

Mass composition of ultra-high-energy cosmic rays at the Pierre Auger Observatory

Thomas Fitoussi^{a,*} for the Pierre Auger Collaboration^b

^aKarlsruhe Institute of Technology  Institute for Astroparticle Physics, Karlsruhe, Germany

^bObservatorio Pierre Auger, Av. San Martín Norte 304, 5613 Malargüe, Argentina

Full author list: https://www.auger.org/archive/authors_2024_11.html

E-mail: spokespersons@auger.org

For the past 20 years, the Pierre Auger Observatory has collected the largest dataset of ultra-high-energy cosmic rays (UHECRs) ever achieved using a hybrid detector. The study of this dataset has led to numerous unexpected discoveries that enhance our understanding of the origins of UHECRs. One of the key points in this study is their mass composition. In this work, we will present the most recent results regarding the mass composition of UHECRs at the Pierre Auger Observatory. In particular, we will focus on the measurement of the depth of the maximum of air-shower profiles, denoted as X_{\max} . This determination has been achieved through both direct measurements from the Fluorescence Detector data and the application of machine learning for estimating X_{\max} on an event-by-event basis using the Surface Detector data. The latter has allowed us to extend the measurement to energies up to 100 EeV and indicates a correlation between changes in composition and three features of the energy spectrum (ankle, instep, steepening). Moreover, the results provide evidence of a heavy and nearly pure primary beam for energies greater than 50 EeV that is independent of the hadronic interaction model. The implications of these findings for astrophysics and for modelling hadronic interactions will be discussed.

*7th International Symposium on Ultra High Energy Cosmic Rays (UHECR2024)
17-21 November 2024
Malargüe, Mendoza, Argentina*

*Speaker

1. Introduction

The Pierre Auger Observatory began collecting data in 2004 and was fully deployed in 2008. The Observatory consists of an array of 1,660 water-Cherenkov detectors (WCDs) spaced 1,500 m apart, along with two smaller arrays spaced 750 m and 433 m apart, respectively, to probe lower energies. Together, these form the Surface Detector (SD), which covers an area of 3,000 km² [1]. Additionally, 24 fluorescence telescopes, grouped in sets of six, are located at four different sites around the SD, enabling hybrid measurements. Each telescope has a field of view of 30° and an elevation range between 1.5° and 30° [2]. Furthermore, three fluorescence telescopes, aimed at higher elevations (30° to 60°), are deployed at the Coihueco site to study energies below 10¹⁸ eV [3]. These 27 telescopes constitute the Fluorescence Detector (FD) [4]. In 2019, the deployment of an upgrade to the Observatory, dubbed AugerPrime, started [5]. In this proceeding, we will focus on the data collected prior to the upgrade, referred to as Phase I data.

The mass composition of ultra-high-energy cosmic rays (UHECRs) is important for understanding their origins. The best available proxy for a mass estimate is the maximum of the air-shower development, denoted as X_{\max} [6], which can be measured directly with the FD. Due to the limited duty cycle (approximately 13%), FD data are scarce at the highest energies. In contrast, the SD has a nearly 100% duty cycle but does not have direct access to X_{\max} . Nevertheless, the time distributions of the signal traces measured by each WCD [7] contain information about X_{\max} and allow for its estimation.

In this proceeding, we present X_{\max} measurements from the FD and X_{\max} estimations derived from the SD data using machine learning, along with interpretations of their mass composition.

2. Mass composition from the Fluorescence Detector data

The maximum development of the air shower, denoted as X_{\max} , can be directly observed using the fluorescence telescope. However, not all events meet the required quality standards, necessitating an event selection process. Initially, a minimal selection is applied to reject events recorded during periods of poor conditions, whether due to the electronic acquisition system or atmospheric factors (such as clouds or high aerosol content). Only events that trigger at least one water-Cherenkov detector (WCD) (referred to as hybrid events) and have a reconstructed profile are retained. A quality selection is then performed to ensure that there is no compositional bias from the hybrid detection, that the profile is accurately fitted, and that X_{\max} is within the field of view. Finally, a fiducial selection is applied to guarantee that the complete X_{\max} distribution can be reconstructed. Only events with energies above 10^{17.8} eV are included [8, 9].

