}(D, v)$ estimated from the χ^2 statistics.

40 g/cm², respectively. Due to the additional parameter, we needed to adapt the parameters of the MCMC chain. We increased the sampling length to 10⁶ steps and adjusted the burn-in period to 20000. Although an alternative way to increase the number of steps would be to change how the MCMC proposals are generated, none of the options available in emcee have provided a sufficient improvement.

Fig. 3.14 shows the results of the MCMC fit with a shift in the data X_{\max} scale as an additional fit parameter. The fit was performed with Sibyll 2.3d and EPOS-LHC interaction models to the ICRC19 data set. The normalized posterior distributions are plotted using a color scheme, with the points corresponding to the MAP point estimates. While the posterior distributions for the fractions are unimodal, with the MAP estimates coinciding with the mode of the distribution, the fit for the δX_{\max} can have several local minima, with the MAP not neces-

3. Measurement of Mass Composition

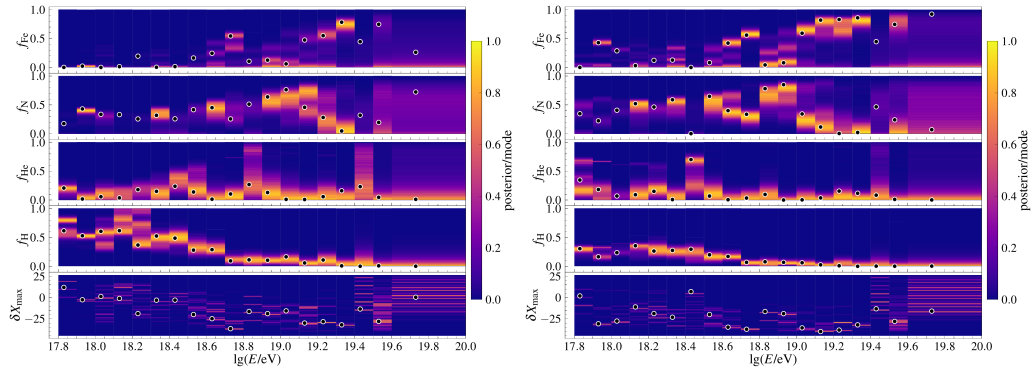


Figure 3.14: The 4-component MCMC fit with a shift in the data X_{\max} scale as an additional fit parameter. The fit is performed with the EPOS-LHC (left) and Sibyll 2.3d (right) hadronic interaction models for the ICRC19 data. The color scale shows the posterior distribution for the composition fractions and the shift δX_{\max} . The posterior distributions are normalized. The black dots are the point estimates of the fitted quantities (MAP).

sarily equal to the mode of the distribution. Allowing for the shift in the mass composition fit results in a noticeable change in the predicted fractions, with a significant increase in the fraction of N and Fe nuclei for both interaction models. Consequently, the He fraction becomes compatible with zero in more than 2/3 of the energy bins, even at energies above $10^{18.6}$ eV, where the He-N mix dominates in a standard fit. The fraction of protons in the composition mix remains almost unchanged.

It is clear that in the EPOS-LHC model, the favoured shift in data tends to be close to or above zero at lower energies and gradually decreases with increasing energy. The shift value reaches approximately -30 g/cm² beyond 10^{19} eV (except for the last energy bin, where the preferred shift is back to the default X_{\max} scale). In the case of the fit with the Sibyll 2.3d model, the average shift is about -25 g/cm², except for a few energy bins with a jump in the fit results, without showing any particular trend with the evolution of the energy.

Fig. 3.15, shows the results of the fit to the most recent ICRC23 data with Sibyll 2.3d interaction model. The trends in the evolution of the predicted composition and the shift in the X_{\max} scale are very similar to those derived from the fit to the ICRC19 data, indicating an increase in the contribution from heavier nuclei. At lower energies, because the fit favors the near-zero or positive shift in the X_{\max} scale, the fitted composition is either consistent with the results of the fit with a

3. Measurement of Mass Composition




Figure 3.15: Mass composition fit to the ICRC23 data with a shift in the X_{\max} scale using the Sibyll 2.3d model. The color scale shows the posterior distribution for the compositional fractions and the shift δX_{\max} normalized to the corresponding modes. The black dots are the point estimates of the fitted quantities (MAP).

fixed X_{\max} scale or has an increased contribution from the lighter nuclei, i.e., H and He. With increasing energy, the shift in the X_{\max} scale decreases so that the corresponding composition is dominated first by N and then by Fe nuclei after the δX_{\max} reaches a value of $\approx -30 \text{ g/cm}^2$ at the highest energies.

It is worth noting that estimating the X_{\max} scale energy-by-energy may not be the optimal approach, as it could lead to inconsistent shifts at different energies. As an alternative method to the energy-by-energy fitting of the X_{\max} scale, a scan over a given range can be performed, followed by summing the χ^2 over the energies to estimate the best-fit single shift in the X_{\max} scale. In this scenario, we obtain a shift of $-4.5^{+1}_{-1.4} \text{ g/cm}^2$ and $-12.7^{+1.8}_{-0.8} \text{ g/cm}^2$ for the fit to the ICRC23 data with the EPOS-LHC and Sibyll 2.3d interaction models, respectively. For the ICRC19 data, the results are similar, with a shift of $-6.5^{+0.6}_{-2.2} \text{ g/cm}^2$ for the EPOS-LHC models and $-13.7^{+1.9}_{-0.8} \text{ g/cm}^2$ for the Sibyll 2.3d model. In addition, the Sibyll 2.3d fit has a second, local, minimum 2.2σ away from the global minimum, predicting a shift in the data X_{\max} scale of $-28.3^{+0.7}_{-2} \text{ g/cm}^2$. The second minimum value of the shift in the X_{\max} scale agrees well with the results obtained

in [166] for the Sibyll 2.3 hadronic interaction model. There is no second minimum observed with the EPOS-LHC interaction model.

3.3 Results

In this section, we report the cosmic-ray mass composition estimated from the full Phase I data of the Pierre Auger Observatory and discuss the interpretations of the obtained results. Furthermore, we investigate additional potential applications of the MCMC fractions fit procedure.

3.3.1 Composition fractions

In Fig. 3.16, the results of the mass composition fit with the MCMC algorithm is shown for a combination of four particle species, representing four elemental groups, approximately equally spaced in $\ln A$, for the EPOS-LHC and Sibyll 2.3d models for the latest ICRC23 data, including the systematic uncertainty. The fit is performed on the X_{\max} range of 0-2000 g/cm², and a binning of the X_{\max} distribution of 10 g/cm² was chosen. The energy binning is consistent with PRD14, employing a bin width of 0.1 $\lg(E/\text{eV})$ for all energies up to the last integral bin at $> 10^{19.6}$. All events above this energy are merged into one integral bin. For this result, we quote the more precise p-values obtained from fitting mock data, see discussion at the beginning of Sec. 3.2. Both statistical and total uncertainties on the fit fractions are shown. The total uncertainty on the composition fractions includes the statistical uncertainty from the MCMC posterior distributions and the effect of the systematic uncertainty in the measurement of the X_{\max} scale. For the fits for the individual X_{\max} distributions and the values of the fitted fractions, see Appendix B.

Overall, the cosmic-ray mass composition can be described as a mix of H, He, and N nuclei at lower energies and a He-N mix at the higher energies above the ankle in the cosmic-ray spectrum. Although the qualitative behavior for both interaction models is the same, one can also see the significant dependence of the choice of the interaction model on the individual fractions. On average, the Sibyll 2.3d model results in a He fraction that is $\approx 20\%$ larger at lower energies and an increase in the fraction of N nuclei at higher energies compared to EPOS-LHC. The proton fraction obtained with EPOS-LHC reaches up to 70% around $10^{18} - 10^{18.2}$ eV and then drops to less than 20% above $10^{18.7}$ eV. The Sibyll 2.3d

3. Measurement of Mass Composition




Figure 3.16: The MCMC fit to the ICRC23 data with four elemental mass groups: H , He , N and Fe (top four panels). The error bars denote statistical (inner cap) and total (outer cap) uncertainties. The bottom panel shows the p -values of the fit. The fractions are shown for the following hadronic interaction models: Sibyll 2.3d (blue squares) and EPOS-LHC (red circles). Error bars for the fractions denote statistical and total uncertainties. Error bars for the p -values correspond to the variations in the fit fractions within the X_{\max} scale uncertainty range.

predicts a smaller proton fraction over the energy range considered, with a near-zero contribution at the higher energies. The amount of iron in the cosmic-ray mix is consistent with zero within uncertainties at all energies. Within the energy range observed, the data is compatible with a cycle from H to He to N , referred to in the literature as Peter’s cycle (see [167] for further discussion in the astrophysical context).

Similarly, as in the previous section, we can calculate the first two moments of the X_{\max} distribution, mean $\langle X_{\max} \rangle$, and standard deviation $\sigma(X_{\max})$, from the

3. Measurement of Mass Composition




Figure 3.17: The first two moments of the X_{\max} distribution, mean $\langle X_{\max} \rangle$ (left), and standard deviation $\sigma(X_{\max})$ (right) obtained from the MCMC-fitted fractions for the ICRC23 data. The predictions from the composition fit with Sibyll 2.3d and EPOS-LHC hadronic interaction models are compared to the data moments. The uncertainties on the derived moments are the statistical uncertainties calculated from the posterior distributions for the moments.

obtained fractions. Fig. 3.17 compares the moments of the data and the moments derived from the fractions fitted with EPOS-LHC and Sibyll 2.3d interaction models. Overall, the predicted moments follow a similar trend in evolution with energy and are in good agreement with the data within the statistical uncertainty limits. A difference between the data and the calculated $\sigma(X_{\max})$ is observed at several energy bins, where there are fluctuations in the data, while the $\sigma(X_{\max})$ derived from the fractions evolves rather monotonically with energy. Around $10^{18.6}$ eV, a slight increase in $\sigma(X_{\max})$ is observed in both data calculated moments, with the latter being less pronounced, despite the overall trend indicating a decrease in X_{\max} fluctuations with an increase in energy. Conversely, around $10^{18.9}$ eV and $10^{19.1}$ eV, the $\sigma(X_{\max})$ derived from the fractions fit with Sibyll 2.3d has higher values than the observed standard deviation. The predictions from EPOS-LHC, although also lower than data values, agree with them within the uncertainty limits. It is worth noting that at both $10^{18.6}$ eV and $10^{18.9}$ eV, the fit quality is poor, especially for the Sibyll 2.3d interaction model, with p-values below 10^{-3} .

3.3.2 Further applications of the MCMC mass composition fit

This section explores several examples highlighting the advantages of the MCMC inference approach, such as the straightforward estimation of uncertainties from the sampled posterior distribution and their propagation into any other distribution derived from the estimated fractions.

Fraction Limits

One example of the MCMC applications is the estimation of the upper bounds of the proton/photon fraction. This can be done directly from the posterior distribution by selecting the quantile of interest. Fig. 3.18 shows statistical 95% c.l. upper limits on the proton and photon fractions. The obtained results also show good agreement with the limits derived using a Neyman construction in [168]. At the lowest energies, the proton fraction in both the EPOS-LHC and Sibyll 2.3d models is about 1/3 of the mix, so the upper limits are quite high, reaching up to $\approx 60\%$ and 40% , respectively. At the energies of around 10^{19} eV, the maximum possible amount of protons in the Sibyll 2.3d model drops to less than 10% for the Sibyll 2.3d and less than 20% for EPOS-LHC but increases again at the highest energies. In general, as is also seen from the fitted fractions, the EPOS-LHC model always predicts more protons in the mix. The upper limit on the photon fraction, shown in Fig. 3.18 (right), is at a maximum of $\approx 10\%$ at the highest energies and is around 2% around the ankle in the cosmic-rays spectrum. Note that since we have limited statistics in the data, there is always a minimum possible value for the photon fraction [169]. The sensitivity to the photons should improve with an increase in the number of events.

Rigidity

Another example of the application of the MCMC posterior distribution for the fractions, and an error propagation can be the estimation of the average cosmic-ray rigidity similar to how the X_{\max} moments are obtained by using the MCMC samples f_i in the expression for the rigidity $R = \sum_i f_i \cdot E_i / Z_i$, where E_i is energy and Z_i is a species atomic number. The average rigidity calculated in such a way is shown in Fig. 3.19 as a function of energy for three hadronic interaction models for the ICRC 2019 data set. It is promising to see that the rigidity is increasing over the full energy range, i.e., that the charge increases slower than the energy, and thus, the deflection in cosmic magnetic fields decreases with energy. Since

3. Measurement of Mass Composition




Figure 3.18: The upper 95% statistical differential limits for the proton (left, for Sibyll 2.3d and EPOS-LHC models) and photon fractions (right, for Sibyll 2.3d model) for the fit to ICRC23 data. The photon limits are derived from the five-component fractions fit (γ , H , He , N , and Fe).

the magnetic deflection angle is proportional to the inverse of the rigidity, but in general $1/\langle x \rangle \neq \langle 1/x \rangle$, we also calculate the average of $1/R$ from the MCMC samples and display the inverse of that average (i.e., $1/\langle 1/R \rangle$) in the right panel of Fig. 3.19. As can be seen, the trend with energy is similar to that of $\langle R \rangle(E)$, but the absolute values change substantially. For a typical deflection angle of $\langle \theta_{\text{def}} \rangle = 3^\circ \times 60 \text{EV} \times \langle 1/R \rangle$ in the coherent Galactic magnetic field [170], the value of $\langle R^{-1} \rangle^{-1} \sim 10^{18.8} \text{ eV}$ derived for the EPOS-LHC fractions at the highest energies corresponds to $\langle \theta_{\text{def}} \rangle = 28^\circ$.

Nevertheless, such an average deflection angle is unlikely to be a very meaningful quantity for a mixed composition. For example, a mixture of protons and iron may have a large average deflection angle dominated by the nearly isotropic iron. However, the protons in this sample will still point back to regions close to the sources. Given a model of the Galactic magnetic field, the MCMC samples of the fractions can be used in future studies to derive a posterior of the distribution of deflection angles, from which angular quantiles can be derived that give a better representation of the cosmic ray arrival directions than the mean.

3. Measurement of Mass Composition

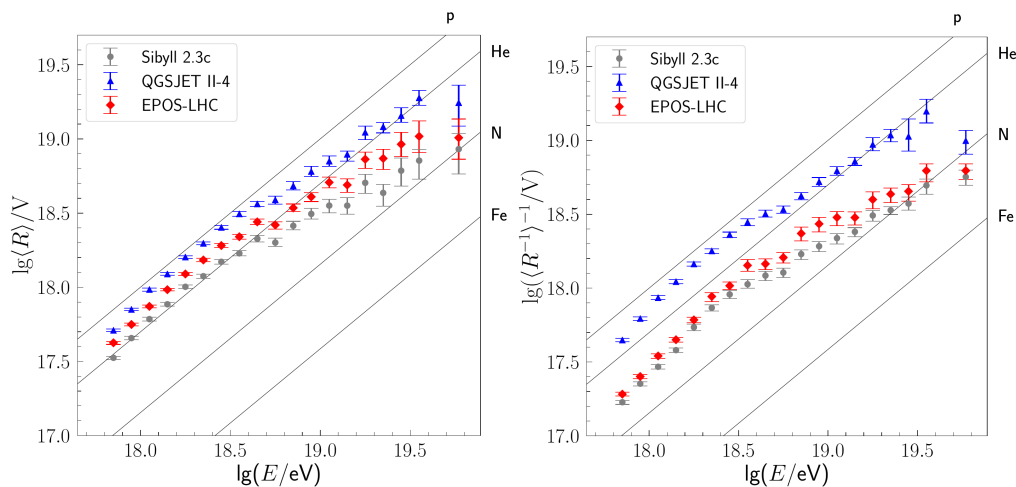


Figure 3.19: Average cosmic-ray rigidity $\langle R \rangle$ (left) and rigidity computed from inverse $\langle R^{-1} \rangle^{-1}$ (right) as a function of logarithm energy for the following hadronic interaction models: Sibyll 2.3c (gray circles), QGSJETII-04 (blue triangles), EPOS-LHC (red diamonds). The calculations were done for the PRD14 data. Only statistical uncertainties are shown.

Fractions from on/off Galactic plane

A previous study [171, 172] showed evidence for difference in the X_{\max} distribution from different parts of the sky, in particular from the region around the Galactic plane and the complementary region. With the fraction fit we can investigate which elemental group is responsible for this difference.

We split the data into two parts: one subsample consisting of events with a Galactic latitude, $|b| \leq 30^\circ$, referred to as on-plane sample, and another subsample with $|b| \geq 30^\circ$, referred to as the off-plane sample. Fig. 3.20 shows the fractions estimated from the on- and off-plane regions compared to the estimation from the full data. Since there is no statistically significant difference in the acceptance parameterization for on- and off-plane data, and the resolution parameterizations agree within the uncertainty [172], we use the same acceptance and resolution parameterizations for the two subsamples as for the full-sky analysis. The differences seen in the fitted on- and off-plane fractions are similar to the difference in the first two moments for the two regions of the plane (see [172] for the comparison of the moments of the X_{\max} distribution from on- and off-plane regions and further discussion of the astrophysical implications). While at lower energies the

3. Measurement of Mass Composition




Figure 3.20: The comparison of the composition fractions from on- (blue circles) and off-plane regions (violet diamonds). The fit is done with the Sibyll 2.3d interaction model to the ICRC23 data. The bottom panel shows the quality of the fit.

splitting of the data into two subsamples does not affect the results, at energies $> 10^{18.7}$ eV, there is an increase in the heavier component, i.e., in the N fractions, from the on-plane sample, reaching almost 100% at energies $> 10^{19}$ eV. In contrast, the fit to the full-sky, as was shown in Fig. 5.26, and off-plane data suggests a He fraction of up to 40%. A fluctuation in the proton fraction for the off-plane region at $\approx 10^{19.2}$ is also consistent with the data moments. Note that within the overall uncertainty limits, the proportions calculated for two regions of the plane agree with each other and with the full-sky results, except for a few energy bins where the deviation of one or another particle species exceeds the uncertainty.




Figure 3.21: Examples of the difference between the generated and reconstructed X_{\max} , $X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}}$. The differences are shown for two different energies and X_{\max} ranges.

3.4 X_{\max} -dependent X_{\max} bias

One of the factors sometimes omitted in the X_{\max} -based analyses, but which can indeed affect the interpretations of the measurements, is the existence of the non-constant reconstruction bias in the X_{\max} scale, which depends on the X_{\max} value itself. While both detector and reconstruction biases are estimated and corrected, the potential dependence of the reconstruction bias on X_{\max} is often overlooked. In this section, we attempt to estimate the X_{\max} -dependent X_{\max} bias and evaluate its effect on the composition measurements.

We estimate the bias using the extensive Monte Carlo simulations produced with air shower simulation programs for the description of the physics processes in the EAS and the detector responses [173]. As the simulations include all the reconstruction steps except the atmospheric effects, we can compare the true generated X_{\max} , X_{\max}^{gen} , with the corresponding reconstructed value X_{\max}^{rec} and evaluate the dependence in the difference on the X_{\max} . Fig. 3.21 shows an example of the difference between the generated and reconstructed X_{\max} . To ensure sufficient statistics in the X_{\max} distribution, we increased the energy binning to $0.2 \lg E$ while keeping the last integral energy bin unchanged, and we divided the X_{\max} range into 20 g/cm^2 intervals. For each such interval, we estimated a mean $\mu(X_{\max})$, and a

3. Measurement of Mass Composition




Figure 3.22: Examples of the evolution of the mean $\mu(X_{\max})$ (top) and standard deviation $\sigma(X_{\max})$ (bottom) of the difference between the generated and reconstructed X_{\max} scale with the generated (true) X_{\max} . The differences are shown for two energy ranges. The dashed line corresponds to the fit with a cubic/quadratic function to the data points. The solid line shows a fit obtained from interpolating the bias parameterization as a function of energy.

standard deviation $\sigma(X_{\max})$ of the difference $X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}}$.

The change in the difference with the evolution in X_{\max} is illustrated in Fig. 3.22 for two energy bins. As the MC simulations include four particle species, H, He, N, and Fe, in equal proportions (25% of the dataset each), we verified that

3. Measurement of Mass Composition

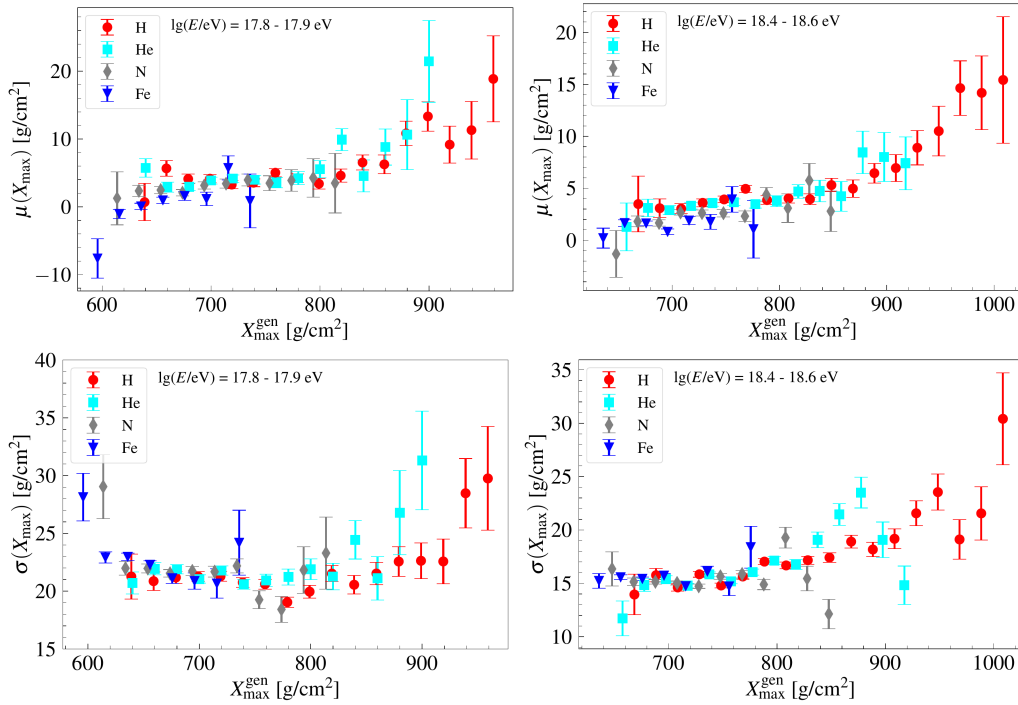


Figure 3.23: Examples of the evolution of the mean $\mu(X_{\max})$ (top) and standard deviation $\sigma(X_{\max})$ (bottom) of the difference between the generated and reconstructed X_{\max} scale with the generated (true) X_{\max} . The differences are shown for two energy ranges. Each particle species is shown separately.

we get the same trend in the mean and standard deviation trend seen in the simulations including all elements does not differ from the predictions for individual species (see Fig. 3.23).

