

# SPECTRUM OF LIGHT COSMIC RAYS IN THE PEV ENERGY RANGE MEASURED WITH ICECUBE

Zur Erlangung des akademischen Grades eines  
**Doktors der Naturwissenschaften (Dr. rer. nat.)**  
von der KIT-Fakultät für Physik des  
Karlsruher Instituts für Technologie (KIT)  
genehmigte

## **Dissertation**

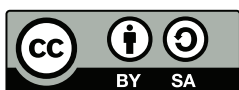
von  
Julian Philipp Saffer (M.Sc.)  
aus Kalchreuth

Referent: Prof. Dr. Ralph Engel

Korreferentin: Prof. Dr. Kathrin Valerius

Betreuer: Dr. Andreas Haungs

Tag der mündlichen Prüfung: 4. Juli 2025



Dieses Werk ist lizenziert unter einer Creative Commons Namensnennung –  
Weitergabe unter gleichen Bedingungen 4.0 International Lizenz (CC BY-SA 4.0):  
<https://creativecommons.org/licenses/by-sa/4.0/deed.de>

## ABSTRACT

### **Spectrum of Light Cosmic Rays in the PeV Energy Range Measured with IceCube**

The all-particle energy spectrum of cosmic rays spans over 11 orders of magnitude and has been measured by numerous experiments. Probing cosmic radiation above a few hundred TeV is only possible via the indirect approach of detecting extensive air showers. With this method, mass reconstructions are generally only capable of assigning a so-called mass group to an individual event, or they are even limited to statistical inference such as reporting mass group fractions or an average mass. The IceCube Neutrino Observatory consists of a  $1\text{ km}^3$  array of optical instruments deep in the ice of Antarctica as well as an arrangement of sensors at the surface for the detection of particle cascades, which are initiated by cosmic rays reaching Earth. Selecting events that are registered in both parts of this particle telescope in coincidence opens the possibility of reconstructing the energy and mass of the primary cosmic ray.

This analysis provides a reconstruction of the combined spectrum of protons and helium nuclei between 200 TeV and 10 PeV using IceCube data. It is the first measurement of the light cosmic-ray flux in the energy range 300 TeV–1.4 PeV using post-LHC hadronic interaction models to interpret the measurements. This energy range is of particular interest as it marks the transition from direct detection of cosmic radiation to indirect measurements with air showers. The PeV energy range exhibits a deficiency in composition measurements. Determining the spectrum of light cosmic rays contributes in constraining models of the astrophysical origin of those highly energetic particles.

In the framework of this analysis, a large database of event simulations has been created, including three of the latest hadronic interaction models and five different primary particle types. A neural network has been constructed, which is trained and applied to air-shower event data. It takes information from the surface and the in-ice detector as well as from a hybrid shower reconstruction and returns predictions of primary energy and type.

The obtained light cosmic-ray spectrum is in good agreement with astrophysical models that claim a power law with a spectral index of approximately  $-2.7$ . Furthermore, it is in agreement with the assumption of protons and helium nuclei dominating the cosmic-ray flux below the knee. Minor differences are observed between the results based on the three hadronic interaction models that have been compared. A gradual deviation from the unbroken power law towards a softer spectrum, that is, a soft knee-like structure, is indicated between 2 PeV and 8 PeV, depending on the hadronic interaction model.



# ZUSAMMENFASSUNG

## **Messung des Spektrums der leichten kosmischen Strahlung im PeV-Energiebereich mit IceCube**

Das Gesamtenergiespektrum der kosmischen Strahlung erstreckt sich über 11 Größenordnungen und wurde von zahlreichen Experimenten gemessen. Die Untersuchung der kosmischen Strahlung oberhalb einiger hundert TeV ist nur indirekt möglich, indem ausgedehnte Luftschauer detektiert werden. Bei dieser Methode sind Rekonstruktionen der Masse generell nur in der Lage, einem einzelnen Event eine sogenannte Massengruppe zuzuordnen, oder sie beschränken sich auf eine statistische Auswertung, indem Massengruppenanteile oder die Durchschnittsmasse bestimmt werden. Das IceCube Neutrino-Observatorium besteht aus einem Array optischer Instrumente innerhalb eines Volumens von  $1 \text{ km}^3$  tief im Eis der Antarktis, sowie einer Anordnung von Detektoren an der Oberfläche für den Nachweis von Teilchenkaskaden, die durch kosmische Strahlung, welche die Erde erreicht, ausgelöst werden. Die Selektion von Ereignissen, die in beiden Teilen dieses Teilchenteleskops gleichzeitig registriert werden, eröffnet die Möglichkeit, die Energie und die Masse der Primärteilchen der kosmischen Strahlung zu rekonstruieren.

Diese Analyse bietet eine Rekonstruktion des kombinierten Spektrums von Protonen und Heliumkernen zwischen 200 TeV und 10 PeV unter Verwendung von IceCube-Daten. Es ist die erste Messung des Flusses der leichten kosmischen Strahlung im Energiebereich von 300 TeV bis 1.4 PeV, bei der post-LHC hadronische Interaktionsmodelle für die Interpretation der Messungen verwendet werden. Dieser Energiebereich ist von besonderem Interesse, da er den Übergang vom direkten Nachweis der kosmischen Strahlung zu indirekten Messungen mit Luftschauern markiert. Der PeV-Energiebereich weist ein Defizit an Messungen der Zusammensetzung kosmischer Strahlung auf. Die Bestimmung des Spektrums der leichten kosmischen Strahlung trägt dazu bei, die Modelle für den astrophysikalischen Ursprung dieser hochenergetischen Teilchen einzuschränken.

Im Rahmen dieser Analyse wurde eine große Datenbank mit Ereignissimulationen erstellt, die drei der neuesten hadronischen Wechselwirkungsmodelle und fünf verschiedene Primärteilchentypen umfasst. Ein neuronales Netz wurde konstruiert, trainiert und auf Luftschauerdaten angewendet. Es erhält Informationen von der Oberfläche und dem Detektor im Eis sowie von einer hybriden Schauer-Rekonstruktion und liefert Vorhersagen zu Primärenergie und -typ.

Das erhaltene Spektrum der leichten kosmischen Strahlung steht in guter Übereinstimmung mit astrophysikalischen Modellen, welche ein Potenzgesetz mit einem Spektralindex von etwa  $-2,7$  annehmen. Außerdem stimmt es mit der Annahme überein, dass Protonen und Heliumkerne den kosmischen Strahlungsfluss unterhalb des Knies dominieren. Zwischen den Ergebnissen, die auf den drei verglichenen hadronischen Wechselwirkungsmodellen basieren, werden geringfügige Unterschiede beobachtet. Eine allmähliche Abweichung vom ungebrochenen Potenzgesetz in Richtung eines weicheren Spektrums, d.h. einer runden knieähnlichen Struktur, zeigt sich zwischen 2 PeV und 8 PeV, je nach hadronischem Modell.



# CONTENTS

<b>Abstract</b>	<b>i</b>
<b>Zusammenfassung</b>	<b>iii</b>
<b>Contents</b>	<b>v</b>
<b>Introduction</b>	<b>1</b>
<b>Chapter 1: Cosmic Rays</b>	<b>5</b>
1.1 Energy Spectrum	5
1.2 Mass Composition	7
1.3 Sources, Acceleration and Propagation	8
1.4 Direct Detection	12
<b>Chapter 2: Extensive Air Showers</b>	<b>13</b>
2.1 Air-Shower Development	13
2.2 Air-Shower Components	17
2.3 Detection Methods	20
2.4 Air-Shower Experiments	22
<b>Chapter 3: Machine Learning</b>	<b>25</b>
3.1 Introduction to Machine Learning	25
3.2 Neural Networks	27
3.3 Model Training	30
<b>Chapter 4: The IceCube Neutrino Observatory</b>	<b>35</b>
4.1 Detector Components	35
4.1.1 In-Ice Array	36
4.1.2 IceTop	38
4.2 IceCube Online System	39
4.2.1 Triggering	40
4.2.2 Filtering	43
<b>Chapter 5: Event Simulation</b>	<b>45</b>
5.1 CORSIKA	45
5.1.1 Air-Shower Simulation with CORSIKA	45
5.1.2 Simulation Production	47
5.2 Detector Simulation	49
5.3 Injection of In-Ice Background	51
5.4 Environmental Effects	54
5.4.1 Atmosphere	54
5.4.2 Snow	57
5.4.3 Ice Model	58
<b>Chapter 6: Offline Processing and Reconstruction</b>	<b>61</b>
6.1 Level 2	61
6.2 Level 3	63
6.3 SLC Calibration	66
6.3.1 Waveform Digitization	66
6.3.2 Charge Calibration Using HLCs	67
6.3.3 ATWD Crossover Points	68
6.4 Level 4	71
6.4.1 Coincident Event Selection and Pulse Cleaning	71
6.4.2 Hive splitting	72

## CONTENTS

6.4.3	In-Ice Track Reconstruction . . . . .	73
6.4.4	Background Tagging . . . . .	74
6.4.5	Connecting COGs . . . . .	75
6.4.6	RockBottom . . . . .	75
6.4.7	Background Rejection . . . . .	76
6.4.8	Containment . . . . .	78
6.4.9	In-Ice Energy Loss . . . . .	80
6.5	Monte Carlo Verification . . . . .	81
6.5.1	Event Rates . . . . .	81
6.5.2	Pulse Series . . . . .	82
<b>Chapter 7: Quality Cuts . . . . .</b>		<b>87</b>
7.1	Passed IceTop Filters . . . . .	87
7.2	In-Ice Charge . . . . .	88
7.3	RockBottom Convergence . . . . .	89
7.4	Containment . . . . .	89
7.5	Shower Size . . . . .	90
7.6	Survival Rates . . . . .	94
7.7	Data / Monte Carlo Comparison . . . . .	95
<b>Chapter 8: Light Spectrum Reconstruction . . . . .</b>		<b>99</b>
8.1	The Neural Network Model . . . . .	99
8.1.1	Model Inputs . . . . .	99
8.1.2	Network Architecture . . . . .	103
8.1.3	Training . . . . .	104
8.1.4	Model Evaluation . . . . .	105
8.2	Light Fraction Calibration . . . . .	109
8.2.1	Method 1: agnostic . . . . .	109
8.2.2	Method 2: GSF-informed . . . . .	112
8.3	Efficiency Correction . . . . .	113
8.4	Pressure Correction . . . . .	117
8.5	Environmental Variation . . . . .	119
8.6	Recovery Tests . . . . .	119
<b>Chapter 9: Results . . . . .</b>		<b>123</b>
9.1	Light Cosmic-Ray Fraction . . . . .	124
9.2	Differential Light Cosmic-Ray Flux . . . . .	124
9.3	Comparison of Hadronic Interaction Models . . . . .	129
9.4	Comparison with Composition Models . . . . .	130
9.5	Comparison with Other Experiments . . . . .	132
<b>Chapter 10: Conclusion . . . . .</b>		<b>135</b>
<b>Appendix A: Addendum Monte Carlo Verification . . . . .</b>		<b>139</b>
<b>Appendix B: Light Fraction Calibration . . . . .</b>		<b>143</b>
<b>Appendix C: Light Flux with QGSJet-II.04 and EPOS-LHC . . . . .</b>		<b>155</b>
<b>Appendix D: Confidence Intervals for Efficiency Ratios . . . . .</b>		<b>161</b>
<b>Bibliography . . . . .</b>		<b>165</b>
<b>Acknowledgments . . . . .</b>		<b>175</b>
<b>List of Acronyms and Initialisms . . . . .</b>		<b>177</b>
<b>List of Symbols and Constants . . . . .</b>		<b>181</b>

# INTRODUCTION

*“If you can stand books, then you might like this.”*

— Diane Morgan as *Philomena Cunk*

The Universe is permeated by high-energy particles. Hadronic cosmic rays as well as elementary particles such as charged leptons, neutrinos and gamma rays reach energies of literally astronomical magnitude, making them propagate through galaxies and intergalactic space at relativistic speed. With the discovery of cosmic radiation in the 1910s [1], the foundation for astroparticle physics has been laid. Over the past decades, this field of research — also known as high-energy physics — has been evolving to a highly sophisticated scientific discipline. In recent years, the advent of multi-messenger astronomy opened up the possibility of observing distant astrophysical objects not just in the wide electromagnetic spectrum, but also probing them by the means of high-energy neutrinos [2] and gravitational waves [3]. Astrophysical neutrinos are generated by cosmic rays and also the majority of the cosmic gamma radiation originate from hadronic processes. Therefore, it is important to study all those messenger particles as well as gravitational waves. Observations with several detectors in coincidence help shed light on the origin of cosmic rays and the violent processes in their sources required for them to obtain energies far beyond what is possible to achieve with particle accelerators on Earth.

The IceCube Neutrino Observatory is an exceptional particle detector, sensitive not only to astrophysical neutrinos but also to hadronic cosmic radiation and gamma rays. IceCube can detect cosmic rays indirectly by measuring the distribution of secondary particles that are created by the primary hadron or nucleus as it enters the atmosphere. The square-kilometer surface array IceTop detects the electromagnetic component of these air showers as well as low-energy muons. The volumetric array of optical sensors embedded in the Antarctic ice below measures the amount of Cherenkov light emitted by high-energy muons that propagate up to a few kilometers through the glacier. When an air-shower event is registered by both IceTop and the in-ice detector, the energy as well as the type of the primary cosmic ray can be estimated. However, the calorimetric air-shower measurement is prone to uncertainties much larger than those with direct detection devices, which are commonly used at lower energies. Furthermore, the distinction of primary mass groups suffers from large fluctuations during the shower development.

In the range from several hundred TeV to a few PeV primary energy, a lack of composition measurements is pointed out. At higher energies, air-shower observatories have measured the mass components of cosmic rays, whereas satellite and balloon experiments identify primary hadrons directly up to about  $10^{14.5}$  eV. The objective of this work is to close this gap for the light flux component of protons and helium nuclei with data taken by IceCube.

## Thesis Overview

- Cosmic rays are introduced in [chapter 1](#), illustrating their wide and steep energy spectrum and the changing mass composition. Possible sources of cosmic radiation are described on the basis of models for acceleration and propagation of charged particles in magnetic fields.
- The formation and detection of extensive air showers is the subject of [chapter 2](#). The Heitler and Heitler–Matthews models are explained in detail. Particle distributions of the distinct shower components are discussed and illustrated using simulations. Different experiments using various types of air-shower detectors are reviewed.
- The theory part is concluded with an introduction to machine learning in [chapter 3](#). After a general overview, the focus is put on neural networks, including their architecture and training concept.
- The IceCube Neutrino Observatory is presented in [chapter 4](#). The hardware section provides details of the detector components, especially the optical in-ice array and the IceTop stations at the surface. This is followed by an elaboration of the online computing system, which is responsible for triggering detector readout and creating event filters.
- The process of Monte Carlo simulation is covered in [chapter 5](#) and starts with the simulation of air showers with CORSIKA. Subsequently, the simulation setup for the detector response is developed, introducing the injection of in-ice muon background. Finally, the effects of the atmospheric condition, snow height and the assumed ice model are discussed.
- The data processing levels 2 and 3 as well as the newly developed level 4 are the subject of [chapter 6](#). This includes calibrating and cleaning of pulses, shower reconstruction methods and the comparison of Monte Carlo distributions with measured data.
- Quality cuts have been defined specifically for this analysis. They are motivated and demonstrated one by one in [chapter 7](#).
- The analysis method is described in [chapter 8](#). The chosen architecture and inputs to the neural network model are presented, followed by the calibration of predicted primary energy and mass. Two methods for the estimation of the proton+helium fraction are introduced. The systematic uncertainty on the reconstructed light cosmic-ray flux comprises five contributions.
- The results of applying the developed analysis to one year of IceCube data is presented in [chapter 9](#). After a cross-check with previously published all-particle spectra by IceCube, the reconstructed light flux is compared to models of cosmic-ray composition as well as other experiments that measured the flux of protons and helium nuclei. Eventually, a comparison of hadronic interaction models is provided.

- In [chapter 10](#), the thesis is concluded with a summary of the analysis and the obtained results, highlighting the contribution of this work in the context of cosmic-ray composition around the spectral knee. Lastly, possible future enhancements of this analysis are discussed in an outlook.



# CHAPTER 1

## COSMIC RAYS

*“You people pride yourself on your roots, right? Well, this is as rooty as it gets.”*

— Jamie Zvirzdin, *Subatomic Writing*

Cosmic rays (CRs) are highly energetic charged particles moving through space at relativistic speeds. They are mostly composed of fully ionized atomic nuclei but also electrons and antiparticles contribute to the flux of cosmic radiation. At low energy, the mass fractions of hadronic CRs can be broken down to approximately 90% free protons (i.e., hydrogen nuclei), about 9%  $\alpha$ -particles (helium nuclei) and a remaining small portion of heavier nuclei, reaching up at least to the chemical element iron<sup>1</sup> [4]. Wider definitions extend this group of particles to electrons,  $\gamma$ -rays and even neutrinos. However, in this work, the term *cosmic rays* refers to the classical notion comprising hadronic CRs only. Depending on their origin, cosmic rays can reach energies from a few GeV up to hundreds of EeV [5, 6], which equals tens of Joules.

Since their discovery by Hess in 1912 [1], cosmic rays have been studied for over a century, leading to groundbreaking discoveries in particle physics (like the first detection of the positron [7], the muon [8], the pion [9] and the kaon [10]) and allowing for new insights into our Galaxy and beyond. But, as always, when the boundaries of the known are expanded, CR research also opens up new questions about the Universe that still await to be answered.

After introducing the cosmic-ray energy spectrum in [section 1.1](#), one of the not yet fully understood aspects of the spectrum and the main objective of this work, the elemental composition of CRs, is discussed in [section 1.2](#). Potential sources of cosmic radiation are listed in [section 1.3](#), which also covers the acceleration of CRs and their propagation through space. Finally, direct detection methods are presented in [section 1.4](#).

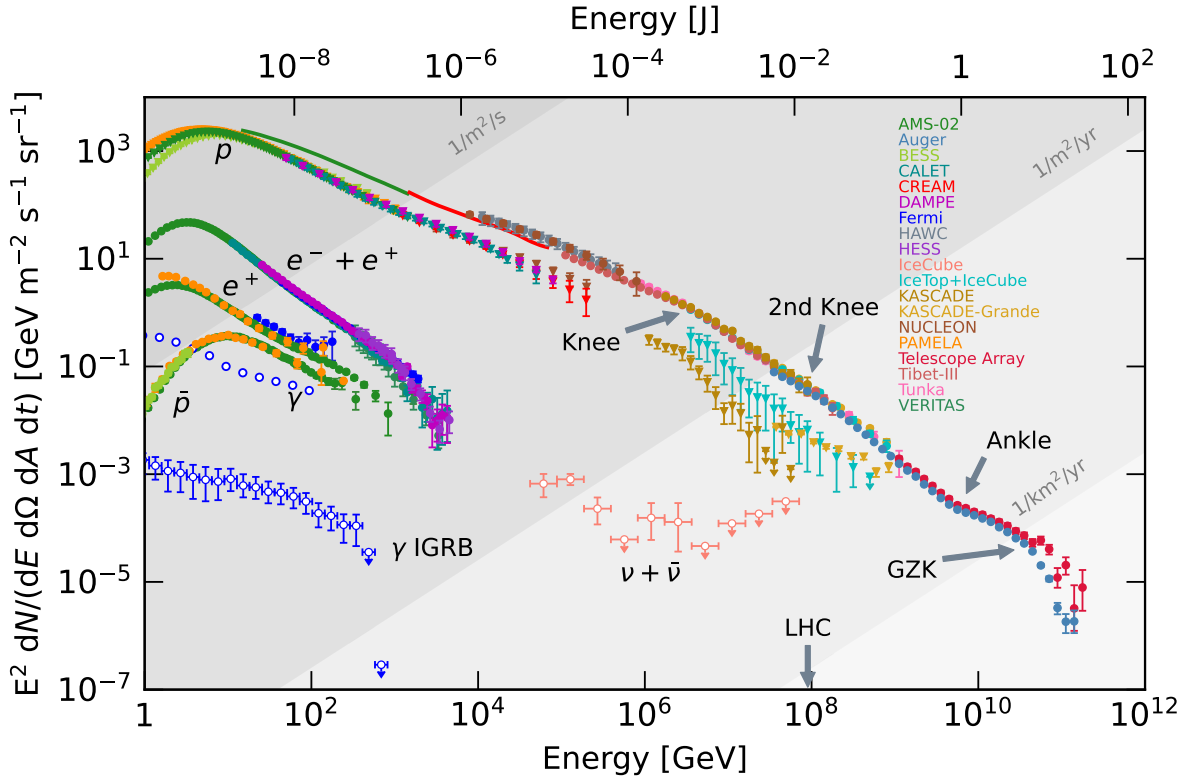
### 1.1 Energy Spectrum

The energies of cosmic rays span an enormous range, starting at the barely relativistic GeV regime and ending at ultra-high energies of about a hundred EeV, which is almost 12 orders of magnitude. No human-made particle accelerator is capable of bringing hadrons to such extreme energies. [Figure 1.1](#) illustrates the spectrum of cosmic rays and reveals its steep fall. Low-energy cosmic rays make up the bulk of the flux with several particles/m<sup>2</sup>/s. With increasing energy the abundance drops rapidly to only a few per square meter and year at PeV energies to less than one per km<sup>2</sup> in a year towards the high-energy end of the spectrum. The differential CR flux

$$J = \frac{d\Phi}{dE} = \frac{d^4N}{dE d\Omega dA dt} \quad (1.1)$$

---

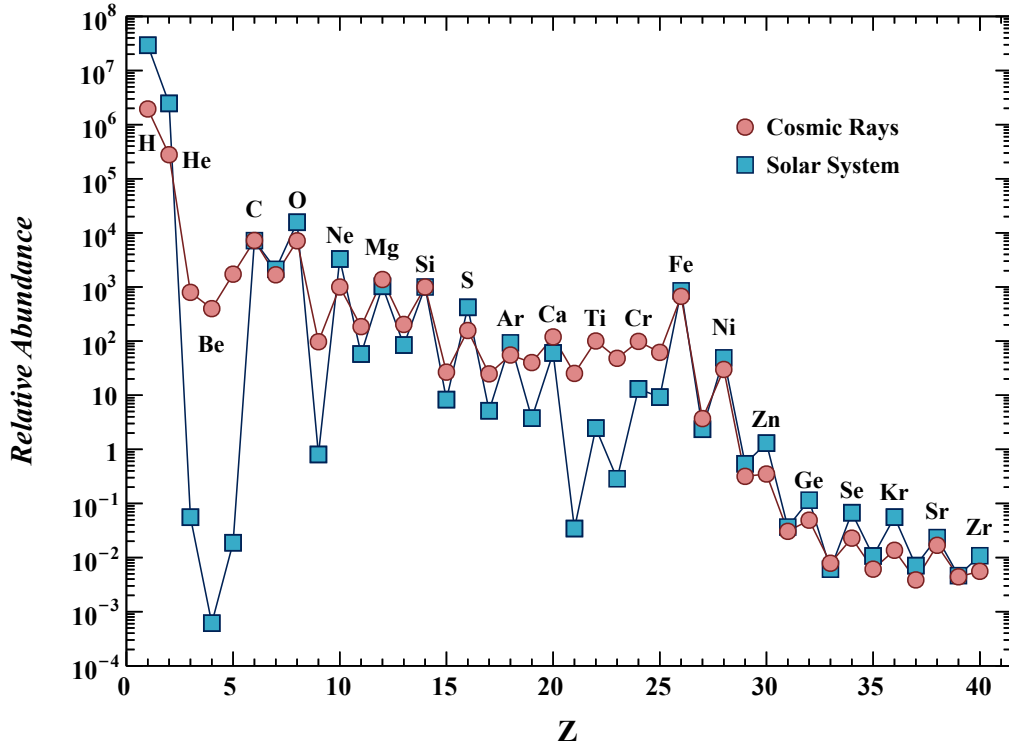
<sup>1</sup>Ultraheavy cosmic rays with atomic number  $Z > 26$  do exist, although their flux is highly suppressed.



**Figure 1.1:** The energy spectrum of cosmic rays, including electrons, positrons,  $\gamma$ -rays and neutrinos as measured by numerous experiments. Proton spectra are plotted with triangles and solid lines represent the all-particle spectra. The indicated rates labeling the shaded areas assume a  $4\pi$  field of view. The maximum energy achieved in proton–proton collisions at the large hadron collider (LHC) is marked at the bottom. Plot adopted from Ref. [11] and modified.

is the number of cosmic-ray primaries  $N$  per energy interval, solid angle, area and time. For better visualization,  $J$  is commonly multiplied by a power of the energy (in Fig. 1.1 by  $E^2$ ). This facilitates the identification of features, specifically changes in the spectral slope.

At the lowest energies ( $\lesssim 10$  GeV), the spectrum is dominated by what is called solar modulation [12]. The Sun not only emits GeV protons but also a plasma, the solar wind, which suppresses the flux of incoming GeV CRs from outside the Solar System. Additionally, coronal mass ejections (CMEs) can contribute temporarily to the flux of GeV cosmic rays but also lead to a flux reduction a few days after the CME, a phenomenon called *Forbush decrease* [4]. The flux of GeV cosmic rays is thus generally time-dependent. With higher energy, the CR spectrum can be described as a power law  $\propto E^{-\gamma}$  with the differential spectral index  $\gamma \approx 2.7$ . At about 3 PeV the spectrum softens; that is, it gets steeper, and at ca. 100 PeV it steepens once more, arriving at a slope of approximately 3.1. These two bends in the spectrum are called the *knee* and *second knee*, respectively. At about 5 EeV the spectrum gets harder again, continuing with a  $\gamma$  of roughly 2.6. This feature is dubbed the *ankle*, which is eventually followed by the spectral cutoff above 100 EeV. A possible explanation for the origin of these structures will be given in section 1.3.

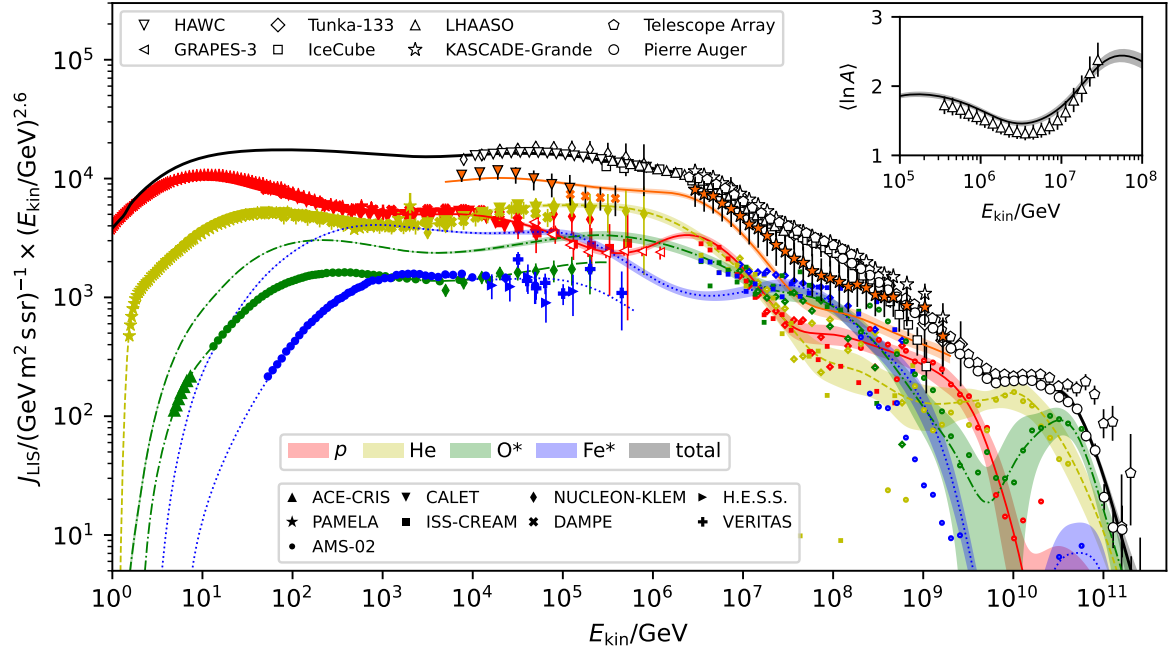


**Figure 1.2:** Relative abundance of chemical elements in cosmic rays and the Solar System. The values are normalized for silicon to be at  $10^3$ . Plot adopted from Ref. [13] (Fig. 30.2) and modified.