In this analysis, it is essential to account for the detector response to correct for small residual biases in the X_{\max} distributions. This correction encompasses the X_{\max} -dependent combined efficiencies for event triggering, event reconstruction, and event selection (referred to as "acceptance"), as well as the reconstruction biases in energy and X_{\max} , and the average X_{\max} resolution as a function of energy. To compute these factors, we generated air showers using CONEX [10] with the Sybill2.3d [11] hadronic model for protons and iron nuclei. The simulation of the atmosphere, along with the detector response and reconstruction, was conducted using the Auger Offline software [12] to account for the real-time state of the SD and FD [13].

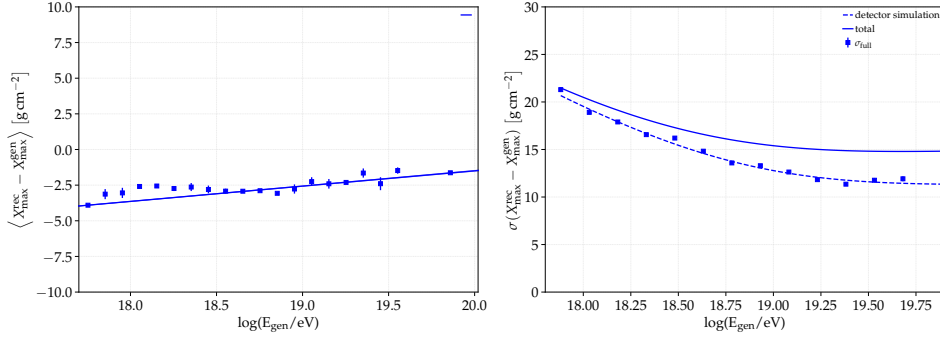


Figure 1: Energy dependency of X_{\max} bias and resolution. See text for more information.

Simulations are selected in the same manner as the data. To compute the resolution, the simulated X_{\max} distributions are also reweighted to match the energy and X_{\max} distributions of the data. This approach allows for the calculation of end-to-end residual biases and X_{\max} resolution. Fig. 1 displays the X_{\max} bias (on the left) and the X_{\max} resolution (on the right). These correspond to the mean and standard deviation of the $\Delta X_{\max} = X_{\max}^{\text{rec}} - X_{\max}^{\text{gen}}$ distributions, respectively, as a function of energy. The bias decreases with increasing energy, remaining between -3 g/cm^2 and -1 g/cm^2 , which is significantly lower than the systematic uncertainties (approximately $\pm 10 \text{ g/cm}^2$). The standard deviation also decreases with energy, from 22 g/cm^2 to 12 g/cm^2 (dashed line). The total X_{\max} resolution is represented in Fig. 1 (right) by the solid line. It includes uncorrelated uncertainties from the atmosphere (such as the precision of aerosol content estimation and molecular scattering) and the alignment of the telescopes. The systematic uncertainties of the X_{\max} scale and resolution remain unchanged from [8], except for the VAOD uncertainties, which have been updated [14] leading to an increase of the $\langle X_{\max} \rangle$ by 1 g/cm^2 to 4 g/cm^2 .

A total of 75 210 events passed all selection criteria. First, we compute the raw X_{\max} distributions after applying quality and fiducial field of view cuts. Next, X_{\max} moments are estimated from these distributions, taking into account the X_{\max} acceptance using the Λ_η method described in [8]. Finally, the X_{\max} resolution is subtracted in quadrature from the computed $\sigma(X_{\max})$ values. The obtained X_{\max} moments are shown in Fig. 2. The results previously obtained by the Pierre Auger Collaboration [8, 15, 16] are confirmed with a dataset that is 60% larger. As can be seen, there is a noticeable break in the $\langle X_{\max} \rangle$ evolution at approximately $10^{18.4} \text{ eV}$, illustrated in Fig. 2 by a broken line (dashed line). The $\langle X_{\max} \rangle$ elongation rate fitted before the break is approximately $\sim 81 \text{ g/cm}^2/\text{decade}$, while after the break it is $\sim 28 \text{ g/cm}^2/\text{decade}$. The comparison with the elongation rate expected for a constant composition ($\sim 60 \text{ g/cm}^2/\text{decade}$) indicates that the composition becomes lighter from $10^{17.8} \text{ eV}$ to $10^{18.3} \text{ eV}$ before transitioning to a heavier composition as the energy increases. Simultaneously, the change in slope of the $\sigma(X_{\max})$ evolution suggests a mixed composition below $10^{18.3} \text{ eV}$ and a purer composition at higher energies.