Both the mean and standard deviation of the difference between the generated and reconstructed X_{\max} show a strong dependence on X_{\max} , increasing as X_{\max} increases. This effect could lead to an X_{\max} deeper by up to 20 g/cm^2 in the tail of the X_{\max} distribution, especially at the lowest energies as the X_{\max} is, in general, larger for low energies and decreases at high energies. At the same time, small X_{\max} values can become even smaller when the difference becomes negative. This is observed, for example, for the energy around $10^{17.8}$ eV. It is important to note that, on average, there is an almost flat region of approximately 80-100 g/cm^2 in the central X_{\max} range, indicating that the bulk of the X_{\max} distribution does have a nearly constant bias.

3. Measurement of Mass Composition




Figure 3.24: X_{\max} -dependent X_{\max} bias parameterization as a function of energy for the $\mu(X_{\max})$ interpolated with a cubic function with parameters par0 , par1 , par2 and par3 .

The observed in $\mu(X_{\max})$ and $\sigma(X_{\max})$ X_{\max} -dependent X_{\max} bias can be parametrized as a function of X_{\max} . For now, we have decided to use a cubic function for $\mu(X_{\max})$ and a quadratic function for $\sigma(X_{\max})$ as they gave the best quality of fit. For X_{\max} values exceeding the MC range, we use a linear extrapolation as we do not expect the bias to continue to grow cubically above $\approx 1200 \text{ g/cm}^2$.

Although the same functional dependencies apply to all energies, the quantitative and qualitative behavior varies from energy to energy, so extending the

3. Measurement of Mass Composition




Figure 3.25: X_{\max} -dependent X_{\max} bias parameterization as a function of energy for the $\sigma(X_{\max})$ interpolated with a quadratic function with parameters $par1$, $par2$, and $par3$.

parameterization over the whole energy range is not entirely trivial. It may eliminate the possible energy-by-energy fluctuations but could lead to less accurate results if the fitting function is not properly selected. It was found that the energy dependence of the bias parameterization for $\mu(X_{\max})$ and $\sigma(X_{\max})$ can be described well with a cubic function:

$$\begin{aligned}\mu(X_{\max}) &= par_{0,\mu}X_{\max}^3 + par_{1,\mu}X_{\max}^2 + par_{2,\mu}X_{\max} + par_{3,\mu}, \\ \sigma(X_{\max}) &= par_{0,\sigma}X_{\max}^3 + par_{1,\sigma}X_{\max}^2 + par_{2,\sigma}X_{\max} + par_{3,\sigma}.\end{aligned}\tag{3.14}$$

For the corresponding numerical values of the parameterizations, see [Appendix A](#).

Fig. 3.24 illustrates the dependence of the bias parameterization as a function of energy for the mean difference between the generated and reconstructed X_{\max} , and Fig. 3.25 shows the dependence on the energy for the $\sigma(X_{\max})$ parameterization. In the latter case, since we use the quadratic function for the $\sigma(X_{\max})$ dependence on X_{\max} , the coefficient of the cubic term, $par0$, is equal to 0, so it is not shown. To verify that the functional dependence of the bias parameterization on energy gives the correct $\mu(X_{\max})$ and $\sigma(X_{\max})$ over the whole X_{\max} range, we verified that the function with the fit parameters obtained from the energy dependence reproduces the fitted points and the fit itself well, as demonstrated in Fig. 3.22.

As we do not know the true values in the measurements, we cannot correct the data directly. Instead, we implement the correction in the X_{\max} distribution templates by modifying the double Gaussian detector resolution function. For

3. Measurement of Mass Composition

this, we re-calculate the standard deviations of the Gaussians and shift each of the Gaussians by $\mu(X_{\max})$, e.g. we would have now:

$$R(X - X_{\max}) = fG(\mu(X_{\max}), \sigma_1^{\text{corr}}) + (1 - f)G(\mu(X_{\max}), \sigma_2^{\text{corr}}). \quad (3.15)$$

For simplicity, we assume that f depends only on energy. We also assume that the ratio between σ_1 and σ_2 does not evolve with X_{\max} either so that we can relate $\sigma(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ to the resolution of each Gaussian in the detector resolution function.

In principle, we could also fit the double Gaussian function to $X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}}$ instead of taking the full width of the difference, but the corresponding fit is unstable, leading to significant fluctuations in the results, even for an unbinned maximum likelihood fit. Therefore, as a first approximation, we use the full width and the mean of the distribution of the difference between the generated and reconstructed X_{\max} instead of directly fitting the resolution function to it. A comparison of the X_{\max} distribution template for H nuclei with and without implemented bias correction is shown in Fig. 3.26.

With implemented correction for the X_{\max} -dependent X_{\max} bias, we evaluated its effect on the mass composition estimation. As can be seen from Fig. 3.27, the fitted mix of particle species almost does not change independently of the bias correction, which would be expected assuming that the fit is dominated by the central part of the X_{\max} distribution with a constant X_{\max} bias. The only minor differences, mainly negligible within the uncertainty limits, are seen in the predicted composition, indicating a slightly lighter composition over the whole energy range, with an increase in the proton component at the lower energies and He component at $> 10^{18.6}$ eV.




Figure 3.26: The X_{\max} distribution templates with and without X_{\max} -dependent X_{\max} bias correction for proton primary at $10^{17.8}$ eV.

3. Measurement of Mass Composition




Figure 3.27: The comparison between the fitted fractions with and without the correction for the X_{\max} -dependent X_{\max} -bias. Only statistical uncertainties are shown. The quality of fit (lower panel) is accessed via the Poissonian likelihood.

3. Measurement of Mass Composition

Chapter 4

Measurement of the proton-proton cross section from the tail of the X_{\max} distribution

The standard method for measuring the interaction cross section from ultrahigh-energy cosmic-ray data is based on the fit to the tail of the X_{\max} distribution and the subsequent conversion of the derived value to the proton-air interaction cross sections with air shower simulations. This is possible because the exponential tail of the X_{\max} distribution has a slope proportional to the nucleon-air interaction length. So even though we cannot directly probe the depth of the first interaction point, which is directly related to the interaction length of the protons in the atmosphere and thus to the proton-air cross section, we can still infer $\sigma_{p\text{-Air}}$ from the distribution of the depth of the air shower maximum [174].

This approach for measuring the proton-air cross section from the X_{\max} distribution tail was initially applied in analyzing the Fly's Eye data [174, 176] and was adopted by other air shower experiments later, such as the HiRes Project [177], Pierre Auger Observatory [178, 175] and Telescope Array [179, 180].

4.1 Method

The fit to the tail of the X_{\max} distribution is typically done using a certain fraction of events in the tail. In the Telescope Array (TA) Collaboration measurements, the start of the exponential fit is determined as the minimum stable X_{\max} value that maximizes the number of events in the tail of the distribution, which is estimated

4. p - p cross section from the tail of the X_{\max} distribution

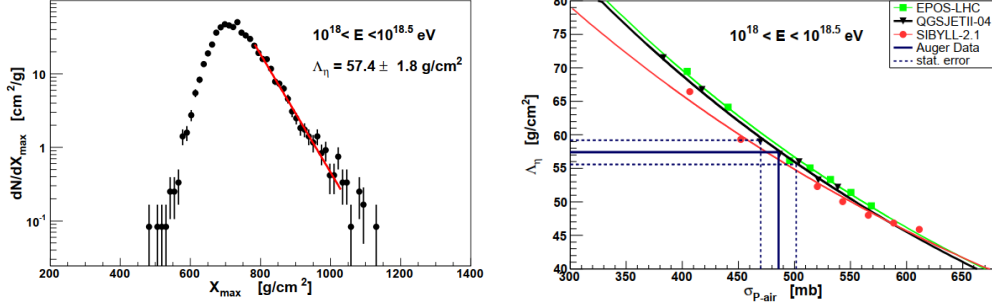


Figure 4.1: The unbinned likelihood fit to the tail of the X_{\max} distribution (left) and the conversion to the proton-air interaction cross section for the energy range of 10^{18} - $10^{18.5}$ eV. Figure from [175].

from the data as $X_i = \langle X_{\max} \rangle + 40$ g/cm². In the Pierre Auger Observatory analysis, the fit range is set to the tail fraction $\eta=0.2$, meaning that 20% of the deepest showers are considered with the fit starting at the corresponding depth. Here, we will adopt the approach of the Pierre Auger Collaboration, where, additionally, a selection of events for cross section analysis is conducted to guarantee an unbiased X_{\max} distribution across the entire depth range of observed X_{\max} values as follows. We will refer to this selection as xsecFoV in the following.

Firstly, a fiducial event selection (see Sec. 2.5) is performed for the range of X_{\max} values containing 99.8% of the X_{\max} distribution. The X_{\max} value corresponding to $\eta=0.2$ is then found from the obtained subset. The range of the tail to fit is thus defined by the beginning of the 20% tail, X_{20} , and the upper limit of the 99.8% of the X_{\max} distribution, $X_{99.8}$. Only those events with geometries allowing the complete observation of the fit range are selected to maximize the statistics of the unbiased distribution in the tail. Then, the exponential tail is fitted with:

$$f(z) = \frac{1}{k} e^{-\frac{z}{\lambda_\eta}}, \quad (4.1)$$

Here, $z = X_{\max} - X_{20}$, k is a normalization factor defined over the $[X_{20}, X_{99.8}]$ interval, and λ_η represents the slope of the exponential profile. The conversion of λ_η to proton-air cross sections is carried out through air shower simulations with an energy-dependent scaling factor included to modify the original models. Note that the conversion is subject to uncertainties due to extrapolations of hadronic properties from low-energy data and the assumptions made in this context. Additionally, the experimental uncertainties may also have a non-negligible effect

on the estimation of the proton-air interaction cross section. The potential presence of helium in the data is considered as the dominating systematic uncertainty, assuming a He fraction of less than 25%.

More substantial amounts of helium could result in a significant overestimation of the proton-air cross section by more than 80 mb if helium constitutes more than half of the composition mix, see Fig. 5.12 below. In Fig. 4.1, an example of the fit to the tail of the X_{\max} distribution alongside the conversion curve for the proton-air cross section is shown for three different interaction models. The corresponding cross section measurements are shown in Fig. 1.5 (top), with the proton-proton cross sections being calculated from the proton-air cross sections with the Glauber formalism and taking into account the corresponding uncertainty.

In the case of a standard event selection used for the mass composition estimation, to which we will refer further in the text as compFoV selection, the damping in the exponential tail due to the detector acceptance should be taken into account in the fit [34]:

$$f(z) = \frac{1}{k} e^{-\frac{z}{\lambda_\eta}} \begin{cases} 1 & z < z_{\text{acc}} \\ e^{-\frac{z-z_{\text{acc}}}{\lambda}} & \text{otherwise,} \end{cases} \quad (4.2)$$

where $z = X_{\max} - X_{\text{start,tail}}$, $z_{\text{acc}} = x_2 - X_{\text{start,tail}}$ and k is the normalization:

$$k = \lambda_\eta (1 + \exp(-\frac{z_{\text{acc}}}{\lambda_\eta})) \left[\frac{\lambda_2}{\lambda_2 + \lambda_\eta} - 1 \right], \quad (4.3)$$

where λ_2 and x_2 are the detector acceptance parameterizations. The starting point of the fit (X_{20}) is defined in the same way as for the unbiased selection in the tail, and there is no upper limit on the fit range, e.g. the fit is performed to




Figure 4.2: An example of the unbinned likelihood fit to the tail of the simulated X_{\max} distribution at $10^{17.8}$ eV with acceptance damping using a broken exponential profile (blue line). Additionally, a line for $\exp(-z/\lambda_\eta)$ is shown (in green).

4. p - p cross section from the tail of the X_{\max} distribution

infinity. Fig. 4.2 shows an example of this "broken exponential" fit to the tail of the simulated X_{\max} distribution.

4.2 Bias and uncertainties

As we established the existence of X_{\max} -dependent X_{\max} reconstruction bias, the λ_η measurements have to include the corresponding correction. While this bias had nearly no impact on the composition measurements, the correction provided for the fit to the tail of the X_{\max} distribution is quite significant as the bias is the greatest at the large X_{\max} values, i.e., in the tail of the X_{\max} distribution. The bias correction results in an increase in λ_η for both the $xsecFoV$ and $compFoV$ selections, thereby improving the agreement between them as well as with the previous measurements. Note that the effect is slightly different for the two event selections considered, as the $compFoV$ selection includes larger X_{\max} values due to the absence of the upper limit on the X_{\max} values, and the parameterization of the bias correction for $xsecFoV$ is different from the one for the $compFoV$. The fits to the simulated exponential distribution were used to obtain the corrections for




Figure 4.3: Bias correction for the λ_η fit. Left: The difference between the fit results to 1000 simulated data sets with (in green) and without (in pink) bias for $\lambda_\eta=60$ g/cm². Right: an example of the relation between λ_η measured from the biased distribution to the unbiased distribution for $10^{18.3}$ - $10^{18.4}$ eV.

the X_{\max} -dependent X_{\max} bias. We simulated 1000 X_{\max} distributions following single and broken exponential distributions in the tail for the xsecFoV and compFoV selections, respectively, and fitted them with the corresponding exponential profiles. Then, we adjusted each entry in the distribution based on the bias parameterization and re-fitted the biased distributions accordingly. As bias evolves with X_{\max} , the distribution entries shift unevenly, leading to a modification in the slope of the exponential tail of the X_{\max} distribution. We use the same X_{\max} ranges for the fits to the simulated data as we would use for the fits to actual data at each energy. We performed a fit to the biased and unbiased simulated distributions for a wide range of λ_η values to construct a function connecting the biased and unbiased λ_η measurements. We then corrected the measurements for the λ_η for the X_{\max} -dependent X_{\max} bias using the obtained correlation function. Fig. 4.3 shows the effect of the bias in the X_{\max} distribution for the λ_η measurements and provides an example of the correlation function. There are two different event selections shown in the Figure. One is a selection used previously for the cross section analysis by the Pierre Auger Observatory [178], xsecFoV, and another one is a selection used for mass composition estimation, compFoV. Fig. 4.4 compares the fit to the tail of the X_{\max} distribution with and without the correction for the X_{\max} dependent bias. The effect of an increase in λ_η value has a greater significance for the compFoV as it covers a broader range of X_{\max} values, as seen in the figure.

In addition to the statistical uncertainty from the fitting, there is also a systematic uncertainty for the λ_η fit with compFoV related to the detector acceptance parameterization. To estimate the impact of the acceptance-related uncertainty on the measured λ_η , we repeat the fit for four different combinations of δx_2 and $\delta \lambda_2$ uncertainties: both with a positive sign, both with a negative sign and one with a positive sign and another with a negative sign. The combinations giving the largest positive and negative deviations, on average within 1-2 g/cm^2 , from the fitted λ_η are then reported as the systematic uncertainty.

4.3 Results

The previous proton-proton cross section measurements from the Pierre Auger Observatory utilize the data recorded between December 2004 and September 2010. Since then, more than ten additional years of operation have been added, and there is now nearly four times as much data available, see discussion in Sec. 2.5.

4. p - p cross section from the tail of the X_{\max} distribution

Fig. 4.5 shows the results of the fit to the tail of the X_{\max} distribution for full Phase I data of the Pierre Auger Observatory. The measurements are corrected for the X_{\max} -dependent X_{\max} bias. The results obtained agree with each other and with the measurements of [175], referred to as ICRC2015 in the figure. The statistical uncertainties in the compFoV are smaller than those in the xsecFoV because fewer events in the tail of the X_{\max} distribution are passing the latter selection. Also, the fit uncertainties are much larger at the highest energies due to the smaller statistics in the data. The two datasets share less than 20 % of the common events, so they can be considered almost independent.

In general, λ_η decreases with energy, except for energies below 10^{18} eV, where there is a slight increase. The fluctuations in λ_η seen at $\approx 10^{18.6}$ and $10^{19.1}$ eV correspond to the fluctuations in the first two moments of the X_{\max} distribution.

Since the compFoV gives a smaller statistical uncertainty than the xsecFoV, we use it to estimate the proton-proton interaction cross sections. By modifying the proton-proton cross sections in the simulations, we can derive the function for converting λ_η into σ_{pp} similar to what is shown for proton-air interactions in Fig. 4.1.

Fig. 4.6 shows proton-proton inelastic cross section obtained from the fit to the tail of the X_{\max} distribution. The average estimated cross sections align with the Sibyll 2.3d extrapolation within the uncertainty limits, except for the interval

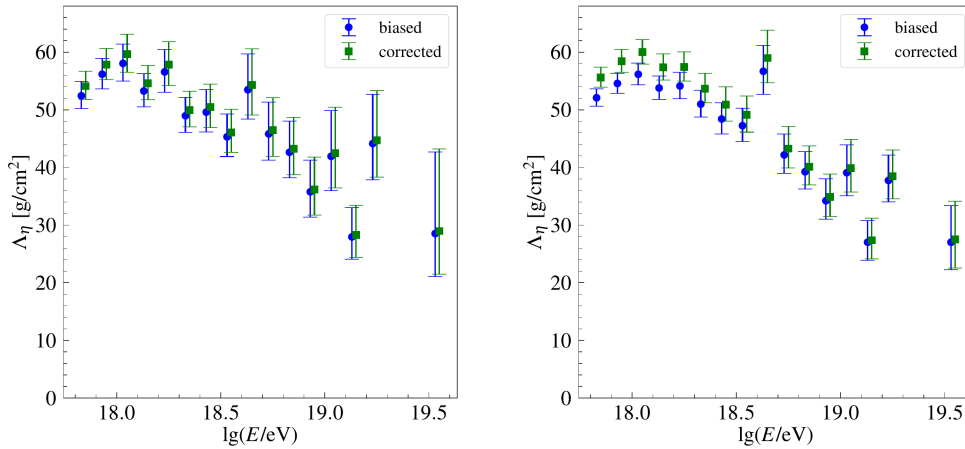


Figure 4.4: Comparison of λ_η measurement with and without the X_{\max} -dependent X_{\max} bias correction for the xsecFoV (left) and compFoV (right) event selections.




Figure 4.5: The measured λ_η from the ICRC23 data. Both $xsecFoV$ and $compFoV$ selections are shown. For the latter, in addition to the statistical uncertainty (inner cap), the total uncertainty, including the systematic uncertainty from detector acceptance and resolution, is shown (outer cap). The results obtained in this work are compared to the previous measurements [175].

of $10^{18.4}$ - $10^{18.6}$ eV where the large cross sections are measured, which corresponds to a decreasing trend in λ_η . Note that above $10^{18.3}$ eV, the He fraction estimated from the data reaches up to 50%, which could introduce significant bias and reduce the accuracy of the λ_η fit for cross section measurements, as previously discussed. The range shown is limited by the energies for which a sufficient amount of protons is present in the composition mix, see Fig. 3.16.

We can also estimate the λ_η corresponding to the prediction from the mass composition fit. We construct the X_{\max} distributions using the Gumbel function under the assumption of the composition estimated in the previous section from the ICRC23 data and fit the 20% tail of the distribution. In this example, we assume an ideal detector response, so the detector resolution and acceptance effects are not accounted for in the simulated X_{\max} distributions. Figure 4.7 compares the

4. p - p cross section from the tail of the X_{max} distribution




Figure 4.6: Proton-proton inelastic cross section derived from the fit to the tail of the X_{max} distribution (λ_η) compared to the previous Auger results and to measurements from the accelerator experiments.

predicted from the composition values of λ_η with measurements from the Pierre Auger Observatory and Telescope Array. The fitted fractions have been smoothed with a Gaussian kernel, resulting in a smooth transition from one energy bin to another. In addition to the λ_η corresponding to the measured composition mix, the fit to the pure proton simulations is shown. Results are shown for the EPOS-LHC and Sibyll 2.3d hadronic interaction models.

As can be seen, the difference between the λ_η obtained from the X_{max} distribution simulated under the mixed composition assumption and that of a pure proton composition is small in the energy range studied by the Pierre Auger Observatory. However, there is a significant change between the two scenarios at the higher energies studied by the Telescope Array. Consequently, the closeness of the Telescope Array measurements to the λ_η values for proton-initiated air showers [181]

4. p - p cross section from the tail of the X_{max} distribution




Figure 4.7: Comparison of the expectations for a pure proton composition (indicated by dashed lines) and a mixed composition (indicated by solid lines) of λ_η with measurements from the Pierre Auger Observatory [175] and the Telescope Array [179, 180]. In addition, the λ_η values derived in this study from the latest Pierre Auger Observatory data are included in the comparison. Only statistical uncertainties are given.

is unexpected. However, considering the statistical uncertainties associated with the measurements, the Telescope Array measurements are also consistent with predictions for a mixed composition.