## 1.2 Mass Composition

The types of nuclei in the cosmic radiation are not uniformly distributed. The relative abundance of chemical elements in low-energy cosmic rays is shown in Figure 1.2 alongside the quantities for the Solar System. Both distributions show a similar trend, such as the general decrease in abundance with heavier mass and nuclei that have an even number of protons  $Z$  being more abundant due to their higher stability. However, two clear differences are visible as well. The lightest elements, hydrogen and helium, are less abundant in CRs than in the Solar System material. This could be because of a different composition at the source itself or the relatively high energy needed to ionize hydrogen and helium atoms. The other apparent deviation between both graphs is that the mass groups Li–B and Sc–Mn are largely suppressed in the Solar System while they are much more abundant in cosmic rays. It is widely accepted that these nuclei are products of spallation (i.e., fragmentation) processes resulting from collisions of cosmic rays in the CNO group and iron primaries with the interstellar medium (ISM) [4].

It can be seen in Figure 1.1 that the contribution of protons to the all-particle spectrum is not constant over the observed energy range. Similarly, the flux fractions of heavier nuclei change with energy, too. Measurements of four mass groups by a variety of experiments are presented over a wide energy range in Figure 1.3 together with spline fits to the individual spectra. Depending on the energy, different particle types are dominating the CR flux. The gap in the mass group spectra at around  $10^6$  GeV is a consequence of the fundamentally different detection methods used for cosmic rays above and below the knee (→ section 1.4 and



**Figure 1.3:** All-particle cosmic-ray spectrum (black) with the flux from individual mass groups (red: proton, yellow: helium, orange: proton+helium, green: oxygen, blue: iron) together with spline fits (GSF model) to the respective elements. The inset plot shows the mean logarithmic mass  $\langle \ln(A) \rangle$  as a function of energy. Plot from Ref. [18].

section 2.3). The elemental composition of the cosmic radiation holds convoluted information on the types of sources and hence the mechanism needed to accelerate charged particles to highly relativistic speeds. Furthermore, the conditions during propagation to Earth influence the mixture of cosmic-ray primary hadrons.

There are several approaches to model the CR mass composition such as *poly-gonato*<sup>2</sup> by Hörandel [14], H3a/H4a<sup>3</sup> by Gaisser [15] and GST by Gaisser, Stanev and Tilav [16]. They all have in common that their structural features are primarily linked to the nuclear charge of the individual flux components and they attempt to explain observed features using broken power laws. This is contrasted with the *Global Spline Fit (GSF)* by Dembinski et al. [17] (Fig. 1.3). The purpose of GSF is to only describe the observed cosmic-ray spectra with cubic B-splines. Taking into account the statistical and systematic uncertainties of numerous direct detectors and indirect observatories, a best-fit model with a covariance matrix has been obtained.

### 1.3 Sources, Acceleration and Propagation

Interstellar space is permeated by magnetic fields with a strength up to a few  $\mu\text{G}$  in our Galaxy while extragalactic magnetic fields are typically weaker than  $1\text{ nG}$ <sup>4</sup> [20]. Because of their electric charge, CRs are deflected in these magnetic fields because of the Lorentz force.

<sup>2</sup>Greek for *many knees*.

<sup>3</sup>Hillas model with three source populations or a fourth alternative proton-only population.

<sup>4</sup>For comparison: Earth's magnetic field has a strength of a few hundred mG [19].

### 1.3. SOURCES, ACCELERATION AND PROPAGATION

The resistance against this magnetic deflection is called *rigidity*

$$R = \frac{pc}{Ze} \quad (1.2)$$

with  $pc$  being the particle's kinetic energy<sup>5</sup> and  $Ze$  its charge. It can be seen from Eq. (1.2) that rigidity of cosmic rays increases with their energy and decreases with their atomic number. It is for that reason that only ultra-high-energy cosmic rays (**UHECRs**,  $E \gtrsim 10^9$  GeV) when they arrive at Earth point approximately back to their source as long as it is not farther away than  $O(100)$  Mpc.

#### Neutral Particles

One objective of multi-messenger astronomy is therefore the identification of cosmic-ray sources via the detection of neutral particles, which are expected to be produced by the charged nuclei. Photons do not experience the Lorenz force, which enables  $\gamma$ -ray astronomers to determine their direction of origin in the sky. However, at high energies ( $\sim 1$  PeV), gamma rays have an increasing probability of pair production

$$\gamma_{\text{HE}} + \gamma_{\text{CMB}} \rightarrow e^+ + e^-$$

via interaction with photons of the *cosmic microwave background* (**CMB**) and the *cosmic infrared background*, which limits their range to the galaxy of their production [21].

Since neutrinos have an extremely small interaction cross section, they can travel cosmological distances without any deflection and hardly any absorption. This makes neutrinos excellent messenger particles in the endeavor to reveal the sources of high-energy cosmic rays, although it also makes their detection very challenging.

#### Fermi Acceleration

The first to bring supernovae as potential sources of cosmic rays into play were Baade and Zwicky [22]. It is now believed that not the supernova explosion itself but rather its heir, the supernova remnant (**SNR**), is responsible for accelerating the bulk of high-energy CRs that reach Earth. Fermi published a first attempt to describe the acceleration process of cosmic rays in a moving cloud of plasma [23]. As a charged particle enters the cloud, it diffuses via collisionless elastic scattering in the magnetic field of the plasma, which leads on average to a net gain in energy after leaving the cloud again. It turned out, however, that this process is not sufficiently efficient to describe the observed energy density of CRs. Since the average relative energy gain per encounter<sup>6</sup>  $\xi$  is proportional to the square of the non-relativistic speed of the plasma cloud  $v^2 \ll c^2$ , it is retrospectively called *second-order Fermi acceleration*. A higher efficiency can be reached by *first-order Fermi acceleration*, also called *diffusive shock acceleration*, which takes place at shock fronts between a fast high-density plasma and a low-density gas environment, like a SNR expanding into the ambient ISM. The probability distribution of the entrance and exit angle to the plasma region changes when considering a

<sup>5</sup>Since astroparticle physics takes place at highly relativistic speeds, the approximation  $E = E_{\text{kin}} = pc$  is used here, dropping the negligible rest mass.

<sup>6</sup>After  $n$  encounters, a particle with an initial energy  $E_0$  has reached  $E_n = E_0(1 + \xi)^n$ .

shock front, eventually leading to  $\xi \propto v$ ; that is, on average a charged particle gains a larger amount of energy when encountering a shock front compared to the second-order process in plasma clouds. Following the elaboration in Ref. [4], one can show that the Fermi acceleration mechanism predicts a power-law energy spectrum with an index  $\gamma = 2$  at the source. During propagation, the energy spectrum becomes softer due to energy-dependent diffusion as well as the decay and spallation of nuclei. There are models such as the description of the Galaxy as a “leaky box” that can explain the changing spectral index to match the observation on Earth.

### Maximal Energy

Having the acceleration via magnetic fields at the core of the Fermi mechanism implies a maximal energy achievable by a finitely extended accelerator. The gyroradius of a charged nucleus moving in a magnetic field  $B$  is

$$r_G = \frac{R}{B}.$$

If  $r_G$  exceeds the dimension of the acceleration site, the particle can no longer be contained inside and eventually escapes. The maximum reachable rigidity in an accelerator of size  $r$  and magnetic field strength  $B$  is thus

$$R_{\max} = Br.$$

Substituting in Eq. (1.2) results in the maximum energy

$$E_{\max}^* \sim ZeBr \tag{1.3}$$

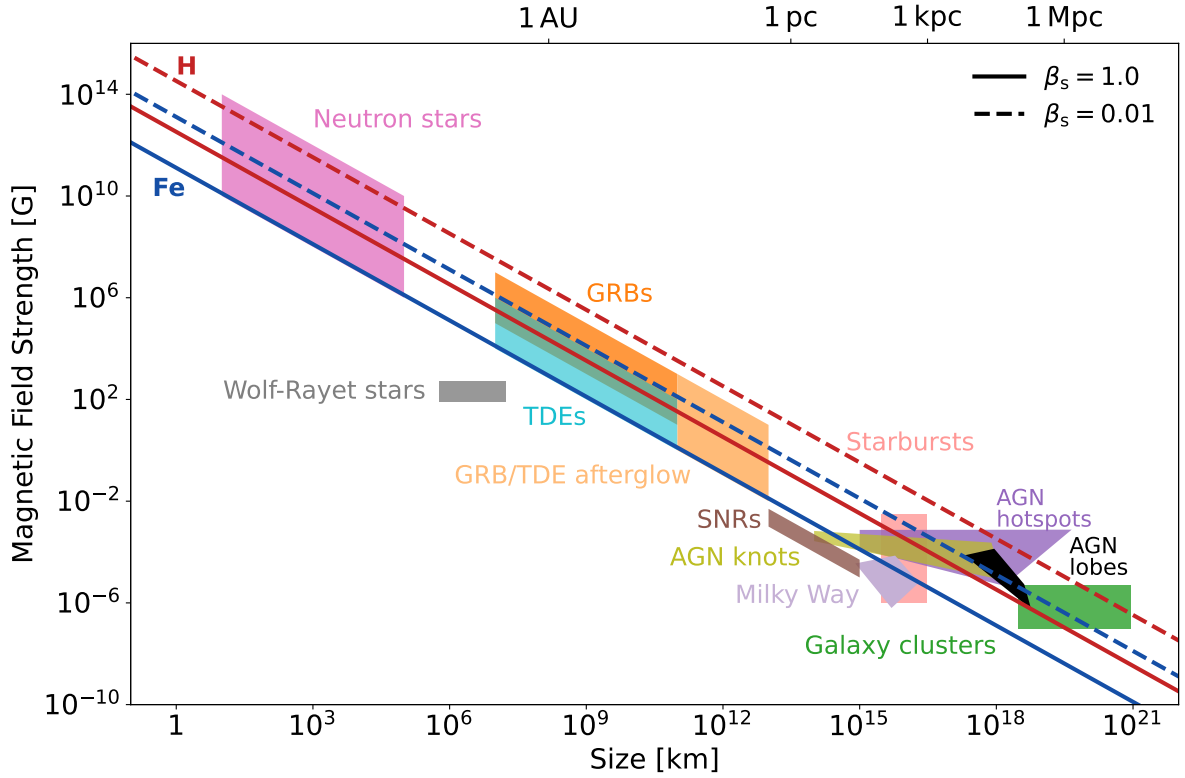
that a CR source of size  $r$  with a magnetic field of strength  $B$  can accelerate to. However, considering the diffusive shock acceleration mechanism, a correction factor for this general relation (1.3) emerges. The diffusion length

$$\lambda_D = D/\beta_s c$$

in the acceleration region is required to be much shorter than the shock radius, while  $\lambda_D$  cannot be smaller than the particle’s gyroradius, where  $D$  is the diffusion coefficient and  $\beta_s$  the shock velocity [24]. This inequality leads to a maximal energy achievable via diffusive shock acceleration of

$$E_{\max} \sim Ze\beta_s Br. \tag{1.4}$$

This relation is illustrated by the so-called Hillas plot [25] in Figure 1.4, which gives an overview of various cosmic-ray source candidates. It becomes clear from that diagram how a certain kind of accelerator can produce heavy nuclei with higher energy than it can for light nuclei like protons. This charge-dependent — and thus implicitly mass-dependent —  $E_{\max}$  (1.4) of a specific source population could be a possible explanation for the changing mass fractions in the cosmic-ray energy spectrum. Assuming that SNRs can accelerate protons to 1 PeV, they should be able to produce iron nuclei with energies up to 26 PeV. Such a sequence of mass-dependent cutoffs has been first proposed in Ref. [26] and is known as a *Peters cycle*.



**Figure 1.4:** A Hillas plot depicts potential sources of high-energy cosmic rays in the parameter space of object size and magnetic field strength. Diagonal lines represent the minimum conditions to accelerate protons (red) and iron nuclei (blue) to an energy of 100 EeV. Solid lines correspond to fast shock velocities, whereas dashed lines illustrate the scenario of slower shock fronts. Plot inspired by Ref. [27] (Fig. 10), courtesy of F. Oikonomou.

### Extragalactic Sources

While the estimated combined luminosity of supernovae in the Galaxy could sustain most of the flux of CRs that is measured at Earth, there are indications of other types of cosmic accelerators besides supernova remnants. The dimension of and the magnetic field inside SNRs results in a maximum rigidity of a few PV, which covers the spectrum only up to the second knee. Furthermore, even if another source type in the Milky Way was able to produce CRs with rigidities orders of magnitude higher than SNRs can, their gyroradius would exceed the size of the Galaxy. Since there is no evidence of nearby point sources of UHECRs, this indicates that cosmic rays beyond the second knee are of extragalactic origin. Among the source candidates for UHECRs are neutron stars, gamma-ray bursts (GRBs) and active galactic nuclei (AGNi) [28] ( $\rightarrow$  Fig. 1.4).

### Spectral Cutoff

The Peters cycle for Galactic sources gives an explanation for the position of the knee (proton cutoff) and the second knee (iron cutoff). Traditionally, the ankle is believed to mark the energy where extragalactic CRs become dominant [29] but there are indications that this transition might already take place at lower energies [30]. The end of the spectrum could be caused by another Peters cycle acting on the much more powerful UHECR accelerators. It has,

however, been predicted by Greisen [31], Zatsepin and Kuzmin [32] (GZK) that protons above  $\sim 6 \cdot 10^{19}$  eV will interact with low-energy CMB photons and lose energy via the production of pions<sup>7</sup>:

$$p_{\text{HE}} + \gamma_{\text{CMB}} \rightarrow \Delta^+ \rightarrow \begin{cases} p + \pi^0 \\ n + \pi^+ \end{cases}$$

Therefore, even if there are astrophysical objects with the ability to produce cosmic rays above this GZK limit, they would not be able to reach Earth with their initial energy. The predicted GZK cutoff energy coincides with the observed sharp drop of the energy spectrum ( $\rightarrow$  Figs. 1.1 and 1.3). Greisen also addresses the case of heavier cosmic rays, although there are competing effects like the giant dipole resonance, that might explain the spectral cutoff, too. Furthermore, heavy nuclei with ultra-high energy are prone to photodisintegration processes [31].

#### 1.4 Direct Detection

The high flux of cosmic rays below a few hundred TeV primary energy allows for their direct detection at high altitude with relatively small devices. Instruments such as calorimeters and magnetic spectrometers are carried to the top of the atmosphere by balloons as part of long-duration measurement campaigns, or they orbit Earth aboard satellites and on the International Space Station. Because of their excellent energy, charge and mass resolution, detectors like PAMELA, AMS-02 and CREAM can even distinguish between individual isotopes and identify antimatter. It is thanks to these direct observations that there is a broad understanding of the cosmic-ray spectrum up to about 100 TeV.

With higher energy, the detection rate of air- and spaceborne experiments (with their geometric acceptance limited to a few  $\text{m}^2 \text{sr}$ ) quickly drops below a few CRs per day. It is therefore necessary to measure the cosmic radiation over a much larger area, which is only possible in an indirect manner on the ground. The concept of air-shower experiments will be covered in the next chapter.

#### Key Messages

*The energy range of cosmic radiation spans multiple orders of magnitude. The spectral features and changes in the chemical composition can be described by elaborate models, although a full explanation has not been found yet. Measurements of the combined proton and helium flux are missing below and around 1 PeV primary energy, although there are proton spectra and measurements of the mean logarithmic atomic mass. The goal of this work is to provide a measurement of the H+He spectrum in that energy range with IceCube.*

---

<sup>7</sup>The neutron subsequently decays via  $n \rightarrow p + e^- + \bar{\nu}_e$ .

# CHAPTER 2

## EXTENSIVE AIR SHOWERS

*“Twinkle twinkle little star,  
Supernova, au revoir!  
You got so big, too big perhaps,  
Electron capture, core collapse.  
Twinkle twinkle former star,  
A black hole’s all that you now are.”*

— Henry Reich, «Astronomically Correct Twinkle Twinkle»

Primary [CRs](#) arriving at Earth generally will not reach the ground except in very high terrain. They will instead interact with the atmosphere and produce a cascade of secondary particles in a sequence of numerous interactions. Because of the steeply falling energy spectrum of cosmic rays, the rate of secondary particles detected on the ground (or below) is dominated by events with low initial energy.

The foundation for the discovery of these particle showers was laid by Kolhörster and Bothe, who connected spatially separated Geiger-Müller counters to allow for the registration of temporal coincidences [\[33\]](#). This pioneering setup was then developed further in the early 1930s by Rossi, who observed secondary shower particles on the ground in vertical as well as horizontal coincidence arrangements [\[34\]](#). Later, Auger and his group extended the spatial separation of the detectors up to 300 m, thereby providing access to larger air-shower cascades initiated by cosmic rays exceeding 1 PeV primary energy [\[35\]](#).

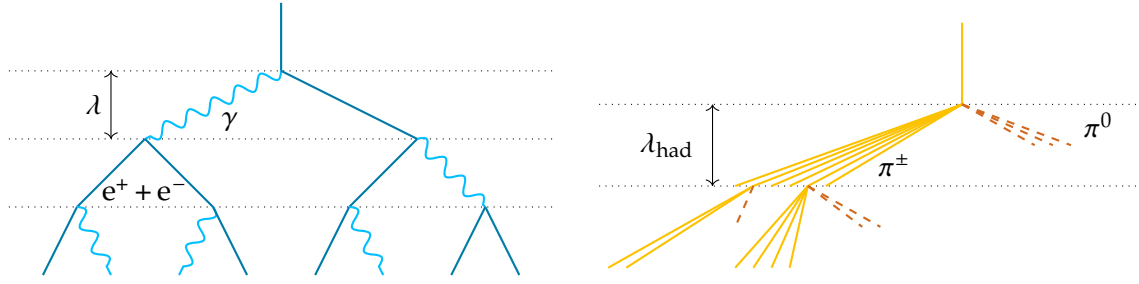
This chapter covers the production of *extensive air showers* ([EASs](#)) in [section 2.1](#) and the particle species constituting them ([section 2.2](#)). After that, [section 2.3](#) gives an overview of different approaches for air-shower detection and some ground-based experiments are introduced with their respective properties in [section 2.4](#).

### 2.1 Air-Shower Development

The first interaction of an incoming cosmic ray with the atmosphere typically occurs at a height  $h$  of about 20–30 km above sea level. Since the rate of interaction and energy loss depends on the amount of traversed matter, it is customary, however, to instead consider the slant depth

$$X(h) = \int_h^\infty \frac{\rho_{\text{air}}(h')}{\cos(\theta)} dh',$$

where  $\theta$  denotes the zenith angle and  $\rho_{\text{air}}$  is the density of air, which can be approximated as an exponential function of height (more details in [subsection 5.4.1](#)). Vertical air showers are thus initiated at an atmospheric slant depth of around 25 g/cm<sup>2</sup>.



**Figure 2.1:** *Left:* Schematic of the Heitler model for an electron-induced shower. At each step of length  $\lambda$ , either a bremsstrahlung photon is emitted (wave lines) or an  $e^\pm$  pair is created (straight lines). *Right:* Heitler–Matthews model for a hadron-induced shower. At each step of length  $\lambda_{\text{had}}$ , multiple pions are produced:  $\frac{2}{3}$  charged pions (solid lines) that interact further after another interaction length and  $\frac{1}{3}$  neutral pions (dashed lines), which decay immediately into photon pairs (not included). Only a subset of the pions is shown.

### Heitler Model

A simple model for describing the formation of electromagnetic particle showers has been conceptualized by Heitler [36]. According to that model, particles interact after a fixed amount of traversed material, the collision length  $\lambda$ , where  $\gamma$ -rays decay into  $e^\pm$  pairs and  $e^\pm$  leptons emit bremsstrahlung photons. The secondary particles resulting from this interact themselves after traveling for another collision length. This cascading sequence leads to a doubling of the particle multiplicity

$$N(X) = 2^{X/\lambda} = e^{X/X_0} \quad (2.1)$$

with each step, where the radiation length  $X_0 = \frac{\lambda}{\ln(2)}$  is introduced, which is about  $37 \text{ g/cm}^2$  in air. A schematic sketch of the first few steps for an electron-induced shower is depicted in Figure 2.1. One more simplification of the Heitler model is that it assumes an equal allocation of energy for every interaction. Thus, the energy per particle drops as

$$E(X) = \frac{E_0}{N(X)} = E_0 \cdot 2^{-X/\lambda}$$

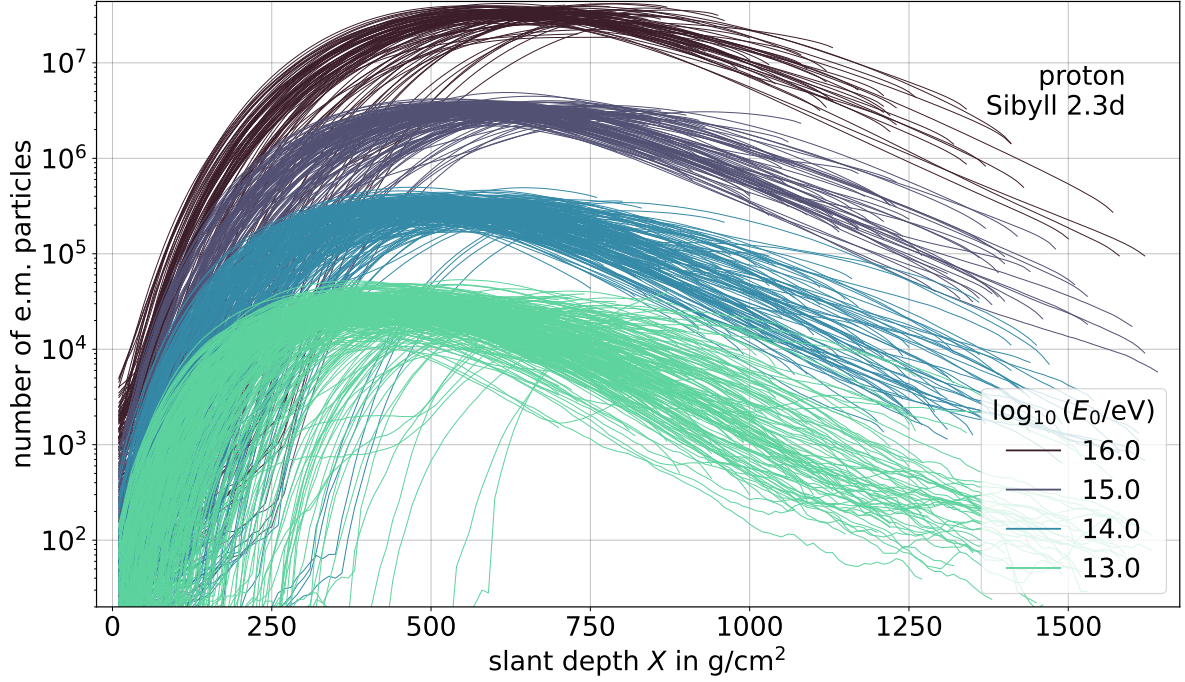
for a shower initiated by a particle with initial energy  $E_0$ . The shower development comes to an end as soon as  $E(X)$  falls below the critical energy  $E_c$ , at which no further splitting occurs as the particles dominantly lose energy via ionization or are absorbed. From this stopping condition, the maximum number of particles

$$N_{\text{max}}(E_0) = \frac{E_0}{E_c} \quad (2.2)$$

can be deduced. When plugging this in Eq. (2.1), it follows that the position of maximal particle multiplicity, the *shower maximum*, lies at

$$X_{\text{max}}(E_0) = X_0 \cdot \ln\left(\frac{E_0}{E_c}\right), \quad (2.3)$$

after which the shower starts to die out. Even though the predictions from the Heitler model do not match the measured quantities in cosmic-ray-induced EASs, the proportionalities  $N_{\text{max}} \propto E_0$  (Eq. (2.2)) and  $X_{\text{max}} \propto \ln(E_0)$  (Eq. (2.3)) have been proven valid through observation. A visualization of these relations with simulated EASs is provided in Figure 2.2.



**Figure 2.2:** Longitudinal shower profiles of CORSIKA simulations ( $\rightarrow$  [section 5.1](#)). The multiplicity of electromagnetic (e.m.) particles in proton-induced air showers of different primary energy is plotted as a function of slant depth. The simplified Heitler model only describes the average behavior, whereas the particle interactions and decays in air showers are of stochastic nature, leading to fluctuations in particle multiplicities.

### Heitler–Matthews Model

In order to describe EASs more realistically, the electromagnetic Heitler model has been extended to hadronic showers by Matthews [37]. Analogous to the Heitler model, it is assumed that the primary hadron undergoes the initial interaction at a specific slant depth, the hadronic interaction length  $\lambda_{\text{had}}$ , and that subsequent hadronic interactions take place at each integer multiple of  $\lambda_{\text{had}}$  of traversed matter. Each interaction produces more hadrons,  $\frac{2}{3}$  being charged mesons, the remaining third comprising neutral mesons. For simplicity, the Heitler–Matthews model deals only with pions as the exclusive type of meson being created. A schematic sketch of the Heitler–Matthews model is given in Figure 2.1. The amount of energy that goes into  $\pi^\pm$  mesons is hence

$$E_{\text{had}}(X) = \left(\frac{2}{3}\right)^{X/\lambda_{\text{had}}} \cdot E_0,$$

whereas the remaining energy that is transferred to  $\pi^0$  is converted into electromagnetic particles as neutral pions immediately decay into pairs of  $\gamma$ -rays, which in turn initiate electromagnetic sub-showers as described by Heitler. For hadronic showers, the depth of shower maximum can be approximated by only considering the first interaction, which yields

$$X_{\text{max}}^{\text{had}}(E_0) \approx \lambda_{\text{had}} + X_0 \cdot \ln\left(\frac{E_0}{3n_{\pi^\pm}E_c}\right),$$

where  $n_{\pi^\pm}$  is the number of charged pions resulting from a single interaction. Similarly to the Heitler model, the production of mesons also ceases, namely when  $E_{\text{had}}(X)$  drops below a

threshold energy  $E_c^\pi$ . At this point, the remaining  $\pi^\pm$ 's decay into muons. For that reason, the predicted muon multiplicity  $N_\mu$  is identical to the eventual number of charged pions, which results in

$$N_\mu(E_0) = N_{\pi^\pm} = \left( \frac{E_0}{E_c^\pi} \right)^\alpha, \quad (2.4)$$

where the exponent  $\alpha = \ln(n_{\pi^\pm}) / \ln\left(\frac{3}{2}n_{\pi^\pm}\right)$  takes values roughly between 0.8 and 0.95.