The X_{\max} moments can be converted [19] to the mean and variance of the logarithm of the mass number $\ln A$ using the following equations:

$$\langle X_{\max} \rangle = \langle X_{\max} \rangle_p + f_E \langle \ln A \rangle \quad (1)$$

$$\sigma^2(X_{\max}) = \langle \sigma_{\text{sh}}^2 \rangle + f_E^2 \sigma^2(\ln A) \quad (2)$$

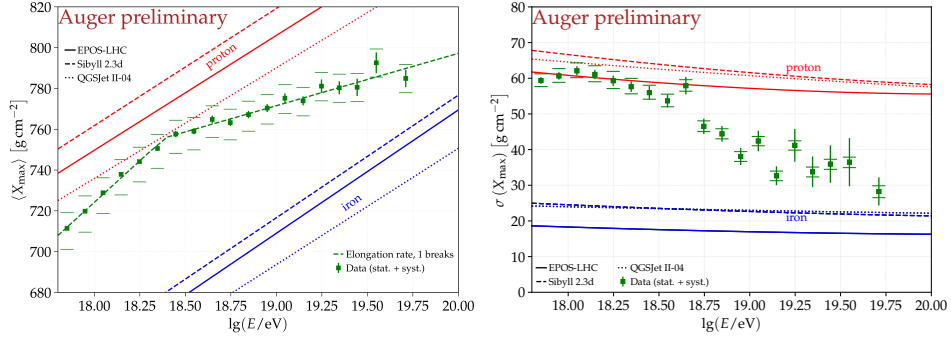


Figure 2: Mean $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$ of the X_{\max} distributions versus energy with the $\langle X_{\max} \rangle$ broken line fit. As comparison predictions for proton and iron nuclei from hadronic models QGJSJet II-04 [17], EPOS-LHC [18] and Sybill2.3d [11] are also displayed. (x-bin centers slightly shifted for better visibility)

where $\langle X_{\max} \rangle_p$ is the mean X_{\max} for protons, σ_{sh}^2 represents the average shower-to-shower fluctuations, and f_E is an energy-dependent parameter. All three factors are parameterized for different hadronic models. The results for the hadronic models QGJSJet II-04 [17], EPOS-LHC [18], and Sibyll 2.3d [11] are shown in Fig. 3. It can be noted that the interpretation with QGJSJet II-04 results in negative values for $\sigma^2(\ln A)$, which is not physically plausible. For the other two models, we corroborate our previous findings of a $\sigma^2(\ln A)$ close to zero, indicating that a less mixed cosmic-ray composition arrives at Earth above 10^{19} eV.

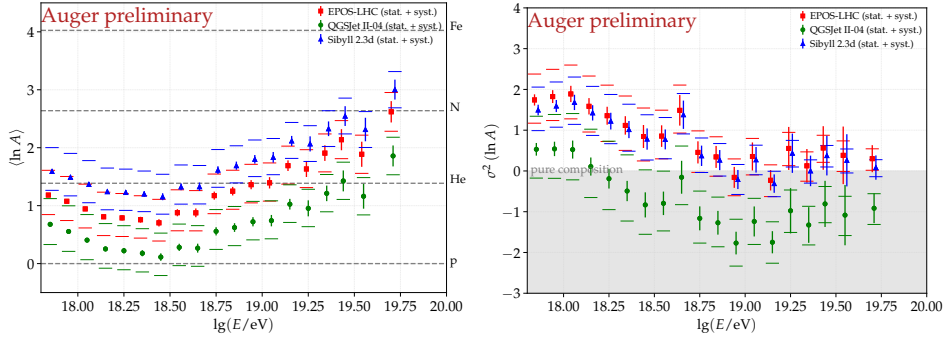


Figure 3: Mean and variance of the logarithm of the mass number $\ln A$ inferred from the data using hadronic models QGJSJet II-04 [17], EPOS-LHC [18] and Sybill2.3d [11].