4. *p-p* cross section from the tail of the X_{\max} distribution

Chapter 5

Combined estimation of the cosmic-ray mass composition and proton-proton inelastic cross section

The mass composition of cosmic rays and the cross sections of particle interactions are crucial for understanding the origin and nature of ultrahigh-energy particles and the fundamental properties of their interactions. Both quantities can be estimated from the X_{\max} data measured by the fluorescence detector, but their separate estimation is considerably influenced by the assumptions needed in the corresponding analyses. The mass composition estimations rely on the predictions for air shower characteristics from the hadronic interaction models. However, as discussed in the Sec. 1.3.1, these predictions can vary significantly between different interaction models, particularly for $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$, as shown in Fig. 1.6. The extrapolations of the measurements obtained in the collider experiments at the LHC energies are needed to interpret cosmic ray data at ultrahigh energies. One of the main challenges is the poorly understood uncertainties associated with extrapolation, so the existing interaction models may not fully cover the phase space of interaction characteristics. The construction of hadronic interaction models can benefit from studies of the interaction properties of ultrahigh-energy cosmic rays. However, testing the interaction models extensively with the air shower simulations may entail making some assumptions about the mass composition. For instance, the measurements of the cross sections for proton-air or proton-proton interactions, which are the focus of this study, are typically made under the assumption of the proton-dominated tail of the X_{\max} distribution. Any potential He-contamination effect is then considered as the systematic uncertainty (see, for

example, [178]).

The chapter is structured as follows: first, we discuss the modifications to the Sibyll interaction model, in the spirit of the f_{19} approach for the proton-air interactions covered in Sec. 1.3.1, but applied to the proton-proton interactions. Then, we present a new approach for the combined measurement of the cosmic ray mass composition and the proton-proton interaction cross sections and evaluate the performance of the method. Lastly, we will apply the new method to the actual data and discuss the implications of the obtained results for the physics of cosmic rays.

Some of the results discussed in this chapter were presented at the International Cosmic Ray Conferences in 2021 and 2023. For the corresponding proceedings, see [182] and [183], respectively.

5.1 Rescaling of the proton-proton interaction cross section

In this work, we follow the approach introduced in [102] to modify the interaction properties in the existing interaction models. However, instead of rescaling the proton-air interactions, as was done previously, and propagating those modifications to proton-proton interactions, we directly modify the latter.

The original proton-proton cross section is modified by the energy-dependent factor $f(E)$:

$$\sigma_{\text{mod}}^{\text{pp}} = \sigma_{\text{orig}}^{\text{pp}} f(E), \quad (5.1)$$

5. Combined mass composition & cross section fit




Figure 5.1: The effect of the rescaling in σ^{pp} on the σ^{A-air} and the relation between them. Left: The modified σ^{A-air} normalized to the original Sibyll 2.3d cross section for proton, helium, nitrogen, and iron nuclei. The x-axis is shown as energy per nucleon E/A . Right: The relationship between the proton-proton and nucleus-air interaction cross sections. The simulations were performed with a modified Sibyll 2.3d interaction model in CONEX.

where the rescaling factor has a form of:

$$f(E) = 1 + H(E - E_0)(f_{lgE_1} - 1) \frac{\lg(E/E_0)}{\lg(E_1/E_0)}, \quad (5.2)$$

with $H(E - E_0)$ being a Heaviside step function, E denotes the energy of interest, E_1 is an energy at which $f(E) = f_{lgE_1}$ and E_0 is a threshold energy corresponding to the onset of the modifications in the cross sections. As E_0 , we use the center-of-mass energy at the LHC of $\approx 10^{16.95}$ eV, as of time of writing this thesis (see [Appendix C](#) for the effect of the choice of E_0). Below this energy, the post-LHC version of the hadronic interaction models are tuned to the collider experiment data. We keep the reference energy, E_1 , equal to 10^{19} eV as in the original analysis. The modifications are implemented for the Sibyll 2.3d model in the CONEX air shower simulation program, where the cascade equations are incorporated below the threshold energy, and the full Monte Carlo simulations are performed above it. The σ^{pp} predicted by the Sibyll 2.3d model is 80 mb at the threshold energy and 110 mb at 10^{19} eV. The hadron-nucleon cross sections ($\pi+p$, $K+p$) are also rescaled by the f_{19}^{pp} factor. Given the input cross section for proton-proton

5. Combined mass composition & cross section fit




Figure 5.2: The algorithm for generating the X_{\max} distribution for the modified cross section.

interaction, a self-consistent rescaling of any nucleus-nucleus cross sections can be obtained using the Glauber calculation, see Sec. 1.3.3.

Fig. 5.1 illustrates the impact of the modified proton-proton interaction cross sections on nucleus-air interaction cross sections and the corresponding relationship between them. The x -axis is shown as energy per nucleon E/A since the onset of modifications for each nucleon $A > 1$ scales as atomic mass A according to the superposition model. The modified cross sections in Fig 5.1 (left) are normalized to the original Sibyll 2.3d cross sections. The effect of σ^{pp} rescaling is shown for three f_{19} values, 0.5, 2.0, and 3.0, for H, He, N, and Fe nuclei. Since the simulations were conducted at 10^{19} eV, which surpasses the E/A threshold for the modification of iron interactions, all four particle species are affected by the rescaling. The most significant changes in nucleus-air cross sections are present for protons. The heavier the nuclei, the smaller the deviations from the original cross section as anticipated. The relation between the proton-proton and nucleus-air interaction cross sections is shown in Fig. 5.1 (right), which can be used for a straightforward conversion between the quantities. The central question of interest here is how the introduced modification in the proton-proton cross section would impact the predicted X_{\max} distributions as the measured X_{\max} is sensitive to the interaction properties of cosmic rays and is commonly used to estimate the interactions of ultrahigh-energy cosmic rays. Fig. 5.2 shows a schematic depiction of the generation of the X_{\max} distributions for the modified proton-proton interaction cross section. First, the $f_{19} E_1$ factor is introduced for the energy of interest $E = E_1$, and the proton-proton cross section is rescaled. Then, the proton-air cross section, or any nucleus-air cross section in general, is modified according to Glauber theory. Finally, CONEX air shower program generates the X_{\max} distribution predicted by the modified interaction model.

The changes in the shape of the X_{\max} distribution with the rescaling factor f^{pp} are shown in Fig. 5.3. As can be seen, increasing the rescaling factor results in a narrower and shallower distribution. This can also be seen in Fig. 5.4, where the effect of rescaling the proton-proton cross section on the X_{\max} distribution is shown separately for the first two moments of the X_{\max} distribution, the mean, $\langle X_{\max} \rangle$, and the standard deviation, $\sigma(X_{\max})$.

Qualitatively, it is equivalent to that discussed in Sec. 1.3.2 for rescaling the proton-air interaction cross sections. As the primary particle mass increases, there are fewer changes in the moments of the X_{\max} distribution, with interactions of iron nuclei almost unaffected. Moreover, as the rescaling factor values increase, the magnitude of the changes decreases. At large rescaling values, $f_{19}^{\text{pp}} = 2.5$ and above, there is almost no sensitivity of X_{\max} to further increases in f_{19}^{pp} .

Fig. 5.4 illustrates a change in the depth of the first interaction point X_1 and the difference between X_{\max} and X_1 for different values of the rescaling factor. This difference is highly relevant for the measurement of the particle interaction cross section in air, see Sec. 1.2.2, as it provides a mean for retrieving the X_1 distribution from the measured X_{\max} and air shower simulations as shown in Fig. 1.4, as it is not possible to measure the depth of the first interaction directly.

5.2 Fitting algorithm

Now that we have a tool for estimating the fractions of cosmic ray nuclei and a method for modifying the proton-proton cross section, we can combine them to obtain an independent and simultaneous measurement of the cosmic ray mass composition and the proton-proton interaction cross sections.

The basis of the approach lies in a standard fit of the mass composition with model predictions constructed from the modified hadronic interaction properties, particularly the particle interaction cross section. This is done by rescaling the proton-proton production cross section, with the subsequent self-consistent modification of any nucleus-nucleus cross section via Glauber theory, as discussed in the previous sections.

In general, the combined mass composition and cross section fitting algorithm




Figure 5.3: Proton X_{\max} distributions for modified σ^{pp} at 10^{19} eV. The color scheme shows the change in the rescaling factor.

5. Combined mass composition & cross section fit




Figure 5.4: Top: the effect of the rescaling in σ^{PP} on the first two moments of the X_{max} distribution, mean $\langle X_{max} \rangle$ (left) and standard deviation $\sigma(X_{max})$ (right). Bottom: the effect of the rescaling in σ^{PP} on the point of the first interaction X_1 (left) and the difference between X_{max} and X_1 (right). The simulations were performed with a modified Sibyll 2.3d interaction model in the CONEX for proton, helium, nitrogen, and iron nuclei at 10^{19} eV. The averages of 5000 simulated showers are shown.

can be broken down into the following steps. First, we generate the X_{\max} distribution templates as predicted by the Sibyll 2.3d interaction model with modified interaction cross sections for closely spaced, with a step size of 0.1, f_{19}^{PP} rescaling factor values (refer to Fig. 5.2 for an illustrative representation of the steps involved in generating the X_{\max} distribution for the modified cross section). Then, with each generated template, i.e., for each value of the rescaling factor, we perform a standard mass composition fit. Given that the X_{\max} scale is not well-defined by the existing interaction models, as discussed in Sec. 1.3.1, we also account for the corresponding uncertainty in the fit by shifting the data by some δX_{\max} . Therefore, ultimately, we obtain a discrete scan over a range of f_{19}^{PP} and δX_{\max} values with the estimate of the best-fit mass composition for each combination of them. We then can access the goodness of fit for each combination of the shift in the X_{\max} scale and the cross section rescaling factor by calculating the deviance of the fit as defined by the logarithm of the Poissonian likelihood adopted in the maximum-likelihood fit, which is approximately χ^2 -distributed.

By comparing the resulting χ^2 values, we can find which combination of f_{19}^{PP} and δX_{\max} gives the best quality of the fit and determine the corresponding composition. In this way, we obtain the best-fit estimate of the proton-proton cross section rescaling factor f_{19}^{PP} , the shift in the X_{\max} scale δX_{\max} , and the cosmic ray mass composition at ultrahigh energies. We use the contour of

$$n_{\sigma} = \sqrt{\chi^2 - \chi_{\min}^2} \quad (5.3)$$

at $n_{\sigma} = 1$ to obtain the one-dimensional 68% CL uncertainties on δX_{\max} and f_{19}^{PP} . The χ_{\min}^2 is a minimal value of the χ^2 of the fit. When one of the parameters is fitted in the one-dimensional case, $\chi^2 + 1$ is given as the statistical uncertainty. Fig. 5.5 shows an example of the 1D fit to the simulated data with only a scan over f_{19}^{PP} considered for one energy bin. From this example, one can see which combination of the nuclei fractions and the f_{19}^{PP} gives the best fit quality, e.g., the smallest χ^2 . Note that instead of f_{19}^{PP} the x-axis shows $f_{17.8}^{\text{PP}}$, e.g. the rescaling factor value at that particular energy of $10^{17.8}$ eV calculated according to Eq. 5.2. Since the energy of $10^{17.8}$ eV in the example fit is below the threshold for the modifications of the nitrogen and iron interaction cross section, the variations in the proton-proton interaction cross sections do not affect the EAS variables for the heavier elements.

As established with simulated data, the fitting the cross section and X_{\max} scale at each energy is an underconstrained problem leading to ambiguous results (multiple minima). However, for a mixed composition evolving with energy, these

5. Combined mass composition & cross section fit




Figure 5.5: An example of the combined mass composition and cross section fit. The four upper panels display the estimated fractions as a function of the rescaling factor. The bottom panel shows the χ^2 of the fit. The fit is shown for the data simulated at $10^{17.8}$ eV with 95% H and 5% He.

ambiguities can be resolved, since only one X_{\max} shift fits at all energies, see discussion in Sec. 3.2.4 and Sec. 5.4. Therefore, we assume that both the factor f_{19}^{PP} and the shift δX_{\max} are constant across energies and consider all energies simultaneously to estimate the best-fit values of the fitted variables. To do this, we perform composition fits for all energies, scanning over f_{19}^{PP} and δX_{\max} , we sum the χ^2 values over the entire energy range, thus finding the f^{PP} and δX_{\max} that minimize the total χ^2 for the entire data set.

Since conducting full simulations of the X_{\max} distributions with air shower codes, with sufficient statistics to avoid random fluctuations, could be time-consuming, especially if simulations are needed for a broad range of cross sectional modifications, we utilize the Gumbel function parameterization for the X_{\max} distributions instead. The Generalized Gumbel distribution has a form of [159]:

$$G(z) = \frac{1}{\sigma} \frac{\lambda^\lambda}{\Gamma(\lambda)} (e^{-\lambda z - \lambda e^{-z}}), \quad z = \frac{x - \mu}{\sigma} \quad (5.4)$$

where $\Gamma(\lambda)$ denotes the Gamma-function, and λ , μ , and σ are free fit parameters.

The functional dependence of the shape parameters of the generalized Gumbel

5. Combined mass composition & cross section fit

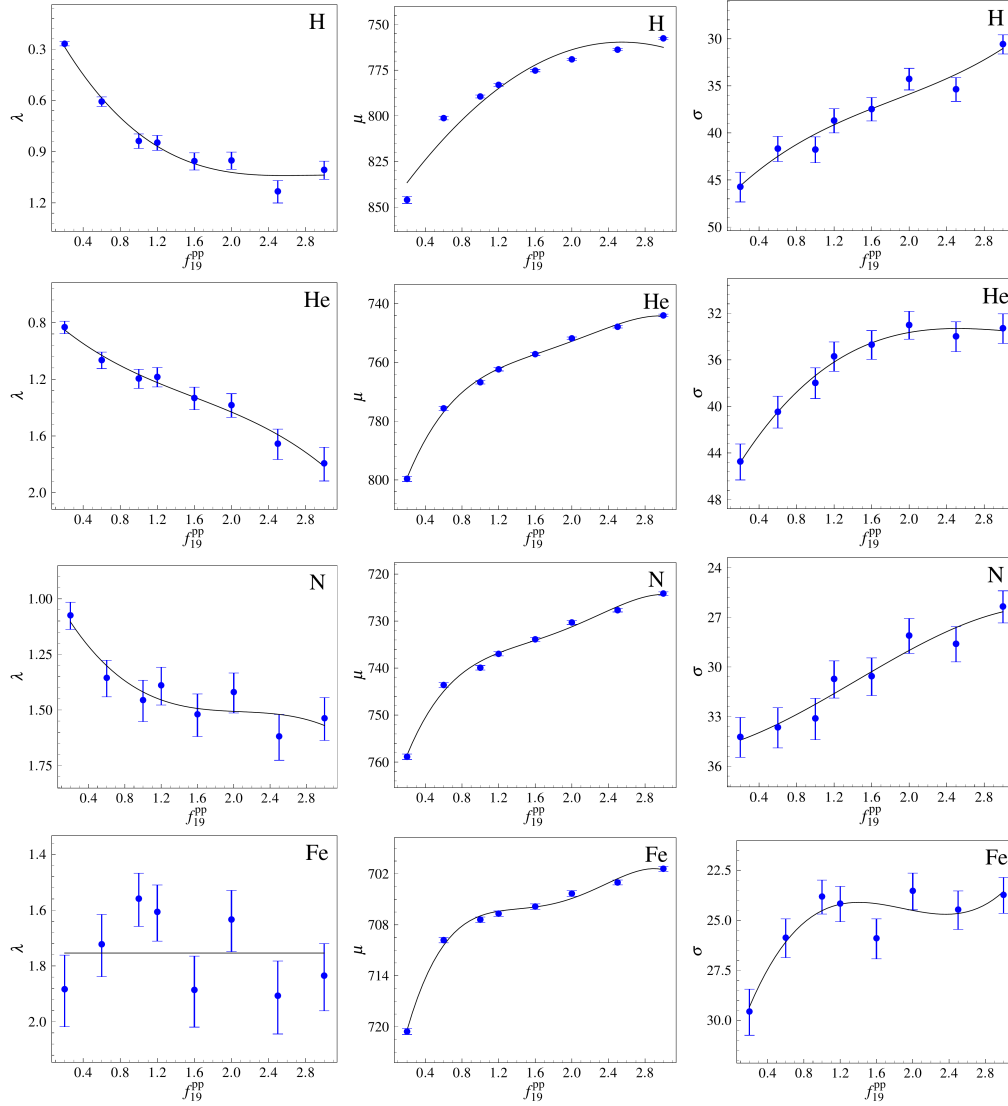


Figure 5.6: An example of the parameterization of the Gumbel function shape parameters, λ (left), μ (center) and σ (right) as a function of the rescaling factor f_{19}^{pp} at 10^{19} eV is shown for H (first row), He (second row), N (third row), and Fe (last row).

distribution on the scaling factor is obtained by fitting them to the X_{\max} distributions simulated with CONEX with limited statistics for a wide range of modified proton-proton cross sections. An example of the dependence of the Gumbel function parameterization on the rescaling factor at 10^{19} eV for proton, helium, nitrogen, and iron nuclei is shown in Fig. 5.6. To obtain the interpolation function, we fitted the evolution of the shape parameters with the rescaling factor using a polynomial function up to the 5th order and evaluated the quality of the fit. For each energy, particle species, and each parameter of the Gumbel distribution, we selected the function that gives the best p-value.

5.3 Effect of the cross section modification on the mass composition fit

Firstly, before reporting the final results, we evaluate the effect the change in the proton-proton interaction cross section has on mass composition. As an example, we performed a fit to the ICRC23 data for four values of f_{19}^{PP} factor: 0.6, 0.8, 1.2, 1.4. We keep the X_{\max} scale at the default value. Fig. 5.7 illustrates the expected deviation in the cross section corresponding to the given rescaling factor values. The figure also includes measurements from accelerator experiments and estimates derived from cosmic ray data recorded by the Pierre Auger Observatory. The latter covers, within its uncertainty, the range of rescaling factor values between 0.8 and 1.2.




Fig. 5.8 shows the corresponding change in composition with the change in interaction cross section. The uncertainties on the fitted fractions are only shown for the fit to the model predictions with the original Sibyll 2.3d cross section, as the same data are used in the fit for each value of the rescaling factor. The error bars are, therefore, very similar. At lower ener-

Figure 5.7: Comparison of the model extrapolations for σ^{PP} with different values of f_{19}^{PP} rescaling factor: 0.6, 0.8, 1.0, 1.2, and 1.4.

gies, where the composition is a mixture of H, He, and N nuclei, the change in the predicted fractions is observed for the lighter elements, H and He. The larger proton-proton interaction cross section, corresponding to the increase in f_{19}^{PP} , leads to larger proton fraction in the fitted mix, thus reducing the contribution of He.

For example, increasing the rescaling factor from 0.8 to 1.2 reduces the He fraction from ≈ 0.7 to 0.4, while the H fraction increases by almost the same amount. Up to energies where no protons are observed in the data (less than 5% in the mixture), the N fraction remains unaffected by cross section modifications. Above $\approx 10^{18.7}$ eV, the change in the fitted composition fractions for He and N nuclei is observed with the same trend as for H and He nuclei at the lower energies, with lighter nuclei becoming more dominant for the larger proton-proton cross section. The difference from the fit with the original Sibyll 2.3d cross section increases with energy. The iron fraction remains stable at almost all energies, except at the highest energies where its contribution is non-zero, and the increase in f_{19}^{PP} leads to the decrease in the Fe fraction. In general, an increase in the rescaling factor, and thus a larger proton-proton interaction cross section, results in a lighter composition over the whole energy range. The overall trend in the change in fractions with energy following Peter's cycle remains unaffected.

Fig. 5.9 shows two first moments of the X_{max} distribution, $\langle X_{\text{max}} \rangle$ and $\sigma(X_{\text{max}})$, derived from the fitted fractions, as well as the slope λ_{η} determined from the fit to the tail of the X_{max} distribution, compared to the data. Throughout the energy range, the mean X_{max} remains insensitive to changes in the interaction cross section and agrees well with the data within the uncertainty. The standard deviation $\sigma(X_{\text{max}})$ also does not show much dependence on the assumed value of the rescaling factor, the differences between them being within a few g/cm^2 . The second moment of the X_{max} distribution is smaller for larger cross sections, and the difference between the $\sigma(X_{\text{max}})$ of the modified Sibyll 2.3d model and the original one increases with energy, reaching up to 10% deviation for the 20% larger rescaling factor above $10^{18.5}$ eV. This behavior in $\sigma(X_{\text{max}})$ is consistent with a lighter composition obtained in the combined fit for the larger cross section since both an increase in the scaling factor and an increase in the fraction of lighter nuclei would have the same effect on the X_{max} distribution, narrowing it. Generally, there is good agreement between the $\sigma(X_{\text{max}})$ calculated using the fitted fractions and the data for all the rescaling factor values considered, except for a few energies where the original Sibyll 2.3d model also provides a poor description of the data.

To study the dependence of the λ_{η} derived from the estimated composition on the modified cross section, we constructed the X_{max} distributions using the CONEX simulations and the fitted fractions. We then fit the 20% tail of the distri-

5. Combined mass composition & cross section fit




Figure 5.8: The variation in the measured mass composition depending on the rescaling factor f_{19}^{pp} . Four top panels: composition fractions. Bottom panel: the quality of the fit accessed by a deviance $P_{\chi^2}(D, \nu)$.

bution to estimate the λ_η . The dependence of λ_η on the rescaling factor is similar to that observed for $\sigma(X_{\text{max}})$. As expected, with an increase in the rescaling factor (and therefore for larger proton-proton cross sections) the value of λ_η becomes smaller. While there is little dependence on f_{pp}^{19} at the lower energies, the difference between the results obtained under the modified cross section assumption increases with energy. As can be seen, the results of the fit to the tail of the X_{max} distribution corresponding to the fitted composition agree well with the data for the 20% deviation in the value of the rescaling factor.

5.4 Performance of the method

Before proceeding to the data analysis with a method for the combined measurement of the proton-proton interaction cross section and cosmic-ray mass compo-

5. Combined mass composition & cross section fit




Figure 5.9: The variation in the first two moments of the X_{max} distribution, mean $\langle X_{max} \rangle$ (left) and standard deviation $\sigma(X_{max})$ (center), as well as the slope λ_η from the fit to the tail of the X_{max} distribution (right), derived from the fitted composition fractions as a function of the rescaling factor f_{19}^{PP} . The calculations are compared to the data of the Pierre Auger Observatory.

sition, we verify how well the algorithm performs.