### From Nucleons to Nuclei

The Heitler–Matthews model does not stop at describing hadronic showers as a single class of events with uniform behavior. When treating composite cosmic-ray nuclei of atomic mass  $A$  and energy  $E_0$  as *superposition* of  $A$  independent nucleons with an energy  $E_0/A$ <sup>1</sup>, the mass-dependent longitudinal profile of hadronic showers can be described qualitatively. The linear relation in Eq. (2.2) results in the number of electromagnetic particles at the shower maximum being basically independent of the primary mass  $A$ :

$$N_{\max}^A(E_0) = A \cdot N_{\max}(E_0/A) = N_{\max}(E_0)$$

However, the depth of the electromagnetic shower maximum becomes shallower because of the lower energy per nucleon:

$$\begin{aligned} X_{\max}^A(E_0) &= X_{\max}(E_0/A) \\ &= X_{\max}^H(E_0) - D_e \cdot \ln(A), \end{aligned}$$

where  $X_{\max}^H$  is the depth of shower maximum for proton primaries and  $D_e \approx 28 \text{ g/cm}^2$  is known as the elongation rate. This means that the shower maximum of iron events is about  $110 \text{ g/cm}^2$  higher than for a proton event with the same  $E_0$  ( $\rightarrow$  Fig. 2.3, left). Lastly, the power-law structure of Eq. (2.4) leads to the muon number increase with  $A$  as

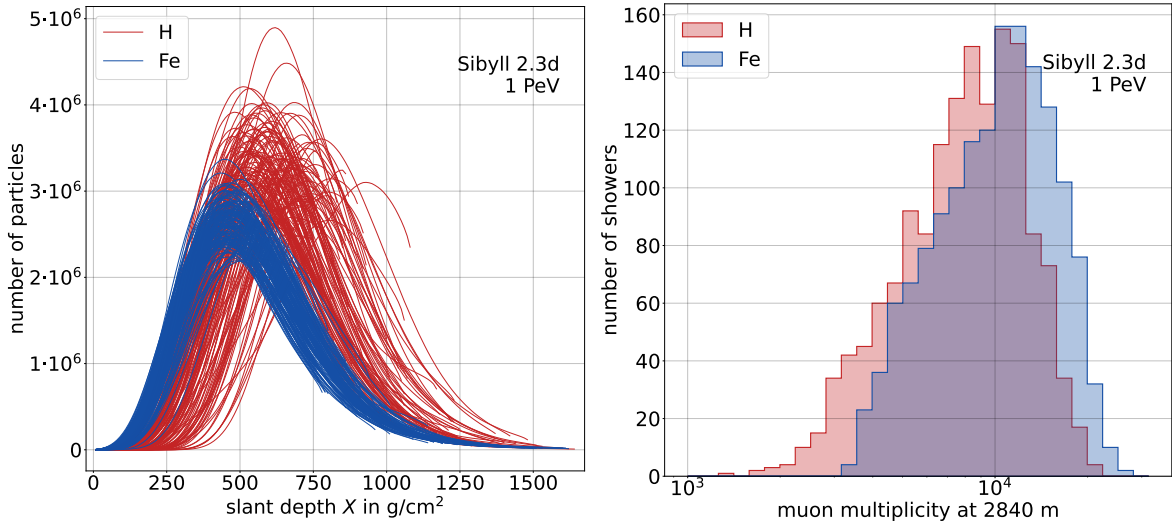
$$N_\mu^A(E_0) = A^{1-\alpha} \cdot \left( \frac{E_0}{E_c^\pi} \right)^\alpha.$$

Consequently, assuming  $\alpha \approx 0.92$ , the expected muon multiplicity in iron-induced showers ( $A = 56$ ) with initial energy  $E_0$  is about 40% higher than  $N_\mu^H(E_0)$  ( $\rightarrow$  Fig. 2.3, right). Having a superposition of many small electromagnetic cascades, the variance in the longitudinal shower profile for heavy primaries is reduced compared to proton events, which are subject to larger fluctuations.

Apart from these effects of the primary mass  $A$  on the shower development, the point of first interaction is also dependent on  $E_0$ . Since the hadronic interaction cross section increases with energy and mass,  $\lambda_{\text{had}}$  becomes smaller, thus high-energy EASs or those initiated by heavy primaries not only develop farther down through the atmosphere but also tend to start higher up than low-energy showers.

---

<sup>1</sup>This is a valid approximation since the binding energy of nucleons is negligible compared to the energy involved in air-shower interactions.



**Figure 2.3:** Difference between air showers induced by protons (red) and iron nuclei (blue). The distributions are from CORSIKA simulations with primary energy set between 1 PeV and 1.26 PeV. **Left:** Longitudinal shower profiles for proton and iron CRs. Note how fluctuations decrease for heavier primaries as predicted by the superposition model. **Right:** Number of muons at the altitude of 2840 m above sea level. Heavier primaries systematically produce more muons of lower energy.

## 2.2 Air-Shower Components

The previous section introduced the most abundant particle types in air showers. Here, an overview of the individual components present in EASs is given.

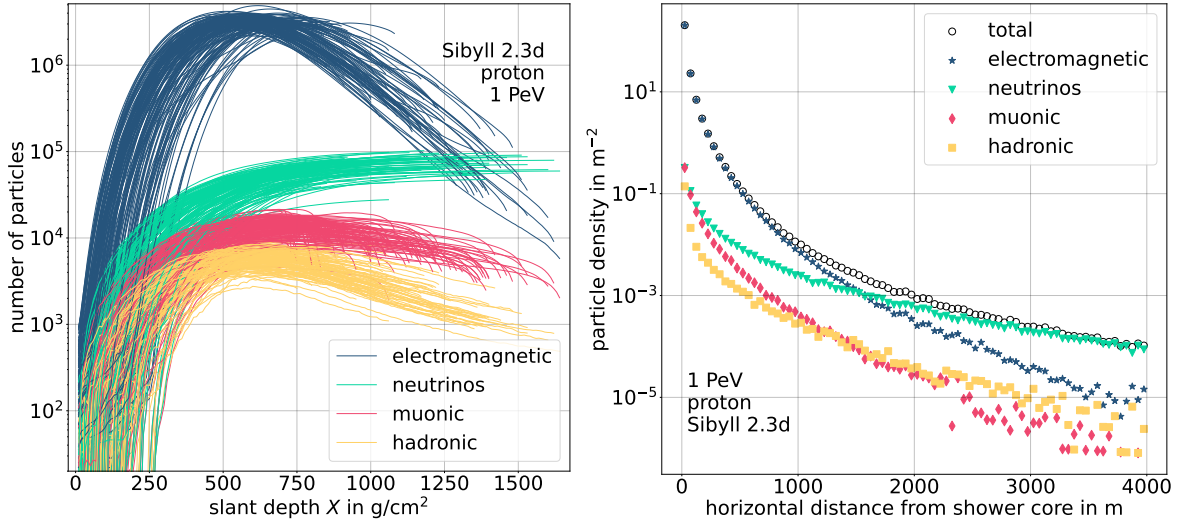
### Hadrons

Hadrons only make up a small fraction of the number of particles in hadron-induced showers. As the mesons that are primarily created in hadronic interactions decay after a few nanoseconds (in the case of charged mesons) or even faster (neutral mesons), their energy eventually feeds into the muonic and mostly the electromagnetic shower component. Despite the Heitler–Matthews model only considering pions, heavier mesons like kaons are also produced in hadron-air interactions. While the decays of  $\pi^\pm$  and  $K^\pm$  produce the *conventional* muon flux, the rare and almost instantaneous decays of charmed hadrons and light vector mesons are responsible for the *prompt* muon component.

Secondary hadrons are produced with a transverse momentum of a few hundred MeV [4]. The angle of low-energy hadrons with the shower axis (the extended track of the incident primary CR) is therefore on the order of several degrees. This results in a flat lateral distribution of low-energy muons and neutrons (the decay products of those laterally separated hadrons) on the ground ( $\rightarrow$  red and yellow markers in the right plot of Figure 2.4). Very-high-energy hadrons, on the other hand, remain close to the shower axis and can be detected on the ground only close to the center of the shower footprint.

### Electromagnetic Particles

The majority of shower particles are of the electromagnetic type, i.e., electrons, positrons and  $\gamma$ -rays. They are produced in sub-showers, which have their origin in the decay of



**Figure 2.4:** *Left:* Longitudinal profiles of the four shower components for simulated proton events with 1.00–1.26 PeV primary energy. *Right:* Lateral distribution of the four components at 2840 m elevation for a single vertical proton shower with  $E_0 = 1$  PeV. Only particles exceeding a species-dependent energy threshold ( $\rightarrow$  subsection 5.1.2) are included.

neutral mesons. In the previous section it has been shown that as the shower development progresses, an increasing fraction of the initial primary energy  $E_0$  is transferred into those e.m. showers. Hence, reconstructions of  $E_0$  rely primarily on the detection of  $e^\pm$  and  $\gamma$ -rays at ground level. Air-shower observatories are located between altitudes of 4410 m above sea level (LHAASO, the *Large High Altitude Air Shower Observatory*, vertical atmospheric depth<sup>2</sup>  $\Xi$  of 600  $\text{g}/\text{cm}^2$ ) and 110 m (KASCADE, the *Karlsruhe Shower Core and Array Detector*,  $\Xi = 1022 \text{ g}/\text{cm}^2$ ). The point of maximum particle multiplicity ( $X_{\text{max}}$ ) of PeV air showers lies in the range of 400–600  $\text{g}/\text{cm}^2$  slant depth.

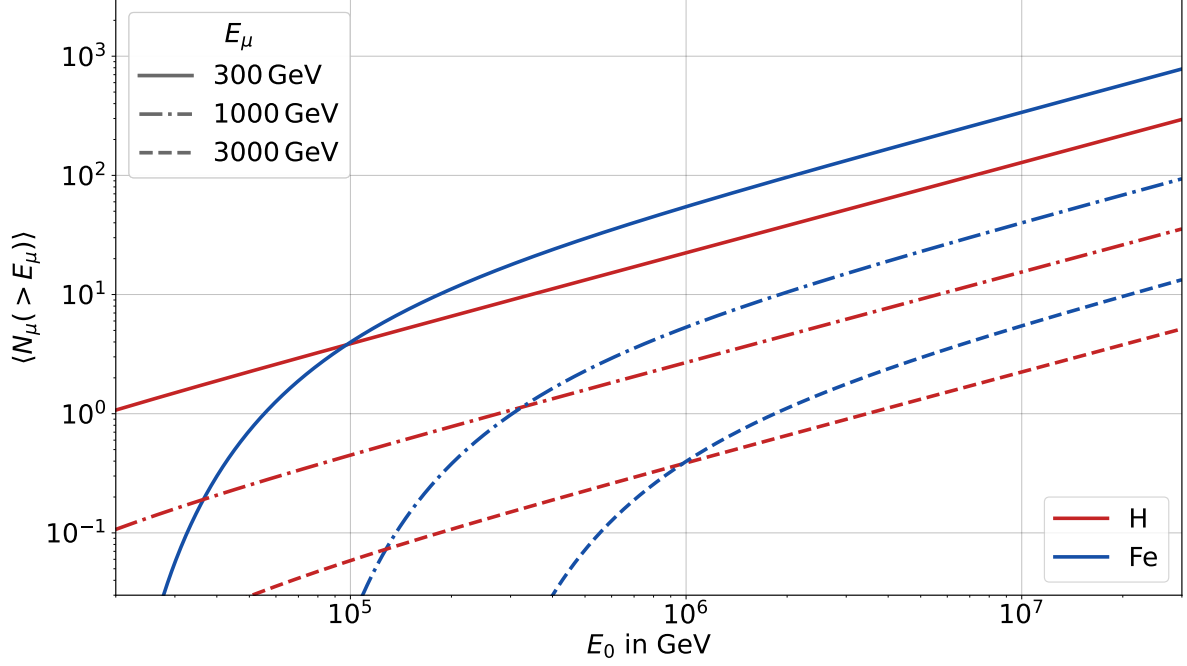
### Muons

Air-shower muons (and anti-muons) originate from the decay of charged mesons. In contrast to the electromagnetic shower component, the maximum muon multiplicity in EASs is reached rather late ( $\rightarrow$  red lines in the left plot of Fig. 2.4) as the muon lifetime is sufficiently long due to Lorentz boosting for many muons to reach the ground before decaying. Apart from that, muons in atmospheric cascades can be considered minimally ionizing, leading to a relatively low energy loss rate compared to  $e^\pm$ . Because of their higher mass, muons from meson decay are created with a larger transversal momentum component compared to electrons and positrons. They are therefore more prominent at large lateral distances on the ground than close to the shower core.

Some experiments are only sensitive to muons above a certain minimum energy  $E_\mu$ . Their average number can be estimated via

$$\langle N_\mu (> E_\mu) \rangle \approx A \frac{\kappa_m}{E_\mu \cos(\theta)} \left( \frac{E_0}{A E_\mu} \right)^{\gamma_\mu - 1} \left( 1 - \frac{A E_\mu}{E_0} \right)^{5.25} \quad (2.5)$$

<sup>2</sup>Vertical atmospheric depth, also known as grammage, is the special case of slant depth  $X$  for  $\theta = 0^\circ$ .



**Figure 2.5:** The number of high-energy muons above  $E_\mu$  in vertical air showers according to Eq. (2.5). Generally, heavier primaries produce more muons than light ones. However, with the ratio  $E_0/E_\mu$  decreasing, muon multiplicity becomes mass-independent and eventually flips.

with  $\kappa_m = 14.5 \text{ GeV}$  and  $\gamma_\mu = 1.757$  [38]. For example, a vertical 1 PeV iron shower is expected to produce 65 muons above a threshold of 273 GeV. In Figure 2.5,  $\langle N_\mu(> E_\mu) \rangle$  is plotted for vertical proton- as well as iron-induced showers and three different muon energy thresholds. In first approximation, iron showers have 2.66 times higher  $N_\mu$  than proton events<sup>3</sup>. However, the last term in Eq. (2.5) leads to a decrease in muon multiplicity, which is stronger for heavier primaries. Since the energy of a cosmic ray is partitioned between all its nucleons, the expected number of muons produced rapidly falls as  $E_0/A$  approaches the muon energy. As a consequence, the discrimination power of the high-energy muon multiplicity for light and heavy cosmic rays ceases at low primary energy.

### Neutrinos

The first three air-shower components discussed so far — electromagnetic, muonic and hadronic — can be detected at the surface with dedicated devices (more on methods of detecting secondary shower particles in the following section). Some of the energy, however, goes into the (almost) invisible neutrino component. Both the decay of charged mesons and the decay of muons are weak interactions that are always accompanied by the production of neutrinos:

$$\begin{aligned}\pi^- &\rightarrow \mu^- + \bar{\nu}_\mu \\ \mu^- &\rightarrow e^- + \nu_\mu + \bar{\nu}_e\end{aligned}$$

<sup>3</sup>This is in contrast to the ~40% excess predicted by the Heitler–Matthews model, illustrating the large variations between models of air-shower development.

Since neutrinos can only interact weakly themselves, their interaction cross section with nuclei is extremely small. Hence, the neutrino component of EASs does not die out but rather grows steadily as hadrons and muons decay ( $\rightarrow$  green lines in the left plot of Fig. 2.4). They hardly even get absorbed by Earth, which is why these *atmospheric neutrinos* are present everywhere on the globe coming from all directions. This makes atmospheric neutrinos the main background for the identification of *astrophysical neutrinos*.

### 2.3 Detection Methods

There are various ways to detect cosmic-ray air showers. Depending on the particle type of interest, the targeted energy range of both the primary and the secondary particles, the observatory location as well as the planned analyses, one may choose a certain kind of detector or build a hybrid array using a combination thereof. What follows is an overview of commonly used detection techniques for EASs.

#### Scintillators

Scintillation detectors are mostly sensitive to the  $e^\pm$  particles in EASs but can also detect single muons penetrating the device. Scintillators are often operated as pure muon detectors, which requires adding a shielding layer on top. The covering material absorbs electromagnetic particles, and with the choice of material type and thickness, the energy threshold for the muon component can be adjusted.

#### Underground Detectors

Taking the idea of shielding scintillators against electromagnetic particles even further, detectors can be deployed deep below the surface. With several meters of soil or water between atmosphere and detector, it can be guaranteed that only highly energetic muons (or the products of neutrino interactions) will be seen.

#### Water- / Ice-Cherenkov Tanks

When relativistic charged particles pass through a dielectric medium with refractive index  $n$ , they emit Cherenkov radiation as long as their velocity exceeds  $c/n$ . Water, both in its liquid and solid state, has a refractive index of about 1.35 [39], making it a suitable yet inexpensive detection medium with a Cherenkov angle of  $42^\circ$ . Light sensors like photomultiplier tubes ( $\rightarrow$  section 4.1) can be installed inside tanks filled with water or ice to record the flashes of Cherenkov light from shower particles passing through. Besides the accumulated charge within the readout time of a Cherenkov tank, the fine temporal structure, that is, the pulse shape, holds useful information about the particle type. Unless the tank has a spherical shape, the path the particle took through the detector volume has an influence on the recorded signal. For example, in a cylindrical Cherenkov tank, a vertically through-passing muon has a shorter track length than a diagonally traversing one, but not as short as an inclined muon, which only penetrates the top peripheral area of the tank close to the edge, also known as *corner clippers*.

Particle detectors such as Cherenkov tanks and scintillators are usually arranged in regular grids on the ground. From timing information of particles hitting individual detectors and the lateral distribution of deposited charge, both the arrival direction and the energy of the primary cosmic ray can be reconstructed. The distribution of charge deposited per particle differs for e.m. particles (with a broad continuous spectrum) and muons, which can be considered minimum ionizing particles (MIPs) and hence stand out by a characteristic charge deposit value.

### Radio Antennas

The electrons and positrons in air showers emit electromagnetic radiation in the radio spectrum. The reasons for this are the deflection of  $e^\pm$  in Earth's magnetic field (geomagnetic effect) as well as the charge excess that builds up along the shower axis caused by ionization of the air [40]. Antennas can be installed to measure those radio pulses in the MHz-to-GHz range. Since the detection of radio signals from air showers is sensitive to the weather, the measurement uptime with this method is about 90% or higher, depending on the site.

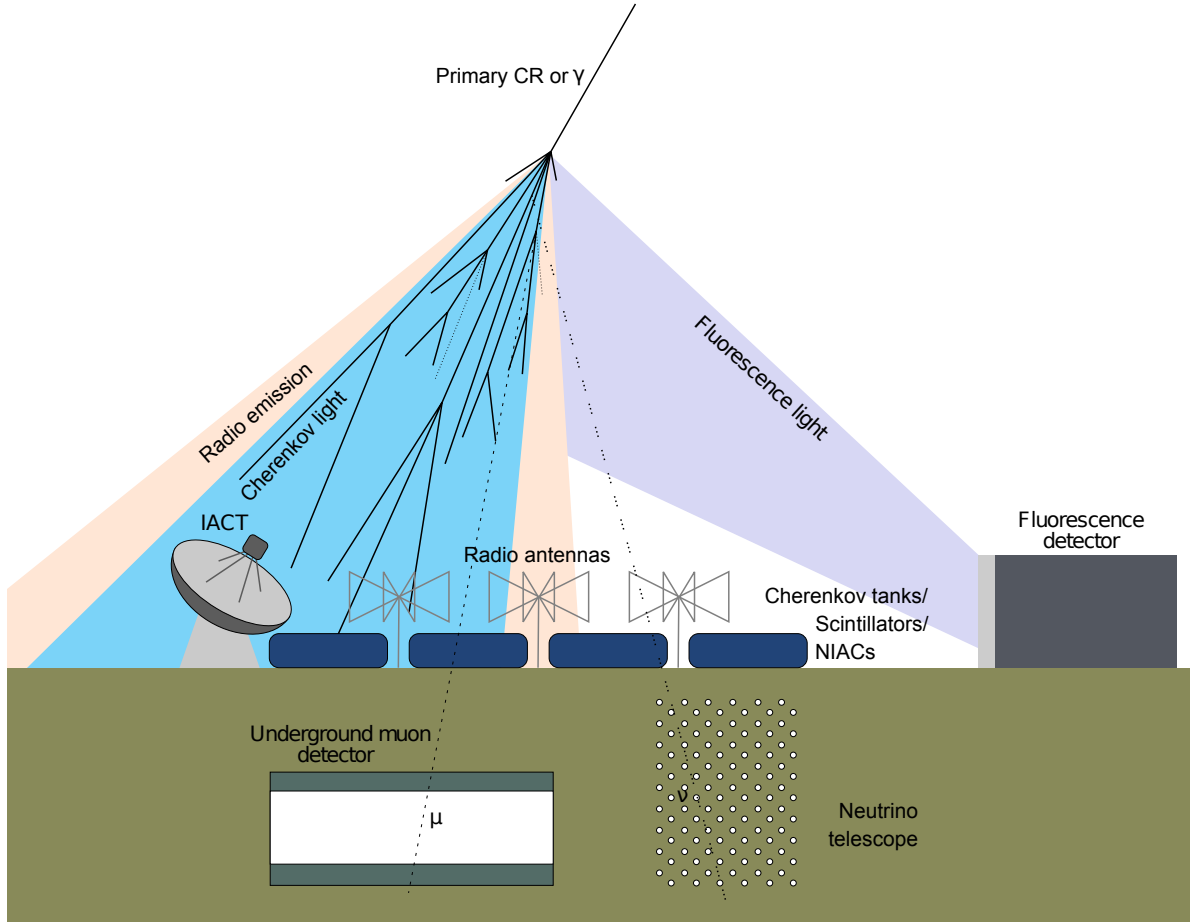
### Fluorescence Telescopes

As charged air-shower particles traverse the atmosphere, they excite nitrogen molecules in the air. Their immediate de-excitation is accompanied by the emission of fluorescence light in the ultraviolet (UV) part of the electromagnetic spectrum [41]. Fluorescence telescopes can image those traces of UV light in the sky from outside the shower footprint, enabling an accurate estimation of  $X_{\max}$  from the longitudinal shower profile.

### Air-Cherenkov Telescopes

Alternatively to the fluorescence emission, the Cherenkov light produced by the shower particles themselves in air can be observed directly as well inside the shower footprint. The refractive index of air changes with altitude as a consequence of the atmospheric density profile and is barely greater than unity. The resulting Cherenkov angle lies between  $1.36^\circ$  at sea level and only  $0.07^\circ$  above 40 km [4]. Non-imaging air-Cherenkov telescopes (NIACs) are arranged in regular grids similar to scintillator and Cherenkov tank arrays and can be used for calorimetric measurements based on the amount of collected light. Imaging air-Cherenkov telescopes (IACTs) on the other side are able to see the shower development in the atmosphere using high-resolution cameras. Since these cameras can only point in one direction at a time, the field of view of IACTs is rather limited. Similar to fluorescence telescopes, IACTs and NIACs can operate only during clear moonless nights, drastically reducing their uptime to about 10–15%. In comparison, Cherenkov tanks and scintillators have a duty cycle of basically 100%. Moreover, fluorescence telescopes and IACTs have in common that by deploying a stereoscopic imaging setup with many telescopes it is possible to reconstruct the incident direction of the primary cosmic ray, improve the energy reconstruction and distinguish between gamma-ray-induced and hadronic air showers.

Figure 2.6 illustrates how the detector types described above probe different parts and properties of air showers.



**Figure 2.6:** A schematic overview of the detector types for EAS particles and their emitted radiation as discussed in the main text. Deep muon detectors and neutrino telescopes can be deployed both under soil and rock as well as inside large volumes of water or ice. Adopted from Ref. [28] (Fig. 55) and modified.

## 2.4 Air-Shower Experiments

Two observatories dedicated to the detection of extensive air showers have already been mentioned: LHAASO in Sichuan, China, is a hybrid experiment at a vertical atmospheric depth of  $600 \text{ g/cm}^2$ , comprising telescopes for air-Cherenkov and fluorescence light, water-Cherenkov detectors as well as a  $1 \text{ km}^2$  array of shielded scintillators and underground muon detectors [42]. The former KASCADE experiment used an array of various types of scintillators with and without shielding and a hadron calorimeter to achieve sensitivity to the distinct shower components [43]. The “Grande” extension added more unshielded and farther separated scintillators, increasing the area from  $0.04 \text{ km}^2$  to almost  $0.5 \text{ km}^2$  [44]. KASCADE(-Grande) operated from 1996 to 2013 on the site of Forschungszentrum Karlsruhe (today: Campus North of KIT, the Karlsruhe Institute of Technology,  $\Xi = 1022 \text{ g/cm}^2$ ) in Germany.

The *High Altitude Water Cherenkov Experiment* (HAWC) in Puebla, Mexico, uses an array of water-Cherenkov tanks 4100 m above sea level ( $\Xi = 637 \text{ g/cm}^2$ ) for the detection of secondary air-shower particles [45]. Due to their dense instrumentation at high altitude, observatories like HAWC, LHAASO and the *Astrophysical Radiation with Ground-based Observatory* in Yang-

bajing, Tibet, (ARGO-YBJ,  $\Xi = 606 \text{ g/cm}^2$ ) [46] are sensitive at primary energies as low as 100 TeV. UHECR experiments such as the *Pierre Auger Observatory* in Mendoza, Argentina [47], and the *Telescope Array* (TA) in Utah, USA [48], on the other hand, are located at much lower altitudes of 1400 m above sea level ( $\sim 875 \text{ g/cm}^2$ ) and use an arrangement with 1.2–1.5 km spacing between surface detector stations together with fluorescence telescopes and radio antennas. This allows those largest observatories to detect air showers with primary energy beyond several EeV.

IceCube combines a deep underground muon detector with a square-kilometer ice-Cherenkov surface array ( $\rightarrow$  section 4.1) at 2830 m, which corresponds to about  $695 \text{ g/cm}^2$ . This puts the southernmost EAS detector in the central energy range around the spectral knee, making it a pivotal experiment to bridge the gap between low-energy and high-energy observations.

### Key Messages

*As their primary energy exceeds about 1 PeV, cosmic rays can only be detected indirectly on the ground through the extensive air showers they produce in the atmosphere. The various species of particles that are created in those cascades require specialized detector types. In this analysis, ice-Cherenkov tanks are used to detect low-energy muons and the electromagnetic shower component. High-energy muons are observed by an underground array of light sensors.*



# CHAPTER 3

## MACHINE LEARNING

*“The first rule of communication is to shut up and listen [...]”*

— David Spiegelhalter, *The Art of Statistics: Learning from Data*

For the analysis presented, a machine learning technique is used to identify air showers induced by protons or helium nuclei and to estimate their primary energy. After a general introduction into the realm of machine learning in [section 3.1](#), the focus is put on neural networks ([section 3.2](#)), including their architecture and the training process ([section 3.3](#)).

### 3.1 Introduction to Machine Learning

The notion of machine learning (ML) refers to a branch of computer science that provides powerful tools for statistical data analysis. In a general scope, the purpose of ML is to find (that is, to *learn*) the function  $f$  that maps predictors  $x_i \in X$  (also called inputs or features) to their respective labels  $y_i \in Y$  (also referred to as targets). In this sense, machine learning can be considered the extension of simple linear fitting but with more complex and diverse domains ( $X$ ), codomains ( $Y$ ) as well as relations between them. In order to achieve this, ML encompasses a wide variety of methods, which can be classified, for example, by the technique used, their field of application or the data format needed. Among the most popular tools are tree-based models (random forests and boosted decision trees), support vector machines and the “k nearest neighbors” algorithm [49]. The functionality of neural networks, which are the model of choice for this analysis, will be described in more detail in the following.