The X_{\max} distributions are also utilized to fit the composition fractions of protons, helium, nitrogen, and iron. The fitting was performed using the Markov Chain Monte Carlo (MCMC) method [20], which offers several advantages over frequentist inference (see [21]). In Fig. 4, the mass composition fit is presented for a combination of four particle species: proton (H), helium (He), nitrogen (N), and iron (Fe), representing four elemental groups that are approximately equally spaced in $\ln A$. The total uncertainty in the composition fractions includes the statistical uncertainty from the MCMC posterior distributions and the impact of systematic uncertainty on the X_{\max} scale, evaluated by fitting the data with a consistently varied shift of X_{\max} within the scale uncertainty. The trends observed in the evolution of cosmic ray composition with energy are consistent with our previous results presented in [15, 22]. Minor differences from the previous results in the

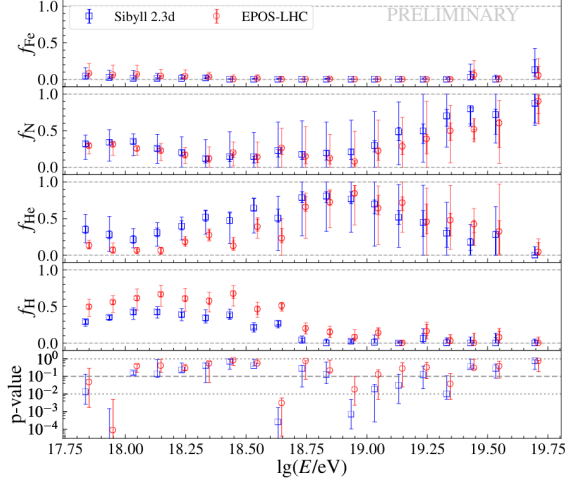


Figure 4: The mass composition fit for four elemental mass groups (top four panels). The bottom panel shows the p -values of the fit.

individual mass groups are likely attributed to the larger dataset (which includes more observation years, changes in event reconstruction, and updates to the aerosol model) and the use of the most recent version of the Sibyll interaction model, which predicts slightly shallower showers than its predecessor [11] (~ 20 g/cm² for protons). Although the qualitative behavior remains the same, there is a significant dependence of the choice of interaction model on the individual fractions. The proton fraction obtained with EPOS-LHC reaches up to 70% around $10^{18.0}$ - $10^{18.2}$ eV and then drops to less than 20% above $10^{18.7}$ eV. The Sibyll 2.3d model predicts a smaller proton fraction across the considered energy range, with a near-zero contribution at higher energies. On average, the Sibyll 2.3d interaction model results in a helium fraction that is approximately 20% larger at lower energies. The amount of iron in the cosmic-ray mix is consistent with zero within uncertainties at all energies. Within the observed energy range, the data are compatible with a successive extinction of the proton and helium components.

3. Mass composition from the Surface Detector

To reconstruct X_{\max} using the Surface Detector (SD), we first ensure a high-quality selection by applying the standard SD selection criteria for near-vertical events [23]. This involves rejecting events below $10^{18.5}$ eV and those with zenith angles greater than 60° , while accepting only events where each station in the hexagon surrounding the station with the largest signal is operational. Additionally, we apply a fiducial selection derived from simulations to include only events that fall within a zenith range where the X_{\max} estimation bias of the Deep Neural Network (DNN) for both protons and iron nuclei is less than 10 g/cm². This cut is energy-dependent and eliminates events with geometries that are challenging the X_{\max} reconstruction with SD. After applying these selection criteria, our dataset consists of 48 824 events recorded with the SD array of the Pierre Auger Observatory, spanning from January 1, 2004, to August 31, 2018.