We start with a simple case of the fit at one energy bin with a fixed X_{max} scale. As the primary source of uncertainty in the estimation of the particle cross sections from the fit to the tail of the X_{max} distribution is the possible presence of He in the tail of the X_{max} distribution, we simulated data with different helium contamination, in particular we had a look at 0%, 5%, 25%, 40%, 50%, 60%, 75%, 95% and 100% of He in the data.

We also looked at several different rescaling factor values, e.g., we did tests on the original Sibyll 2.3d proton-proton cross section and the modified values. The total number of H and He events is always set to 6000 to avoid possible effects of varied sample size. This is similar to the number of events in the Pierre Auger Observatory data in the energy bins around 10^{18} eV. The data and the X_{max} distribution templates were generated using the parameterization of the generalized Gumbel distribution as a function of the scaling factor, with a binning of 1 g/cm². The selected energy range is $10^{17.8}$ - $10^{17.9}$ eV. For each composition and cross section scenario, we generated 100 data realizations and evaluated the fit performance based on the average results.

In the case of the simulated data mainly being dominated by one particle species, the χ^2 as a function of the scaling factor has a parabolic shape with a prominent minimum, as shown in Fig. 5.5, and correspondingly, the best-fit interaction cross sections are well-defined. However, if the composition is more mixed, the fit can have two minima or a very flat χ^2 even in a simple scenario

5. Combined mass composition & cross section fit




Figure 5.10: The χ^2 contour of the combined mass composition and cross section fit in a two-component H-He mix: proton-dominated composition with 5% of He (left), equal H-He mix (right), and He-dominated composition with 20% of H (bottom). The color bar shows how much the χ^2 deviates from the minimum value in units of sigma. The actual simulated parameters are shown in a blue dashed line.

involving two particle species. Fig. 5.10 shows an example of the dependence of the resulting χ^2 profile shape on the composition in the fitted simulated data for three scenarios: almost pure proton composition with 5% He fraction, equal H-He mix, and a He-dominated composition with 20% H.

As can be seen, the scenario with an equal amount of H and He in the simulated data can have two minima - one that correctly reproduces the input parameters and a second one where the predicted cross section is much smaller, approximately 1/4 of the simulated value ($f_{17.8}^{\text{PP}} \approx 0.25$), and almost pure He composition. This degeneracy is also seen in Fig. 5.11, where the two X_{max} distribution fits are compared. Although there is a difference in the fitted composition, this is compensated for by the significantly modified σ^{PP} , so there is no difference in the predicted X_{max} distributions as both of them describe the data equally well. For one energy bin, the only way to distinguish these several minima is by theoretical prejudice since the proton-proton cross section does not fall to a quarter of its value at LHC energies within a decade in beam energy (factor of 3 in center-of-mass energy). If more energy bins of data are present, and the composition varies with energy, then the extreme fit of the second minimum can also be excluded, as it will not be viable for other composition fractions. Note that this is a rather extreme example, and such scenarios are rarely observed in actual data. In the case of the He-dominated composition scenario, as shown in Fig. 5.10 (bottom), there is only one minimum in the χ^2 profile, but a very broad $\chi^2 + 1$ contour, resulting in the greater uncertainty in the estimated cross section and mass composition.

In Fig. 5.12, the results of the proton-proton cross section measurement from a combined fit are compared with those obtained from the fit to the tail of the X_{max} distribution, as discussed in Chap.5. The fit is performed on simulated data with varying helium fraction in the mix and different scaling factor values. The left side of the figure shows the bias in the estimated proton-proton cross section $\sigma_{\text{est}}^{\text{PP}}$ relative to the simulated $\sigma_{\text{sim}}^{\text{PP}}$. As can be seen, the proton-proton cross sections obtained from the combined fit are compatible with the standard approach for the




Figure 5.11: The fitted X_{max} distribution for the combined mass composition and cross section fit with an equal H-He mix in the simulated data. The shown X_{max} distributions correspond to two minima in the fit: one at H:He=1:1 with $f_{17.8}^{\text{PP}}=1.0$ and another one H:He=0:1 with $f_{17.8}^{\text{PP}} \approx 0.25$.

5. Combined mass composition & cross section fit

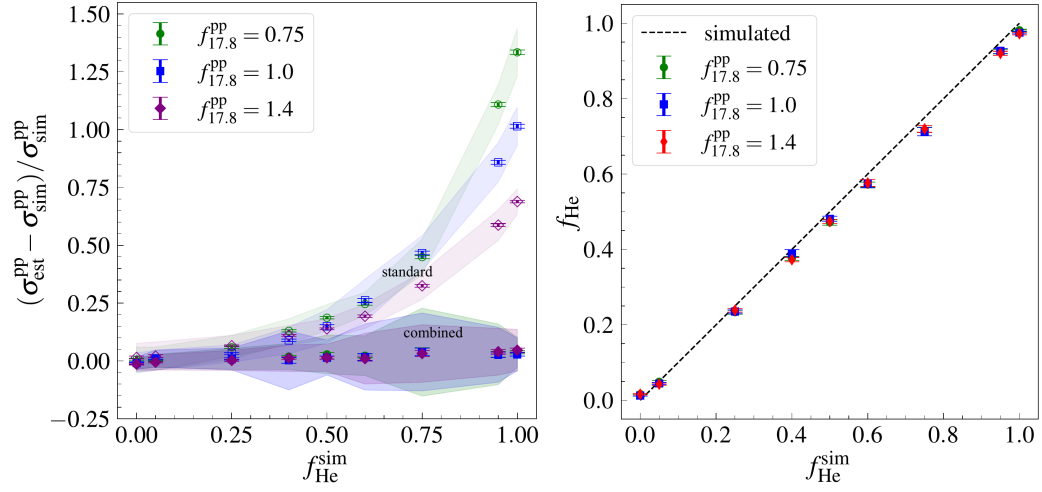


Figure 5.12: The performance of the combined fit as a function of He fraction for a two-component mix of proton and helium. Left: the bias of the fitted proton-proton interaction cross sections obtained from the combined fit (filled symbols) and the fit to the tail of the X_{max} distribution (open symbols) as a function of the He fraction. Right: the fitted He fraction as a function of the simulated He fraction. The averages of the fit to 100 simulated data realizations at $10^{17.8}$ - $10^{17.9}$ eV are shown. Error bars correspond to the uncertainty on the mean, and the shaded area shows a standard deviation.

proton-dominated composition. Furthermore, the combined mass composition and cross section fit not only reproduces the input parameters well for the large proton fraction in the simulated data but also yields nearly unbiased σ^{pp} even with an arbitrarily high contribution from helium in the mix, where the X_{max} distribution tail fit provides significantly overestimated values. Therefore, the combined fit does not require an almost pure proton composition in the tail of the X_{max} distribution and can be used to estimate proton-proton cross sections even with a significant amount of helium. The combined fit also performs well with respect to the composition and fully reproduces the simulated fractions, as shown on the left. The results are independent of the choice of the energy.

As outlined in the description of the fitting algorithm, in addition to scanning the rescaling factor, we estimate the best-fit shift in the X_{max} scale, δX_{max} . Therefore, as a further elementary test of the performance of the method, we fit the simulated data set with varied δX_{max} values. Figure 5.13 illustrates a fit to the proton-dominated simulated data - pure proton composition and a H-He mix

with 25% He, shifted by some δX_{\max} , scanned in a range between -40 and 40 g/cm^2 . Note that this is an extreme case, as the presence of N or Fe nuclei will make the range of possible shifts smaller. Again, we display the corresponding bias instead of showing the cross sections. There are noticeable changes in the predicted composition and cross section depending on the value of the shift. In particular, for the pure proton composition, the larger shift, heavier composition, and smaller cross sections are expected to compensate for the change in the X_{\max} scale. However, it should be noted that there is very little sensitivity in the quality of the fit to negative shifts as the χ^2 value remains nearly the same independently of the δX_{\max} value. However, large positive shifts in data are entirely disfavored by the χ^2 . With the increase in the shift value, the cross section rapidly decreases. However, this is insufficient to reproduce the data, as no further changes in the composition are possible, having reached the 100% proton fraction. This pattern remains the same as the composition becomes more mixed with an increase in the He fraction. Fig. 5.13 (bottom) shows the extreme case of a pure He composition in the simulated data. Again, at significant negative shifts in the X_{\max} scale, the fit prefers a well-mixed heavier composition and smaller cross sections. Nevertheless, as the X_{\max} scale approaches the default value, the deviation from the input value becomes smaller, eventually leading to a positive bias. Around $\delta X_{\max} = 0$ g/cm^2 , the cross section gradually gets smaller, and the fitted composition reaches the simulated 100% He fraction. At approximately 15 g/cm^2 , there is a phase transition from pure He to pure H composition, leading to a corresponding change in the proton-proton cross section to almost twice the input value. Similar to the pure proton composition, the changes in χ^2 are small for negative shift values, and all shift values above $\delta X_{\max} = 0$ g/cm^2 are disfavored by large χ^2 . The sharp transition from pure He to pure proton composition is also reflected in the χ^2 dropping to smaller values.

As a final test of the fit performance at individual energies, we simulate 100 data realizations with a composition similar to that observed in the cosmic ray data and for the composition close to the one derived in Chap. 3. We perform a fit at each energy bin, adjusting the number of events to match the observed data. The X_{\max} scale is fixed at its default value. Fig. 5.14 shows the bias in the estimated cross section and the fitted composition. Both quantities are well reproduced by the fit with little or no bias over the entire considered energy range independently of the simulated composition as it is different in different energy bins.

After preliminary tests of the method's performance at individual energies, we estimate a single rescaling factor and a shift in the X_{\max} scale δX_{\max} for all energies. We use the same procedure to test the algorithm, simulating the 100

5. Combined mass composition & cross section fit

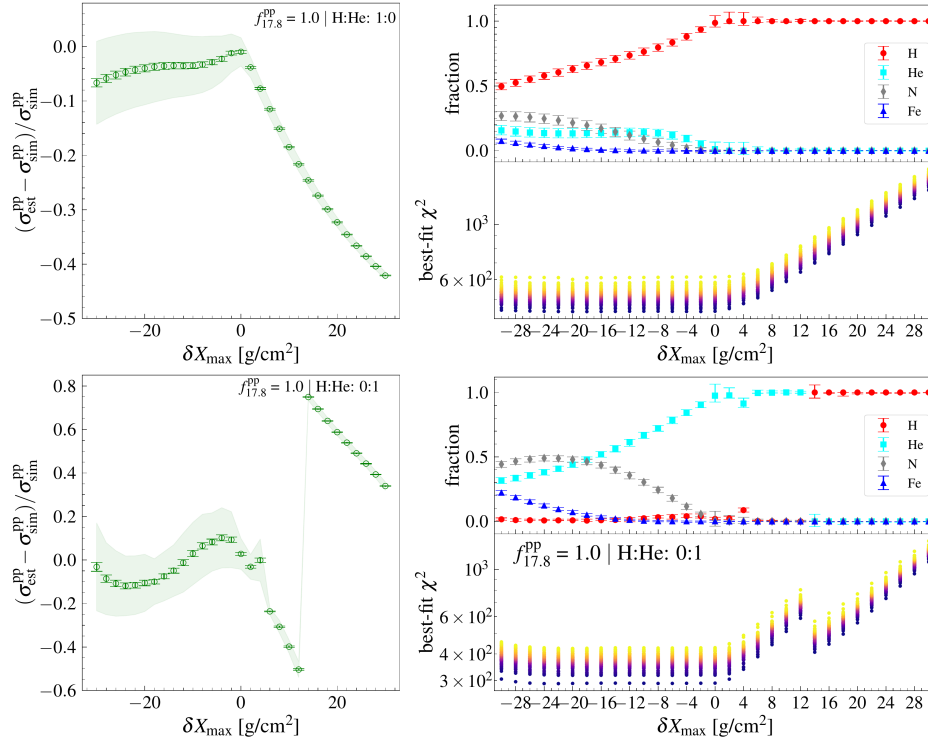


Figure 5.13: The performance of the combined fit with a shift in the X_{\max} scale δX_{\max} . The fit is done to data simulated at $10^{17.8}$ - $10^{17.9}$ eV with a pure proton composition (top) and a pure He (bottom). Left: The bias in the estimated proton-proton cross section. Right: The fitted composition and the χ^2 of the fit. The averages of 100 fits are shown for the fractions and cross sections. For the χ^2 , all fits are shown using a color scheme, where each color represents one of 100 fits.

X_{\max} distribution realizations for different composition scenarios and performing the fit with a scan over a range of f_{19}^{PP} and δX_{\max} values. We first generate data with the same parameterization of the Gumbel function as in the templates under the assumption of an ideal detector response. The number of events per energy bin equals the corresponding numbers in the actual data. In Fig. 5.15, the results of the combined fit are shown for several different composition scenarios in the simulated data with the ratios of H to He to N to Fe of 0.25:0.25:0.25:0.25, 0.15:0.6:0.15:0.1, 0.15:0.2:0.6:0.05, 0.6:0.2:0.1:0.1. The fit to the simulated data with AugerMix2023, which denotes realizations using the fractions derived in Chap. 3, is also shown. Note that while the other composition scenarios have

5. Combined mass composition & cross section fit




Figure 5.14: The performance of the combined mass composition and cross section fit for several energy bins. The fit was performed to the simulated data with a composition close to the observed composition in cosmic-ray data. Left: the bias in the estimated proton-proton cross section. Right: The fitted fractions.

the same fractions at all energies, in the latter case, the composition varies from energy to energy as observed in data. Depending on the simulated composition mix, there is variation in the bias observed in the measured cross section. Among the composition mixes considered, it is largest for the He-dominated composition with $\approx 10\%$ deviation towards the increase of the cross section value, and almost zero for the H- and N-dominated scenarios with 20% He. Generally, the bias does not exceed 8% for all the cases considered. For the composition measured in the actual data using unmodified Sibyll 2.3d model, the deviation in the fitted cross section from the input value at the level of 4%. The level of the bias also has some minor dependence on the presence of the δX_{max} shift in the fit. The estimated shift is negative and less than 0.1 g/cm^2 for all composition scenarios considered except the AugerMix2023 case with $\delta X_{\text{max}} = -2.56 \pm 0.16 \text{ g/cm}^2$. Thus, since the shifts obtained are compatible with zero or have minor deviations from zero, there is little to no difference in the composition and cross section predictions between the fits with a free and fixed X_{max} scale. A slight reduction in the bias of less than 2% is observed (see Figure 5.15), which could also be related to the more straightforward way of estimating the best-fit combination of variables, as with fixed $\delta X_{\text{max}} = 0$ the parabolic approximation/interpolation around the minimum

5. Combined mass composition & cross section fit




Figure 5.15: The estimated proton-proton cross sections obtained by fitting the data realizations simulated with different composition scenarios. Left: with a fixed X_{\max} scale. Right: with fitted shift in the X_{\max} scale δX_{\max} . The shaded area corresponds to the standard deviation, while the dashed lines mark the estimated value.

is needed only in one dimension. The plot for a fixed X_{\max} scale includes two additional composition scenarios: pure proton composition and equal H-He mix. As observed, the fit is slightly biased towards smaller cross section values for the pure proton composition and larger cross section values for the H-He mix, with an average deviation in the value of the rescaling factor f_{19}^{PP} from unity of ≈ 0.05 . Note that fitting a pure composition with a varied shift δX_{\max} makes no sense due to some degeneracy between f_{19}^{PP} and δX_{\max} , which results in all parameters combinations being equally preferred.

Fig. 5.16 compares the fitted and simulated fractions for different composition scenarios, constant in energy, and Fig. 5.17 shows the fit result for the data simulated with AugerMix2023 fractions, i.e. variable with energy composition. Independently of the assumed composition, the fit recovers the input fractions well with nearly no bias in the results, except for the highest energies where the deviation of up to 10% in the fitted fractions is observed in the He-dominated case.

In the tests above we used data simulated with the Gumbel function parameterized as a function of the f_{19}^{PP} scaling factor. However, there could be some bias in the data fit originating from the accuracy of the parameterization. To check this, we generated data assuming the same fractions of nuclei and number of events

measured in the composition-only fit for the original Sibyll 2.3d model using the Gumbel function parameterization and the CONEX air shower simulations program. In the case of the fit to the simulations with the Gumbel function parameterization without the dependence in the rescaling factor for the original Sibyll 2.3d model (see [159] for the parameterization), we get a bias of $\approx 11\%$ in the estimated value of the rescaling factor, e.g., $f_{19}^{\text{PP}} = 1.1$, and the X_{max} scale has a small bias of $\delta X_{\text{max}} = -2.85 \text{ g/cm}^2$. The fitted shift in the X_{max} scale remains consistent with that obtained from tests using data simulated by parameterizing the Gumbel function for the modified cross section. In the case of the simulated X_{max} distribution with CONEX simulations, the deviation of the fitted f_{19}^{PP} from the input increases, reaching a deviation of up to 24% in the estimated cross section at 10^{19} eV ($f_{19}^{\text{PP}}=1.24$). The shift in the X_{max} scale remains nearly unchanged compared to the test with Gumbel distribution. Assuming an equal number of events in each energy bin in the simulated data, the deviation from the 'true' simulated values decreases to $f_{19}^{\text{PP}} = 1.15$ and $\delta X_{\text{max}} = -1.98 \text{ g/cm}^2$. Furthermore, introducing a change in composition from that observed in the data to an equal mixture of all four nuclei further reduces the bias. The difference between the simulated and input values of the rescaling factor becomes less than 10% ($f_{19}^{\text{PP}} = 1.08$), and the shift in the X_{max} scale is compatible zero ($\delta X_{\text{max}} = -0.24 \text{ g/cm}^2$) as simulated. Furthermore, the issue might not be related to the precision of the parameterization used here, but to the accuracy with which the Gumbel function is capable to describe air shower simulations, see e.g.[184] for a discussion of different X_{max} parameterizations and their quality in the tail of the X_{max} distribution.

In addition to the tests conducted with Sibyll 2.3d simulated data, we also examined the fit performance using data simulated under the assumption of different interaction models, particularly QGSJETII-04. This model was chosen because, among all available hadronic interaction models with different predictions for air shower observables, QGSJETII-04 is the most optimal for fits with the Sibyll 2.3d templates. Due to the different approaches of treating the nuclear fragmentation, we cannot fully reproduce the X_{max} distribution shape predicted by EPOS-LHC, even with the modified Sibyll 2.3d model, see [185] for a discussion of an update of the treatment of nuclear fragmentation in the next version of the model. In the case of the Sibyll 2.1 MC generator, we cannot recover the slope of the extrapolated proton-proton interaction cross section, as it is not tuned to the most recent LHC data. The estimated proton-proton cross section from the fit to the QGSJETII-04 data simulated with AugerMix2023 composition is compared to the original QGSJETII-04 values in Fig. 5.18 (left). The results are shown for a fixed X_{max} scale and a varied δX_{max} . The best-fit shift obtained with a free X_{max}

5. Combined mass composition & cross section fit




Figure 5.16: The estimated fractions of particle species obtained by fitting the data realizations simulated with different composition scenarios with a composition constant across the energies. The input fractions are shown as dashed lines of the same color as the fitted composition.

scale is $14.2 \pm 2 \text{ g/cm}^2$, which is close to the average difference in the predicted X_{max} scale between the QGSJETII-04 and Sibyll 2.3d models. The observed deviation from the simulated values is minimal, approximately $\pm 5\%$ at 10^{19} eV , in both cases, with a larger cross section predicted in the case of the allowed shift in the X_{max} scale and a smaller cross section when $\delta X_{\text{max}} = 0$, compared to the original QGSJETII-04 extrapolations. The fitted fractions of particle species deviate from the simulated values towards a heavier composition (see Fig. 5.18 (right)). The predicted composition is significantly heavier with the X_{max} scale fixed, with up to 50% iron at the highest energies and no more than 30% He and 20% H at all energies.

In contrast, the simulated composition has up to half the mix being protons at the lowest energies and an almost pure He composition around $10^{18.7}$ - 10^{19} eV .

5. Combined mass composition & cross section fit




Figure 5.17: The estimated fractions of particle species obtained by fitting the data realizations simulated with AugerMix2023 composition.

With the X_{\max} scale set free, the composition estimate becomes more accurate, with tiny discrepancies of less than a few percent between the simulated and fitted fractions for H and Fe over the whole energy range, as well as for He and N at the lowest and most of the highest energies. In the intermediate energy range, the estimated composition also has an enhanced contribution from nitrogen as in the fit with $\delta X_{\max} = 0$, but to a much lesser extent. Thus, this example demonstrates that leaving the X_{\max} scale free in the combined mass composition and cross section fit leads to a more accurate estimation of the variables.

Finally, we tested the performance of the method with end-to-end Monte Carlo simulations of the EAS, including the detector response and all the physical process in the light emission and transmission and other environmental factors, like the night-sky background. Since the existing simulations have limited statistics and the number of events follows the cosmic ray spectrum with an index of $\alpha = 2.2$, there are fewer events at the lowest energies compared to the intermediate

5. Combined mass composition & cross section fit




Figure 5.18: The results of the fit to the data simulated with QGSJETII-04 interaction model with AugerMix2023 composition. Left: estimated proton-proton interaction cross section. The shaded area corresponds to the standard deviation, while the dashed lines mark an uncertainty on the estimated cross section.. Right: fitted mass composition. The fit is shown for a fixed X_{max} scale and the best-fit $\delta X_{max} = 14.2 \pm 2 \text{ g/cm}^2$.

energy range. However, the actual data has the highest number of events at lower energies of ≈ 10000 , which limits the ability to generate more than five fully independent datasets from the full MC simulations assuming the AugerMix2023 composition. Since increasing the number of simulations to 100 independent datasets is time-consuming, the existing MC simulations are resampled.