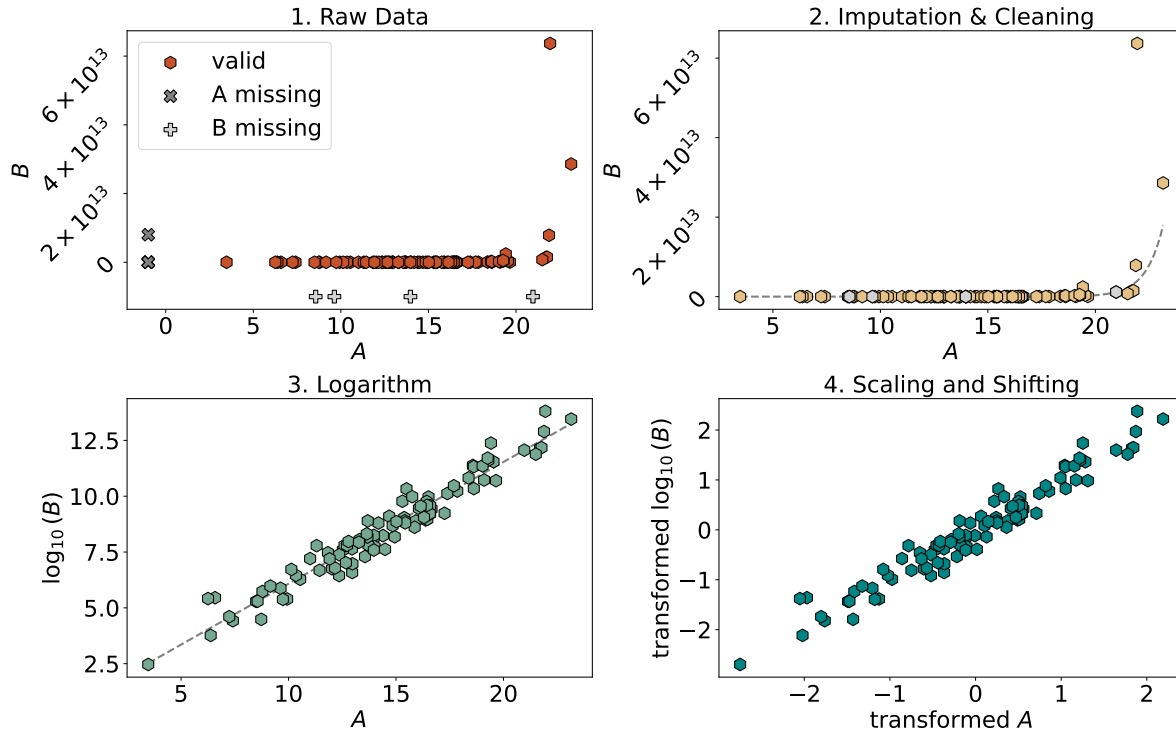
#### Training Supervision

The process of adapting any ML model to its desired task is called *training* and requires a set of exemplary data to learn from. Typically, this training data contains both the features and the corresponding target values for each sample in the dataset. However, one discipline of machine learning aims to learn patterns from unlabeled data, so-called unsupervised learning, which includes clustering and anomaly detection techniques. The presented cosmic-ray analysis, however, uses a *supervised* approach with Monte Carlo simulations (→ [chapter 5](#)) as labeled training data.

#### Regression vs. Classification

Supervised machine learning tasks usually fall into one of two categories: *Regression* denotes the prediction of continuous values such as physical quantities like temperature or energy. Assigning one of  $n \geq 2$  categories to an input sample on the other side is called *classification*. One common approach to turn the generally continuous output value of a ML model into a classification statement is to feed the raw value into a logistic sigmoid function

$$\sigma_{\log}(x) = \frac{1}{1 + e^{-x}}, \quad (3.1)$$



**Figure 3.1:** Visualization of a possible preprocessing procedure applied to generated fake data (1) described by two features, A and B. The raw data has samples with missing information, which is plotted on a default value of A and B, respectively. Elements with a missing A value are omitted, whereas cases of unknown B obtain replacement values estimated via a fit (2). As B spans many orders of magnitude, its logarithm is calculated (3). In the final step (4), mean is set to zero and standard deviation is scaled to one via linear transformations.

which returns an output  $0 < \sigma_{\log}(x) < 1$ . Thresholds on this modified value can then be defined to differentiate between classes. This is particularly useful for binary classification problems that only require a single threshold. In the case of  $n \geq 3$  classes, it is rather advisable to let the model return  $n$  output values and select the class with the highest score as the prediction.

### Preprocessing

Depending on the format of data and the type of ML model, a *preprocessing* step of the inputs and outputs is generally necessary. This includes storing the data — both features and targets — in the shape required by the model architecture (e.g. vectors of strings, tensors filled with numbers or single integers representing discrete categories). In the process of data taking it can happen that certain bits of information are missing or corrupt for some samples. It is crucial to handle those feature values with the necessary care. Depending on the circumstances, faulty values can be imputed using similar events with a complete set of features or replaced by dummy values that are not expected to appear in correct data.

Numerical data can range over many orders of magnitude following all kinds of distributions. Even though machine learning models are generally quite flexible in terms of the magnitude of the numbers they are fed and supposed to return, their performance and the time it takes

to train them often benefit from a rescaling. In particular, it is favorable to apply a sequence of (usually linear) transformations to the values that roughly yields a Gaussian-like distribution centered at zero. An example of a sequence of preprocessing steps on generated data points is shown in Figure 3.1.

### Training, Testing and Validating

Similarly to basic fitting of algebraic functions, the training of ML models is susceptible to overfitting, or *overtraining*. When the model adapts to the training samples too much (for example, over an excessive number of training iterations), it gives overly accurate predictions on the samples it was trained on but performs poorly on data of the same type that has not been used for training. It has thus become standard practice to permanently split the available data into two or even three datasets.

After the model has been optimized using the *training* set, the accuracy<sup>1</sup> can be analyzed by applying the model to the separate *test* dataset. If the model is performing significantly worse on the test set than on the training set, this is indicative of over- or undertraining. A *validation* set can be utilized with some ML techniques to monitor the model performance in real time during the training process. While a perfectly learning model would yield a similarly high accuracy for both training and validation data<sup>2</sup>, an overtraining model continues improving the training accuracy, whereas the validation performance suffers.

### Hyperparameters

All machine learning models are specified with two fundamentally different sets of parameters. The *learnable parameters* (such as slopes, thresholds or weights) are gradually adjusted during the training process, whereas *hyperparameters* (such as number of decision trees, size of convolution kernels or number of samples processed at once) are specified by the user in order to configure the model's architecture and control the training procedure. A well-considered choice of hyperparameters can often be used to counteract overtraining.

## 3.2 Neural Networks

Artificial *neural networks* (NNs) are the type of machine learning model that has received by far the most attention in the past decades. Their broad field of applications made NNs a popular tool for all sorts of data analysis applications and also an essential component of services in every-day life that use *artificial intelligence*.

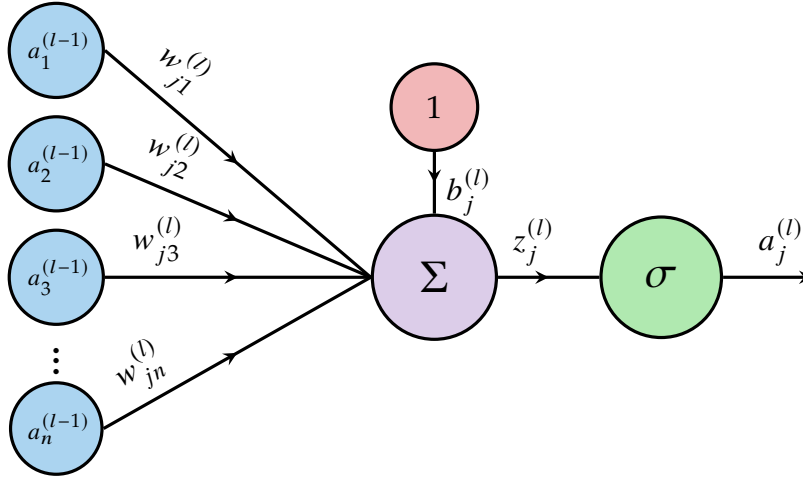
### The Perceptron

The concept of neural networks is based on the idea of imitating the structure of biological neural systems, such as those found in the brain. The smallest unit of a NN is the neuron, in the context of machine learning also called *perceptron* or simply node. A neuron takes several inputs  $a_k$  — coming from  $n$  other neurons — and returns a single output  $z$ , which is

---

<sup>1</sup>Even though the term “accuracy” has a very specific meaning in machine learning, in this context it is used to refer to any kind of performance metric.

<sup>2</sup>Ideally, the training accuracy should be slightly higher than the validation accuracy.



**Figure 3.2:** Schematic sketch of a single perceptron with  $n$  inputs  $a_k^{(l-1)}$ , the weights  $w_{jk}^{(l)}$  as well as bias  $b_j^{(l)}$  which are combined in a weighted sum. The result  $z_j^{(l)}$  is eventually fed into an activation function  $\sigma$  which returns the activation value of the perceptron  $a_j^{(l)}$ . The meaning behind the subscript  $j$  and the superscripts are covered in the discussion of backpropagation in the main text.

a weighted sum of its inputs

$$z = b + \sum_{k=1}^n w_k a_k, \quad (3.2)$$

where the  $w_k$  are the weight factors of the perceptron. The bias value  $b$  is added to provide direct control of the neuron's *activation*, i.e., its effect on other neurons that are connected to it. In this context, the symbol  $\vartheta$  is used<sup>3</sup> to refer to all weights and biases as the free parameters of a network collectively.

### Activation

In analogy to biological neurons which serve as a model, artificial perceptrons only send out a relevant signal if they receive sufficiently strong inputs. This non-linear activation can be simulated by applying a function  $\sigma$  to the raw perceptron output  $z$ . Typical choices for such an *activation function* are sigmoid<sup>4</sup>, hyperbolic tangent ( $\tanh$ ) and the so-called *rectified linear unit*

$$\text{ReLU}(z) = \max(0, z) = \begin{cases} 0 & , z < 0 \\ x & , z \geq 0 \end{cases}$$

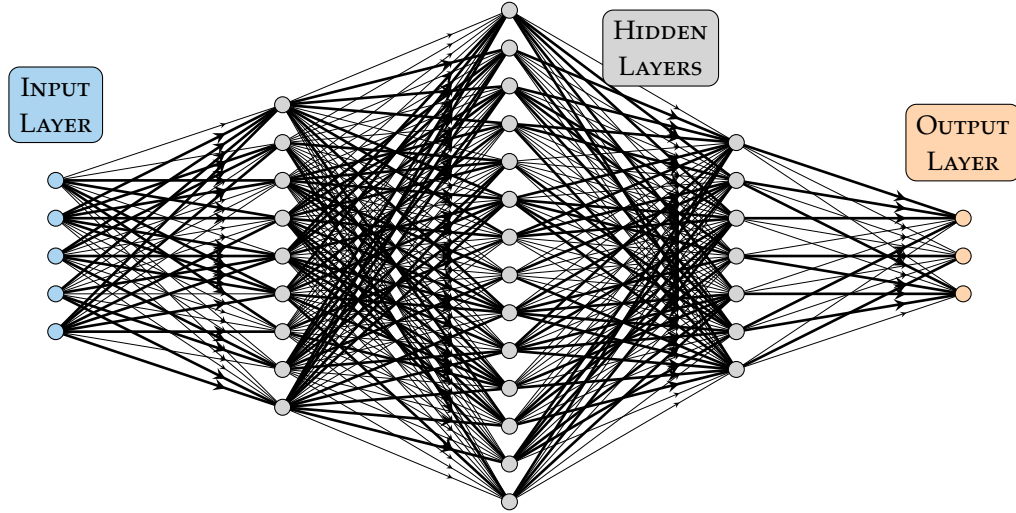
as well as various derivatives thereof. Figure 3.2 shows the layout of a single perceptron with the activation function attached to it.

### The Multilayer Perceptron

Since a single neuron is rather limited in its capabilities, many nodes — each with an individual set of weights and bias values — are arranged in what is called a *layer*. Despite

<sup>3</sup>Not to be confused with the symbol for the zenith angle  $\theta$ .

<sup>4</sup>In the early years of neural networks, logistic sigmoid (Eq. (3.1)) has been the standard activation function, which is why the letter  $\sigma$  nowadays represents activation functions in general. In order to avoid confusion, the subscript “log” is added when the sigmoid function is referred to.



**Figure 3.3:** Architecture of a fully-connected multilayer perceptron with five input features (blue), three hidden layers (gray) and three output targets (orange). The gray and orange nodes represent individual neurons without explicitly showing their biases and the activation. The varying thickness of connecting arrows represents the weight value the neurons have for each of their inputs. This specific network configuration would be referred to as a 5-9-14-7-3 network.

not being connected to each other, the nodes of a specific layer share their inputs. Thus, a network layer comprising  $k$  neurons can be thought of as transforming the  $n$  inputs into  $k$  output values.

In order to achieve a higher model complexity for more sophisticated tasks, many layers are combined into a large network of neurons, a *multilayer perceptron* (MLP)<sup>5</sup>. The idea behind this concept is that layers hold higher-level information on their inputs in an abstract representation. Stacking several layers is thought to enable reach an arbitrary level of complexity in order to solve basically any data analysis problem. In the standard configuration of a *fully-connected* MLP, the outputs of one layer are used as the inputs of the subsequent layer. While the first layer takes the preprocessed features from the data as input, the output of the last layer is interpreted as the network's prediction. The output layer can have multiple nodes in the case of multi-class classification or when several quantities are predicted at once. Layers in between input and output are dubbed *hidden layers* as there is no interaction of the user with this part of the network. Figure 3.3 gives an example for a rather small fully-connected MLP architecture. The amount of free parameters of such a simple architecture grows rapidly with the size of the network. For a network size of  $L$  layers and  $N_l$  nodes in the  $l^{\text{th}}$  layer, the number of weight and bias values turns out to be

$$N_{\theta} = \sum_{l=1}^{L-1} (N_l + 1) N_{l+1}.$$

The 5-9-14-7-3 network ( $L = 5$ ) in Figure 3.3 thus contains 323 free parameters.

<sup>5</sup>In this work, the terms (neural) network, multilayer perceptron and model are used interchangeably.

Besides the MLP, the concept of neural networks has been advanced to create more sophisticated and specialized model architectures. One example for this are *convolutional neural networks* (CNNs), which are widely used in the domain of image recognition.

### 3.3 Model Training

When a neural network architecture has been specified in code and a model instance has been created, all weights and biases are initialized as random values<sup>6</sup>. As a consequence, the network's response to input samples is not correlated with the true target values.

#### Metrics and Loss Functions

Several types of metrics can be used to quantify the performance of a network. Popular choices for regression problems are the *mean squared error* (MSE)

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (f_{\theta}(x_i) - y_i)^2$$

and the *mean absolute error* (MAE)

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |f_{\theta}(x_i) - y_i|,$$

where the mean is taken over  $n$  samples and  $f_{\theta}$  describes the output of a model with free parameters  $\theta$ . In order to abbreviate the notation, the network response to the input  $x_i$  can also be denoted  $\hat{y}_i = f_{\theta}(x_i)$ . A less common metric in ML is the *coefficient of determination*

$$R^2 = \frac{\sum_{i=1}^n (\hat{y}_i - \bar{y})^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

with  $\bar{y} = \sum_{i=1}^n y_i / n$  the target average. Maximizing  $R^2$  is equivalent to minimizing the *fraction of variance unexplained* (FVU) [50]

$$\text{FVU} = 1 - R^2 = \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}, \quad (3.3)$$

which can also be expressed as  $\text{MSE}(f_{\theta}) / \text{Var}(y)$ . Hence, the FVU loss considers the regression error in relation to the target variance.

For classification tasks, *cross entropy* is often used, which, for binary classification ( $y_i \in [0, 1]$ ) in particular, takes the form

$$\text{BCE} = \frac{-1}{n} \sum_{i=1}^n \left( y_i \cdot \log(\sigma_{\log}(\hat{y}_i)) + (1 - y_i) \cdot \log(1 - \sigma_{\log}(\hat{y}_i)) \right). \quad (3.4)$$

When there are two classes (positive and negative), one can also consult

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

---

<sup>6</sup>By default, initial values usually are drawn from a normal distribution with vanishing mean and unit variance.

as well as

$$\text{recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

where **TP** refers to correctly classified true positives, **FP** labels false positives, that is, negative samples that are predicted as positive, and vice versa **FN**, which are mistakenly classified as members of the negative group.

One such metric needs to be selected for each label and if necessary modified in such a way that good agreement of the model output  $f_{\theta}(x_i)$  and the target value  $y_i$  results in a small metric quantity. Bad predictions on the other hand have to be penalized with large values. The chosen metric (or combination of metrics) is called the *loss* or *cost function* of the network. Training a NN means finding the set of parameters that minimizes the cost:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} C(f_{\theta}(x), y)$$

Due to the extremely high dimension of the parameter space, this is not feasible analytically. Instead, an algorithm called *stochastic gradient descent* (**SGD**) is applied.

### Stochastic Gradient Descent

Even though it is not possible to determine the global minimum of  $C(f_{\theta}(x), y)$  directly, a local minimum can be found by iteratively updating the vector of randomly initialized parameters  $\theta$  in small steps that lead to a reduction of the loss. With the partial derivatives with respect to all weights and biases  $\nabla_{\theta} C = \left( \frac{\partial C}{\partial w_{11}^{(2)}}, \dots, \frac{\partial C}{\partial b_{N_L}^{(L)}} \right)$  and a small hyperparameter  $\eta$ , the parameter update turns out to be

$$\theta \rightarrow \theta - \Delta\theta = \theta - \eta \nabla_{\theta} C. \quad (3.5)$$

The factor  $\eta$  is called the *learning rate* and needs to be chosen carefully. A too large value can lead to an unstable training with the risk of overshooting local minima, whereas an overly small learning rate causes the training to slow down and potentially even cease before a good minimum is reached. Eq. (3.5) demonstrates that the step size  $\Delta\theta$  also scales with the loss gradient itself. This means that when the training begins and  $C$  is far from its minimum, larger steps are made; that is, the optimization generally progresses quickly. As the local minimum is approached, the gradient  $\nabla_{\theta} C$  flattens, so only minor adjustments are made to the parameters in the final stage of training.

In practice, the training data is partitioned randomly into several *batches*, which are processed after each other. This random division into smaller sets accelerates the training and makes the gradient descent algorithm *stochastic*. The  $N_{\theta}$ -dimensional gradients for all samples in the batch are averaged and the minimization step (3.5) is performed before the next batch is loaded into memory. Following this procedure, all training samples are seen once by the SGD algorithm and help making a few small steps towards the local minimum of the loss function.

As a rule, a single run of the described steps is not sufficient to reach the minimum of  $C$  but rather tens or hundreds of iterations are necessary to achieve a satisfactory model. One

repetition is called a training *epoch*. In practice, the gradient descent is overseen by a so-called optimizer such as the *adaptive moment estimation* algorithm “Adam” [51]. By keeping track of the direction of  $\nabla_{\theta} C$  in previous epochs and adding a corresponding momentum term to Eq. (3.5), Adam emulates the downward motion of a ball on an undulating surface due to gravity. Similarly to the ball surging past bumps in the 2-dimensional ground, the loss  $C$  is supposed to decrease towards the local minimum, unperturbed by fluctuations in the  $N_{\theta}$ -dimensional loss surface.

### Backpropagation

It may first seem impractical to determine the gradient  $\nabla_{\theta} C$  of the cost function with respect to all the network parameters, especially for deep neural networks with a large number of layers  $L$ . During a *forward pass* the initial activation provided by the input features  $x_i$  propagates through the network, culminating in the eventual activation of the output nodes  $\hat{y}_i$ , which serve as model predictions. The *backpropagation* technique regards this forward-directed information flow in reverse — from the output back to the first hidden layer [52].

For networks with more than one loss component, the total loss

$$C = \sum_{j=1}^{N_C} C_j(\hat{y}, y) \quad (3.6)$$

is a sum over all defined cost functions  $C_j$ <sup>7</sup>. The partial derivatives of  $C$  with respect to all network parameters  $\theta$  can be obtained via repeated application of the chain rule. Starting with the output layer  $L$ , the influence of the weight  $w_{jk}^{(L)}$  (which connects the  $j^{\text{th}}$  neuron of layer  $L$  with the  $k^{\text{th}}$  node in the preceding layer) on the loss  $C$  can be identified to be

$$\begin{aligned} \frac{\partial C}{\partial w_{jk}^{(L)}} &= \frac{\partial C}{\partial a_j^{(L)}} \frac{\partial a_j^{(L)}}{\partial z_j^{(L)}} \frac{\partial z_j^{(L)}}{\partial w_{jk}^{(L)}} \\ &= \frac{\partial C}{\partial a_j^{(L)}} \sigma' \left( z_j^{(L)} \right) a_k^{(L-1)}, \end{aligned} \quad (3.7)$$

where in the special case of the last layer, the activation of the output node  $a_j^{(L)}$  corresponds to the prediction value  $\hat{y}_j$ . Moreover, for regression tasks, the activation function in the final layer commonly equals the identity, i.e.,  $\sigma(z) = z$ , simplifying the second derivative to  $\sigma' \left( z_j^{(L)} \right) = 1$ . The first two terms in Eq. (3.7) demonstrate that both the cost function  $C$  and the activation function  $\sigma$  need to be differentiable. The third factor  $a_k^{(L-1)}$  follows directly from Eq. (3.2). Similarly, the partial derivatives of  $C$  with respect to the  $j^{\text{th}}$  bias value in the

---

<sup>7</sup>The number of loss components  $N_C$  typically coincides with the number of output nodes  $N_L$ . It is, however, possible to construct multiple losses for some targets or to have several output nodes covered by a single loss. Furthermore, it is possible to introduce summation weights into Eq. (3.6), which has been omitted in this derivation for the sake of simplicity.

last layer emerge as

$$\begin{aligned}\frac{\partial C}{\partial b_j^{(L)}} &= \frac{\partial C}{\partial a_j^{(L)}} \frac{\partial a_j^{(L)}}{\partial z_j^{(L)}} \overbrace{\frac{\partial z_j^{(L)}}{\partial b_j^{(L)}}}^{=1} \\ &= \frac{\partial C}{\partial a_j^{(L)}} \sigma' \left( z_j^{(L)} \right).\end{aligned}$$

Finally, in order to propagate the loss derivatives backwards through the network, it is crucial to also determine the contribution of the activations in the second-to-last layer to the cost gradient, yielding

$$\begin{aligned}\frac{\partial C}{\partial a_k^{(L-1)}} &= \sum_{j=1}^{N_L} \frac{\partial C}{\partial a_j^{(L)}} \frac{\partial a_j^{(L)}}{\partial z_j^{(L)}} \frac{\partial z_j^{(L)}}{\partial a_k^{(L-1)}} \\ &= \sum_{j=1}^{N_L} \frac{\partial C}{\partial a_j^{(L)}} \sigma' \left( z_j^{(L)} \right) w_{jk}^{(L)},\end{aligned}$$

where again the third term  $w_{jk}^{(L)}$  arises from Eq. (3.2). The summation over all nodes  $j$  in the  $L^{\text{th}}$  layer is a consequence of the total loss  $C$  having contributions from all output nodes [53]. The backpropagation can be continued further towards the first hidden layer, adding a factor  $\sigma' \left( z^{(l)} \right) \cdot w^{(l)}$  and another summation for each layer<sup>8</sup>:

$$\frac{\partial C}{\partial a_m^{(L-q)}} = \sum_{j_{L-q+1}=1}^{N_{L-q+1}} \cdots \sum_{j_L=1}^{N_L} \frac{\partial C}{\partial a_{j_L}^{(L)}} \sigma' \left( z_{j_L}^{(L)} \right) w_{j_L j_{L-1}}^{(L)} \cdots \sigma' \left( z_{j_{L-q+1}}^{(L-q+1)} \right) w_{j_{L-q+1} m}^{(L-q+1)}$$

In practice, all these calculations that constitute both the forward pass and the backpropagation are performed efficiently using tensor objects. Given a batch size of  $n_{\text{sample}}$ , the weights are stored in matrices of dimension  $N_l \times N_{l-1}$ , biases are kept in vectors of length  $N_l$  and the resulting activation takes the form of  $n_{\text{sample}} \times N_l$  matrices. The forward pass  $f_{\theta}$  can thus be expressed as a sequence of alternating matrix multiplications and element-wise vector additions.

### Batch Normalization

The same positive effect of normalized and centered input distributions that has been mentioned in the context of preprocessing also applies for the inputs of each layer inside a neural network. Faster training progress can be achieved by preventing the occurrence of vanishing gradients when the distribution of node activation values is normalized for each batch of training data. This can be achieved by introducing dedicated *batch normalization* layers between the linear layers of the network [54].

<sup>8</sup>For example:  $\frac{\partial C}{\partial a_m^{(L-3)}} = \sum_{l=1}^{N_{L-2}} \sum_{k=1}^{N_{L-1}} \sum_{j=1}^{N_L} \frac{\partial C}{\partial a_j^{(L)}} \sigma' \left( z_j^{(L)} \right) w_{jk}^{(L)} \sigma' \left( z_k^{(L-1)} \right) w_{kl}^{(L-1)} \sigma' \left( z_l^{(L-2)} \right) w_{lm}^{(L-2)}$

## CHAPTER 3. MACHINE LEARNING

### **Key Messages**

*Neural networks are powerful and versatile machine learning tools. However, great care is required when configuring a model, including its architecture, the selection of loss functions and the assembly of preprocessed input features. This analysis utilizes a multilayer perceptron for the reconstruction of the primary energy and mass of cosmic rays, involving the compilation of regression and classification losses.*

# CHAPTER 4

## THE ICECUBE NEUTRINO OBSERVATORY

*“There’s a South Pole,” said Christopher Robin, “and I expect there’s an East Pole and a West Pole, though people don’t like talking about them.”*

— Alan Alexander Milne, *Winnie-the-Pooh*

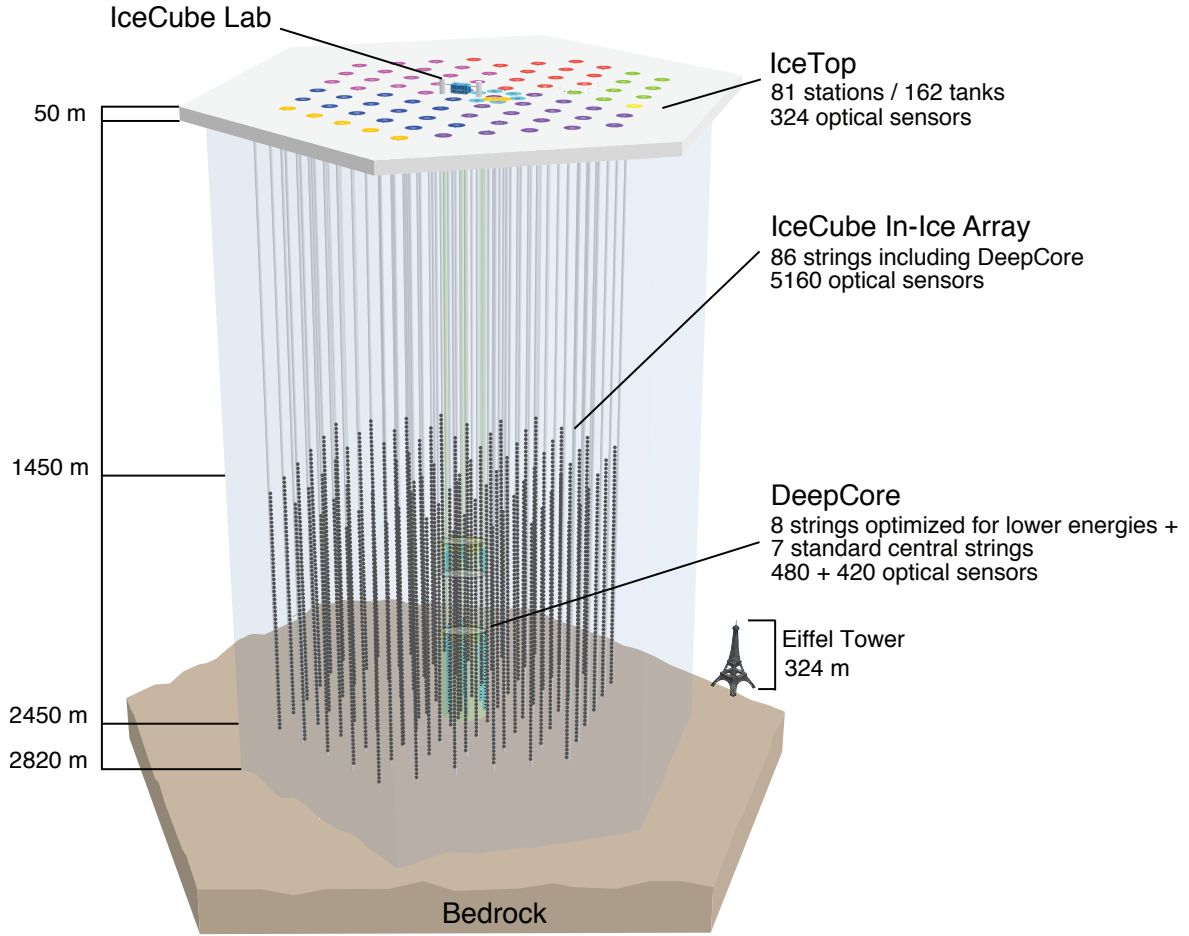
The Antarctic continent provides a unique environment for scientific research in a variety of fields. This includes astronomy and astroparticle physics for which the geographic South Pole is a valuable location for long-term observations as it has the entire southern sky overhead continuously. For example, the *South Pole Telescope* (SPT) takes measurements of the CMB [55] and contributed as part of the Event Horizon Telescope collaboration to the first images taken of black holes [56]. Other benefits that Antarctica offers are the stable temperature in the stratosphere and the circumpolar wind during the austral summer, which enable long-duration balloon missions like the *Antarctic Impulsive Transient Antenna* (ANITA) [57].