To reconstruct the depth of the shower maximum, X_{\max} , we utilize the time-dependent particle footprint measured at each triggered WCD, employing deep neural networks. The arrival times,

along with the measured signal traces, serve as input for the algorithm. In the first step, the signal traces are analyzed and characterized into features using a Long Short-Term Memory (LSTM) network [24]. Subsequently, a Convolutional Neural Network (CNN) leverages these station-wise characteristics and the arrival times. This is achieved by employing hexagonal convolutions to analyze the spatial and time-dependent information, taking advantage of the triangular structure of the surface array.

The network is trained using a comprehensive air-shower library simulated with CORSIKA [25] and Offline [12], employing the hadronic interaction model EPOS-LHC [18]. For the training data, we utilized a composition with equal fractions of protons, helium, oxygen, and iron. During the training of the neural network, data augmentation techniques were applied to simulate realistic variations in operating conditions, such as malfunctioning PMTs or broken stations, as well as different states of the electronics.

When applied to simulations using the hadronic interaction models QGSJet II-04 [17] and Sibyll 2.3d [11], which were not used during training, our studies found a bias in $\langle X_{\max} \rangle$ amounting to up to -15 g/cm^2 . In contrast, the estimations of $\sigma(X_{\max})$ and the method's resolution do not exhibit a strong dependency on the interaction models.

To eliminate the dependency on the interaction model used during training, we calibrate the X_{\max} reconstruction of the DNN to the X_{\max} measurements from the FD. The application to hybrid data is illustrated in Fig. 5. A strong correlation between the reconstructions is evident in the left figure, yielding a Pearson correlation coefficient of $\rho = 0.7$. This result aligns well with the correlation expected from simulation studies, assuming the same energy spectrum and a composition scenario inferred from the FD data. However, an offset of approximately -31 g/cm^2 between the SD and FD reconstructions is visible in the right figure. This offset is larger than the one predicted by simulation studies, which anticipated biases of up to approximately -15 g/cm^2 . This substantial bias observed in the data reflects discrepancies between the current generation of hadronic interaction models and the data, particularly in the modeling of the muonic component and minor remaining differences in the simulated detector response, including atmospheric models. Our simulation studies predicted that any bias would be approximately independent of energy. As shown in the right figure, no significant deviations from this expectation are apparent in the data. Therefore, we perform an energy-independent calibration of the SD X_{\max} , although our tests using several energy-dependent models found no significant impact on the results. Any potential energy dependence is accounted for as a systematic uncertainty.

Since the resolution of the neural network (see Fig. 5, middle) and $\sigma(X_{\max})$ do not exhibit a strong dependence on the hadronic interaction model, we do not perform a calibration of the second moment. This conclusion is based on the observation that fluctuations in X_{\max} generally vary little between the models. Nevertheless, the systematic differences of $\leq 5 \text{ gcm}$ found in the reconstruction of $\sigma(X_{\max})$ when applying the network to various compositions and interaction models are propagated into the systematic uncertainties of the $\sigma(X_{\max})$ measurement.

The X_{\max} reconstruction using the SD data allows us to conduct high-statistics studies of the energy evolution of the first and second moments of the X_{\max} distributions. This increase in statistics by a factor of 10 compared to previous analyses based on the FD significantly enhances our ability to investigate changes in the $\langle X_{\max} \rangle$ elongation rate. The elongation rates measured by the FD and the SD agree well (see Fig. 6). A characteristic structure can be identified when fitting more complex