The results obtained demonstrate the existence of some bias in the estimator. The average fitted cross section deviates from the input value towards a larger cross section with the fitted rescaling of 1.29 ± 0.009 when the X_{max} scale is not fixed. The corresponding shift in the data is $\delta X_{max} = -8.6 \pm 0.2 \text{ g/cm}^2$. If we assume $\delta X_{max} = 0$, the bias in the fit is reduced to $\approx 25\%$ at 10^{19} eV ($f_{19}^{pp} = 1.25 \pm 0.006$). The fitted mass composition is heavier than simulated if the X_{max} scale is left free, and otherwise lighter if it is assumed default. The primary change is evident in the He and N fractions, where the amount of N increases correspondingly for the best-fit δX_{max} , while the contribution from He decreases by the same amount. The opposite picture, with an increase in the He fraction and a decrease in the N fraction, respectively, is present for $\delta X_{max} = 0$. The amount

5. Combined mass composition & cross section fit




Figure 5.19: The proton-proton inelastic cross section obtained from the fit to data simulated with Sibyll 2.3d interaction model including detector response. Left: fit with a fixed X_{max} scale. Right: fit with a scan in δX_{max} . Additionally, the effect of the X_{max} -dependent X_{max} bias correction is shown. The shaded area corresponds to the standard deviation, while the dashed lines mark an uncertainty on the estimated cross section.

of He from the fit with a fixed X_{max} scale is larger than the input fraction by $\approx 10\%$ for almost all energies, except in the small range between 10^{18} and $10^{18.5}$ eV where the difference is a few percent. This results in a composition around the ankle that is almost purely He, as opposed to the simulated H-He mix dominated by He. There is also a slight increase in the proton fraction compared to the simulations. For the $\delta X_{\text{max}} \neq 0$ fit, the difference between the fitted and simulated He fractions, and thus, N fraction, is about 10% at the lowest energies and rises to about 25% for energies above $10^{18.5}$ eV. Furthermore, an increased presence of iron is observed in the lowest and highest energy bins.

However, as the full MC simulations are subject to the exact reconstruction as the data, the X_{max} -dependent X_{max} bias should be considered. As was shown previously, we found minimal effect on composition measurements and, conversely, quite a change in the cross section measured from the fit to the tail of the X_{max} distribution. Therefore, a comparable bias contribution is anticipated in the combined mass composition and interaction cross section estimation. The bias was corrected in the same manner as for the composition fit, with the X_{max} templates constructing the model predictions modified accordingly.

5. Combined mass composition & cross section fit




Figure 5.20: The mass composition fractions obtained from the fit to the data simulated with the Sibyll 2.3d interaction model, including the detector response. The results of the fit with (blue squares) and without (red circles) the shift in the X_{max} scale are compared with the simulated composition (green diamonds). The effects of the X_{max} -dependent X_{max} bias correction are also shown (black triangles).

Fig. 5.19 shows the change in the estimated cross section for the fit with both a free and a fixed X_{max} scale when the X_{max} -dependent X_{max} bias is taken into account. Correcting for the X_{max} -dependent X_{max} bias significantly reduces the deviation in the estimated cross section, resulting in a 10 % smaller cross section at 10^{19} eV, corresponding to the rescaling factor value of $f_{19}^{\text{pp}} = 1.19 \pm 0.01$, compared to the fit without bias correction. Prior to the correction, the rescaling factor has a value of $f_{19}^{\text{pp}} = 1.29 \pm 0.009$. The X_{max} scale shift estimated from the fit with the bias correction becomes larger and is $\delta X_{\text{max}} = -13.8 \pm 0.54$ g/cm². In the case of a fixed X_{max} scale, the remaining deviation from the simulated cross section after correction for the X_{max} -dependent X_{max} bias is even smaller and amounts to less than 5 %.

5. Combined mass composition & cross section fit

fit, the fractions of the particle species are not sensitive to the bias correction (see Fig. 5.20). Note that the studies of the X_{\max} -dependent X_{\max} bias performed in this work are only valid for the Pierre Auger Observatory data presented at ICRC23, as different datasets use different reconstructions, and quantitatively, there will be differences in the significance of the bias.

'data'	f_{pp}^{19} , %	δX_{\max} , g/cm ²
Gumbel distr. as a func. of f_{pp}^{19}	+4	-2.6
Original Gumbel distr. (no f_{pp}^{19} dependence)	+11	-2.9
CONEX air shower simulations	+24	-2.9
End-to-end Monte Carlo sim. (with CONEX)	+19	-13.8
QGSJETII-04 Gumbel distr.	+7	≈ -9

Table 5.1: The sources of bias in the combined mass composition and cross section fit for the rescaling factor f_{19}^{pp} and shift in the X_{\max} scale δX_{\max} , in the simulated data we assume the composition derived in Sec 3.2, $\delta X_{\max} = 0$ and the number of events as in the ICRC23 data.

In summary, we can conclude that the combined mass composition and cross section method works well. Using data generated with the Gumbel function parameterization, we have verified that for a single energy, the estimator remains unbiased for any He fraction. We also did not get any significant deviation from the input parameters, assuming the composition mix observed in data. Note that we may have several minima in a fit for extreme cases of a pure composition due to some existing degeneracy between the f_{pp}^{19} and fractions. If this is the case, we choose the cross section that is closest to the original Sibyll 2.3d value, as there is no physical justification for the estimated value to be twice as large as the current extrapolations from the LHC measurements. The fit also recovers the input parameters well, with a bias of less than 10%, when estimating a single rescaling factor across all energies, with some minor dependence on the assumed composition. Suppose the X_{\max} scale is left free. In that case, there is a slightly larger deviation from the simulated value, but only within a few percent, which could also be related to the accuracy of finding the minimum in the 2D case. We also ran a test with a different interaction model as input, and the method allowed us to estimate both the cross section and a shift in the X_{\max} scale, as well as the composition, quite accurately. The test with QGSJETII-04 as input also showed that leaving the X_{\max} scale gives a better estimate of the proton-proton interaction cross section and mass composition. However, it generally leads to a larger statistical uncertainty in the results.

The largest source of bias, leading to an increase in the rescaling factor of up to 20% and thus in the estimated cross section with a deviation ranging from nearly 8% at the lowest energies to 24% at the highest energy, is the parameterization of the Gumbel distribution as a function of the rescaling factor f_{19}^{PP} . This bias is evident in tests on data simulated with a standard Gumbel parameterization, CONEX simulations, and full Monte Carlo simulations. All sources of bias, except for the X_{max} -dependent X_{max} bias, as the correction for it is implemented in the fit itself, are listed in Tab. 5.1.

5.5 Results

In this section, we apply the combined mass composition and cross section fit to the data of the Pierre Auger Observatory and discuss the results obtained. We implement the correction for the X_{max} -dependent X_{max} bias as described above for the most recent data. We also implement the correction for bias seen in the full Monte Carlo simulations for the final results by subtracting half of it from the estimated fit parameters.

As we leave the X_{max} scale free, we no longer include the X_{max} -related systematic uncertainty, as we did for the fixed X_{max} scale. However, since we consider a constant shift in the X_{max} scale across all energies, we estimate the energy-dependent X_{max} systematics instead. We repeat the fitting process 100 times using the data with varying shifts in the X_{max} scale at each energy obtained from the covariance matrix of the X_{max} systematics. We then estimate the best-fit f_{19}^{PP} , δX_{max} , and composition for each iteration. In such a way, we obtain the distribution of the estimated variables, the 16% and 84% quantiles of which can then be quoted as the corresponding systematic bounds. Another source of uncertainty in the fit results is the energy scale systematic uncertainty, for which we adopt the 14% estimated in [133]. We refit the data twice to account for this, shifting the energies up and down accordingly. In addition, we include the systematic uncertainty from the detector acceptance and resolution parameterizations by taking the combinations of uncertainties for each parameter that give the extreme acceptance values and the upper and lower 1- σ bounds on the resolution.

Fig. 5.21 compares the results of the combined mass composition and cross sectional fit for several different data: PRD14 data, ICRC19 data, ICRC19 data with the same reconstruction as in PRD14 (denoted GH in the plot), and ICRC23 data. The results are shown for the free X_{max} scale, and the corresponding best-fit shifts are listed in Table 5.2. As can be seen, there is quite a difference between

5. Combined mass composition & cross section fit

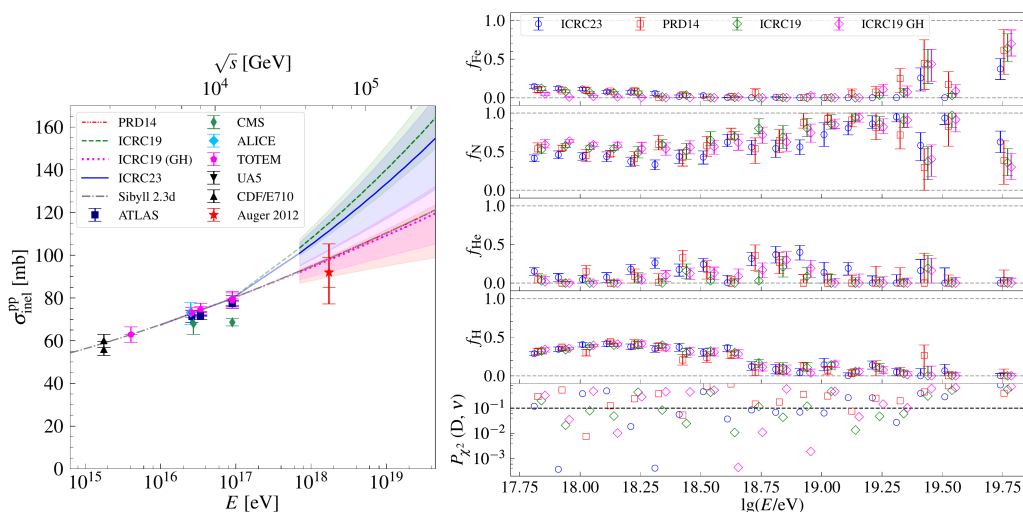


Figure 5.21: The comparison of the proton-proton inelastic cross section $\sigma_{\text{inel}}^{\text{pp}}$ (left) and the composition fractions (right) from the combined mass composition and cross section fits to the different datasets of the Pierre Auger Observatory. The shaded area corresponds to the statistical uncertainty on the estimated cross section.

	PRD14	ICRC19	ICRC19 GH	ICRC23
f_{19}^{pp}	$1.0^{+0.02}_{-0.13}$	$1.28^{+0.09}_{-0.09}$	$1.0^{+0.08}_{-0.01}$	$1.21^{+0.11}_{-0.15}$
$\delta X_{\text{max}}, \text{g/cm}^2$	$-8.7^{+1.64}_{-0.24}$	$-13.7^{+0.81}_{-1.00}$	$-9.8^{+0.8}_{-0.8}$	$-11.54^{+0.53}_{-1.08}$

Table 5.2: The rescaling factor f_{19}^{pp} and the shift in the X_{max} scale δX_{max} measured from the different datasets of the Pierre Auger Observatory. Statistical uncertainties are shown.

the results. The highest cross section is observed in the ICRC19 data, while both the PRD14 data and the ICRC19 + GH data, which is reconstructed in the same way as the PRD14 data, agree with the Sibyll 2.3d extrapolations of the LHC measurements. Note that neither fit has a correction for the X_{max} -dependent X_{max} bias as we can estimate it only for the most recent data. The cross sections estimated with a combined fit from the older PRD14 data are in reasonably good agreement within the uncertainty with the result of the fit to the tail of the X_{max} distribution reported by the Pierre Auger Collaboration in 2012.

In Fig. 5.22, the proton-proton cross section estimated from the ICRC23 data is compared with two scenarios: one using a default fixed X_{max} scale and the other

5. Combined mass composition & cross section fit




Figure 5.22: The comparison of the proton-proton cross section σ^{PP} estimated from the combined mass composition and cross section fit to the ICRC23 data with the default X_{max} scale and a free X_{max} scale. Additionally, the effect of the correction for the X_{max} -dependent X_{max} bias is considered. After the bias correction, the results are the same for fit with $\delta X_{\text{max}} = 0$ and for free δX_{max} , so only the latter is shown. The shaded area corresponds to the statistical uncertainty on the estimated cross section.

incorporating a fitted shift in X_{max} . The effect of the bias correction is also shown for each case. From the biased data, we obtain a significant deviation of more than 20% at 10^{19} eV in the estimated cross section from any hadronic interaction model for both fixed and fitted X_{max} scales, with the rescaling factor values of $1.37^{+0.04}_{-0.16}$ and $1.22^{+0.10}_{-0.15}$ accordingly. In the latter case, the best-fit X_{max} scale is shifted by $-11.5^{+0.54}_{-0.10}$ g/cm² compared to the prediction from the original Sibyll 2.3d interaction model. After correcting for the X_{max} -dependent X_{max} bias, the fitted f_{19}^{PP} decreases, leading to a smaller interaction cross section. The effect of the bias

5. Combined mass composition & cross section fit




Figure 5.23: The comparison of the composition fractions estimated from the combined mass composition and cross section fit to the ICRC23 data with the default X_{\max} scale and a free X_{\max} scale. Additionally, the effect of the correction for the X_{\max} -dependent X_{\max} bias is considered. After the bias correction, the results are the same for fit with $\delta X_{\max} = 0$ and for a free δX_{\max} , therefore displaying only the latter.

correction is greater for the fit with a fixed X_{\max} scale, which gives the same cross section as the fit with a free X_{\max} scale after the bias correction, corresponding to a rescaling factor of $1.16^{+0.08}_{-0.09}$. The estimated shift in the X_{\max} scale from the fit with correction for the X_{\max} -dependent X_{\max} bias is $\delta X_{\max} = -15.8^{+0.5}_{-0.6} \text{ g/cm}^2$.

The corresponding mass composition is shown in Fig. 5.23. The fitted X_{\max} distributions can be found in Appendix C. As discussed above, the presence of the X_{\max} -dependent X_{\max} bias has a negligible effect on the bulk of the X_{\max} distribution, and, therefore, the fitted composition is nearly insensitive to it, with no change in the fit quality. The choice between keeping the X_{\max} scale fixed or let-

5. Combined mass composition & cross section fit




Figure 5.24: The combined mass composition and cross section fit to the data with an energy-dependent shift in X_{\max} scale. Left: the estimated rescaling factor value from the fit with a fixed (in green) and free (in blue) X_{\max} scale. Right: the estimated shift in X_{\max} scale δX_{\max} . The median of the obtained distribution is shown as a solid line, and the one standard deviation interval is shown in the dashed line.

ting it free does matter, however, in particular at energies above $10^{18.6}$ eV. The effect is the same as when fitting to simulated data, resulting in a predicted lighter composition over the whole energy range. At the lowest energies, the changes in composition are small, remaining below 5% for He, N, and Fe nuclei and below 10% for protons. The largest difference in the estimated fractions is present in the energy range between an ankle and the end of the cosmic-ray spectrum. While the fit with a free X_{\max} scale favors the nitrogen-dominated composition, the fit with a fixed X_{\max} scale results in an almost pure He composition between $10^{18.6}$ and $10^{19.1}$ eV. In addition, at both the lowest and highest energies, leaving the X_{\max} scale free results in an increased contribution from iron.

As discussed for the fit performance tests using the Monte Carlo simulations, leaving the X_{\max} scale free is beneficial for recovering the X_{\max} scale if it is not equal to the predictions from the existing hadronic interaction models. Furthermore, it is apparent from analyzing the fits with the energy-dependent shift in X_{\max} data (see Fig. 5.24) that releasing the X_{\max} scale results in a sizable reduction of the uncertainty on the estimated rescaling factor from nearly 10% to 2%.

As a final step, we correct the obtained cross section for the bias observed in the Monte Carlo simulations. It is not entirely straightforward how to proceed with this bias correction since the extent of the bias depends on the composition of the data and the presence of the X_{\max} shift, which are both measured quantities. We subtract the full bias seen in the end-to-end simulation of the measurement from the obtained results, thus reducing the estimated proton-proton interaction cross section by 19% at 10^{19} eV. The rescaling factor corrected for the X_{\max} -dependent X_{\max} bias will then be equal to $f_{\text{pp}}^{19} = 0.96$.

For the final results, all sources of systematic uncertainty listed at the beginning of this section are considered as well as the uncertainty of the correction for the Monte Carlo bias, and the systematic uncertainty in the the Glauber calculations, taken from [178]. The systematic uncertainty additionally includes the bias obtained by fitting the simulated QGSJETII-04 data with the modified Sibyll 2.3d model predictions. The simulated data for the QGSJETII-04 test was obtained using the Gumbel function. Therefore, the uncertainty component was corrected for the bias observed in the fit to the original Gumbel function parameterization with the modified one. We also incorporate the uncertainty of the end-to-end Monte Carlo bias correction as the uncertainty of the estimated rescaling factor obtained from fitting the simulated data. Table 5.3 summarizes the sources of systematic uncertainty. The parameterization of detector acceptance is the most significant factor contributing to the systematic uncertainty in the cross section estimation. Another significant contribution comes from the uncertainty in the Glauber calculations. On the other hand, the energy scale uncertainty does not impact the rescaling factor value but remains the primary source of uncertainty in the estimation of X_{\max} scale. Fig. 5.25 shows a comparison of the interaction cross section estimated in this work to the model predictions, accelerator data and the previous measurement of the Pierre Auger Observatory. The fitted rescaling factor value is

$$f_{19}^{\text{pp}} = 0.96_{-0.08}^{+0.09}(\text{stat})_{-0.15}^{+0.10}(\text{sys}) . \quad (5.5)$$

After correcting for the X_{\max} -dependent X_{\max} bias and the bias observed in the Monte Carlo simulations, the measured cross section is in agreement, within the uncertainty, with all of the hadronic model extrapolations and the previous measurement from the cosmic-ray data. Most of the uncertainty in the obtained results is due to statistical uncertainty from the fit itself. A contribution from systematic effects, such as the uncertainty in energy estimation or detector parameterization, is minor. The estimated uncertainty permits substantial variation in the development of the proton-proton cross section with energy. This includes a scenario in

5. Combined mass composition & cross section fit




Figure 5.25: The estimated inelastic cross section for proton-proton collisions estimated from the data of the Pierre Auger Observatory using the combined mass composition and cross section fit. The results are compared to predictions from the existing hadronic interaction models, accelerator data, and the previous measurement of the Pierre Auger Observatory. The darker shaded area corresponds to the statistical uncertainty from the fit, and the lighter area corresponds to the total uncertainty.

5. Combined mass composition & cross section fit




Figure 5.26: The estimated composition fractions from the Pierre Auger Observatory data using the combined mass composition and cross section fit (blue circles) compared to the results of the composition-only fit (green squares). Both the statistical uncertainties (inner cap) and the total uncertainties, including systematics, (outer cap) are shown.

which the cross section remains nearly constant over the energies, significantly lower than predicted by interaction models and an increase of up to 20% at the highest energies. The proton-proton interaction cross section resulting from the combined fit of mass composition and cross section is consistent with the measurements derived from the fit to the tail of the X_{\max} distribution discussed earlier.

The estimated shift in the X_{\max} scale with the statistical uncertainty from the fit is equal to $\delta X_{\max} = -2.58 \pm 0.54 \text{ g/cm}^2$. The systematic uncertainty of the X_{\max} scale comprises the same components as the cross section, except the Glauber cal-

5. Combined mass composition & cross section fit

culations. The most significant contribution comes from the fit with the energy-dependent X_{\max} systematic, equal to $\pm 10.0 \text{ g/cm}^2$. Each other systematic uncertainty factor contributes to only a couple of g/cm^2 . The total systematic uncertainty is $+12.6 \text{ g/cm}^2$ for the upper limit and -11.0 g/cm^2 for the lower limit on the cross section estimation. The result for the shift in the X_{\max} scale is:

$$\delta X_{\max} = -2.58 \pm 0.54(\text{stat}) \begin{matrix} +12.6 \\ -11.0 \end{matrix}(\text{sys}) \text{ g/cm}^2. \quad (5.6)$$

The estimated mass composition from the combined mass composition and cross section fit is heavier than the predictions from the composition-only fit, though both measurements agree within the total uncertainty, consisting of statistical and systematic uncertainty. The systematic uncertainty in both fits includes the corresponding uncertainty from the acceptance and resolution parameterization and an energy scale uncertainty. Additionally, the composition-only fit includes the X_{\max} scale uncertainty. Since it is the primary source of uncertainty in the standard composition estimation, the resulting total uncertainty exceeds the one for the fractions estimated from the combined mass composition and cross section fit. Just like for the measured proton-proton interaction cross section and the shift in the X_{\max} scale, the fitted fractions were adjusted for the bias observed in the Monte Carlo simulations fits. It should be noted that this correction is only approximate as it has some minor dependence on the composition, and the applied bias correction was obtained for the fractions in the composition-only fit, which results in a larger contribution from the lighter elements compared to the combined estimation. While the proton fraction remains unaffected and only a minor increase in the Fe fraction is observed, there is a significant difference between the

	$f_{19}^{\text{pp}}, \%$	$\delta X_{\max}, \text{g/cm}^2$
Monte Carlo correction	± 1.25	0.5
Energy scale	-3.1	$\begin{matrix} +6.4 \\ -3.8 \end{matrix}$
Detector acceptance	$\begin{matrix} +6.3 \\ -8.3 \end{matrix}$	± 0.5
Detector resolution	$\begin{matrix} +2.1 \\ -10.4 \end{matrix}$	± 0.5
Energy-dependent X_{\max} systematics	± 1.46	± 10
QGSJETII-04 sim. bias	-2.1	3
Glauber calculations	± 5.7	-
Quadratic sum	$\begin{matrix} +10.4 \\ -15.9 \end{matrix}$	$\begin{matrix} +12.6 \\ -11.0 \end{matrix}$

Table 5.3: The sources of the systematic uncertainty in the estimated rescaling factor and the shift of the X_{\max} scale.