After the *Antarctic Muon And Neutrino Detector Array* (AMANDA) proved that it is possible to detect high-energy neutrinos in a big volume of ice [58], plans for an even larger successor — the *IceCube Neutrino Observatory* — took shape. Starting in 2005, the IceCube collaboration built a cubic-kilometer neutrino telescope at the South Pole in close proximity to the Amundsen–Scott South Pole Station. IceCube uses the vast mass of clear glacial ice as detection medium for high-energetic particles.

Despite its name, the IceCube Neutrino Observatory is not only contributing to a large variety of research topics in neutrino astronomy and neutrino physics. In addition to revealing the astrophysical origin of high-energetic neutrinos as the declared goal of IceCube, cosmic-ray measurements are another major field of interest for IceCube researchers. Furthermore, the IceCube collaboration investigates physics beyond the standard model and impacts Earth science fields like atmospheric and glaciological science. For this multitude of transdisciplinary sciences, IceCube consists of various detector components.

### 4.1 Detector Components

IceCube is a multi-detector astroparticle observatory, primarily comprising the deep *in-ice array* and the surface array *IceTop* (Fig. 4.1). Both are connected to the *IceCube Lab* (ICL), a building in the center of IceTop housing the computing facilities for *data acquisition* (DAQ) as well as *processing* and *filtering* (PnF). The IceCube site is located inside the South Pole station’s so-called *Dark Sector*, which besides ICL also accommodates the *Dark Sector Lab* (DSL) with the *Background Imaging of Cosmic Extragalactic Polarization* experiment (BICEP) and SPT as well as the *Martin A. Pomerantz Observatory* (MAPO) housing the *Keck Array*. The construction of IceCube was finished in 2011. New additional detector types are under development and have already been partially tested at the South Pole, such as the IACT IceAct [59]. A first batch of radio antennas and scintillators is already operated at the experiment site as well.



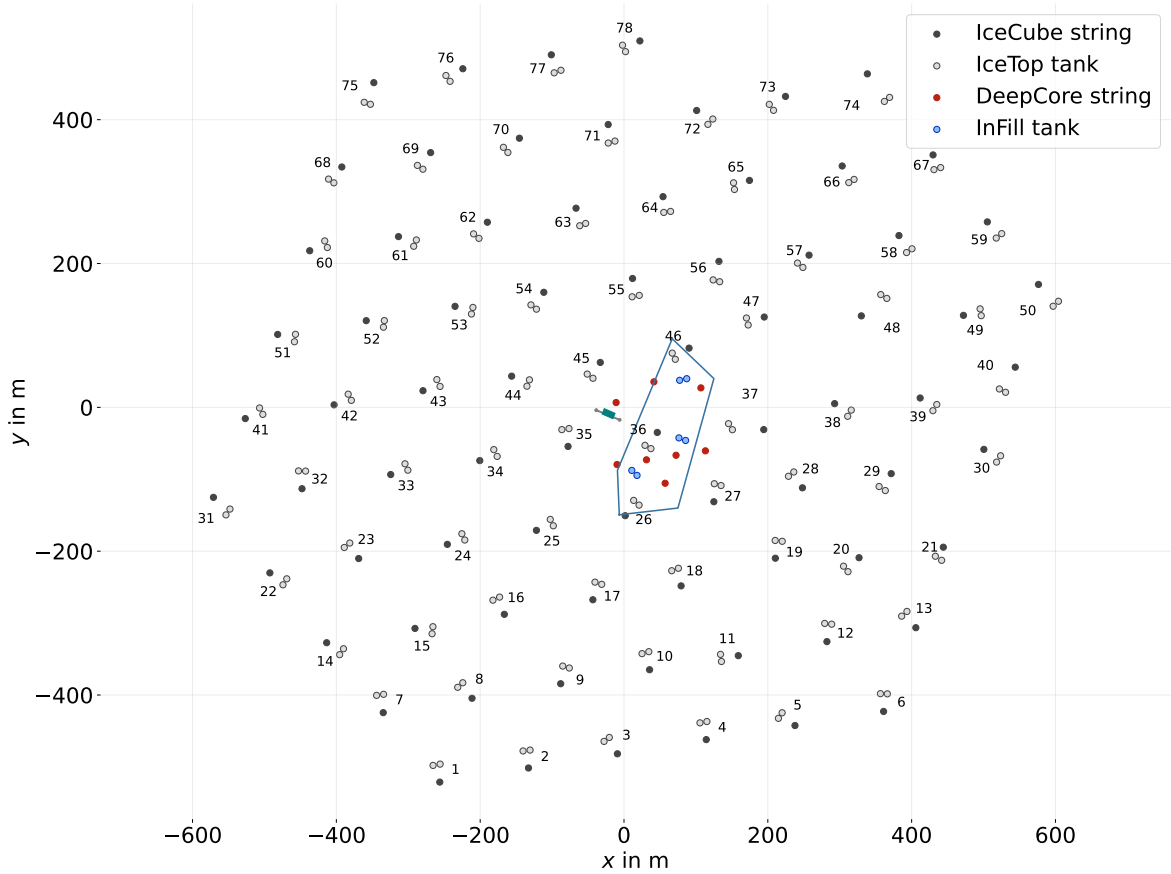
**Figure 4.1:** Schematic of IceCube depicting the strings of the in-ice array and the surface array IceTop, where colors indicate the year of deployment (Fig. 1 in Ref. [62]).

Eventually, those new detectors will constitute the surface array of an even larger future observatory: *IceCube-Gen2*. While IceAct telescopes, scintillators and radio antennas will enhance the surface array, the in-ice array is planned to grow by a factor of 8 in order to increase acceptance for the brightest neutrino events with primary energies beyond a few PeV [60]. Before the construction of Gen2, the *IceCube Upgrade* will be installed during the 2025/26 season and improve the detection of neutrinos with only a few GeV energy, leading to better sensitivity for studying neutrino oscillations [61].

#### 4.1.1 In-Ice Array

At a depth between 1450 and 2450 m below the surface, 5160 light sensors, so-called *digital optical modules* (DOMs) have been deployed in a volume of about  $1 \text{ km}^3$ . They are attached to 86 vertical cables, called *strings*, thus 60 DOMs per string with an inter-DOM spacing of 17 m. 78 of the strings are arranged in an almost regular triangular grid with a horizontal spacing of 125 m (Fig. 4.2). Additionally, eight strings have been deployed in the center of the array in order to lower the energy threshold and hence increasing the sensitivity for dedicated analyses<sup>1</sup>. Together with the surrounding strings, this part of the in-ice array constitutes

<sup>1</sup>It is this concept of a low-energy in-fill array, which is taken to the next level by the IceCube Upgrade.

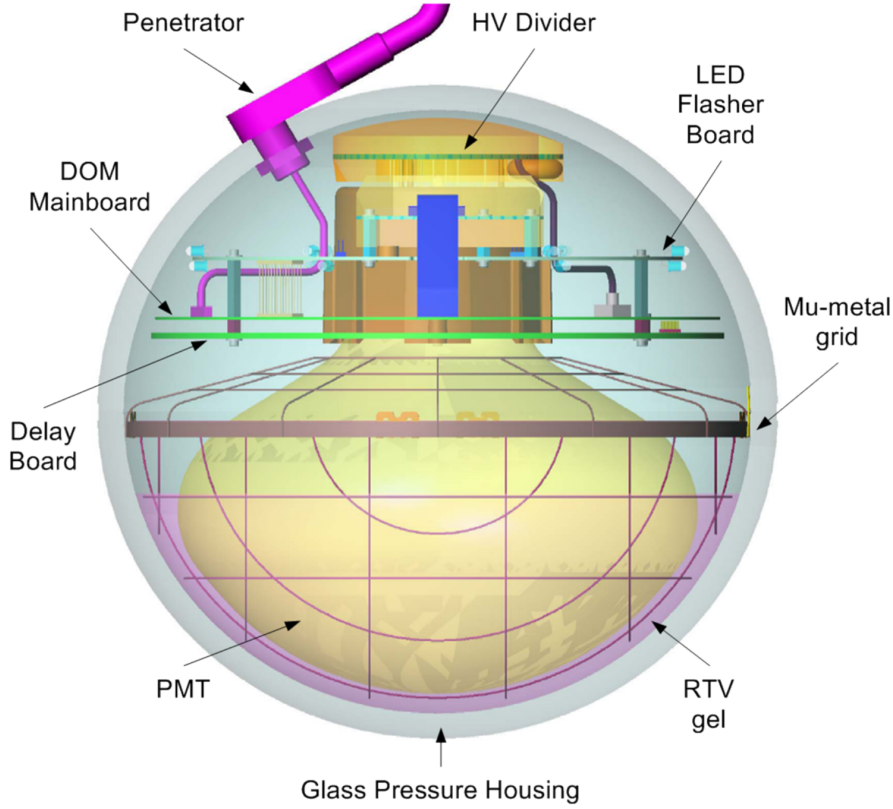


**Figure 4.2:** Layout of all in-ice strings and IceTop tanks with their respective numbers as well as the ICL building in the center. Positions are given in the IceCube coordinate system. Here, only the additional stations 79, 80 and 81 are labeled as “InFill”, whereas in this work the term InFill also includes stations 26, 36 and 46.

the sub-detector *DeepCore*, which is located mainly in the bottom third of the in-ice array ( $\rightarrow$  Fig. 4.1). The horizontal string spacing of *DeepCore* ranges between 42 m and 72 m, the vertical DOM separation between 7 m and 10 m [63]. At a depth of 1975–2100 m below the surface, the deposition of dust locally leads to worse optical properties of the ice [64].

The 3-dimensional IceCube coordinate system has its origin close to the center of the in-ice array at a depth of about 1948 m below the surface and 884 m above sea level. Depth is specified along the z-axis, which is pointing up, away from Earth. The y-axis is defined to align with the prime meridian, pointing towards Greenwich. The x-axis is rotated 90° clockwise with respect to the y-axis, completing a right-handed Cartesian coordinate system. Due to the constant movement of the Antarctic ice shield which IceCube is embedded in, the detector positions relative to the global coordinate system are not stationary. The geographic South Pole is located approximately 1.1 km “south-east” of the center of IceCube.

Each DOM ( $\rightarrow$  Fig. 4.3) is built of a thick glass sphere with a diameter of 33 cm. Inside, a down-facing 25 cm photomultiplier tube (PMT) is optically coupled via a silicone gel to the glass and shielded from Earth’s magnetic field by a mu-metal mesh. The PMT time resolution is about 2 ns [65]. The tube is connected to a mainboard, which digitizes and sends the signal



**Figure 4.3:** Schematic of an IceCube DOM (Fig. 2 in Ref. [66]). The same type of DOM is used for the in-ice array as well as in IceTop.

though a cable to the surface [66]. IceCube DOMs withstand the cold temperature and high pressure in the ice very well. As of 2025, only 79 of the 5160 originally deployed modules stopped operating, most of which (55) failed during their deployment.

A high-energy neutrino heading towards the in-ice array has a very low probability to interact with an atomic nucleus in the surrounding ice or the bedrock and thereby decay into a charged lepton. As a consequence of the highly relativistic speed of the neutrino, the charged lepton travels faster than the effective speed of light in ice ( $\sim 0.75 \times c$ ), causing the emission of Cherenkov light, which can be recorded by the in-ice DOMs.

The hardware that was primarily built for the detection of neutrino events can also be used to study cosmic rays. When the bundle of high-energy muons ( $E_\mu \gtrsim 300 \text{ GeV}$ ) that are produced in an EAS intersects the in-ice array, the emitted Cherenkov light can be recorded in coincidence with the event at the surface. This is crucial for estimating the number of high-energy muons and can be used to improve the reconstruction of the arrival direction.

#### 4.1.2 IceTop

The surface array of IceCube has been constructed to fulfill mainly two purposes: enabling advanced cosmic-ray studies and acting as a veto for the detection of down-going neutrinos. This hybrid detector layout was first tested during the operation of AMANDA with the surface arrays SPASE (South Pole Air Shower Experiment) [67] and VULCAN [68]. IceTop consists of 81 stations on top of almost all IceCube strings. Similarly to the in-ice array, 78 stations are

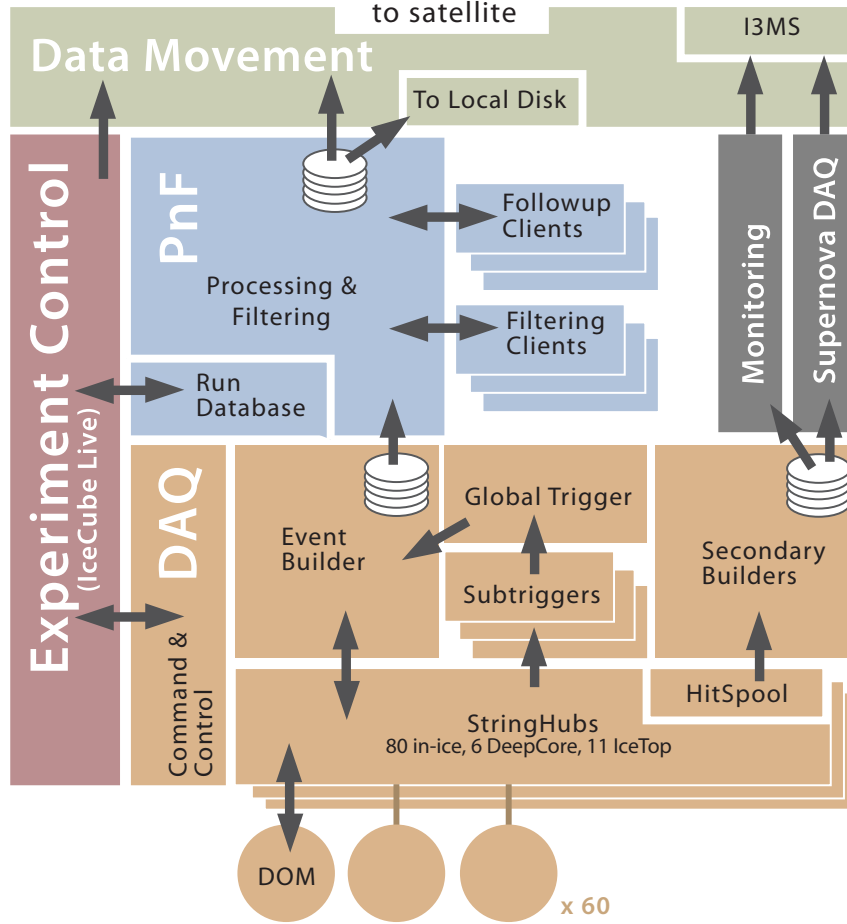
arranged in an almost regular triangular grid, covering an area of  $1 \text{ km}^2$ . Additionally, three more stations are located above the first three DeepCore strings, thus creating a denser *InFill* sub-array at the surface ( $\rightarrow$  Fig. 4.2). In contrast to some previous publications, in this work the term “InFill” covers the stations 26, 36, 46, 79, 80 and 81, which have a low separation between each other, thus not including stations 27 and 37. IceTop stations are not perfectly aligned with the strings below but rather offset by approximately 25 m. Therefore, the surface array (excluding stations 79–81) is less regular with a median station spacing of 124 m, but ranging from 71 m to 193 m. The total size and the spacing between stations determine the range of primary cosmic-ray energy accessible for IceTop, which reaches up to a few EeV. In the main array, the energy threshold lies at about 300 TeV. In the InFill, neighboring stations have a separation just below 50 m, which allows some air showers with primary energy even below 100 TeV to trigger the detector readout ( $\rightarrow$  subsection 4.2.1). IceTop is located 2830 m above sea level, which corresponds to a vertical atmospheric depth  $\Xi$  of roughly  $695 \text{ g/cm}^2$ .

Each IceTop station is a pair of cylindrical ice-Cherenkov tanks, 10 m apart and labeled *A* and *B* respectively. Every tank houses two down-facing DOMs embedded in  $2.3 \text{ m}^3$  of clear ice, one operating in high-gain (HG), the other in low-gain (LG) PMT mode for an increased dynamic range. Similarly to the in-ice array, IceTop utilizes the Cherenkov effect for the detection of relativistic charged particles. Light emitted inside the tank volume is registered by the DOMs [69]. At deployment, the tank lids have been flush with the snow level around. In the following years, snow accumulated on top ( $\rightarrow$  Fig. 5.7), leading to attenuation of the recorded signal, which results in a decreased trigger efficiency at low energies.

In the HG DOMs of IceTop tanks 26A and 67B, one of the three ATWD chips ( $\rightarrow$  subsection 6.3.1), which are needed for signal digitization, failed. As a consequence, the gain mode of all DOMs in those tanks has been switched from high to low and vice versa. In 2007, DOM 39-61 and in 2022 DOM 74-61 (both HG) failed permanently. Their partner DOMs have been set to high-gain mode in order to provide the maximum possible sensitivity to air-shower events.

## 4.2 IceCube Online System

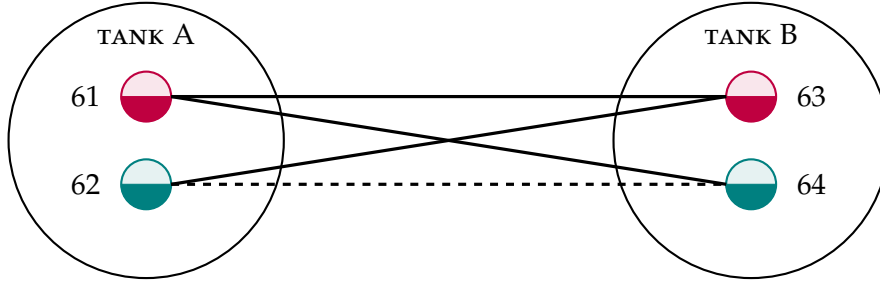
The IceCube online system at the South Pole (Fig. 4.4) is responsible for triggering the detector readout and basic filtering steps before the data can be sent north via satellite and saved for later analysis. Summarized information on all DOMs, triggers, filters and environmental conditions of each data-taking run is made available for monitoring purposes. Typical IceCube runs last eight hours. Once every year, typically in May, the online system is updated, marking the start of a new season (for example, IC86.2018 where “IC86” stands for the 86 IceCube strings in the detector configuration and “2018” indicates the start year of the season). The detector uptime of IceCube is above 99.7% [60]. Recording every electronic pulse in each DOM at any time and sending them north for later examination would be unnecessarily expensive and impractical. In order to get rid of most of the pure noise, event triggers and subsequently physics filters are needed.



**Figure 4.4:** Flow chart of IceCube's triggering (beige) and filtering system (blue) (Fig. 41 in Ref. [62]).

#### 4.2.1 Triggering

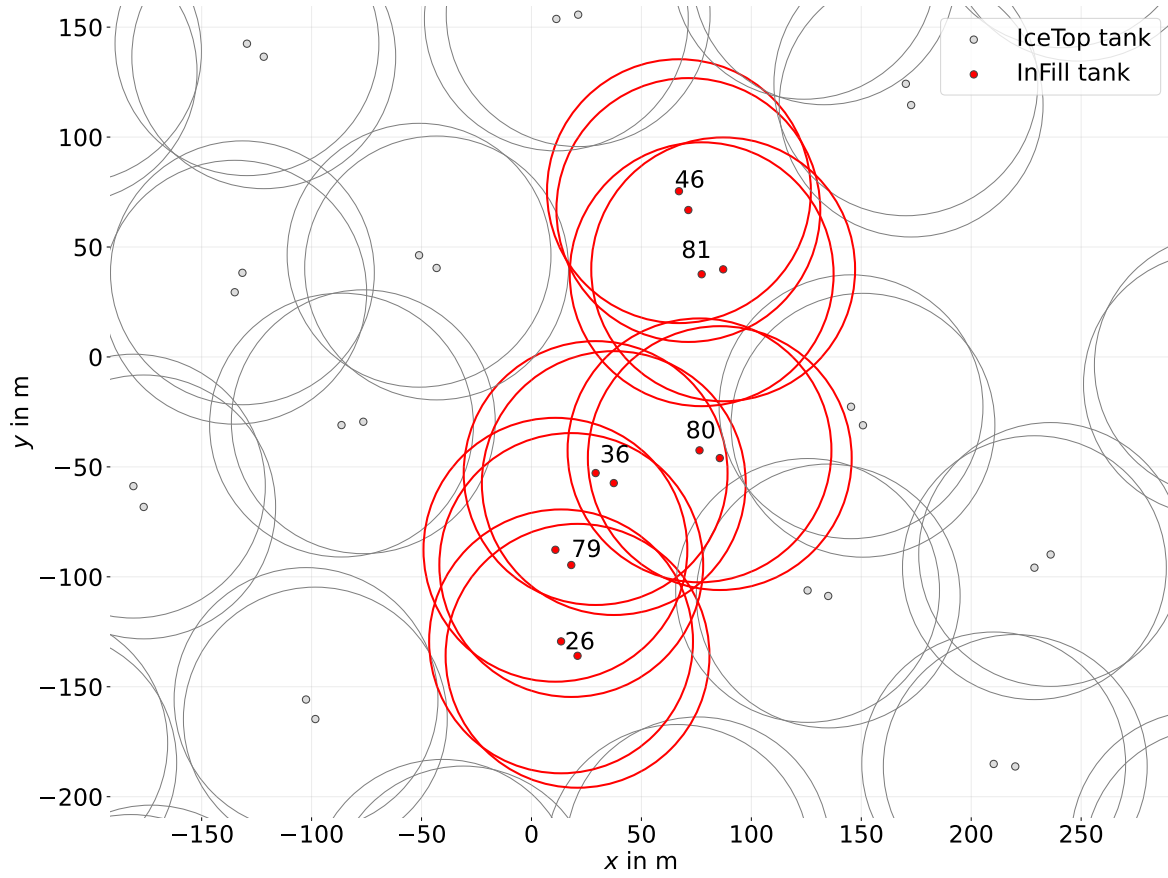
In order for a DOM to digitize the waveform that was recorded with the PMT, the peak pulse charge needs to exceed a predetermined threshold value. This first step towards a recorded pulse is the autonomous *discriminator trigger*, which is defined on each DOM individually. However, the waveform digitization is aborted if the so-called local coincidence (LC) condition is not fulfilled within a  $1\ \mu\text{s}$  time window. For this, triggered DOMs send a signal to neighboring modules. In the in-ice array these are the two DOMs above and below on the same string [62]. When at least one of those four has a discriminator trigger as well, the LC condition is satisfied. For IceTop, only HG DOMs can initiate a LC and they look for modules that passed the discriminator threshold within  $1\ \mu\text{s}$  in the neighboring tank of the same station. For that purpose, cables connect the DOMs of a station (Fig. 4.5). When such a pair of DOMs (HG/HG or HG/LG) in IceTop or a group of neighboring in-ice DOMs exists, their waveforms keep getting digitized and stored in the on-board memory as so-called hard local coincidence (HLC) pulses. In-ice pulses are saved as HLCs when they fulfill the LC condition. The waveform digitization process will be outlined in more detail in section 6.3. Dedicated computers in the ICL, the *DOMHubs*, regularly request the data from all in-ice strings and IceTop stations to check it for potential event-level trigger conditions. There is one DOMHub for each in-ice string and eleven additional DOMHubs for IceTop.



**Figure 4.5:** Sketch of an IceTop station (everything except for the tank separation is to scale). Each tank houses two DOMs, one in HG (red), the other in LG mode (green). DOMs of neighboring tanks are connected via cable. The cable between both LG DOMs (dashed) is only needed in case of a gain switch. DOM numbers are indicated.

On the DOMHubs, a software called *StringHub* takes the HLC data and forms software triggers. There are numerous ways to trigger IceCube. Here, only those triggers relevant for this work are presented. The most fundamental IceTop trigger for analyses with cosmic-ray energies of PeV and above [70] is the simple multiplicity trigger (**SMT**), which is fulfilled with six or more HLC hits within a sliding  $5\ \mu\text{s}$  window without any geometric requirements. This condition can be satisfied by a minimum of three IceTop stations. In order to be sensitive to lower-energy air showers, the IceTop Volume trigger has been implemented in 2016. For it to occur, there need to be four HLC pulses from HG IceTop DOMs within a radius of 60 m around a hit DOM and a time window of 200 ns. Geometrically, this condition can only be fulfilled in the dense InFill ( $\rightarrow$  Fig. 4.6) by the following four station pairs: (26, 79), (36, 79), (36, 80) and (46, 81). Even though unlikely, this trigger condition can also be fulfilled by the station triplets (26, 79, 36) and (79, 36, 80) [71]. The IceTop Volume trigger will be referred to as *2-station trigger* in this work. A SMT trigger is also defined in the in-ice array, requiring at least eight HLC hits within a sliding  $5\ \mu\text{s}$  time window. In analogy to the 2-station trigger, an in-ice volume trigger is initiated when more than three HLC hits occur inside a cylinder of 75 m height and 175 m radius and a  $1\ \mu\text{s}$  time window. Vertical down-going muons can be triggered specifically by a string trigger, which needs five HLCs out of seven adjacent DOMs on a string within  $1.5\ \mu\text{s}$ . For the DeepCore sub-detector, a SMT trigger is build, which requires 3 HLCs in the lower DeepCore DOMs within  $2.5\ \mu\text{s}$ . The rates of the mentioned triggers are given in Table 4.1.