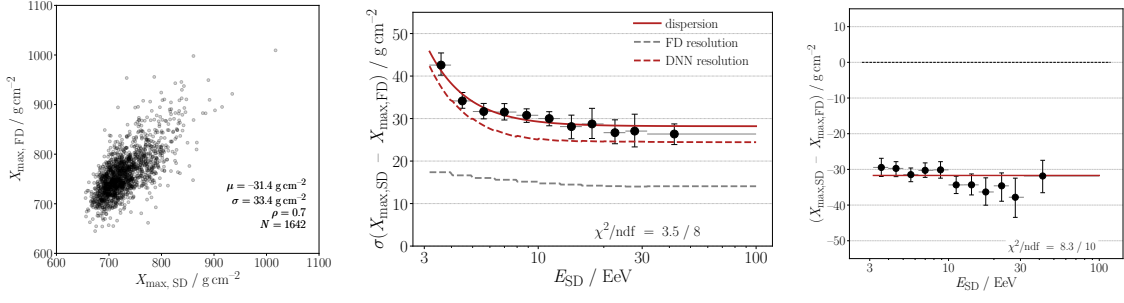


Figure 5: Comparison of X_{\max} reconstructed with FD and SD (left) which gives the DNN resolution (middle) and bias (right).

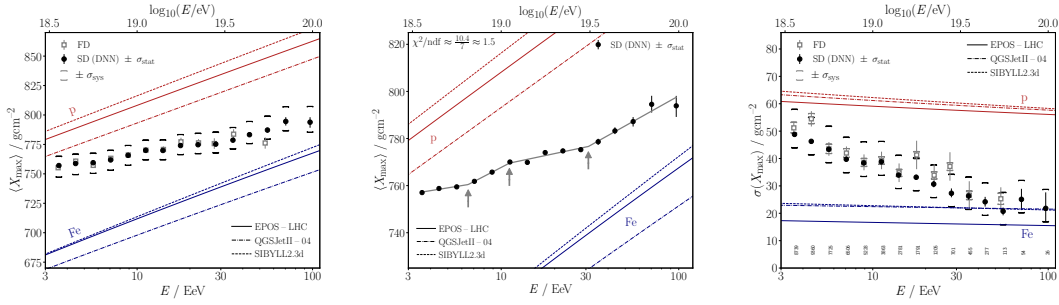


Figure 6: Mean $\langle X_{\max} \rangle$ (left) and $\sigma(X_{\max})$ (right) of the X_{\max} distributions versus energy for SD and FD. Middle: broken-line fit of the SD $\langle X_{\max} \rangle$ elongation rate. As comparison predictions from hadronic models QGJSJet II-04 [17], EPOS-LHC [18] and Sybill2.3d [11] are also displayed.

models, suggesting three breaks in the elongation rate, as presented in Fig. 6 (middle). Within this model, the elongation rate changes from 12 ± 5 g/cm²/decade to 39 ± 9 g/cm²/decade at $E_0 = 6.5 \pm 0.6$ EeV. At an energy of $E_1 = 11 \pm 1.6$ EeV, the elongation rate shifts to 16 ± 3 g/cm²/decade before reaching 42 ± 9 g/cm²/decade above $E_2 = 31 \pm 5$ EeV. The hypothesis of a constant elongation rate can be rejected at 4.4σ .

Similar to the FD X_{\max} moments, the SD X_{\max} moments can be converted into $\ln A$ moments using different hadronic models. When examining the fluctuations in $\ln A$ (bottom row in Fig. 7), a small degree of mass mixing can be identified at energies above 10^{19} eV. While the inferences from EPOS-LHC and Sibyll 2.3d remain within the physically allowed ranges for $\sigma^2(\ln A)$, the negative fluctuation estimates from QGJSJet II-04 disfavor this hadronic interaction model.

4. Conclusion

In this proceeding, we present the latest results on mass composition obtained from the Phase I data of the Pierre Auger Observatory. The X_{\max} reconstruction has been achieved using both the classical FD reconstruction and a new SD reconstruction obtained through DNN analysis [26]. Utilizing the FD data, we confirm previous findings, particularly the transition from a heavier to a lighter composition up to $10^{18.4}$ eV, followed by a shift to a heavier and purer composition. The SD X_{\max} evolution, with higher statistics above $10^{18.5}$ eV, reveals more complex behavior in the