5. Combined mass composition & cross section fit

predicted amount of H and He nuclei for the two fit approaches. The fraction of N nuclei obtained from the combined method constitutes nearly 50% of the mix up to 10^{19} eV and rises above it, resulting in a nearly pure N composition, while the standard fit predicts a more mixed composition at the same energies. The helium fraction remains below 40% throughout the entire energy range, reaching its maximum at around $10^{18.7}$ - 10^{19} eV and is almost zero at the lowest and highest energies. In nearly all energy bins, the quality of the combined fit is better. At the energy of $10^{18.6}$ eV, neither fit performs satisfactorily due to the presence of the event with $X_{\max} > 1200$ g/cm² in the data.

Thus, summarizing what is being discussed, we see that the proton-proton interaction cross section estimated with the combined mass composition and cross section approach introduced in this work agree with the extrapolations in the original Sibyll 2.3d hadronic interaction models as well as with the previous measurement of the Pierre Auger Observatory. The measured shift in the X_{\max} scale is small, negative, and compatible with the default X_{\max} scale within the uncertainty limits. The composition predicted from the combined mass composition and cross section fit is heavier than that measured with the original Sibyll 2.3d model with the default interaction cross sections. The fraction of nitrogen and iron nuclei has increased, reducing the helium fraction, while the contribution from protons has remained unchanged.

5. Combined mass composition & cross section fit

Summary and Outlook

This thesis is dedicated to the combined studies of the cosmic-ray mass composition and hadronic interaction cross sections at ultrahigh energies using the depth of the maximum air shower development measured by the Fluorescence Detector of The Pierre Auger Observatory. A novel strategy has been proposed for the concurrent evaluation of the mass composition and proton-proton interaction cross section, free from the assumptions required in standard independent analyses, such as the validity of the existing hadronic interaction models and a specific proton-dominated composition in the measured data. To perform such a fit, the Sibyll 2.3d hadronic interaction model was modified to obtain self-consistent simulations of nucleus-induced air showers for different values of the inelastic proton-proton interaction cross section. The conversion from the modified proton-proton to the corresponding nucleus-nucleus cross sections is done via the Glauber theory. Using air shower simulations, it is possible to generate model predictions corresponding to the modified interaction properties, which can be used in the standard mass composition fit, allowing one to obtain a mass composition estimate for a varied proton-proton cross section. In addition, the X_{\max} scale was set free to account for both systematic uncertainties of the data and theoretical uncertainties of the properties of particle production in air showers.

The performance of the proposed method was evaluated, and no significant bias in the measured variables was detected, irrespective of the assumed composition mix, contrary to the standard cross-section analysis. It was confirmed that allowing the X_{\max} scale to remain unconstrained leads to enhanced performance if the data has a distinct X_{\max} scale compared to the predictions of the employed interaction model. However, it was discovered that the results of the combined mass composition and cross section fit, as well as the separate cross section analysis, are significantly affected by reconstruction biases. Consequently, the corresponding corrections were derived and applied. The studies with end-to-end Monte Carlo simulations revealed the existence of another additional bias, most likely

5. Summary and Outlook

attributed to the parameterization of the Gumbel function depending on the cross section rescaling factor, and the reported values were corrected for it as well.

The combined mass composition and cross section fit was applied to the full Phase I data of the Pierre Auger Observatory. The outcome of the fit was compared to the predictions of the existing hadronic interaction models, the results of the separate analyses, and the previous measurements from the cosmic-ray data of the Pierre Auger Observatory. The measured cross section for proton-proton interactions, corresponding to the rescaling factor of $f_{19}^{\text{pp}} = 0.96^{+0.13}_{-0.17}$, agrees well with the existing hadronic interaction models and previous analyses. With this analysis, we could decrease the statistical uncertainty of the previous measurement by 20% using the full Phase I dataset. The method developed in this thesis can improve the systematic uncertainty by a factor of 1.3 and 1.8 for the lower and upper uncertainty bands, respectively, with a degree of improvement depending on the assumed helium contamination in the standard analysis. Note that although the measured uncertainty is smaller than previously, it still covers a broad range of cross section modifications, including the logarithmic evolution in cross section with energies and a nearly constant value over the considered energy range.

The estimated shift in the X_{max} scale is compatible with a fixed X_{max} scale but also covers a wide range between -13 g/cm^2 and 10 g/cm^2 compatible with the experimental uncertainty of the X_{max} scale. The composition obtained from the combined approach is heavier than the predictions of the original non-modified Sibyll 2.3d model, with the N fraction reaching more than 80% at energies above 10^{19} eV and a corresponding decrease in He fraction. The proton and iron components of the mix remain stable relative to the cross section modifications.

In addition, separate cross section and mass composition analyses were performed on the full Phase I data, providing an update of the results of the individual analyses with improved methodology and statistics. A Bayesian Markov Chain Monte Carlo approach was used for the mass composition estimation in this thesis. It allows a more straightforward estimation of the uncertainty on the fitted variables and easier marginalization over any quantity derived from the composition fractions, such as cosmic-ray rigidity and upper limits on the proton and photon fractions.

Further studies and improvements in the method could be beneficial for a better understanding of uncertainties in the obtained results and the robustness of the underlying assumptions. For example, as the method relies on Glauber calculations for the conversion between the proton-proton and proton-air cross section, one of the questions of interest is an accurate estimation of systematic uncertainty related to the Glauber formalism. Another issue requiring a more detailed study

could be better quantifying the biases seen in the end-to-end Monte Carlo simulations. Further work of interest could also include studies on the change in the functional dependence of the function used for rescaling the interaction cross section to decrease the related uncertainties that dominates the current systematic uncertainties.

The combined mass composition and cross section fit presented in this thesis provides a new approach for the studies of the cosmic-ray mass composition and the properties of the hadronic interactions. The method shows good performance, and the results obtained from the data fit could be considered reliable.

5. Summary and Outlook

List of Figures

1.1	The cosmic-ray energy spectrum	13
1.2	Average lateral and longitudinal profiles of the vertical air showers	16
1.3	Schematic view of the development of the air shower cascade	18
1.4	Decomposition of the observed X_{\max} distribution	21
1.5	The evolution of the proton-air and proton-proton production cross section with energy	26
1.6	Comparison of model predictions for the moments of the X_{\max} distribution	28
1.7	The effect of the modified hadronic interaction properties on air shower observables for the proton-initiated cascade	32
1.8	The effect of the modified hadronic interaction properties on air shower observables for the iron-initiated cascade	33
2.1	Layout of the Pierre Auger Observatory	36
2.2	The water Cherenkov detector station of the surface detector	39
2.3	The FD site of the Pierre Auger Observatory	40
2.4	The patterns for cosmic-ray tracks identification in FD	42
2.5	Example of a reconstructed shower profile	45
2.6	An example of the longitudinal air shower profile for iron and proton primary	46
2.7	The evolution of the first moments of the X_{\max} distribution as a function of energy from measurements with the fluorescence and surface detectors of the Pierre Auger Observatory and Telescope Array	47
2.8	Resolution and systematic uncertainty of the X_{\max} measurement with Fluorescence Detector of the Pierre Auger Observatory	49
3.1	The detector acceptance as a function of logarithmic energy	60

LIST OF FIGURES

3.2	Example of the MCMC posterior samples for the four-component fit to the PRD14 data	63
3.3	Average fitted fractions for three composition scenarios as a function of the number of detected events	66
3.4	Examples of posterior distributions resulting in biased point estimates for the tests with the simulated data	67
3.5	2–D correlations between the parameters and the distributions of fitted parameters	69
3.6	Coverage probabilities for three different composition scenarios	70
3.7	Comparison of fractions and the quality of the fit for the different binning of the X_{\max} distribution	72
3.8	Comparison of the composition fit with CONEX X_{\max} and Gumbel X_{\max} distribution templates	74
3.9	Comparison of the fractions derived with MCMC to the previous Auger measurements	76
3.10	The first two moments of the X_{\max} distribution from the fit to the PRD14 data	77
3.11	Comparison of the estimated composition fractions for the PRD14, ICRC19 and ICRC23 datasets	79
3.12	Mass composition fit with additional nuclei in the mix (1)	80
3.13	Mass composition fit with additional nuclei in the mix (2)	81
3.14	The 4-component MCMC fit with a shift in the data X_{\max} scale as an additional fit parameter	82
3.15	Mass composition fit to the ICRC23 data with a shift in the X_{\max} scale using the Sibyll 2.3d model	83
3.16	The MCMC fit to the ICRC23 data	85
3.17	The first two moments of the X_{\max} distribution from the fit to the ICRC23 data	86
3.18	The upper limits for the proton and photon fractions	88
3.19	Average cosmic-ray rigidity and its inverse	89
3.20	The comparison of the composition fractions from on- and off-plane regions	90
3.21	X_{\max} -dependent X_{\max} bias: examples of the difference between the generated and reconstructed X_{\max}	91
3.22	X_{\max} -dependent X_{\max} bias: mean and standard deviation of the difference between the generated and reconstructed X_{\max} scale vs energy	92

3.23	X_{\max} -dependent X_{\max} bias: mean and standard deviation of the difference between the generated and reconstructed X_{\max} scale vs energy for each particle species.	93
3.24	X_{\max} -dependent X_{\max} bias parameterization as a function of energy for the $\mu(X_{\max})$	94
3.25	X_{\max} -dependent X_{\max} bias parameterization as a function of energy for the $\sigma(X_{\max})$	95
3.26	The X_{\max} distribution templates with and without X_{\max} -dependent X_{\max} bias correction	96
3.27	The comparison between the fractions fitted with and without the correction for the X_{\max} -dependent X_{\max} -bias	97
4.1	The fit to the tail of the X_{\max} distribution and $\sigma_{p-\text{air}}$ from ICRC2015	100
4.2	An example of the unbinned likelihood fit with a broken exponential profile	101
4.3	Bias correction for the λ_{η} fit	102
4.4	Comparison of λ_{η} measurement with and without bias correction .	104
4.5	The measured λ_{η} from the ICRC23 data with xsecFoV and comp-FoV selections	105
4.6	Proton-proton interaction cross section from the fit to the tail of the X_{\max} distribution	106
4.7	The fit to the tail of the X_{\max} distribution and $\sigma_{p-\text{air}}$ from ICRC2015	107
5.1	The effect of the rescaling in σ^{PP} on $\sigma^{A-\text{air}}$ and the relation between them	111
5.2	The algorithm for generating the X_{\max} distribution for the modified cross-section	112
5.3	Proton X_{\max} distributions for modified σ^{PP}	113
5.4	The effect of the rescaling in σ^{PP} on the first two moments of the X_{\max} distribution and the depth of the first interaction point X_1 . .	114
5.5	An example of the combined mass composition and cross section fit	116
5.6	Gumbel function shape parameters vs f_{19}^{PP}	117
5.7	Comparison of the model extrapolations for σ^{PP} with different values of f_{19}^{PP}	118
5.8	The variation in the measured mass composition depending on the rescaling factor f_{19}^{PP}	120
5.9	The variation in the moments of the X_{\max} distribution and tail fit λ_{η} depending on the rescaling factor f_{19}^{PP}	121

LIST OF FIGURES

5.10	The χ^2 contour of the combined fit for the varied He fraction in the mix	122
5.11	The fitted X_{\max} distribution for the combined fit with an equal H-He mix in the simulated data	123
5.12	The performance of the combined fit for the arbitrary He fraction for one energy bin	124
5.13	The performance of the combined fit with δX_{\max} for one energy bin	126
5.14	The performance of the combined fit for several energy bins	127
5.15	The estimated proton-proton cross sections obtained by fitting the data realizations simulated with different composition scenarios	128
5.16	The estimated fractions of particle species obtained by fitting the data realizations simulated with different composition scenarios	130
5.17	The estimated fractions of particle species obtained by fitting the data realizations simulated with different composition scenarios	131
5.18	The results of the fit to the data simulated with QGSJETII-04 interaction model	132
5.19	The σ^{PP} from the fit to the data simulated with Sibyll 2.3d interaction model including detector response	133
5.20	The fractions fitted to the data simulated with Sibyll 2.3d interaction model including detector response	134
5.21	The comparison of the σ^{PP} and composition fractions from the different data	137
5.22	The σ^{PP} from the fit to the ICRC23 data: the comparison between a fixed and free X_{\max} scale and the correction for X_{\max} bias.	138
5.23	The composition fractions from the fit to the ICRC23 data: the comparison between a fixed and free X_{\max} scale and the correction for X_{\max} bias	139
5.24	The combined mass composition and cross section fit to the data with 100 realizations of the energy-dependent shift in X_{\max} scale	140
5.25	The estimated proton-proton cross section from Pierre Auger Observatory data using the combined mass composition and cross section fit	142
5.26	The estimated composition fractions from the Pierre Auger Observatory data using the combined mass composition and cross section fit	143
B.1	X_{\max} distributions for the fit to the ICRC23 data with the Sibyll 2.3d interaction model	176

B.2	X_{\max} distributions for the fit to the ICRC23 data with the EPOS-LHC interaction model	182
B.3	Comparison of the estimated composition fractions with two versions of the Sibyll MC generator, Sibyll 2.3c and Sibyll 2.3d, for ICRC19 data	183
B.4	Mass composition fit to the ICRC2019 data with a shift in the X_{\max} scale and Si nuclei in the fractions mix	184
C.1	X_{\max} distributions for the fit to the ICRC23 data with the modified Sibyll 2.3d interaction model (best-fit cross section)	189
C.2	The dependence of the estimated cross section on the threshold energy, E_{tr}	190

LIST OF FIGURES

List of Tables

5.1	The sources of bias in the combined mass composition and cross section fit	135
5.2	The f_{19}^{pp} and δX_{max} measured from the different data	137
5.3	The sources of systematic uncertainty in the estimated rescaling factor and the shift of the X_{max} scale	144
A.1	Parametrization of the $\mu(X_{\text{max}}^{\text{gen}} - X_{\text{max}}^{\text{rec}})$ as a function for the compFoV event selection	170
A.2	Parametrization of the $\sigma(X_{\text{max}}^{\text{gen}} - X_{\text{max}}^{\text{rec}})$ as a function of energy for the compFoV event selection	170
A.3	Parametrization of the $\mu(X_{\text{max}}^{\text{gen}} - X_{\text{max}}^{\text{rec}})$ as a function of energy for the xsecFoV event selection	170
A.4	Parametrization of the $\sigma(X_{\text{max}}^{\text{gen}} - X_{\text{max}}^{\text{rec}})$ as a function of energy for the xsecFoV event selection	170
B.1	Fractions from the MCMC fit to the ICRC23 data with the Sibyll 2.3d interaction model	171
B.2	Fractions from the MCMC fit to the ICRC23 data with the EPOS-LHC interaction model	177

LIST OF TABLES

Bibliography

- [1] I. A. Grenier, J. H. Black, and A. W. Strong *Annu. Rev. Astron. Astrophys.* **53** (2015) 199.
- [2] E. G. Zweibel *Phys. of Plasmas* **24** (2017) 055402.
- [3] V. F. Hess *Phys. Z.* **13** (1912) 1084.
- [4] M. Kachelrieß and D. V. Semikoz *Prog. Part. Nucl. Phys.* **109** (2019) 103710.
- [5] G. Kulikov and G. B. Khristiansen *JETP* **35** (1959) 441.
- [6] T. Antoni *et al.*, [KASCADE-Grande Collab.] *Astropart. Phys.* **24** (2005) 1.
- [7] K.-H. Kampert *NPB(PS)* **165** (2007) 294.
- [8] B. Peters *Il Nuovo Cimento* **22** (1961) 800.
- [9] A. D. Erlykin and A. W. Wolfendale *Adv. Space Res.* **27** (2001) 803.
- [10] A. D. Erlykin and S. K. Machavariani *Bull. Lebedev Phys. Inst.* **48** (2021) 31.
- [11] S. I. Nikolsky *NPB(PS)* **39** (1995) 228.
- [12] P. O. Lagage and C. J. Cesarsky *A&A* **125** (1983) 249.
- [13] G. Giacinti, M. Kachelrieß, and D. V. Semikoz *Phys. Rev. D* **90** (2014) 041302.
- [14] H. J. Volk and V. N. Zirakashvili *Astron. Astrophys.* **417** (2004) 807.

BIBLIOGRAPHY

- [15] C. Evoli *The Cosmic-Ray Energy Spectrum* <https://doi.org/10.5281/zenodo.4396125> (2020) . Zenodo 2020.
- [16] M. G. Aartsen *et al.*, [IceCube Collab.] *Phys. Rev. D* **88** (2013) 042004.
- [17] N. Budnev *et al.* *Astropart. Phys.* **117** (2020) 102406.
- [18] W. Apel *et al.*, [KASCADE-Grande Collab.] *Astropart. Phys.* **36** (2012) 183.
- [19] P. Abreu *et al.*, [Pierre Auger Collab.] *EPJ C* **81** (2021) 966.
- [20] M. Nagano *et al.* *J Phys G.* **18** (1992) 423.
- [21] D. Bird *et al.* *ApJ* **424** (1994) 491.
- [22] T. Abu-Zayyad *et al.*, [HiRes-MIA Collab.] *Astropart. Phys.* **16** (2001) 1.
- [23] R. U. Abbasi *et al.*, [Telescope Array Collab.] *ApJ* **865** (2018) 74.
- [24] T. Abu-Zayyad *et al.* [arXiv:1803.07052](https://arxiv.org/abs/1803.07052) [astro-ph.HE].
- [25] W. D. Apel *et al.*, [KASCADE-Grande Collab.] *Phys. Rev. Lett.* **107** (2011) 171104.
- [26] J. Linsley *Conf. Proc. ICRC1963* (1963) 77.
- [27] D. J. Bird *et al.* *Phys. Rev. Lett.* **71** (1993) 3401.
- [28] J. Abraham *et al.*, [Pierre Auger Collab.] *Phys. Lett. B* **685** (2010) 239.
- [29] A. Aab *et al.*, [Pierre Auger Collab.] *Phys. Rev. Lett.* **125** (2020) 121106.
- [30] D. D. Marco and T. Stanev *Phys. Rev. D* **72** (2005) 081301.
- [31] E. Parizot *NPB(PS)* **256-257** (2014) 197.
- [32] A. M. Hillas *J. Phys. G* **31** (2005) R95.
- [33] R. U. Abbasi *et al.*, [HiRes Collab.] *ApJ* **622** (2005) 910.
- [34] A. Aab *et al.*, [Pierre Auger Collab.] *Phys. Rev. D* **90** (2014) 122005.
- [35] V. Berezhinsky, A. Gazizov, and S. Grigorieva *Phys. Rev. D* **74** (2006) 043005.

- [36] R. Aloisio, V. Berezhinsky, and A. Gazizov *Astropart. Phys.* **39** (2012) 129.
- [37] D. Allard, E. Parizot, and A. V. Olinto *Astropart. Phys.* **27** (2007) 61.
- [38] D. Allard, A. V. Olinto, and E. Parizot *A&A* **473** (2007) 59.
- [39] C. Blaksley and E. Parizot *Astropart. Phys.* **35** (2012) 342.
- [40] R. U. Abbasi *et al.*, [HiRes Collab.] *Phys. Rev. Lett.* **92** (2004) 151101.
- [41] N. N. Efimov *et al.* *Proc. Conf. ICCR International Symposium on the Astrophysical Aspects of the Most Energetic Cosmic Rays* (1990) .
- [42] J. Abraham *et al.*, [Pierre Auger Collab.] *Phys. Rev. Lett.* **101** (2008) 061101.
- [43] K. Greisen *Phys. Rev. Lett.* **16** (Apr, 1966) 748.
- [44] G. T. Zatsepin and V. A. Kuz'min *JETP Lett.* **4** (1966) 78.
- [45] J. W. Cronin *NPB(PS)* **97** (2001) 3.
- [46] A. A. Watson *Rep. Prog. Phys* **77** (2014) 036901.
- [47] R. A. Batista *et al.* *Front. astron. space sci.* **6** (2019) 23.
- [48] R. Engel, D. Heck, and T. Pierog *Annu. Rev. Nucl. Part. Sci.* **61** (2011) 467.
- [49] W. Heitler, *The Quantum Theory of Radiation*. Dover Books on Physics Series. Dover Publications, 1984.
- [50] J. F. Carlson and J. R. Oppenheimer *Phys. Rev.* **51** (1937) 220.
- [51] J. Matthews *Astropart. Phys.* **22** (2005) 387.
- [52] B. Rossi and K. Greisen *Rev. Mod. Phys.* **13** (1941) 240.
- [53] J. Alvarez-Muniz *et al.* *Phys. Rev. D* **66** (2002) 033011.
- [54] R. M. Ulrich, *Measurement of the proton air cross section using hybrid data of the Pierre Auger Observatory*. PhD dissertation. Univ. Karlsruhe, 2008.

BIBLIOGRAPHY

- [55] R. Ulrich *et al.* *New J. Phys* **11** no. 6, (2009) 065018.
- [56] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner *Phys. Rev. C* **92** (2015) 034906.
- [57] S. Ostapchenko *Phys. Rev. D* **83** (2011) 014018.
- [58] R. S. Fletcher, T. K. Gaisser, P. Lipari, and T. Stanev *Phys. Rev. D* **50** (1994) 5710.
- [59] F. Riehn *et al.* *Phys. Rev. D* **102** (2020) 063002.
- [60] D. Heck *et al.*, *CORSIKA: a Monte Carlo code to simulate extensive air showers*. Forschungszentrum Karlsruhe Report FZKA 6019. 1998.
- [61] T. Bergmann *et al.* *Astropart. Phys.* **26** (2007) 420.
- [62] A. Ferrari *et al.*, *FLUKA: A multi-particle transport code (Program version 2005)*. Cern, 2005.
- [63] M. Bleicher *et al.* *J Phys G* **25** (1999) 1859.
- [64] V. N. Gribov *JETP* **53** (1967) 654.
- [65] L. Calcagni, C. A. G. Canal, S. J. Sciutto, and T. Tarutina *Phys. Rev. D* **98** (2018) 083003.
- [66] J. Albrecht *et al.* *Astrophys. Space Sci.* **367** (2022) 27.
- [67] S. Ostapchenko *EPJ Web of Conf.* **120** (2016) 04003.
- [68] T. Pierog and K. Werner *NPB(PS)* **196** (2009) 102.
- [69] T. Pierog, B. Guiot, I. Karpenko, and K. Werner *Proc. Conf. 7th International Workshop on Multiple Partonic Interactions at the LHC* (2016) 31.
- [70] S. Ostapchenko *NPB(PS)* **151** (2006) 143.
- [71] S. Roesler, R. Engel, and J. Ranft in *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications*, p. 1033. 2001.
- [72] L. Durand and P. Hong *Phys. Rev. Lett.* **58** (1987) 303.