The DAQ system appends a padding of several microseconds before and after each trigger to form readout windows. All pulses that passed a discriminator trigger and fall into such a readout window are collected. This includes not only HLC hits but also single-hit DOMs, so-called soft local coincidence (**SLC**) pulses. The waveforms of SLCs generally do not get digitized. Instead, only a time stamp and the integrated charge are kept. Moreover, the readout is not limited to the sub-detector that initiated the trigger but takes pulses from the entire detector. To avoid having individual pulses stored in more than one event, overlapping readout windows are merged into a global trigger, which is eventually used to build DAQ events (lower part in Fig. 4.4).



**Figure 4.6:** The inner part of the IceTop array. InFill tanks are shown in red, regular tanks in gray. Around each tank, a circle with 60 m radius is drawn. Only tanks of neighboring stations in the InFill have a separation below 60 m.

**Table 4.1:** IceCube triggers relevant for this work with their respective average rates during the IC86.2018 configuration year.

Trigger Name	Alias	Rate
ICE_TOP_SIMPLE_MULTIPLICITY	IceTopSMT / SMT6	12.7 Hz
ICE_TOP_VOLUME	2-station trigger / IceTopVolume	2.72 Hz
IN_ICE_SIMPLE_MULTIPLICITY	InIceSMT / SMT8	2.13 kHz
IN_ICE_VOLUME	InIceVolume	3.77 kHz
IN_ICE_STRING	InIceString	2.26 kHz
IN_ICE_SIMPLE_MULTIPLICITY_DEEPCORE	DeepCore trigger	0.26 kHz

**Table 4.2:** *IceCube filters relevant for this work with their respective average rates during the IC86.2018 configuration year. The two-digit number at the end indicates the year when the filter was integrated in the DAQ system.*

Filter Name	Alias	Rate
IceTopSTA5_13	5-station filter	1.52 Hz
IceTopSTA3_13	3-station filter	0.68 Hz
IceTop_InFill_STA3_13	3-station InFill filter	0.133 Hz
IceTop_InFill_STA2_17	2-station filter	2.69 Hz
InIceSMT_IceTopCoincidence_13	coincidence filter	0.71 Hz

#### 4.2.2 Filtering

Triggering events using the LC requirement avoids recording pure noise. However, a dedicated filtering stage is necessary to limit the dataset to events that are interesting to analyze.

The DAQ system hands over the triggered events to the PnF system where the waveforms of HLCs are calibrated. By deconvolution with the DOM response to single photons, the waveforms are turned into charge values in units of photoelectrons (PE) and subsequently compressed to a memory-efficient format. The calibration constants for the conversion from voltage to PE for each DOM are stored in files with the geometry, calibration and detector status (GCD) information of each run. Basic reconstructions are applied to the events and filter information is added. Both are necessary to judge whether an event is potentially relevant for analysis or even needs to be transmitted north directly. When it comes to in-ice events, particularly interesting events can trigger an automatic alert via the General Coordinates Network to other observatories around the world in order to start multi-messenger follow-up measurements. All data is stored on tape in the custom i3 file format and transported once a year to the *Wisconsin IceCube Particle Astrophysics Center* (WIPAC) in Madison, Wisconsin, where it is saved on disk [62]. The data from each DAQ readout window is stored in individual Q-frames in the file.

For this work, the relevant filters are briefly presented. Based on the SMT6 trigger, the corresponding 3-station filter has been defined. It requires a minimum of three regular IceTop stations (not including stations 79–81) to have HLC hits. In order to select high-energy showers, the 5-station filter has been introduced, which is passed when five or more stations with HLCs are present in an event. In contrast to that, the energy threshold could be lowered with the 3-station InFill filter, which looks for three stations in the dense InFill region of IceTop<sup>2</sup> [69]. Accessing even lower energies is possible by applying the 2-station filter, which selects all events that have been created from a 2-station trigger. Lastly, a filter has been constructed that checks for coincidences between the in-ice array and IceTop. It requires an InIceSMT trigger, that is, eight in-ice HLC hits and at least one IceTop HLC pulse. The filters are listed with their respective rates in Table 4.2. From it, a total number of filtered events in the order of  $10^8$  per year can be estimated.

<sup>2</sup>In contrast to the 2-station filter, the definition of the 3-station InFill filter also includes stations 27 and 37.

## CHAPTER 4. THE ICECUBE NEUTRINO OBSERVATORY

### **Key Messages**

*IceCube is an extraordinary astroparticle observatory at the South Pole. IceCube comprises several sub-detectors which can be regarded individually but are also designed to operate in conjunction. For the investigation of cosmic-ray mass composition, studying events that are registered by the IceTop air-shower array and in the deep in-ice detector in coincidence is essential. The 2-station trigger enriches the acquired dataset by providing a large amount of low-energy events.*

# CHAPTER 5

## EVENT SIMULATION

*“If you wish to make an apple pie from scratch you must first invent the universe.”*

— Carl Sagan, *Cosmos*

Computer simulations play an indispensable role in modern physics. By comparing measured data with the corresponding simulated values or distributions, insights can be obtained into both the generating model and the true nature of the underlying physics. As the development of air showers is subject to large fluctuations, *Monte Carlo* (MC) simulations are necessary to produce big datasets, covering a wide variety of events. The Monte Carlo method relies on repeated random sampling with the goal of emulating complex processes such as interactions of subatomic particles. MC simulations are also used to tweak calibration procedures, build data processing and event selection pipelines, as well as to test reconstruction methods before they can be applied to experimental data. Furthermore, simulations provide a good way to check competing theoretical or phenomenological models, such as those describing the hadronic interaction in high-energy particle collisions.

The production of MC data in particle physics typically involves two parts: event generation (→ [section 5.1](#)) and simulation of the detector response (→ [section 5.2](#)). For the first time, coincident in-ice background is added to the IceTop event simulation, which is described in [section 5.3](#). The influence of the environmental conditions at the South Pole on the air-shower data taking is discussed in [section 5.4](#).

### 5.1 CORSIKA

For the simulation of extensive air showers, **CORSIKA** (*Cosmic Ray Simulations for KASCADE*) [72] has become the standard tool in the astroparticle physics community after it had been originally developed for the **KASCADE** experiment. The current software version 7 is written in Fortran and receives regular updates. A new C++ framework — CORSIKA 8 — which provides more flexibility and new functionality [73, 74] has been released recently.

#### 5.1.1 Air-Shower Simulation with CORSIKA

Provided with a particle type, the Monte Carlo code injects a primary cosmic ray at the top of Earth’s atmosphere with a random energy and direction sampled from specified distributions. The height of first interaction is sampled randomly as well, given an atmospheric profile, which describes air pressure (atmospheric depth) as a function of altitude (→ [subsection 5.4.1](#)). Following the initial interaction, CORSIKA keeps track of the secondary particles (nuclei, hadrons, muons, electrons and photons<sup>1</sup>) that are created in the cascade as long as their energy exceeds a predefined threshold. Four parts of the program take care of specific processes in the shower development:

---

<sup>1</sup>If needed, also the radiation emitted in the air (Cherenkov, fluorescence and radio) can be simulated.

- decay of unstable particles, tracking, energy losses via ionization, multiple scattering in the atmosphere and deflection in Earth’s magnetic field
- high-energy hadronic interactions with air nuclei
- low-energy hadronic interactions
- interaction and transport of electromagnetic particles ( $e^\pm/\gamma$ )

Secondary particles that reach the specified level of observation are saved with information on their particle type and momentum, as well as location and time relative to injection. Subsequently, this data can be used as input for a detector response simulation, which will be covered in [section 5.2](#). Furthermore, the longitudinal shower profile, that is, the multiplicity of various particle species as a function of slant depth, can also be saved.

Particle interactions can be simulated with various specialized models. For electromagnetic interactions, CORSIKA falls back on the EGS4 (*Electron Gamma Shower*) code [75]. For hadronic interactions, the user can choose from a selection of models. For the description of low-energy interactions, FLUKA (for the German “**F**luktuierende **K**askade”) [76, 77] is the most common option, while for high energies ( $> 80\text{ GeV}$ ), generally EPOS (Energy conserving quantum mechanical multiple scattering approach based on Partons Off-shell remnants Splitting of parton ladders) [78], QGSJet (Quark-Gluon String model with Jets) [79] or Sibyll<sup>2</sup> [80] are used. The prediction of rates and spectra of secondary particles becomes increasingly difficult at energies that exceed those achievable with terrestrial particle accelerators. The extrapolations made by the hadronic interaction models therefore rely on theoretical assumptions tuned by available accelerator data. Tuning of the model parameters, however, is challenging as the detectors at particle accelerator facilities mostly probe collisions with a large momentum transfer (high- $p_T$ ). The development of cosmic-ray showers, on the other hand, is dominated by low- $p_T$  interactions, i.e., small momentum transfer, which cannot be described by perturbative quantum chromodynamics and instead needs to be approximated by Gribov–Regge field theory. An extrapolation to high-energy EASs is further aggravated as the majority of collider experiments investigate proton-proton interactions, whereas collisions of nuclei and hadrons with nitrogen and oxygen are fundamental for air-shower development.

The high-energy hadronic interaction models mentioned above have been tuned to data from the large hadron collider (LHC) at CERN. Their latest versions are thus called *post-LHC* models, but as they follow different approaches in the theoretical description and extrapolation to the regime of high-energy cosmic-ray air showers, the generated particle multiplicities, angular distributions, ratios of charged to neutral particles, etc. generally do not agree with each other. The three most commonly used post-LHC models are currently Sibyll 2.3d [81], QGSJet-II.04 [82] and EPOS-LHC [83]. An overview of the different approaches followed by those hadronic interaction models is provided in Ref. [84].

---

<sup>2</sup>referring to the sibyl oracles in ancient Greece.

### 5.1.2 Simulation Production

CORSIKA offers numerous options that can be activated for air-shower simulation. For the MC production in this work, CORSIKA version 7.7420 has been compiled with the following options<sup>3</sup>:

- CHARM: this option enables explicit treatment of charmed particles and tau leptons
- CURVED: for zenith angles above  $70^\circ$ , the approximation of a flat atmosphere overestimates the slant depth; with this option activated, a more realistic atmospheric profile is mimicked for very inclined showers
- LPM: the Landau–Pomeranchuk–Migdal effect [85, 86] reduces the cross section for bremsstrahlung and pair production at very high energies
- NEUTRINO: the decay of kaons, pions and muons produces (anti-)neutrinos; when selecting this option, CORSIKA keeps track of them without simulating their interactions
- SLANT: this option sets the binning of the longitudinal shower profiles to slant depth intervals instead of vertical altitude bins
- UPWARD: with this option, also secondary particles moving upward are tracked

As hadronic interaction models, FLUKA 2021.2 (for low-energy interactions) and Sibyll 2.3d (for high energy) have been selected. Four datasets have been produced with the primaries being proton (H), helium ( $^4\text{He}$ ), oxygen ( $^{16}\text{O}$ ) and iron ( $^{56}\text{Fe}$ ). The number of simulated showers in an energy bin of width 0.1 in  $\log_{10}(E)$  increases towards lower energy in order to achieve higher statistics in the regime below a few PeV where the trigger efficiency of IceTop drops. For validation purposes, a fifth dataset has been simulated with silicon ( $^{28}\text{Si}$ ) as primary nucleus. For it, the number of showers per energy bin has been scaled down to 20%. Since the focus of this work lies on cosmic rays with an energy just below 1 PeV, the number of showers with the four main primaries has been ramped up to 10 000 per bin for  $5.0 \leq \log_{10}(E/\text{GeV}) < 6.0$  ( $\rightarrow$  Tab. 5.1). A total of 152 500 showers have been produced for each main dataset and 15 000 for the silicon dataset.

Smaller datasets have been produced with the alternative hadronic interaction models EPOS-LHC and QGSJet-II.04. Their size is 10% of the original amount of showers using Sibyll ( $\rightarrow$  right column in Tab. 5.1). In total, 7 500 showers have been simulated for both models each.

At the start of every shower simulation, a steering file provides crucial parameters which include environmental conditions and properties of the cosmic-ray primary. Table 5.2 gives an example for such a CORSIKA steering file. The energy range and slope settings define the minimum and maximum energy and the spectrum to randomly draw a value from.

---

<sup>3</sup>The CURVED, LPM and NEUTRINO options have virtually no effect on the CORSIKA datasets used for the presented analysis, which deals with low-energy, almost vertical cosmic-ray events. They have been activated to ensure consistency with potential extension datasets in the future.

**Table 5.1:** Amount of simulated CORSIKA showers in energy bins of width 0.1 in  $\log_{10}(E)$  for different energy ranges. The numbers for the main datasets with Sibyll 2.3d, the corresponding silicon validation set and the smaller datasets with QGSJet-II.04 and EPOS-LHC are given.

$\log_{10}(E/\text{GeV})$	number of bins	main (Sibyll)	validation (Sibyll)	QGSJet EPOS
4.0–5.0	10	3000	600	300
5.0–5.5	5	10 000	500	250
5.5–6.0	5	10 000	400	200
6.0–6.5	5	1500	300	150
6.5–8.0	15	1000	200	100

**Table 5.2:** Example for a steering file used for the production of CORSIKA datasets. The highlighted lines are described in the text. ‘T’ and ‘F’ decode the true and false Booleans, respectively. For thorough definitions of all available settings,  $\rightarrow$  Ref. [72].

RUNNR	581148				
EVTNR	1				
SEED	3581148	0	0		
SEED	3581149	0	0		
SEED	3581150	0	0		
NSHOW	1				
<b>ERANGE</b>	<b>6.30957E+05</b>	<b>7.94328E+05</b>			
<b>ESLOPE</b>	<b>-1.0</b>				
<b>PRMPAR</b>	<b>1608</b>				
<b>THETAP</b>	<b>0.0</b>	<b>65.0</b>			
<b>PHIP</b>	<b>0.0</b>	<b>359.99</b>			
<b>ECUTS</b>	<b>0.0500</b>	<b>0.0500</b>	<b>0.0100</b>	<b>0.0020</b>	
ELMFLG	F	T			
<b>OBSLEV</b>	<b>2840.E2</b>				
ECTMAP	100				
SIBYLL	T	0			
SIBSIG	T				
SIBCHM	T				
ARRANG	-120.7				
HADFLG	0	1	0	1	0 2
STEPFC	1.0				
DEBUG	F	6	F	1000000	
MUMULT	T				
MUADDI	T				
MAXPRT	0				
<b>MAGNET</b>	<b>16.75</b>	<b>-51.96</b>			
LONGI	T	10.	T	T	
<b>ATMOD</b>	<b>33</b>				

In Table 5.2 an  $E^{-1}$  spectrum ranging from  $10^{5.8}$  to  $10^{5.9}$  GeV is set. The primary particle is encoded by an integer (14 for protons and  $A \cdot 100 + Z$  for nuclei with the mass number  $A$  and the atomic number  $Z$ ). The provided example refers to an oxygen-initiated air shower. The phase space for the arrival direction is set separately for zenith ( $0^\circ \leq \theta \leq 65^\circ$ )<sup>4</sup> and azimuth ( $0^\circ \leq \varphi < 360^\circ$ ). For flat detector arrays such as IceTop, CORSIKA uses a zenith distribution following  $dN/d\theta \propto \sin(\theta) \cos(\theta)$ . The energy thresholds below which particles are no longer tracked have been set to 50 MeV for hadrons and muons, 10 MeV for electrons and 2 MeV for photons. Neutral pions obey the photon cut values instead of the hadronic one. The altitude at which the particle distribution is evaluated and saved, that is, the observation level, has been set to 2840 m above sea level, which is a few meters above the IceTop tanks and thus just above the snow surface during the 2018 season. Environmental conditions are covered by setting the horizontal and vertical component of Earth's magnetic field as well as an atmospheric model. The value of 33 represents the April atmosphere at the South Pole, which will be explained in detail in subsection 5.4.1.

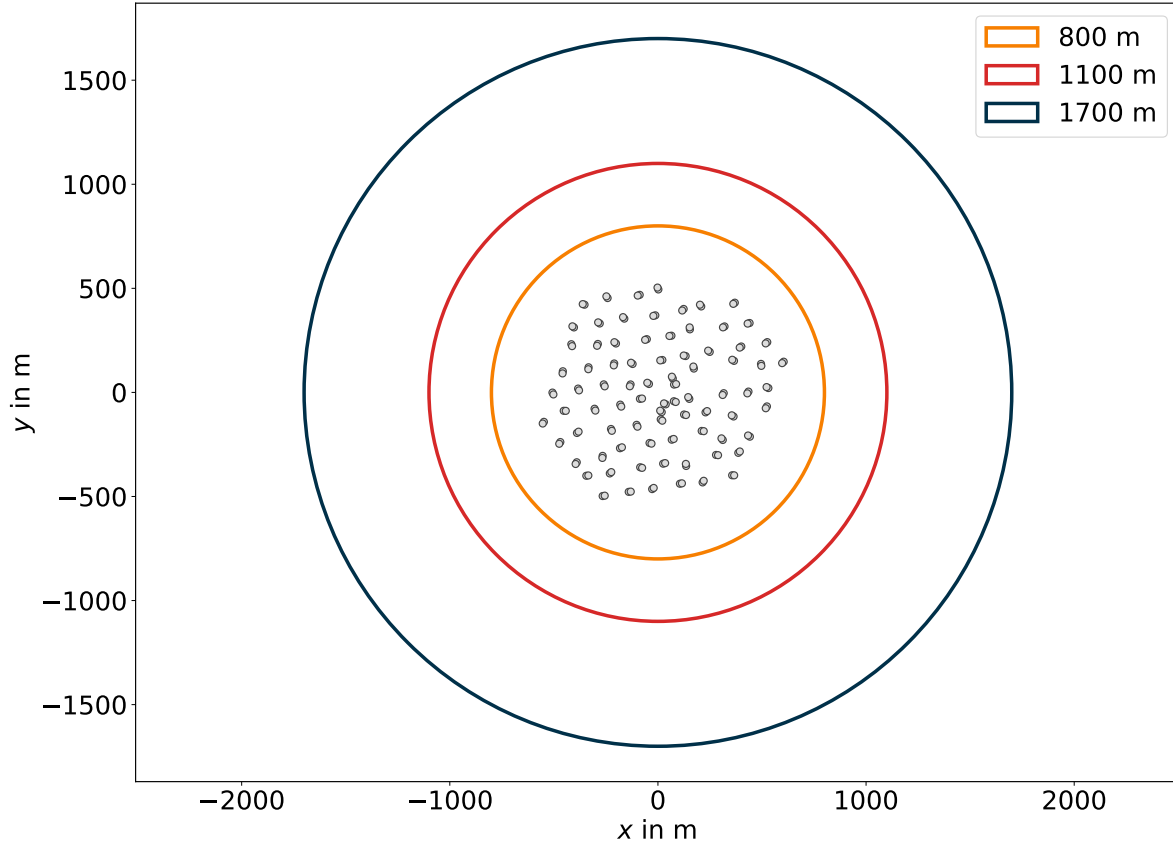
## 5.2 Detector Simulation

In order to complete the Monte Carlo procedure, the detector response of IceTop and the in-ice array to the particles generated by CORSIKA showers must be simulated. The software used for this is part of the IceTray framework. The statistics of the resulting MC datasets can be increased by *resampling*, i.e., injecting each simulated shower multiple times above the detector. The core positions are randomly sampled from a circular area (Fig. 5.1), centered at the midpoint of IceTop and with an energy-dependent radius:

$$r(E) = 800 \text{ m} + (\lfloor \log_{10}(E/\text{GeV}) \rfloor - 5) \cdot (\lfloor \log_{10}(E/\text{GeV}) \rfloor - 4) \cdot 150 \text{ m} \quad (5.1)$$

The resampling area gets larger with each decade in energy as the shower footprint grows, too. The parameterization of Eq. (5.1) has been determined in such a way that air showers centered on the edge of the resampling circle are barely able to trigger IceTop. The default amount of resamples is 100; however, as the trigger efficiency of IceTop almost vanishes below 100 TeV primary energy, the resampling factor has been increased to 400 at those very low energies. The main datasets (H, He, O and Fe) thus are simulated with 24.25 million potential events each (before the trigger condition is applied), the validation set with silicon primaries accordingly has 3.3 million, which sums up to a total of 100.3 million events that are subject to detector simulation in the following step. For QGSJet and EPOS, detector simulation was performed only for the primary energy range  $10^5$ – $10^7$  GeV, which leads to a total number of 350 000 potential events each. Biases that might be introduced by reusing each shower several times are still tolerable. The resampling factors have been chosen to result in only a few events surviving triggering, filtering and subsequent quality cuts on average. Injecting each CORSIKA shower only once, on the other hand, would require a multiple of computing time and immense storage capacities, making this alternative approach currently unfeasible.

<sup>4</sup>All used CORSIKA datasets have a zenith range from  $0^\circ$  to  $65^\circ$  except for the additional showers that fill the statistics of the main Sibyll datasets to 10 000, which are limited to  $45^\circ$ .



**Figure 5.1:** The circles with radius 800 m, 1100 m and 1700 m (from Eq. (5.1)) skirt the resampling areas that are used for the detector simulation of air showers in the energy range 10 TeV–100 PeV. Gray dots represent IceTop tanks.

The simulation of the detector response starts with IceTop. The secondary shower particles are read from the CORSIKA file and only those with a trajectory within 30 cm of an IceTop tank are considered for the further simulation. Since the observation level in the CORSIKA steering file has been set slightly above the snow surface altitude, the particles need to be propagated through air, snow and if necessary detector material (tank walls, ice, glass, etc.). This is done using the Geant4 toolkit [87], keeping track of secondary particles that are created during propagation, as well as simulating energy deposits. When particles pass through a tank, they are potentially visible to the installed DOMs due to Cherenkov light emission. However, the exact simulation of photon propagation is computationally too expensive. Instead, the number of photons seen by the PMT is parameterized based on the number of expected Cherenkov photons, drawn from a Poisson distribution, which in turn depends on the track length inside the tank. The photon arrival time is modeled to follow an exponential distribution with a time constant that depends on the type of reflective coating inside the tank (Tyvek or zirconia). The simulation of several air showers occurring in coincidence is not part of the IceTop simulation at the moment. More details on IceTop simulations can be found in Ref. [69].

The simulation of the in-ice detector response begins with the selection of only those muons in the CORSIKA file with an energy above 273 GeV (recall Eq. (2.5)). This is the energy at

which 0.1% of vertical muons are expected to reach a depth of 1250 m before they decay, i.e., 200 m above the top of the in-ice detector [88]. The muons that exceed this threshold and thus have a chance to be visible to the in-ice DOMs then are propagated down through the ice by the MC software *PROPOSAL* (*Propagator with optimal precision and optimized speed for all leptons*) [89], which also simulates energy losses and the production of secondary particles in the ice. The Cherenkov photons emitted by these high-energy muons and their secondaries are tracked by the OpenCL-based package CLSim [90] while the photons that emerge from stochastic energy losses are treated by PHOTONICS [91]. The propagation of in-ice muons and photons runs on GPUs, in contrast to all other simulation steps for which CPUs are utilized. The optical properties of the South Pole ice have a direct impact on the Cherenkov light propagation. The simulation uses the latest ice model, which not only provides absorption and scattering coefficients, but also describes the tilt of individual ice layers [92]. Since muons from unrelated (often low-energy) showers can pass through the detector at any time, coincident in-ice background is injected within the readout window (→ section 5.3).

Following the simulation of Cherenkov photons reaching the photomultipliers, in-ice PMTs obtain random noise hits, which mimic the effect of radioactive decays in the ice or the detector material. For both detector parts, electronic noise is then added as part of the DOM mainboard simulation. Finally, the triggering system as described in subsection 4.2.1 is simulated. The subsequent filtering and further processing is applied to the Monte Carlo events in the same way as is done on data.

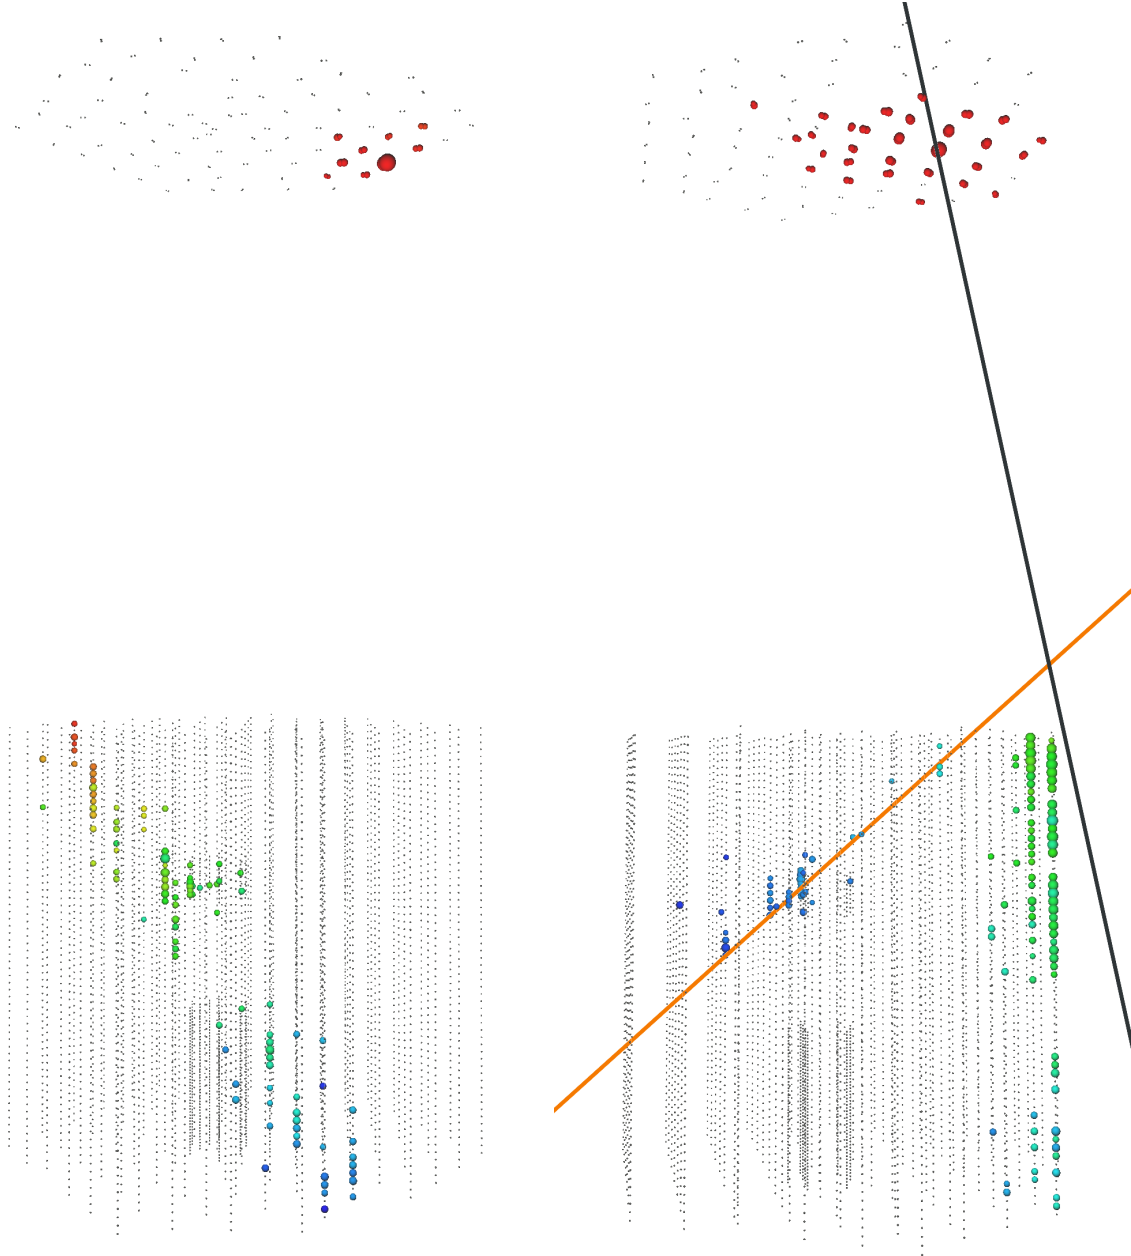
### 5.3 Injection of In-Ice Background

Due to the high flux of low-energetic cosmic rays, the air showers they initiate can produce TeV muons, which have the potential to pass through the in-ice array within the readout window of an IceTop-triggered CR event. These kind of events with coincident muon background are present in data (Fig. 5.2, left), but they have not been part of the detector simulation for air showers in IceCube so far. In analyses of cosmic rays with PeV primary energy and higher, the reconstructed “Laputop” track (→ section 6.2) can be used for additional cleaning of the in-ice pulses, which reliably removes hits from unrelated background muons. In the low-energy regime where showers might trigger only two IceTop stations, this method is not available as Laputop requires a minimum of three stations for a simplified fit and at least five stations for the advanced reconstruction. Hence, an alternative method has to be developed (covered in section 6.4) which demands the inclusion of muon background injection in the simulation chain.