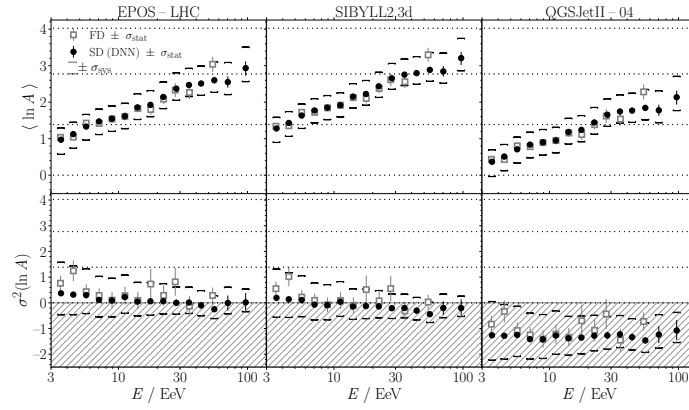


Figure 7: Mean and variance of the logarithm of the mass number $\ln A$ inferred from the SD X_{\max} using different hadronic models.

elongation rate, possibly linked to spectral features [27]. Both measurements have shown large discrepancy to the QGSJet II-04 predictions for $\sigma(X_{\max})$, suggesting not to use this model for air-shower simulation.

References

- [1] I. Allekotte et al., [Pierre Auger Coll. Coll.], *Nucl. Instrum. Meth. A* **586** (2008) 409.
- [2] J. Abraham et al., [Pierre Auger Coll. Coll.], *Nucl. Instrum. Meth. A* **620** (2010) 227.
- [3] T.H.-J. Mathes, [Pierre Auger Coll. Coll.], *EPJ Web Conf.* **3** (2011) 153.
- [4] A. Aab et al., *Nucl. Instrum. Meth. A* **798** (2015) 172.
- [5] A. Aab et al., *arXiv:1604.03637 [astro-ph]* (2016) arXiv: 1604.03637.
- [6] K.-H. Kampert and M. Unger, *Astropart. Phys.* **35** (2012) 660.
- [7] A. Aab et al., *JINST* **16** (2021) P07019.
- [8] A. Aab et al., [Pierre Auger Coll.], *Phys. Rev. D* **90** (2014) 122005.
- [9] T. Fitoussi, [Pierre Auger Coll. Coll.], *EPJ Web Conf.* **444** (2023) 319.
- [10] T. Bergmann et al., *Astropart. Phys.* **26** (2007) 063002.
- [11] F. Riehn et al., *Phys. Rev. D* **102** (2020) 063002.
- [12] S. Argiro et al., *Nucl. Instrum. Methods Phys. Res.* **580** (2007) 1485.
- [13] P. Abreu et al., [Pierre Auger Coll.], *Astropart. Phys.* **34** (2011) 368.
- [14] V.M. Harvey, [Pierre Auger Coll.], *EPJ Web Conf.* **444** (2023) 300.
- [15] J. Bellido, [Pierre Auger Coll.], *EPJ Web Conf.* **301** (2018) 506.
- [16] A. Yushkov, [Pierre Auger Coll.], *EPJ Web Conf.* **358** (2021) 482.
- [17] S. Ostapchenko, *Phys. Rev. D* **81** (2010) 114028.
- [18] T. Pierog et al., *Phys. Rev. C* **92** (2015) 034906.
- [19] A. Aab et al., [Pierre Auger Coll.], *Journal of Cosmology and Astroparticle Physics* **2013** (2013) .
- [20] D.W. Hogg and D. Foreman-Mackey, *ApJS* **236** (2018) 11.
- [21] O. Tkachenko, [Pierre Auger Coll.], *EPJ Web Conf.* **444** (2023) 438.
- [22] A. Aab et al., [Pierre Auger Coll.], *Phys. Rev. D* **90** (2014) 122006.
- [23] A. Aab et al., *JINST* **15** (2020) P10021.
- [24] S. Hochreiter and J. Schmidhuber, *Neural Computation* **9** (1997) 1735.
- [25] D. Heck, J. Knapp, J.N. Capdevielle, G. Schatz and T. Thouw, *CORSIKA: a Monte Carlo code to simulate extensive air showers.* (Feb., 1998).
- [26] A. Abdul Halim et al., *Phys. Rev. D* **111** (2025) 022003.
- [27] A. Abdul Halim et al., *Phys. Rev. Lett.* **134** (2025) 021001.