- [73] E.-J. Ahn *et al.* *Phys. Rev. D* **80** (2009) 094003.
- [74] T. Pierog *PoS ICRC2017* (2018) 1100.
- [75] N. A. Amos *et al.*, [E710 Collab.] *Phys. Rev. Lett.* **68** (1992) 2433.
- [76] F. Abe *et al.*, [CDF Collab.] *Phys. Rev. D* **50** (1994) 2966.
- [77] G. Antchev *et al.*, [TOTEM Collab.] *EPJ C* **80** no. 2, (2020) 91.
- [78] B. Abelev *et al.*, [ALICE Collab.] *EPJ C* **73** no. 6, (2013) 2456.
- [79] M. S. A *et al.*, [CMS Collab.] *JHEP* **07** (2018) 161.
- [80] M. Aaboud *et al.*, [ATLAS Collab.] *Phys. Rev. Lett.* **117** no. 18, (2016) 182002.
- [81] R. Glauber and G. Matthiae *Nuc. Phys. B* **21** (1970) 135.
- [82] A. Aab *et al.*, [Pierre Auger Collab.] *Phys. Rev. D* **90** (2014) 122006.
- [83] A. Yushkov *et al.*, [Pierre Auger, Telescope Array Collab.] *EPJ Web Conf.* **210** (2019) 01009.
- [84] F. Riehn *et al.* *Phys. Rev. D* **102** (2020) 063002.
- [85] A. Aab *et al.*, [Pierre Auger Collab.] *Phys. Rev. D* **91** (2015) 032003.
- [86] A. Aab *et al.*, [Pierre Auger Collab.] *Phys. Rev. Lett.* **117** (2016) 192001.
- [87] R. U. Abbasi *et al.*, [Telescope Array Collab.] *Phys. Rev. D* **98** (Jul, 2018) 022002.
- [88] R. U. Abbasi, [IceCube Collab.] *Phys. Rev. D* **106** (2022) 032010.
- [89] Y. A. Fomin *et al.* *Astropart. Phys.* **92** (2017) 1.
- [90] A. V. Glushkov, K. G. Lebedev, and A. V. Saburov *JETP Lett.* **117** (2023) 257.
- [91] G. R. Farrar and J. D. Allen *EPJ Web of Conf.* **53** (2013) 07007.
- [92] T. Pierog, M. Perlin, and K. Werner *EPJ Web of Conf.* **283** (2023) 05001.

BIBLIOGRAPHY

- [93] T. Pierog *et al.* *PoS ICRC2019* (2020) 387.
- [94] J. Adam *et al.*, [ALICE Collab.] *Nat. Phys.* **13** (2017) 535.
- [95] L. A. Anchordoqui *et al.* *JHEAP* **34** (2022) 19.
- [96] H. Abreu *et al.* *EPJ C* **80** (2020) 61.
- [97] L. A. Anchordoqui *et al.* *Phys. Rep.* **968** (2022) 1.
- [98] D. Soldin *EPJ Web of Conf.* **280** (2023) 04003.
- [99] R. Ulrich, R. Engel, and M. Unger *Phys. Rev. D* **83** (2011) 054026.
- [100] J. Blazek *et al.* *EPJ Web of Conf.* **283** (2023) 05005.
- [101] J. Ebr *et al.* *PoS ICRC2023* (2023) 245.
- [102] R. Ulrich *et al.* *arXiv* **0906.0418** (2009) .
- [103] R. J. Glauber *Phys. Rev.* **100** (1955) 242.
- [104] T. K. Gaisser, R. Engel, and E. Resconi, *Cosmic Rays and Particle Physics*. Cambridge University Press, 2 ed., 2016.
- [105] A. Aab *et al.*, [Pierre Auger Collab.] *NIM-A* **798** (2015) 172.
- [106] I. Allekotte *et al.* *NIM-A* **586** (2008) 409.
- [107] J. Abraham *et al.*, [Pierre Auger Collab.] *NIM-A* **620** (2010) 227.
- [108] A. Aab *et al.*, [Pierre Auger Collab.] *arXiv* **1604.03637** (2016) .
- [109] A. Tripathi *et al.* *NIM-A* **497** (2003) 331.
- [110] J. Abraham *et al.*, [Pierre Auger Collab.] *NIM-A* **613** (2010) 29.
- [111] D. Allard *et al.*, [Pierre Auger Collab.] *arXiv astro-ph/0510320* (2005) .
- [112] T. Suomijärvi, [Pierre Auger Collab.] *J. Phys. Conf. Ser.* (2023) 012010.
- [113] M. T. Dova *Proceedings of the 7th CERN–Latin-American School of High-Energy Physics* (2015) 169.
- [114] X. Bertou *et al.* *NIM-A* **568** (2006) 839.

- [115] J. Abraham *et al.*, [Pierre Auger Collab.] *NIM-A* **523** (2004) 50.
- [116] M. Ambrosio *et al.*, [Pierre Auger Collab.] *NIM-A* **478** (2002) 125.
- [117] R. Knapik *et al.*, [Pierre Auger Collab.] *Proc. Conf. ICRC2007* **4** (2007) 343.
- [118] P. Bauleo *et al.*, [Pierre Auger Collab.] *Proc. Conf. ICRC2005* **8** (2005) 5.
- [119] R. Caruso *et al.*, [Pierre Auger Collab.] *Proc. Conf. ICRC2009* (2009) .
- [120] A. Aab *et al.* *Astropart. Phys.* **95** (2017) 44.
- [121] C. Aramo *et al.*, [Pierre Auger Collab.] *Proc. Conf. ICRC2005* (2005) .
- [122] B. Fick, [Pierre Auger Collab.] *Proc. Conf. ICRC2003* (2003) 449.
- [123] M. Unger *et al.* *NIM-A* **588** (2008) 433.
- [124] T. K. Gaisser and A. M. Hillas *Proc. Conf. ICRC1977* **8** (1977) 353.
- [125] K. Greisen *Annu. Rev. Nucl. Part. Sci.* **10** (1960) 63.
- [126] M. Risse *arXiv astro-ph/0402300* (2004) .
- [127] S. Andringa, R. Conceição, and M. Pimenta *Astropart. Phys.* **34** (2011) 360.
- [128] R. Conceição, S. Andringa, F. Diogo, and M. Pimenta *J. Phys. Conf. Ser.* **632** (2015) 012087.
- [129] J. Linsley *Proc. Conf. ICRC1977* (1977) 89.
- [130] J. Linsley *Proc. Conf. ICRC1983* (1983) 135.
- [131] J. Linsley and A. A. Watson *Phys. Rev. Lett.* **46** (1981) 459.
- [132] K.-H. Kampert and M. Unger *Astropart. Phys.* **35** (2012) 660.
- [133] V. Verzi *et al.*, [Pierre Auger Collab.] *PoS ICRC2013* (2013) 0928.
- [134] B. Dawson, [Pierre Auger Collab.] *PoS ICRC2019* (2020) 231.
- [135] B. Dawson *EPJ Web of Conf.* **144** (2017) 01001.

BIBLIOGRAPHY

- [136] J. Abraham *et al.* *Astropart. Phys.* **32** (2009) 89.
- [137] J. Abraham *et al.*, [Pierre Auger Collab.] *Astropart. Phys.* **33** (2010) 108.
- [138] P. Abreu *et al.*, [Pierre Auger Collab.] *Astropart. Phys.* **35** (2012) 591.
- [139] S. BenZvi *et al.* *NIM-A* **574** (2007) 171.
- [140] A. Aab *et al.*, [Pierre Auger Collab.] *J. Instrum* **16** (2021) P06027.
- [141] A. A. Gioacchino, [Pierre Auger Collab.] *NIM-A* **1044** (2022) 167497.
- [142] A. Aab, [Pierre Auger Collab.] *Phys. Rev. D* **93** (2016) 122005.
- [143] A. Yushkov *et al.*, [Pierre Auger Collab.] *PoS ICRC2019* (2020) 482.
- [144] T. Fitoussi *et al.*, [Pierre Auger Collab.] *PoS ICRC2023* (2023) 319.
- [145] J. Bellido *et al.*, [Pierre Auger Collab.] *PoS ICRC2017* (2017) 506.
- [146] J. Bellido *et al.*, [Pierre Auger Collab.] *PoS ICRC2023* (2023) 211.
- [147] V. M. Harvey *et al.*, [Pierre Auger Collab.] *PoS ICRC2023* (2023) 300.
- [148] A. Condorelli, S. Petrera, F. Salamida, and C. Trimarelli *Internal Note GAP2021-030* (2021) .
- [149] F. James and M. Roos *Comput. Phys. Commun.* **10** (1975) 343.
- [150] D. W. Hogg and D. Foreman-Mackey *ApJS* **236** (2018) 11.
- [151] D. Foreman-Mackey *et al.* *Publ. Astron. Soc. Pac.* **125** (2013) 306.
- [152] J. Salvatier *et al.* *PeerJ Comput. Sci.* **2** (2016) .
- [153] S. Baker and R. D. Cousins *NIM* **221** (1984) 437.
- [154] W. Metzger. *Statistical Methods in Data Analysis*, 2002.
- [155] T. Bergmann *et al.* *Astropart. Phys.* **26** (2007) 420.
- [156] M. E. Muller *Comm. Assoc. Comput. Mach.* **2** (1959) 19.
- [157] H. Dembinski *et al.* *scikit-hep/iminuit: v2.0.0* . Zenodo 2020.

- [158] A. Gelman and D. B. Rubin *Statistical Science* **7** no. 4, (1992) 457.
- [159] M. D. Domenico, M. Settimo, S. Riggi, *et al.* *JCAP* **2013** (2013) 050.
- [160] T. Bister and M. Wirtz *Astrotools*
<https://astro.pages.rwth-aachen.de/astrotools> .
- [161] S. Marafico and O. Deligny *Internal Note GAP2021-021* (2021) .
- [162] A. Aab *et al.*, [The Pierre Auger Collaboration Collab.] *Phys. Rev. D* **102** (2020) 062005.
- [163] A. Aab *et al.* *JCAP* **2017** no. 4, (2017) 038.
- [164] J. Vícha *et al.*, [Pierre Auger Collab.] *SciPost Phys. Proc.* (2023) 026.
- [165] R. Abbasi and G. Thomson *Proc. Conf. UHECR2016* 011015.
- [166] J. Vícha *et al.*, [Pierre Auger Collab.] *EPJ Web of Conf.* **283** (2023) 02012.
- [167] A. A. Halim *et al.*, [Pierre Auger Collab.] *JCAP* **2023** (May, 2023) 024.
- [168] C. Trimarelli, S. Petrera, and F. Salamida *Internal Note GAP2022-030* (2022) .
- [169] J. Abraham *et al.*, [Pierre Auger Collab.] *Astropart. Phys.* **27** (2007) 155.
- [170] R. Jansson and G. R. Farrar *Astrophys. J.* **757** (2012) 14.
- [171] E. Mayotte *et al.*, [Pierre Auger Collab.] *PoS ICRC2021* (2021) 321.
- [172] E. Mayotte, T. Fitoussi, *et al.*, [Pierre Auger Collab.] *EPJ Web of Conf.* **283** (2023) 03003.
- [173] E. Santos *et al.*, [Pierre Auger Collab.] *PoS ICRC2021* (2021) 232.
- [174] R. M. Baltrusaitis *et al.* *Phys. Rev. Lett.* **52** (1984) 1380.
- [175] R. M. Ulrich *et al.*, [Pierre Auger Collab.] *PoS ICRC2015* (2016) 401.
- [176] R. W. Ellsworth, T. K. Gaisser, T. Stanev, and G. B. Yodh *Phys. Rev. D* **26** (1982) 336.
- [177] K. Belov, [HiReS Collab.] *Nuc. Phys. B* **151** (2006) 197.

BIBLIOGRAPHY

- [178] P. Abreu *et al.*, [Pierre Auger Collab.] *Phys. Rev. Lett.* **109** (2012) 062002.
- [179] R. U. Abbasi *et al.*, [Telescope Array Collab.] *Phys. Rev. D* **92** (2015) 032007.
- [180] R. U. Abbasi *et al.*, [Telescope Array Collab.] *Phys. Rev. D* **102** (2020) 062004.
- [181] P. Lipari *Phys. Rev. D* **103** (2021) 103009.
- [182] O. Tkachenko, R. Engel, R. Ulrich, and M. Unger *PoS ICRC2021* (2022) 438.
- [183] O. Tkachenko *et al.*, [Pierre Auger Collab.] *PoS ICRC2023* (2023) 438.
- [184] L. B. Arbeletche and V. de Souza *Astropart. Physics* **116** (2020) 102389.
- [185] T. Pierog and K. Werner *PoS ICRC2023* (2023) 230.

Appendix

A X_{\max} -dependent X_{\max} bias

Parametrization of the X_{\max} -dependent X_{\max} bias as a function of energy

The mean of the difference between the generated and reconstructed X_{\max} , $\mu(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ can be parameterized as a cubic function w.r.t. X_{\max} :

$$\mu(X_{\max}) = \text{par}_{0,\mu} X_{\max}^3 + \text{par}_{1,\mu} X_{\max}^2 + \text{par}_{2,\mu} X_{\max} + \text{par}_{3,\mu}. \quad (\text{A.1})$$

Each of the parameters in the X_{\max} -dependence of the bias can then be parameterized as a function of energy:

$$\text{par}_n = \text{coeff}_{0,\mu} E^3 + \text{coeff}_{1,\mu} E^2 + \text{coeff}_{2,\mu} E + \text{coeff}_{3,\mu}. \quad (\text{A.2})$$

The standard deviation of the distribution of differences $\sigma(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ can be parameterized in a similar way with a quadratic function. Tabs. [A.1](#) and [A.2](#) show the coefficients of the function dependence of the bias parameterization of the X_{\max} -dependent X_{\max} bias on energy for the mean and standard deviation of the difference between the generated and reconstructed X_{\max} , respectively, for the compFoV event selection.

Similarly, the X_{\max} -dependent X_{\max} bias was parameterized as a function of energy for the xsecFoV selection, see Tabs. [A.3](#) and [A.4](#). In this case we use a quadratic polynomial for the dependence on the X_{\max} for the both mean and standard deviation of the distribution of $X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}}$.

parameter	coeff _{0,μ}	coeff _{1,μ}	coeff _{2,μ}	coeff _{3,μ}
par ₀	1.21e-8	-2.42e-7	-4.44e-6	-8.87e-5
par ₁	-2.6e-5	5.2e-4	9.4e-3	- 0.19
par ₂	1.8e-2	-0.36	-6.5	131.6
par ₃	-4	81	1454	-29708

Table A.1: Parametrization of the $\mu(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ as a function of energy for the compFoV event selection.

parameter	coeff _{1,σ}	coeff _{2,σ}	coeff _{3,σ}
par ₁	-1.22e-5	4e-4	- 3e-3
par ₂	1.6e-5	-0.48	3.24
par ₃	-1.58	4	543

Table A.2: Parametrization of the $\sigma(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ as a function of energy for the compFoV event selection.

parameter	coeff _{0,μ}	coeff _{1,μ}	coeff _{2,μ}	coeff _{3,μ}
par ₁	2.9e-6	5.3e-3	-1e-3	1.8e-2
par ₂	-4.4e-3	8.1e-2	1.5	-28.2
par ₃	1.7	-30.3	-574.8	10482

Table A.3: Parametrization of the $\mu(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ as a function of energy for the xsecFoV event selection.

parameter	coeff _{0,σ}	coeff _{1,σ}	coeff _{2,σ}	coeff _{3,σ}
par ₁	-1.8e-6	3.4e-5	6.4e-4	-1.1e-2
par ₂	2.7e-3	-5e-2	-0.94	17
par ₃	-0.9	16.2	312	-5499

Table A.4: Parametrization of the $\sigma(X_{\max}^{\text{gen}} - X_{\max}^{\text{rec}})$ as a function of energy for the xsecFoV event selection.

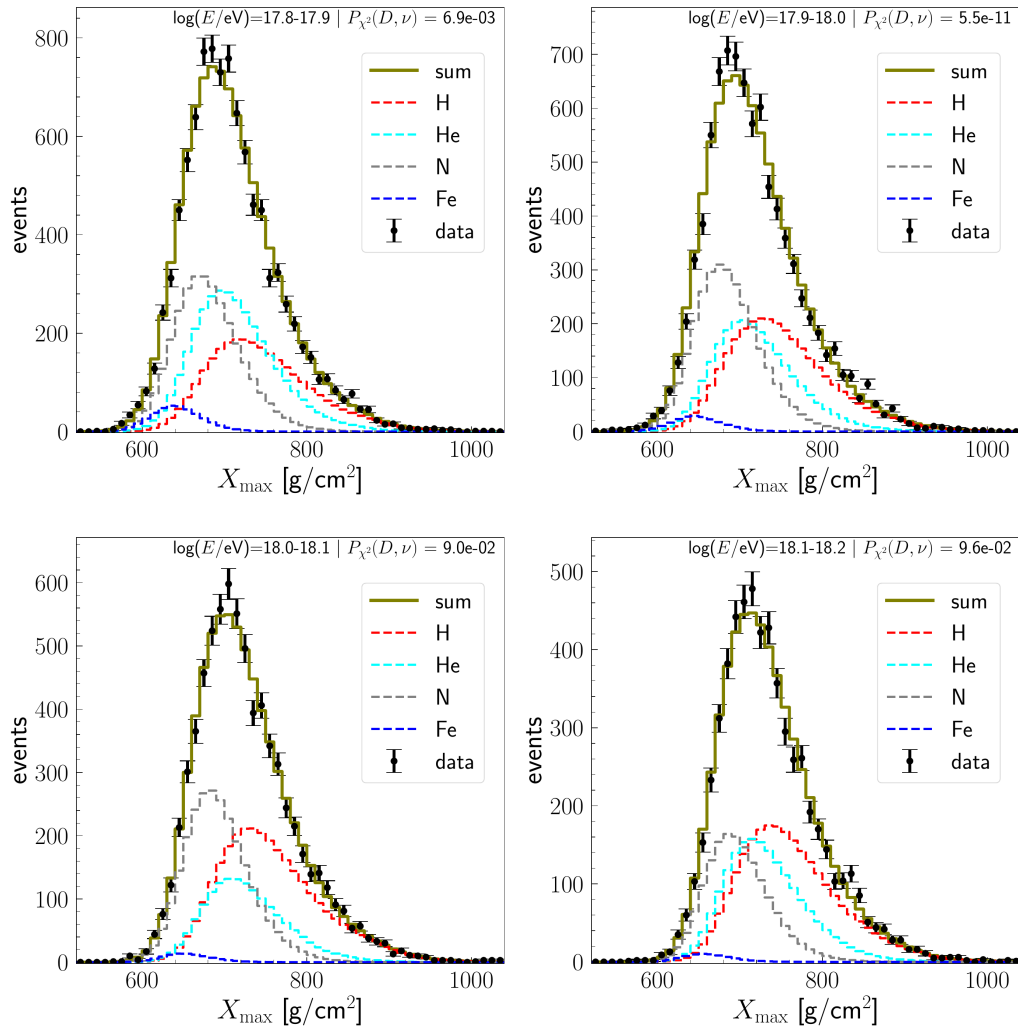
B MCMC mass composition estimation





Mass composition fit to the ICRC23 data with Sibyll 2.3d hadronic interaction model





$\lg E_{\text{low}}/\text{eV}$	$\lg E_{\text{up}}/\text{eV}$	f_{H}	f_{He}	f_{N}	f_{Fe}	p-value
17.8	17.9	$0.29^{+0.025}_{-0.026}$	$0.35^{+0.052}_{-0.049}$	$0.32^{+0.034}_{-0.038}$	$0.042^{+0.011}_{-0.01}$	0.013
17.9	18.0	$0.36^{+0.027}_{-0.029}$	$0.28^{+0.054}_{-0.053}$	$0.34^{+0.035}_{-0.037}$	$0.025^{+0.009}_{-0.009}$	0.00001
18.0	18.1	$0.42^{+0.029}_{-0.031}$	$0.21^{+0.049}_{-0.058}$	$0.35^{+0.038}_{-0.032}$	$0.01^{+0.009}_{-0.009}$	0.16
18.1	18.2	$0.42^{+0.031}_{-0.037}$	$0.3^{+0.066}_{-0.053}$	$0.26^{+0.032}_{-0.043}$	$0.01^{+0.009}_{-0.008}$	0.134
18.2	18.3	$0.39^{+0.03}_{-0.04}$	$0.39^{+0.075}_{-0.05}$	$0.2^{+0.03}_{-0.051}$	$0.01^{+0.01}_{-0.008}$	0.23
18.3	18.4	$0.34^{+0.033}_{-0.048}$	$0.52^{+0.087}_{-0.056}$	$0.12^{+0.033}_{-0.06}$	$0.018^{+0.014}_{-0.009}$	0.41
18.4	18.5	$0.38^{+0.029}_{-0.057}$	$0.47^{+0.106}_{-0.036}$	$0.15^{+0.01}_{-0.066}$	$0.0^{+0.0089}_{-0.0}$	0.72
18.5	18.6	$0.22^{+0.029}_{-0.065}$	$0.64^{+0.131}_{-0.03}$	$0.14^{+0.007}_{-0.086}$	$0.00004^{+0.014}_{-0.00004}$	0.4
18.6	18.7	$0.27^{+0.038}_{-0.06}$	$0.5^{+0.126}_{-0.051}$	$0.23^{+0.019}_{-0.086}$	$0.00001^{+0.011}_{-0.00001}$	0.0003
18.7	18.8	$0.04^{+0.027}_{-0.038}$	$0.79^{+0.082}_{-0.066}$	$0.17^{+0.037}_{-0.069}$	$0.00004^{+0.007}_{-0.00004}$	0.28
18.8	18.9	$0.0012^{+0.045}_{-0.012}$	$0.8^{+0.061}_{-0.086}$	$0.2^{+0.034}_{-0.094}$	$0.0002^{+0.018}_{-0.0002}$	0.12
18.9	19.0	$0.024^{+0.026}_{-0.022}$	$0.77^{+0.097}_{-0.064}$	$0.21^{+0.037}_{-0.103}$	$0.00006^{+0.01}_{-0.00006}$	0.0007
19.0	19.1	$0.0088^{+0.053}_{-0.088}$	$0.7^{+0.097}_{-0.057}$	$0.29^{+0.07}_{-0.083}$	$0.00005^{+0.011}_{-0.00005}$	0.019
19.1	19.2	$0.0004^{+0.027}_{-0.0004}$	$0.52^{+0.091}_{-0.096}$	$0.48^{+0.062}_{-0.11}$	$0.0002^{+0.012}_{-0.0002}$	0.03
19.2	19.3	$0.059^{+0.069}_{-0.059}$	$0.45^{+0.16}_{-0.19}$	$0.49^{+0.1}_{-0.16}$	$0.00002^{+0.02}_{-0.00003}$	0.12
19.3	19.4	$0.002^{+0.07}_{-0.02}$	$0.3^{+0.07}_{-0.19}$	$0.7^{+0.09}_{-0.13}$	$0.0002^{+0.024}_{-0.0002}$	0.01
19.4	19.5	$0.0001^{+0.08}_{-0.0001}$	$0.17^{+0.053}_{-0.17}$	$0.8^{+0.046}_{-0.23}$	$0.03^{+0.029}_{-0.08}$	0.35
19.5	19.6	$0.0008^{+0.13}_{-0.0008}$	$0.28^{+0.05}_{-0.28}$	$0.72^{+0.074}_{-0.27}$	$0.0005^{+0.056}_{-0.0004}$	0.27
19.6	21.0	$0.0009^{+0.054}_{-0.0009}$	$0.00004^{+0.11}_{-0.00004}$	$0.87^{+0.13}_{-0.27}$	$0.13^{+0.14}_{-0.08}$	0.76





Table B.1: Fractions from the MCMC fit to the ICRC23 data with the Sibyll 2.3d interaction model. Only statistical uncertainties are included.