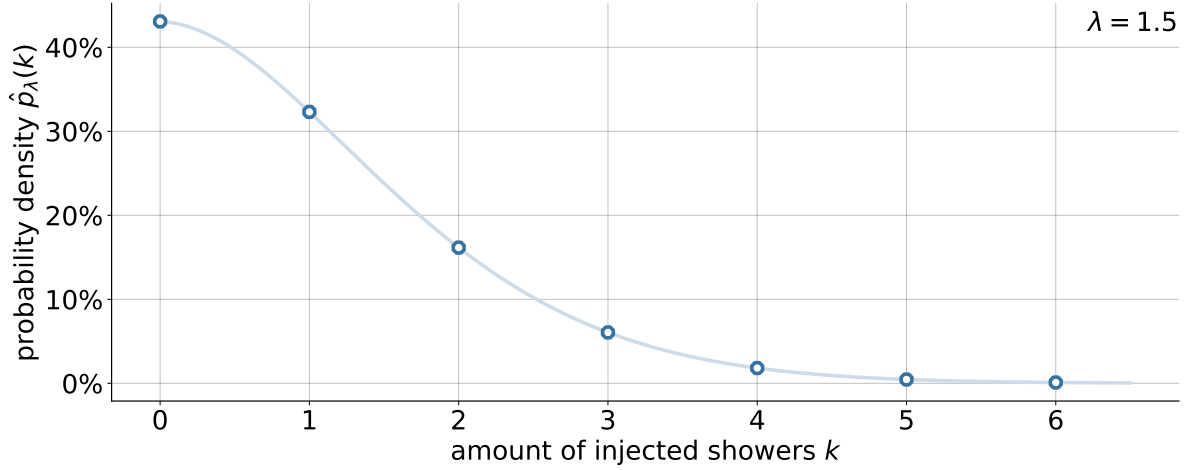
The number  $k$  of coincident muon bundles within a time window  $\Delta t$  around the earliest signal hit follows a Poisson distribution

$$p_{\lambda}(k) = \frac{e^{-\lambda} \lambda^k}{k!}, \quad (5.2)$$

where  $\lambda = R \cdot \Delta t$  is the Poisson mean with  $R$  denoting the muon rate deep in the ice. Since the purpose is to add background muons to already existing CR events, Eq. (5.2) needs to be



**Figure 5.2:** Events at processing Level 3 with background muons. The in-ice pulses underwent cleaning, which will be explained in [section 6.2](#) and — as can be seen — does not remove muon track background reliably. **Left:** Cosmic-ray event recorded on July 21, 2018, with a coincident in-ice muon bundle. **Right:** Simulated Monte Carlo event (helium primary,  $E_0 = 29$  PeV,  $\theta = 14^\circ$ , black line) with injected muon background (proton primary,  $E_0 = 58$  TeV,  $\theta = 54^\circ$ , orange line).



**Figure 5.3:** The modified Poisson distribution (Eq. (5.3)) for  $\lambda = 1.5$ . The continuous version, which uses the Gamma function is added to guide the eye.

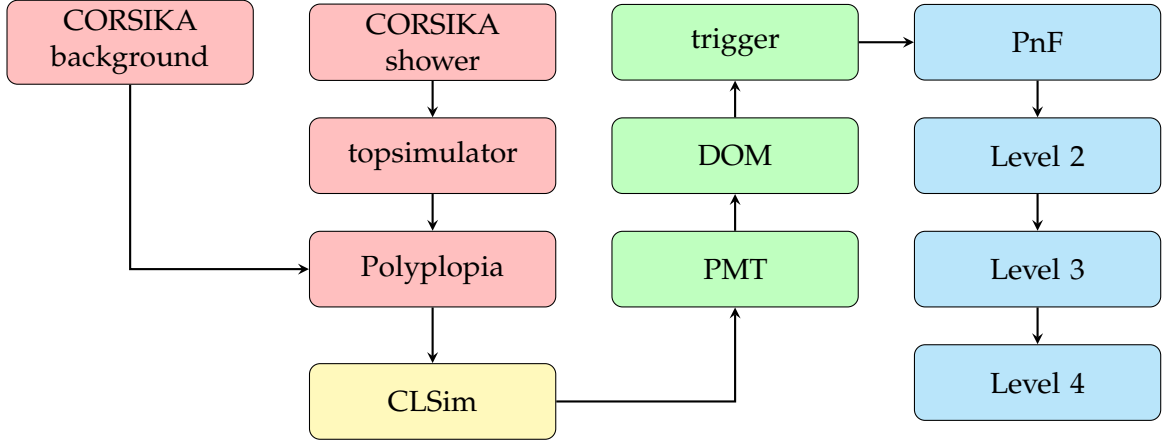
modified in order to account for one slot already being occupied by the primary event:

$$\hat{p}_\lambda(k) = \frac{\lambda^{k+1}}{(e^\lambda - 1)(k+1)!} \quad (5.3)$$

The variable  $k$  now represents the number of additionally injected muon tracks. Since the injection of  $-1$  background tracks (“canceling out” the existing primary event to come to zero events in total) is not possible,  $\hat{p}_\lambda(k)$  was re-normalized via division by  $1 - e^{-\lambda}$ . A time window of  $40\mu\text{s}$  is chosen and the rate of muon tracks inside the in-ice detector is of the order of  $37\text{ kHz}$ , which gives a Poisson mean of about  $1.5$ . The modified Poisson distribution is shown in Figure 5.3.

In practice, a set of CORSIKA background events is generated ( $600\text{ GeV} \leq E \leq 100\text{ EeV}$ , following an  $E^{-2.7}$  spectrum and primaries sampled from the poly-gonato composition model [14]) in parallel with the injection of the main showers into IceTop. In the early stages of detector simulation, all secondary particles (including position, direction, time and energy) and their energy losses in the detector volume are stored in a tree-like structure in the i3 files, the I3MCTree. The IceTray module Polytopia samples effectively from the modified Poisson distribution (5.3) to determine the number of background events to inject. This amount of single muons and muon bundles (it can also be none) is then added to the Monte Carlo primary in the I3MCTree, based on a uniform time distribution within  $\Delta t$ . The simulation of Cherenkov photons from both the main cosmic-ray event and the potentially injected background muons with subsequent detector response follows as described in section 5.2. The complete simulation chain is visualized in Figure 5.4 and an exemplary MC event with injected in-ice background is shown on the right in Figure 5.2.

As part of Level 3 processing (section 6.2), I3MCTree is removed from the i3 files to save disk space. However, information on the injected muon bundles that were visible to IceCube DOMs is kept in the frame. This is necessary so that methods for the suppression of in-ice signal produced by those background muons can be developed ( $\rightarrow$  subsection 6.4.7).



**Figure 5.4:** The detector simulation chain. Generation (red): CORSIKA showers are injected onto IceTop by topsimulator and the high-energy muons are combined with in-ice background by the Polyplopia module. Photon propagation (yellow): CLSim tracks Cherenkov photons to the DOMs. Detector response (green): PMT response, DOM electronics and the trigger system are simulated. Processing and Filtering (blue): filters are applied before Level 2, 3 and 4 ( $\rightarrow$  chapter 6) processing are run.

## 5.4 Environmental Effects

Environmental conditions affect both the development of cosmic-ray air showers and the detection of secondary shower particles. Here, the influence of atmospheric as well as ice properties and the effect of snow covering IceTop tanks will be discussed. Later, in chapter 8, environmental effects on the reconstructed proton and helium flux are studied as part of the systematic uncertainties.

### 5.4.1 Atmosphere

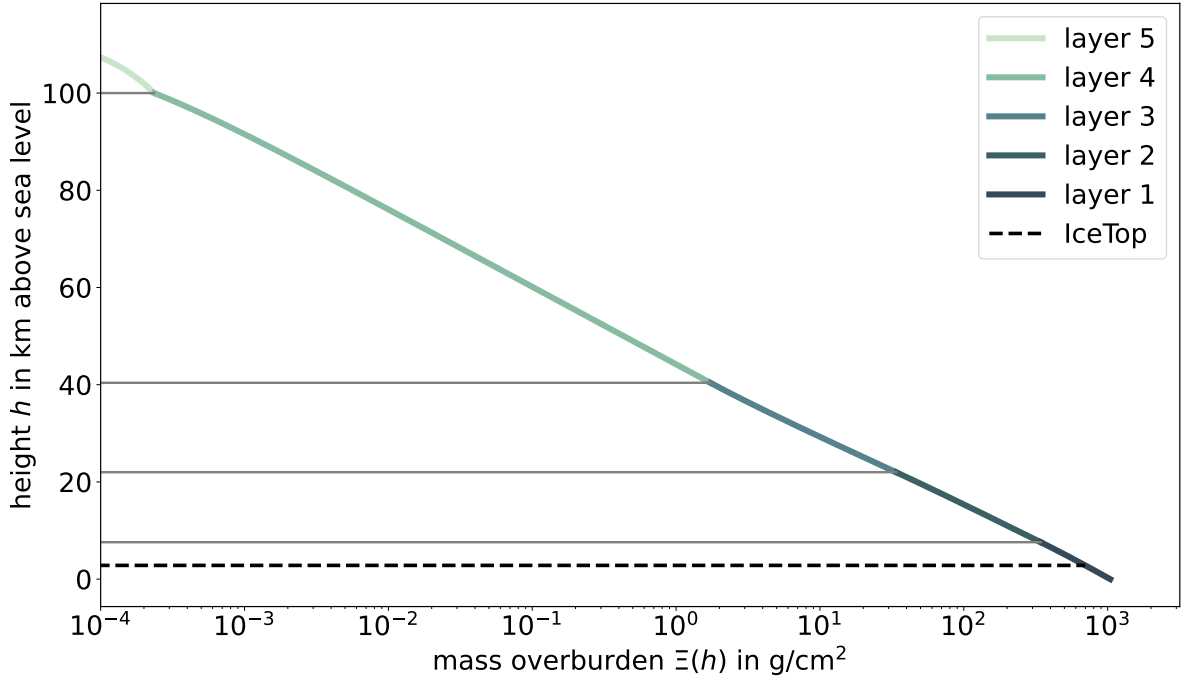
As the development of atmospheric cascades is driven by the interaction of shower particles with nuclei constituting the air molecules, the atmospheric profile plays a significant role in the behavior of EASs. Therefore, it is crucial to provide CORSIKA with an appropriate parameterization of the atmosphere to use in the simulation. The atmosphere is modeled in five continuous layers. The lowest four describe the vertical atmospheric depth (or mass overburden)  $\Xi$  as an exponential function of the height  $h$  above sea level:

$$\Xi(h) = a_i + b_i \cdot e^{-\frac{h}{c_i}}, \quad (5.4)$$

where  $i \in [1, 2, 3, 4]$ . In the top layer, the air pressure decreases linearly until it vanishes at an altitude of 112.8 km:

$$\Xi(h) = a_5 - b_5 \cdot \frac{h}{c_5} \quad (5.5)$$

The parameters  $a_i$ ,  $b_i$  and  $c_i$  are tuned in such a way that  $\Xi(h)$  is differentiable at the transition between layers. In all layers, the volume fractions of  $N_2$ ,  $O_2$  and Ar are set to 78.1%, 21.0% and 0.9%, respectively [72]. CORSIKA provides 41 predefined atmospheric settings for various locations and seasons, among them are parameterizations for all twelve months at the South Pole, which are based on meteorological measurements between 2007 and 2011. For the new MC datasets, the South Pole April atmosphere was chosen because it resembles the yearly



**Figure 5.5:** The atmospheric profile of the South Pole April parameterization. The five colors correspond to the five layers in the model. The altitude of IceTop at 2830 m and the layer boundaries are indicated by horizontal lines.

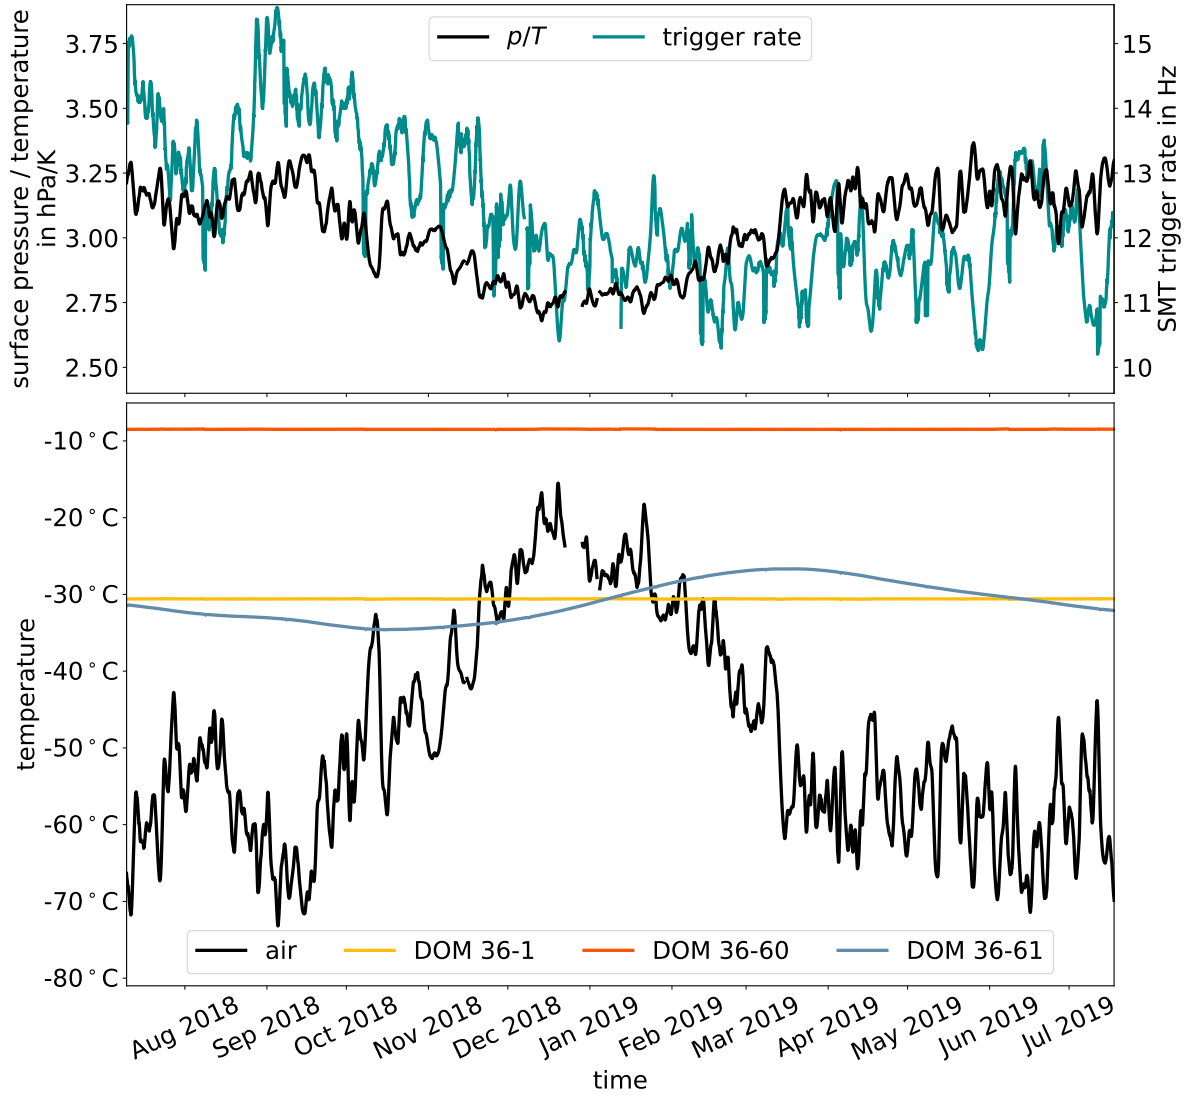
average condition [93]. The corresponding atmospheric parameters for Eqs. (5.4) and (5.5) are specified in Table 53 of Ref. [72]. Figure 5.5 visualizes the exponential nature of atmospheric overburden as a function of height. Evaluating this atmospheric profile at 2830 m above sea level yields an overburden  $\Xi$  of 698.62 g/cm<sup>2</sup>.

Over the course of a year, the atmospheric density varies. Figure 5.6 demonstrates this by plotting the ratio of measured air pressure and temperature at the South Pole for twelve months<sup>5</sup>. The ideal gas law

$$\rho = \frac{p}{R_s T}$$

describes the proportionality between gas density  $\rho$  and the ratio of pressure and temperature  $p/T$  with the specific gas constant  $R_s$ . The air temperature on the South Pole surface ranges between  $-75^\circ\text{C}$  and  $-15^\circ\text{C}$ . During austral summer, the expansion of the relatively warm air leads to a decrease in density. The measured data was taken at the surface, but the entire atmosphere follows this seasonal change with the strength of that effect depending on the altitude. Air density has a large influence on air-shower development, as higher density leads to an increased probability for charged mesons to interact with nuclei in air molecules before they decay. As a result, the count rate of single IceTop DOMs is anticorrelated with atmospheric temperature and varies at a 5% level, whereas the rate of high-energy muons detected in the ice correlates positively with the temperature at a 10% level [94]. The top plot of Figure 5.6 visualizes the correlation between air density and IceTop trigger rate.

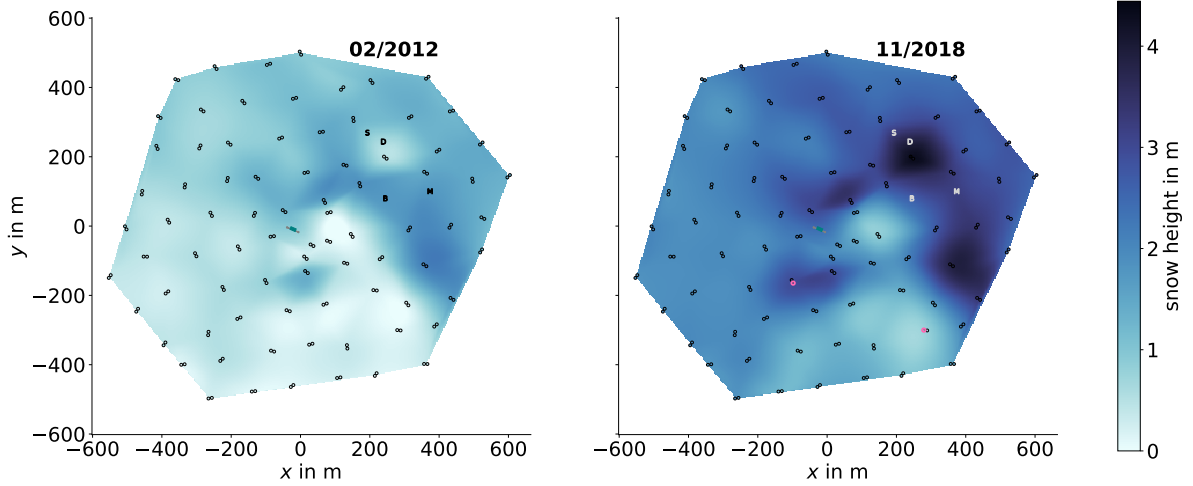
<sup>5</sup>NOAA is the United States' National Oceanic and Atmospheric Administration that provides meteorological data to the IceCube monitoring database.



**Figure 5.6:** *Top:* Ratio of surface air pressure and temperature measured at the South Pole during the IC86.2018 season from July 2018 to July 2019 (black). Under the assumption of an ideal gas, this quantity is proportional to air density. This is contrasted with the IceTopSMT trigger rate (teal). *Bottom:* Temperature measured in DOMs of string/station 36. DOM 60 (dark orange) at a depth of 2450 m has a stable temperature around  $-8.5^{\circ}\text{C}$ , while the DOM at the top of the same string (1) has a temperature of  $-30.6^{\circ}\text{C}$  (yellow). The temperature of the IceTop DOM 61 (blue) is influenced by the seasonal variation of air temperature (black), although delayed by about three months. Rate, pressure and temperature measurements are smoothed by a linear Savitzky–Golay filter with a 24-hour time window. Meteorological data credit: NOAA

Towards the end of the season, the SMT6 rate does not return to the same level as in the previous year due to the accumulation of snow ( $\rightarrow$  subsection 5.4.2).

In contrast to the in-ice DOMs which remain at an almost constant depth-dependent temperature ( $-35^{\circ}\text{C}$  to  $-8^{\circ}\text{C}$ ), the optical modules in IceTop heat up and cool down according to the season (bottom plot in Fig. 5.6). The temporal delay between the air and DOM temperature oscillation as well as the smaller amplitude of the temperature variation of IceTop DOMs demonstrate the insulating property of snow.



**Figure 5.7:** Snow height measurements several years apart. The snow heights between stations has been interpolated with cubic polynomials. The ICL in the center and four more buildings (MAPO, DSL with the SPT and the supplementary building 61) are indicated by their respective initial letter. **Left:** February 2012. **Right:** November 2018. The two tanks with the most extreme natural snow accumulation ( $\rightarrow$  Fig. 5.8) are marked in pink.

#### 5.4.2 Snow

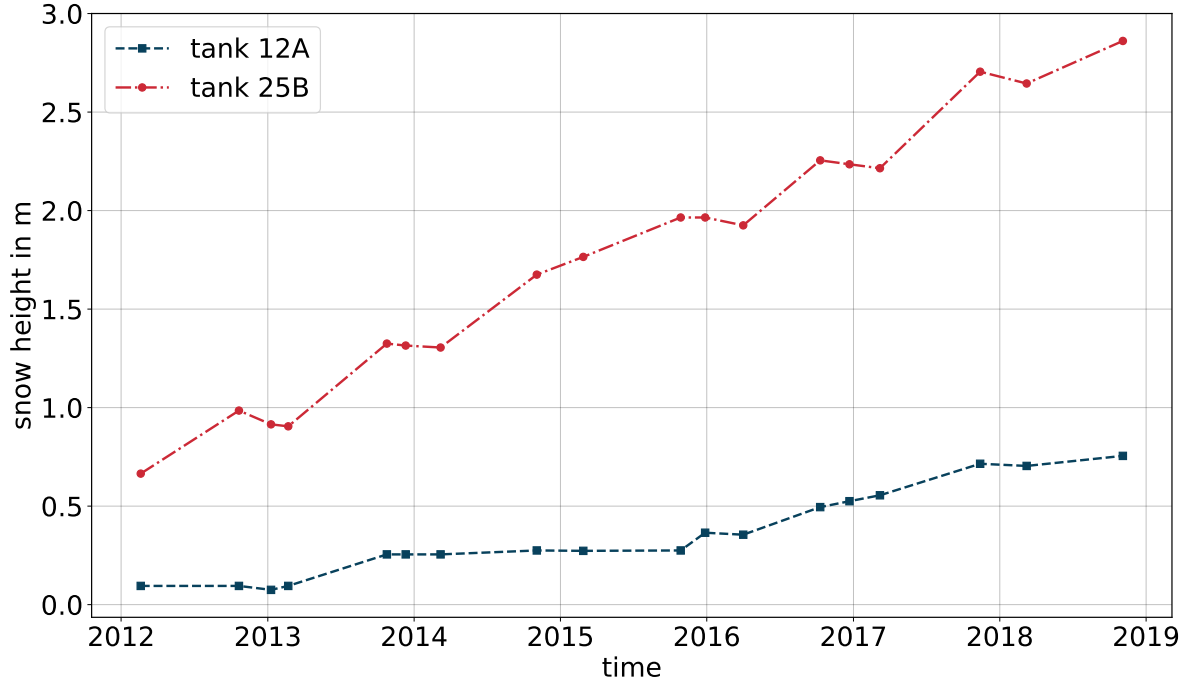
The IceTop tanks were deployed in the snow with their lids on top being level with the surface in order to minimize snow accumulation and temperature fluctuations inside. However, over the years, the strong winds at the South Pole lead to a pile-up of snow above the tanks, whereas the contribution due to precipitation is negligible. The amount of snow drift varies across the array due to a slight slope of the terrain and buildings on site blocking or diverting the wind. The snow height above each IceTop tank is measured at least twice every year. Figure 5.7 gives an impression of the accumulation of snow over time, showing the measured and interpolated snow height for February 2012 and November 2018. On average, the snow height increases by about 21 cm per year; however, it spreads between less than 10 cm and more than 33 cm annually<sup>6</sup> ( $\rightarrow$  Fig. 5.8). For the MC datasets used in this work, the snow heights measured in November 2018 have been used in the IceTop detector simulation.

Snow has a strong attenuation power for the electromagnetic particles in air showers, whereas muons are hardly affected by the snow. For reconstructions, the charge expectation values of tank pulses can be adjusted for this snow effect by estimating the attenuated signal amplitude

$$S_i^{\text{corr}} = S_i^{\text{no snow}} \cdot \exp\left(\frac{-d_i}{\lambda_{\text{snow}} \cos(\theta)}\right),$$

where the subscript  $i$  indicates the tank,  $S_i^{\text{no snow}}$  is the expected charge value without any snow on the tank,  $S_i^{\text{corr}}$  denotes the corrected quantity,  $d_i$  is the vertical snow depth and  $\theta$  is the estimated zenith angle. The effective attenuation length  $\lambda_{\text{snow}}$  for the IC86.2018 season is approximately 2.35 m.

<sup>6</sup>Actually, the largest amount of accumulated snow between 2012 and 2018 has been measured above the tanks of station 57 with a total increase of over 4 m. However, due to its proximity to the DSL, snow levels are changed artificially every year as snow is plowed around the building.



**Figure 5.8:** The minimal snow accumulation between 2012 and 2018 was seen above tank 12A with only 66 cm in total. During that time the largest amount of snow (almost 2.2 m) piled up over tank 25B. The rate of snow accumulation is far from constant and even can be negative temporarily.

However, assuming a constant attenuation length for all hits in all events ignores the much weaker suppression of the muon signal compared to the electromagnetic component. Farther away from the shower axis, the recorded signal becomes dominated by muons, leading to a decreasing effective  $\lambda_{\text{snow}}$  at greater distances. A *radius dependent* (RADE) parameterization that also varies with the shower size  $S_{125}$  has recently been developed [95]. For this purpose, charge distributions have been fitted with a function that combines the electromagnetic and muonic signal components as well as takes into account the number of tanks that did not get triggered.