X_{\max} distributions for the fit to the ICRC23 data with the Sibyll 2.3d hadronic interaction model









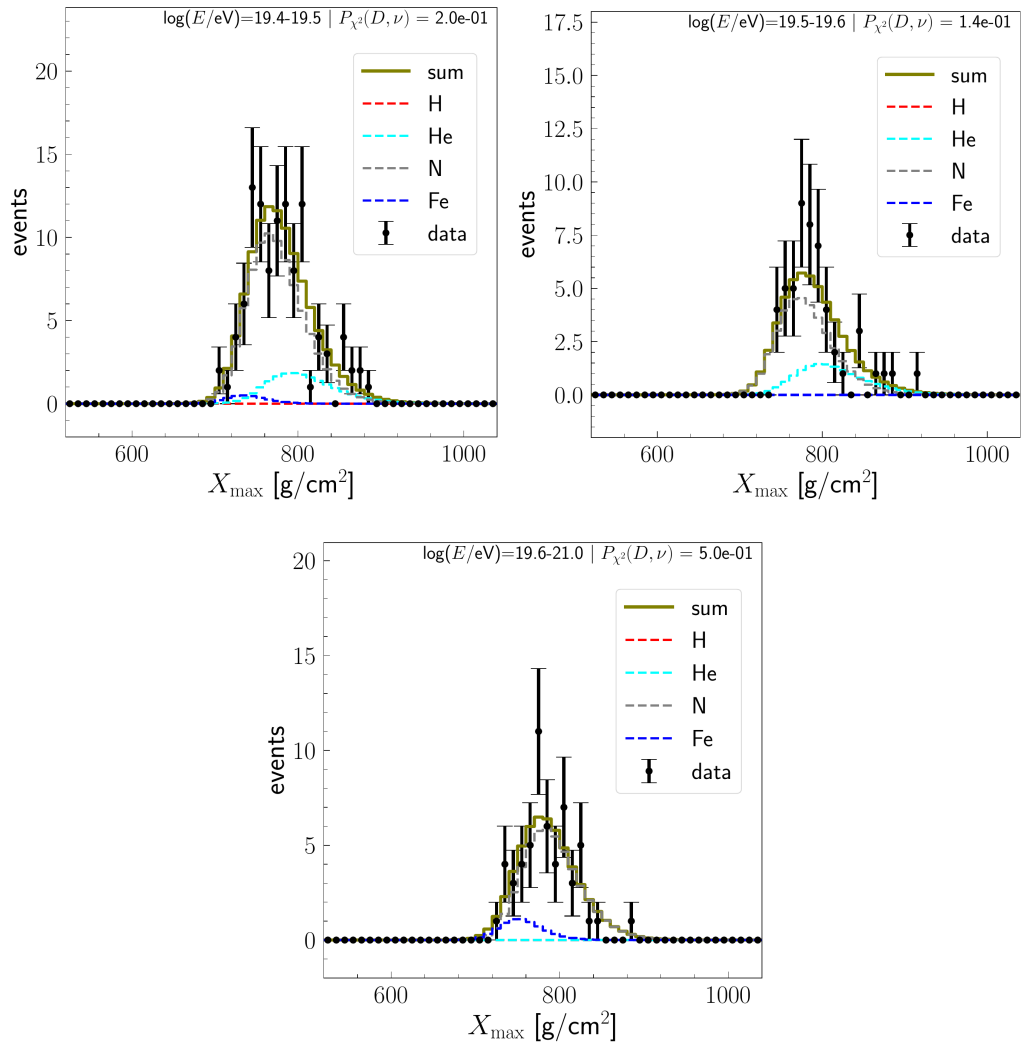



Figure B.1: X_{\max} distributions for the fit to the ICRC23 data with the Sibyll 2.3d interaction model





Mass composition fit to the ICRC23 data with EPOS-LHC hadronic interaction model





$\lg E_{\text{low}}/\text{eV}$	$\lg E_{\text{up}}/\text{eV}$	f_{H}	f_{He}	f_{N}	f_{Fe}	p-value
17.8	17.9	$0.49^{+0.028}_{-0.027}$	$0.13^{+0.46}_{-0.046}$	$0.3^{+0.029}_{-0.029}$	$0.08^{+0.0086}_{-0.0083}$	0.047
17.9	18.0	$0.56^{+0.026}_{-0.0296}$	$0.06^{+0.044}_{-0.042}$	$0.31^{+0.024}_{-0.03}$	$0.061^{+0.009}_{-0.007}$	0.00009
18.0	18.1	$0.61^{+0.028}_{-0.027}$	$0.065^{+0.041}_{-0.047}$	$0.26^{+0.024}_{-0.031}$	$0.07^{+0.009}_{-0.008}$	0.38
18.1	18.2	$0.67^{+0.027}_{-0.033}$	$0.063^{+0.047}_{-0.05}$	$0.23^{+0.024}_{-0.035}$	$0.04^{+0.007}_{-0.01}$	0.4
18.2	18.3	$0.61^{+0.038}_{-0.041}$	$0.18^{+0.061}_{-0.073}$	$0.17^{+0.044}_{-0.035}$	$0.04^{+0.009}_{-0.01}$	0.29
18.3	18.4	$0.58^{+0.039}_{-0.051}$	$0.27^{+0.086}_{-0.063}$	$0.12^{+0.037}_{-0.05}$	$0.035^{+0.012}_{-0.009}$	0.55
18.4	18.5	$0.67^{+0.039}_{-0.065}$	$0.12^{+0.12}_{-0.047}$	$0.2^{+0.013}_{-0.068}$	$0.001^{+0.011}_{-0.001}$	0.82
18.5	18.6	$0.46^{+0.044}_{-0.067}$	$0.39^{+0.12}_{-0.06}$	$0.13^{+0.028}_{-0.077}$	$0.012^{+0.014}_{-0.008}$	0.55
18.6	18.7	$0.5^{+0.049}_{-0.072}$	$0.23^{+0.14}_{-0.06}$	$0.26^{+0.018}_{-0.088}$	$0.0003^{+0.013}_{-0.0003}$	0.003
18.7	18.8	$0.19^{+0.04}_{-0.07}$	$0.66^{+0.14}_{-0.06}$	$0.15^{+0.023}_{-0.1}$	$0.00004^{+0.009}_{-0.00004}$	0.74
18.8	18.9	$0.15^{+0.048}_{-0.074}$	$0.72^{+0.15}_{-0.04}$	$0.12^{+0.087}_{-0.08}$	$0.006^{+0.015}_{-0.006}$	0.21
18.9	19.0	$0.08^{+0.04}_{-0.05}$	$0.84^{+0.1}_{-0.03}$	$0.08^{+0.06}_{-0.09}$	$0.00002^{+0.01}_{-0.00002}$	0.018
19.0	19.1	$0.14^{+0.062}_{-0.078}$	$0.64^{+0.17}_{-0.086}$	$0.22^{+0.04}_{-0.13}$	$0.0^{+0.012}_{-0.0}$	0.12
19.1	19.2	$0.0007^{+0.039}_{-0.0007}$	$0.72^{+0.08}_{-0.1}$	$0.28^{+0.11}_{-0.063}$	$0.0001^{+0.013}_{-0.0001}$	0.27
19.2	19.3	$0.16^{+0.096}_{-0.1}$	$0.45^{+0.24}_{-0.16}$	$0.38^{+0.06}_{-0.2}$	$0.0001^{+0.02}_{-0.0001}$	0.32
19.3	19.4	$0.03^{+0.09}_{-0.03}$	$0.47^{+0.1}_{-0.21}$	$0.5^{+0.1}_{-0.13}$	$0.0001^{+0.023}_{-0.0001}$	0.04
19.4	19.5	$0.0007^{+0.13}_{-0.0007}$	$0.42^{+0.057}_{-0.29}$	$0.52^{+0.15}_{-0.17}$	$0.06^{+0.068}_{-0.052}$	0.3
19.5	19.6	$0.077^{+0.12}_{-0.077}$	$0.32^{+0.15}_{-0.31}$	$0.6^{+0.054}_{-0.33}$	$0.0006^{+0.055}_{-0.0006}$	0.37
19.6	21.0	$0.0003^{+0.082}_{-0.0003}$	$0.044^{+0.13}_{-0.044}$	$0.9^{+0.1}_{-0.31}$	$0.05^{+0.12}_{-0.045}$	0.75





Table B.2: Fractions from the MCMC fit to the ICRC23 data with the EPOS-LHC interaction model. Only statistical uncertainties are included.

X_{\max} distributions for the fit to the ICRC23 data with the EPOS-LHC hadronic interaction model









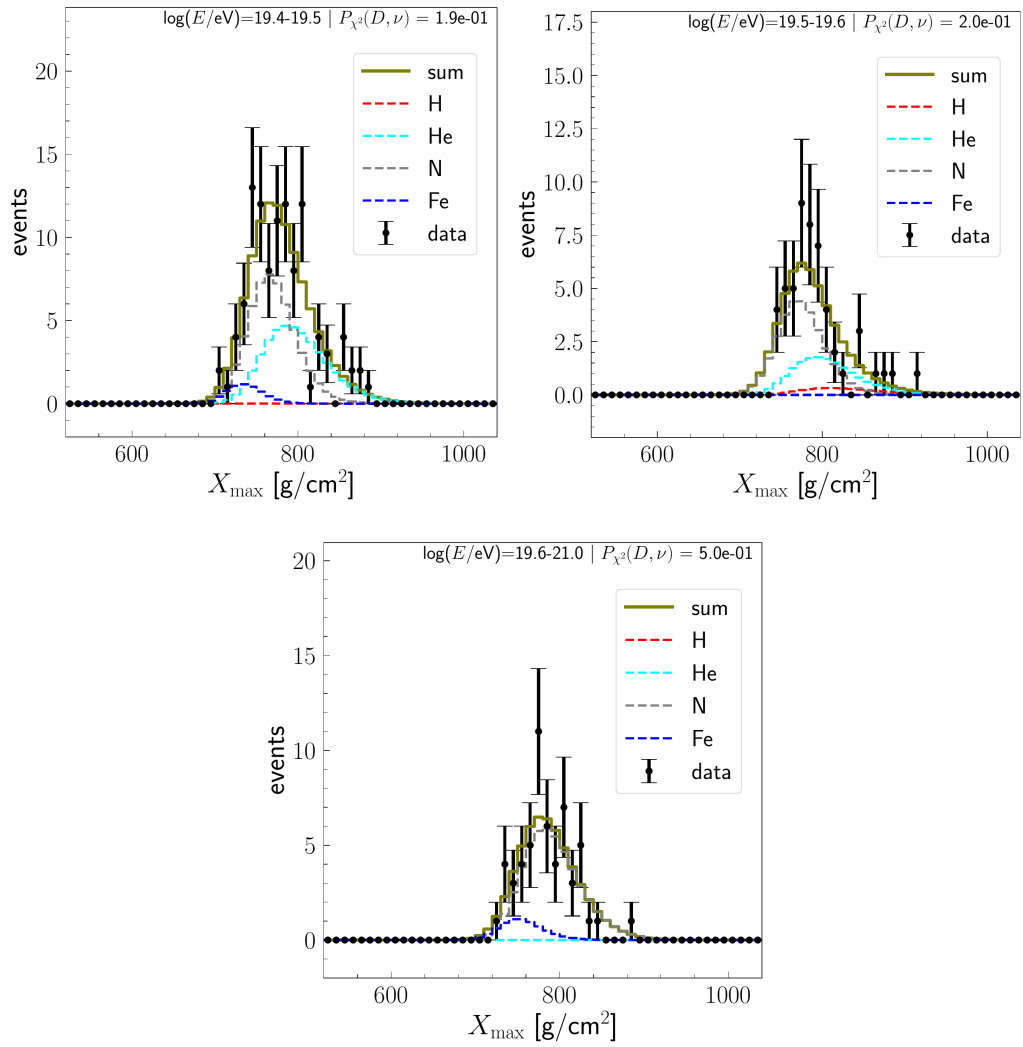


Figure B.2: X_{\max} distributions for the fit to the ICRC23 data with the EPOS-LHC interaction model

Fit with Sibyll 2.3d vs Sibyll 2.3c




Figure B.3: Comparison of the estimated composition fractions with two versions of the Sibyll MC generator, Sibyll 2.3c, which has been used extensively in the past, and the newer version Sibyll 2.3d, for ICRC19 data. Only statistical uncertainties are shown.

Shift in the X_{\max} scale as an additional fit parameter + Si in the composition mix













Figure B.4: Mass composition fit to the ICRC2019 data with a shift in the X_{\max} scale and Si nuclei in the fractions mix. The color scale shows the posterior distribution for the compositional fractions and the shift δX_{\max} normalized to the corresponding modes. The black dots are the point estimates of the fitted quantities (MAP). It can be seen that introducing an additional heavy component into the mix leads to a redistribution of the expected quantities. The δX_{\max} is on average -25 g/cm^2 with much less variation than for the 4-component fit.





C Combined mass composition & cross section fit

X_{\max} distributions for the fit to the ICRC23 data with the Sibyll 2.3d hadronic interaction model with modified proton-proton interaction cross section: best-fit case









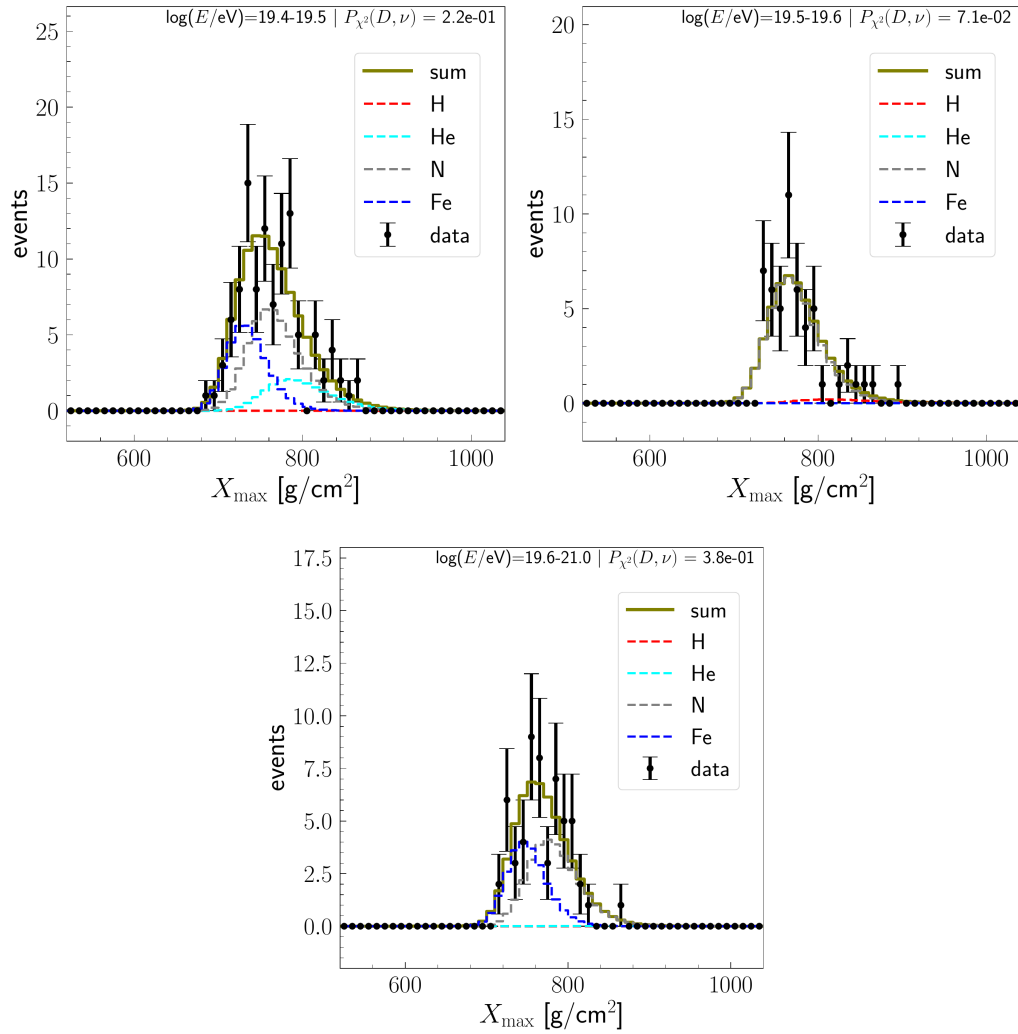


Figure C.1: X_{\max} distributions for the fit to the ICRC23 data with the modified Sibyll 2.3d interaction mode (best-fit cross section).

Energy threshold




Figure C.2: The dependence of the estimated cross section on the threshold energy, E_{tr} . Fit to the ICRC19 data. While there is nearly no difference in the results for the $E_{\text{tr}} < 10^{18}$ eV, the direction of the fit changes completely above it.

Acknowledgment

Staying as a PhD student at the Institute for Astroparticle Physics at the Karlsruhe Institute of Technology was a very nice and fruitful experience and I would like to thank every member of the group for making it like that. I also thank the Doctoral School "Karlsruhe School of Elementary and Astroparticle Physics: Science and Technology (KSETA)" for financial support through the GSSP program of the German Academic Exchange Service (DAAD) for providing me with an opportunity to pursue research in the field of my interest.

I would like to thank my supervisor, Dr. Michael Unger, for all the advice and assistance I received from him while working on my PhD project. The project would never have reached such a depth and extent if he had not been involved. I appreciate Prof. Ralph Engel for providing me with the opportunity to work in the Auger group. The occasional yet highly informative and encouraging meetings with him were essential in addressing and understanding challenges in my analysis. Furthermore, I extend my gratitude to Prof. Ulrich Husemann for accepting the role of the second referee for my thesis.

I would like to express my gratitude to my colleagues at the Pierre Auger Collaboration for their valuable contributions to the development of this work. In particular, I would like to extend a special thanks to Dr. Alexey Yushkov for providing me with Monte Carlo simulations crucial for the completion of the work, as well as for the insightful discussions we had. Further thanks go to Dr. Thomas Fitoussi for the full Phase I data, which became an integral part of the analysis. I also would like to express my thanks to Dr. Darko Veberic and Dr. Markus Roth for their valuable comments on my work. I wish to acknowledge the joint Auger-TA working group for their collaboration and positive environment.

Additionally, I extend my appreciation to my office mates, Katrin Bismark, Luca Deval, and Felix Knapp, for the pleasant and enjoyable moments we shared. Special gratitude goes to Emily Martins for being the friend I truly needed. My deepest thanks are reserved for Tobias Schulz, the love of my life, who consistently stood by me, offering support and comfort whenever necessary. I am also profoundly thankful to my parents for their unwavering belief in me and encouragement.

Last but not least, I would like to thank those who sacrificed their lives and those who are still fighting for the existence of my home and the survival of the Ukrainian nation.

Erklärung

Karlsruhe, den 10.12.2023

Erklärung der selbständigen Anfertigung meiner Dissertationsschrift

Hiermit versichere ich, dass ich die Dissertationsschrift mit dem Titel

**Measurement of the Composition of Cosmic Rays and the Proton-Proton
Inelastic Cross Section at Ultrahigh Energies with the Pierre Auger
Observatory**

selbständig und ohne unerlaubte fremde Hilfe verfasst habe. Dabei habe ich keine anderen, als die von mir angegebenen Hilfsmittel benutzt.

Your Name