#### 5.4.3 Ice Model

A good understanding of the optical ice properties is crucial for an adequate simulation of the detector response and for the development of reconstruction methods. During the operation of AMANDA it has been observed that the absorption and scattering coefficient not only depends on the wavelength of the light but also changes with depth. This is most noticeable in the dust layer at depths 1975–2100 m where the optical properties are significantly worse than in the surrounding *bulk ice* [39]. For the detector deployment, vertical holes with a diameter of about 60 cm were drilled into the ice using hot water which refroze after the string has been installed. This *hole ice* needs to be described distinctly due to the different distribution of bubbles and dust caused by the drilling.

The models of the South Pole ice have been refined continuously, leading to an ever better description of the ice properties. The most recent improvements include an updated description of the light propagation anisotropy based on birefringence [96] and a parameterization

of the tilt of individual ice layers, which results from the glacial ice flowing slowly over the underlying bedrock [92].

### **Key Messages**

*For this analysis, several new Monte Carlo datasets have been produced with CORSIKA, followed by simulating the detector response. The hadronic interaction model Sibyll 2.3d is used in the main dataset, whereas smaller sets have been created with QGSJet-II.04 and EPOS-LHC. For the first time, simulations of cosmic-ray air showers at IceCube include an injected background of high-energy muons.*



# CHAPTER 6

## OFFLINE PROCESSING AND RECONSTRUCTION

*“Excuse me, but Real Programmers use butterflies. They open their hands and let the delicate wings flap once. The disturbances ripple outward, changing the flow of the eddy currents in the upper atmosphere. These cause momentary pockets of higher-pressure air to form, which act as lenses that deflect incoming cosmic rays, focusing them to strike the drive platter and flip the desired bit.”*

— Randall Munroe, *xkcd* #378: ‘Real Programmers’

The data files created by the online [PnF](#) system (→ [section 4.2](#)) are called *Level 1* files. The subsequent offline data processing (*Level 2*, → [section 6.1](#)) uses the computing resources in the north. It performs pulse cleaning and event reconstructions that are of general interest for most analyses in IceCube. It is followed by a dedicated *Level 3* processing (→ [section 6.2](#)), which is individual for each working group within the IceCube collaboration. For cosmic-ray analyses, this includes the charge calibration of [SLCs](#), which is elaborated on in [section 6.3](#). Finally, an even more specialized *Level 4* can be applied. The processing that has been developed and applied for this work will be covered in [section 6.4](#). The scripts used for this offline processing are part of IceCube’s software framework IceTray [97]. In the end, low-level pulse information in simulations and measured data are compared in [section 6.5](#).

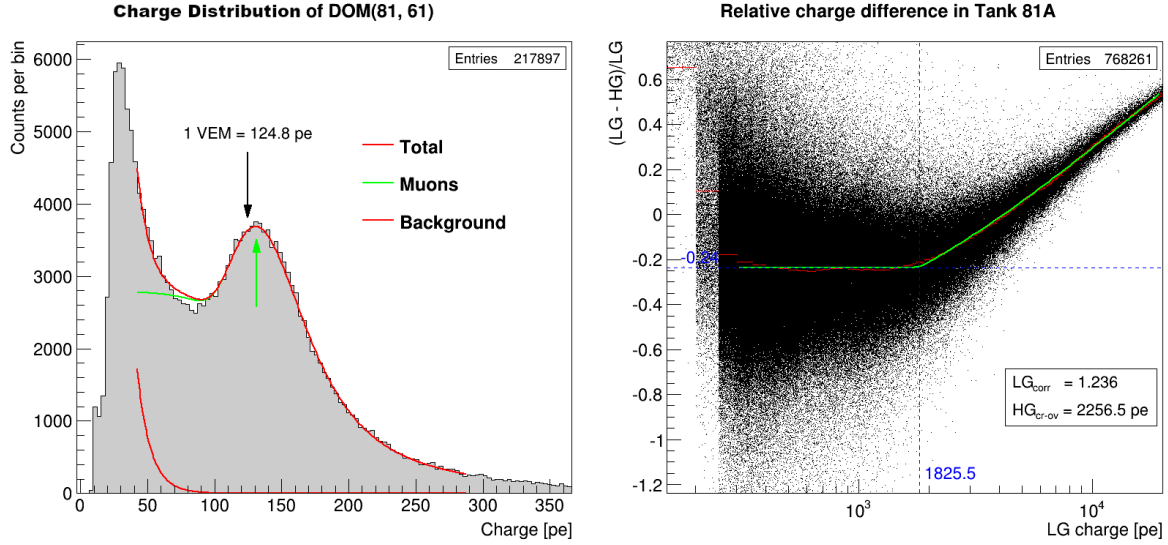
### 6.1 Level 2

Events that pass at least one of the filters in [Table 4.2](#) are stored inside the *i3* files in IceTopSplit P-frames (where ‘P’ stands for Physics) that inherit all information from their original Q-frames. The union of a Q-frame with its associated one or more P-frames is referred to as a frame packet (→ [section 6.4](#)).

The [PMT](#) response to the same light emission differs from [DOM](#) to [DOM](#) and also varies over time. This is especially true for IceTop as over the six years of deployment different hardware was used. For example, tanks installed in the 2005 season have an inner coating with higher reflectivity and feature different pulse characteristics compared to those tanks, which were added later [69]. Also, environmental conditions affect the performance of DOMs in the surface array. In order to account for this, for each [HG](#) DOM every 8192<sup>th</sup> pulse (“minimum bias”) is collected and the resulting charge distributions are fitted biweekly ([Fig. 6.1](#), left). The fit function is the sum of one term describing the contribution of muons and a simple exponential term describing electromagnetic particles:

$$f(x) = p_0 \cdot \left[ \frac{1.85}{p_1} \cdot \frac{1}{\exp\left(\frac{x-p_1}{p_2}\right) + 1} + \text{Landau}(x, p_1, p_2) \right] + p_3 \cdot \exp(p_4 \cdot x)$$

Here,  $\text{Landau}(x, p_1, p_2)$  describes a Landau distribution with a peak at  $p_1$  and a width of  $p_2$ . The peak corresponds to vertical muons that pass through the tank. Since the hits used for the



**Figure 6.1:** *Left:* Charge distribution for a HG DOM in units of photoelectrons after two weeks of data collection. For the purpose of VEM calibration the muon peak is fitted together with a background hypothesis. *Right:* For cross-calibration the relative charge difference between the HG and LG DOM as a function of the LG DOM charge is fitted with a constant offset for low charges and a logarithmic function for charges above the crossover point (Eq. (8) in Ref. [69]). The data was recorded in tank 81A over two weeks in October 2018. Credit: S. Tilav

charge calibration are predominantly caused by low-energy air showers, the detected muons can be considered **MIPs**. Therefore, the amount of detectable light is mainly determined by the distance a muon travels through the tank. **HLC** and **SLC** charges are divided by 95% of this peak position in units of **PE**, so they are transformed into units of vertical equivalent muons (**VEM**). This process is called *VEM calibration*. In the **i3** files, the resulting pulses are accordingly named **VEM** pulses.

While **VEM** calibration can only be performed on the **HG** DOMs as the **LG** threshold lies beyond 1 **VEM**, a cross-calibration is applied to the **LG** DOMs. For this, the relative charge difference  $(Q_{LG} - Q_{HG})/Q_{LG}$  as a function of the low-gain charge  $Q_{LG}$  for all pulses that have been seen by both DOMs in a tank is fitted with two lines, a constant and a logarithmic function. For small **LG** charges, both DOMs behave similarly, they record on average the same charge, except for an offset, which is subtracted as part of the calibration. With increasing charge, the relative deviation increases as the **HG** PMT starts to saturate and hence cannot provide a reliable charge any more. The intersection of both lines is the **HG/LG** crossover point (Fig. 6.1, right).

Next, **HLC** pulses from both DOMs in a tank are merged into so-called tank pulses. Here, three cases are to be distinguished:

- when the charge of the **HG** DOM  $Q_{HG}$  is below the **HG/LG** crossover point,  $Q_{HG}$  is used
- when  $Q_{HG}$  is above the **HG/LG** crossover point; i.e., the **HG** DOM is saturated, and the **LG** DOM in the same tank has a pulse within  $\pm 40$  ns,  $Q_{LG}$  is used

- when the HG DOM is saturated and no matching LG pulse occurred within the 40 ns time window, the charge is set to NaN (“not a number”).

In any case, the tank pulse gets the time stamp from the HG DOM in the tank. SLC pulses are calibrated and merged into tank pulses as part of the Level 3 processing.

A “cluster cleaning” is then applied to the HLC tank pulses to check if more than one cosmic-ray event was recorded inside the DAQ readout window. A maximum time difference between stations is defined as  $\Delta t_{\max} = 200 \text{ ns} + d_{\text{stations}}/c$  with  $d_{\text{stations}}$  the separation of a pair of stations. Stations that have a bigger  $\Delta t$  between each other are considered to belong to separate events and are stored into additional physics frames. Often, P-frames with pulse times later than the main air shower contain *afterpulses* from PMTs that were hit just before. As those hold no information for the analysis of EASs, they are generally discarded, as are noise hits. The cleaned and split pulses get the label CleanedHLCTankPulses.

## 6.2 Level 3

For cosmic-ray analyses, a further processing stage is applied. This includes charge calibration, cleaning of IceTop and in-ice pulses as well as sophisticated air-shower reconstructions [98].

While IceTop HLCs obtained reliable charge values from the deconvolution of the waveforms with the single-photon response during PnF, the charge stamps of SLCs, which are simply the integral of the waveform, are generally less accurate. Level 3 processing starts with calibrating the SLCs of some events. Details on the calibration of SLCs and how it has been extended to all events will be described in section 6.3. The calibrated SLC VEM pulses are then merged into tank pulses, the same way it was done for HLCs in Level 2.

Next is a more rigorous cleaning of the CleanedHLCTankPulses. The so-called *seeded RT* (SRT) algorithm looks for clusters of five HLCs around each HLC DOM within a radius (R) 180 m and a time window (T) of 450 ns. This removes single HLCs, which are usually created from noise or coincident small air showers and survived the “cluster cleaning” step in Level 2 by chance. SRT cleaning is only performed on events with at least five HLCs. The output IceTopHLCSeedRTPulses are thus not available for 2-station events that were triggered by four HLCs in the InFill. Consequently, the following processing steps which all rely on IceTopHLCSeedRTPulses as an input are only performed for events that pass the 3-station or 3-station InFill filter. Figure 6.2 demonstrates how cluster cleaning and SRT cleaning remove distant HLCs in the offline HLC tank pulses and the CleanedHLCTankPulses.

A first estimate for the shower core position is the *center of gravity* (COG). It is determined as

$$\vec{r}_{\text{COG}} = \frac{\sum_i \sqrt{Q_i} \cdot \vec{r}_i}{\sum_i \sqrt{Q_i}}, \quad (6.1)$$

where  $\vec{r}_i$  are tank positions in IceCube detector coordinates and  $i$  runs over the up to seven tanks with the highest charges  $Q_i$ . This ShowerCOG together with a ShowerPlane (the flat shower front, which minimizes the difference between expected pulse times and measured times) is used as a seed for an air-shower reconstruction algorithm, called Laputop.



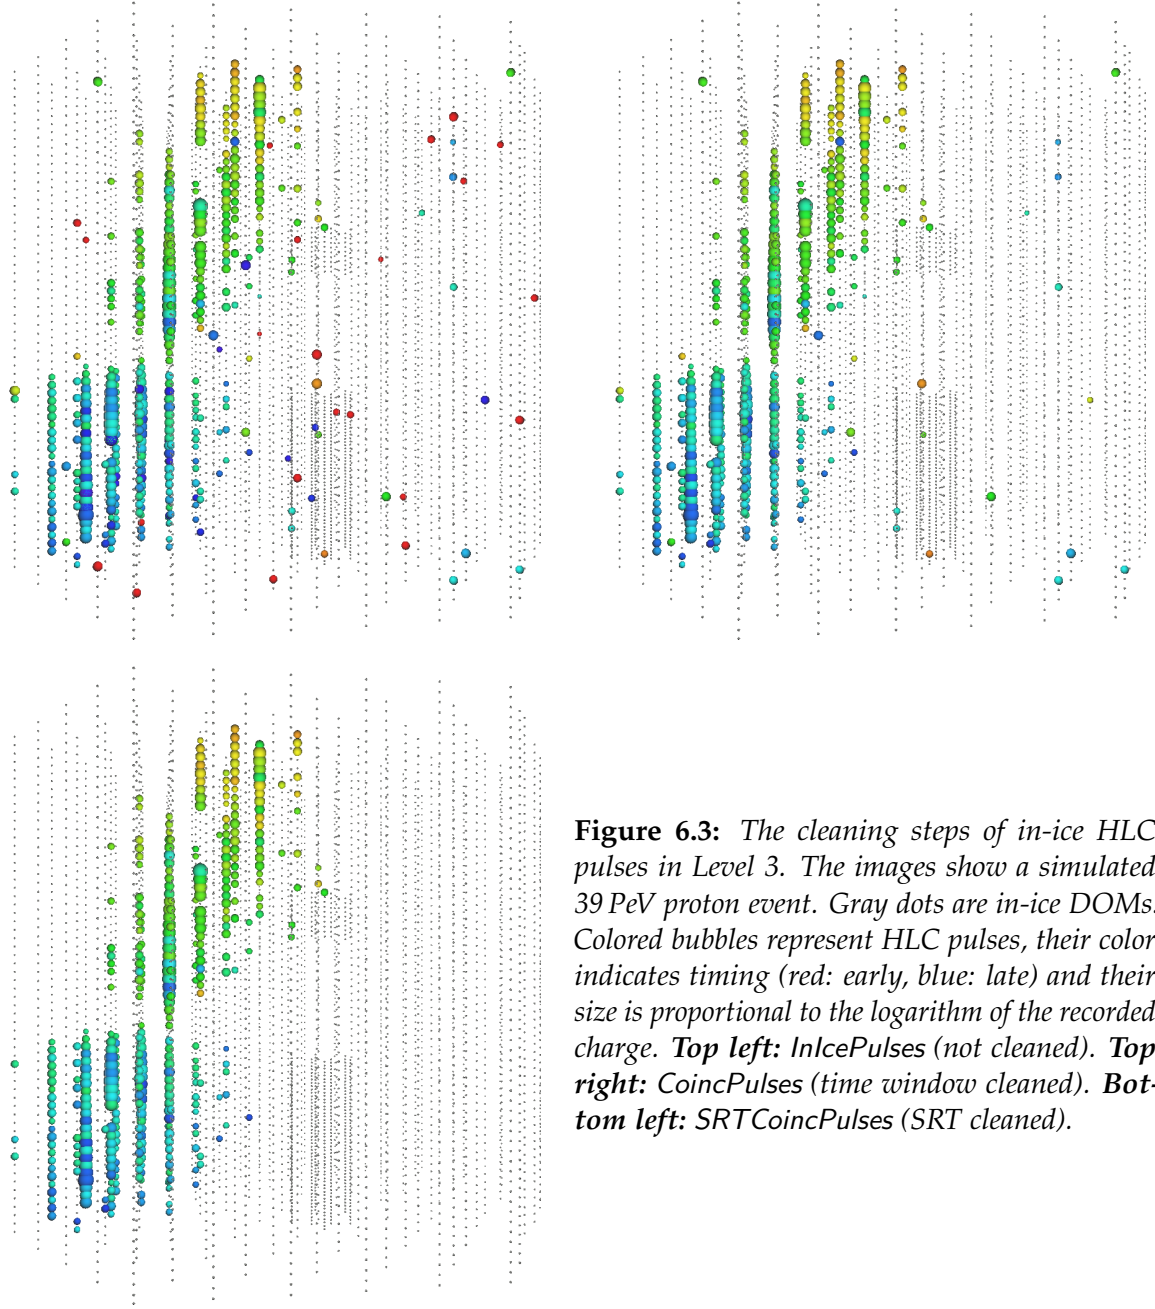
**Figure 6.2:** The cleaning steps of IceTop HLC pulses in Level 2 and 3. The images show a simulated IceTop event of a 37 PeV proton-induced air shower. Gray dots are IceTop tanks. Colored bubbles represent HLC pulses, their color indicates timing (red: early, blue: late) and their size is proportional to the logarithm of the recorded charge. **Top left:** *OfflineIceTopHLCTankPulses* (not cleaned). **Top right:** *CleanedHLC-TankPulses* (cluster cleaned). **Bottom left:** *IceTopHLCSeedRTPulses* (SRT cleaned).

It maximizes a multi-component likelihood to find a good reconstruction of the cosmic-ray event. The charge distribution is parameterized as a *lateral distribution function* (LDF) which describes the signal amplitude  $S$ , that is, the recorded charge in VEM, as a function of the lateral distance  $r$  from the shower axis:

$$S(r) = S_{\text{ref}} \cdot \left( \frac{r}{r_{\text{ref}}} \right)^{-\beta - \kappa \cdot \log_{10} \left( \frac{r}{r_{\text{ref}}} \right)}, \quad (6.2)$$

where the reference distance  $r_{\text{ref}}$  is set to 125 m and the curvature parameter  $\kappa$  is fixed at a value of 0.303. The shower size  $S_{125}$  and the slope  $\beta$  are free parameters of the fit.  $S_{125}$  can be used as a primary energy proxy, as it exhibits a strong correlation with  $E_0$  while not having a significant dependence on the type of primary cosmic-ray particle [70]. The LDF slope  $\beta$  on the other hand is related to the shower age and is therefore a potential composition-sensitive parameter. Apart from the charge likelihood, which is based on this LDF, other contributions to the total likelihood describe pulse timing, stations that were not hit and tanks with saturated DOMs.

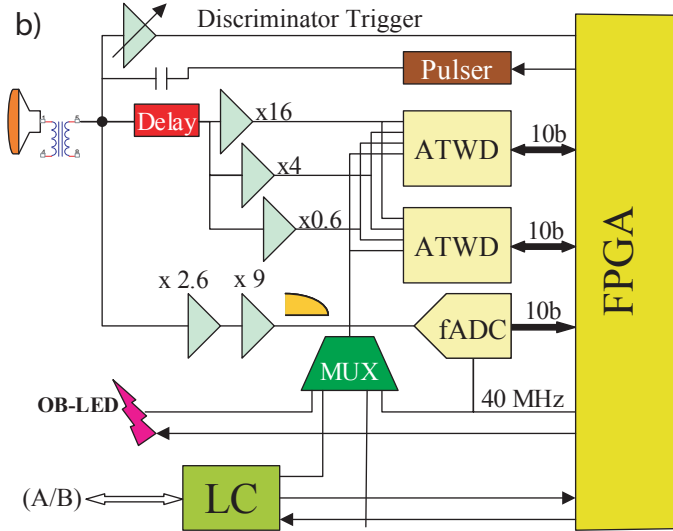
After the treatment of IceTop pulses finished, the in-ice HLC pulses undergo a similar cleaning procedure. Events that have an InIceSMT trigger as well as pass any of the IceTop filters in Table 4.2 receive a *time window cleaning* (TWC) to remove random noise hits and



**Figure 6.3:** The cleaning steps of in-ice HLC pulses in Level 3. The images show a simulated 39 PeV proton event. Gray dots are in-ice DOMs. Colored bubbles represent HLC pulses, their color indicates timing (red: early, blue: late) and their size is proportional to the logarithm of the recorded charge. **Top left:** *InIcePulses* (not cleaned). **Top right:** *CoincPulses* (time window cleaned). **Bottom left:** *SRTCoincPulses* (SRT cleaned).

coincident background at the edge of the trigger window. Only those hits are kept and saved as *CoincPulses* when they lie inside an *InIceSMT* trigger window that needs to start 2–7  $\mu\text{s}$  after the end of the *IceTopSMT* trigger. A padding of  $-300\text{ ns}$  and  $+400\text{ ns}$  is added to the *InIceSMT* trigger window, which, however, must not exceed a length of 6.5  $\mu\text{s}$ .

TWC is followed by SRT cleaning, although the settings for in-ice pulses change with respect to the *IceTop* variables due to the different geometry of the deep *IceCube* array. A cluster needs to have three HLCs within a sphere with 150 m radius and a time window of 1  $\mu\text{s}$  around each HLC in order to survive this step. The cleaned output pulses are labeled *SRTCoincPulses*. Figure 6.3 demonstrates how time window cleaning and SRT cleaning remove noise hits in the *InIcePulses*.



**Figure 6.4:** The IceTop DOM mainboard electronics. A central FPGA integrated circuit receives signals from the discriminator, one of two ATWD chips and the fADC (Fig. 10 in Ref. [69]).

In all the described cleaning steps, the input pulse series does not get deleted, rather the new cleaned version is added to the i3 file. The Level 3 processing for cosmic-ray events involves many more steps that either are not relevant for this work or will be covered by the description of the new Level 4 processing in [section 6.4](#), like the cleaning of IceTop pulses around a reconstructed track and the definition of event containment.

At the end of Level 3 processing, one newly added module is reserved to run only for simulated events. As coincident in-ice background is added as part of the detector simulation now, it will prove useful to have information on the primary particles of those muon bundles available inside a dedicated frame object in the Q-frame.

### 6.3 SLC Calibration

A proper charge calibration of SLC pulses has hitherto been reserved for large IceTop events. While all HLCs undergo a final charge calibration in Level 2, only a subset of SLC pulse series — namely those events that pass the 5-station filter, a random 10% sample of events with a satisfied 3-station (InFill) filter and 1% of coincident events — got a correction of their charge values in Level 3. Low-energy events therefore generally lacked calibrated SLC charges. This section describes the method of SLC calibration and how it has been extended to become applicable to all IceTop events in the upgraded Level 3 processing. For this, it is necessary to cover the signal processing inside the DOM in more detail.

#### 6.3.1 Waveform Digitization

The DOM mainboard electronics initially split the PMT signal into three paths (→ Fig. 6.4):

- A discriminator checks whether the peak pulse charge exceeds the predefined threshold for that DOM. If this is the case, the discriminator trigger is sent to the central *field programmable gate array* (FPGA).
- In parallel, the signal is fed into one of two *analog transient waveform digitizer* (ATWD) chips after being delayed by 75 ns. This gives the integrated FPGA sufficient time

to receive the discriminator trigger so that the ATWD chips only accept input for a passed threshold. The two ATWD chips are used alternately, which reduces the DOM's dead time. Each chip comprises four channels, out of which only three are used (ATWD0, ATWD1 and ATWD2). Those three ATWD channels sample the waveform simultaneously over a time window spanning 427 ns using different gain factors (16, 2 and 0.25). Subsequently, 10-bit Wilkinson *analog-to-digital converters* (ADCs) digitize the analog signal from the ATWD channels, starting with the highest gain (ATWD0). When any of the 128 ATWD bins exceeds a threshold of 768 counts (75% of  $2^{10}$ ) the channel is considered “saturated” and ATWD1 is digitized. Similarly, if ATWD1 reaches saturation as well, the lowest-gain channel ATWD2 is digitized. Eventually, the highest-gain unsaturated channel is passed back to the FPGA.

- In the meantime, a *fast ADC* (fADC) digitizes the amplified PMT output continuously without any discriminator trigger condition. With a lower sampling rate of 40 MHz, the fADC can also process longer signals. For IceTop analyses the fADC output is not relevant.

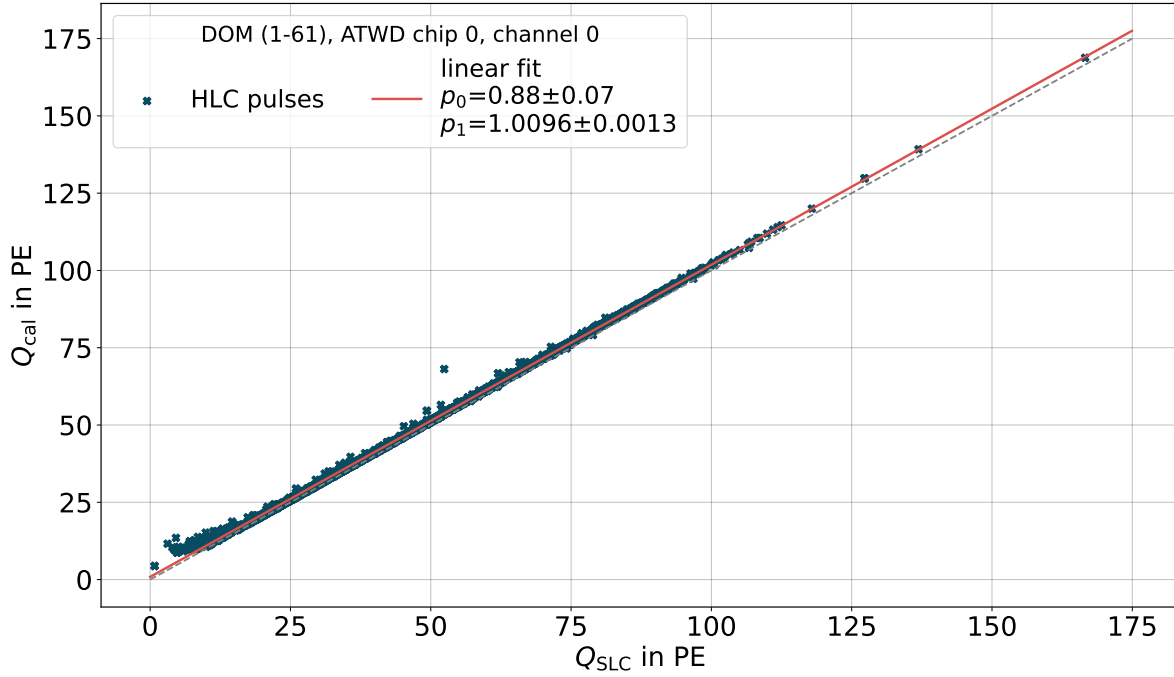
For the conversion from voltage (in units of mV) to charge (in PE), universal calibration constants are used. For more details, see Ref. [69]. The quality of the pulse data to be stored depends on the local coincidence status. HLCs receive their digitized waveforms from the ATWD chip, allowing for an accurate photoelectron calibration. SLCs on the other hand, only get to keep these waveforms under the special conditions listed above. Generally, only the integrated ATWD charge, that is, a charge stamp as well as a time stamp and the pulse width are saved. Being exposed to the surrounding atmosphere, IceTop DOMs are affected by seasonal temperature variations (→ subsection 5.4.1) to some extent, which affects the on-board electronics and thus can lead to deviations of the charge stamp value from the true number of detected photoelectrons.

### 6.3.2 Charge Calibration Using HLCs

Since the waveform-calibrated charges of HLCs can be considered reliable, they are the starting point for the calibration of SLCs. Sending HLC pulses through SLC processing (by flipping their LC bit from 1 to 0) returns the kind of charge stamps that are available for SLCs. There is a linear relation between the correct waveform-calibrated HLC charges  $Q_{\text{cal}}$  and the pretend-to-be-SLC charge stamps  $Q_{\text{SLC}}$ :

$$Q_{\text{cal}} = p_0 + p_1 \cdot Q_{\text{SLC}} \quad (6.3)$$

The values of the parameters  $p_0$  and  $p_1$  are not universal, but vary from DOM to DOM, from ATWD chip to ATWD chip, from channel to channel and over time. Therefore,  $(Q_{\text{cal}}, Q_{\text{SLC}})$  pairs of all HLCs are aggregated in two-week periods to perform the fit (6.3). While the slope  $p_1$  is very close to 1, the intercept  $p_0$  can deviate substantially from 0 PE due to the ATWD baselines varying over time. Figure 6.5 demonstrates the linear relation between  $Q_{\text{SLC}}$  and  $Q_{\text{cal}}$  while Figure 6.6 shows the seasonal variation of the calibration parameters  $p_0$  and  $p_1$ .



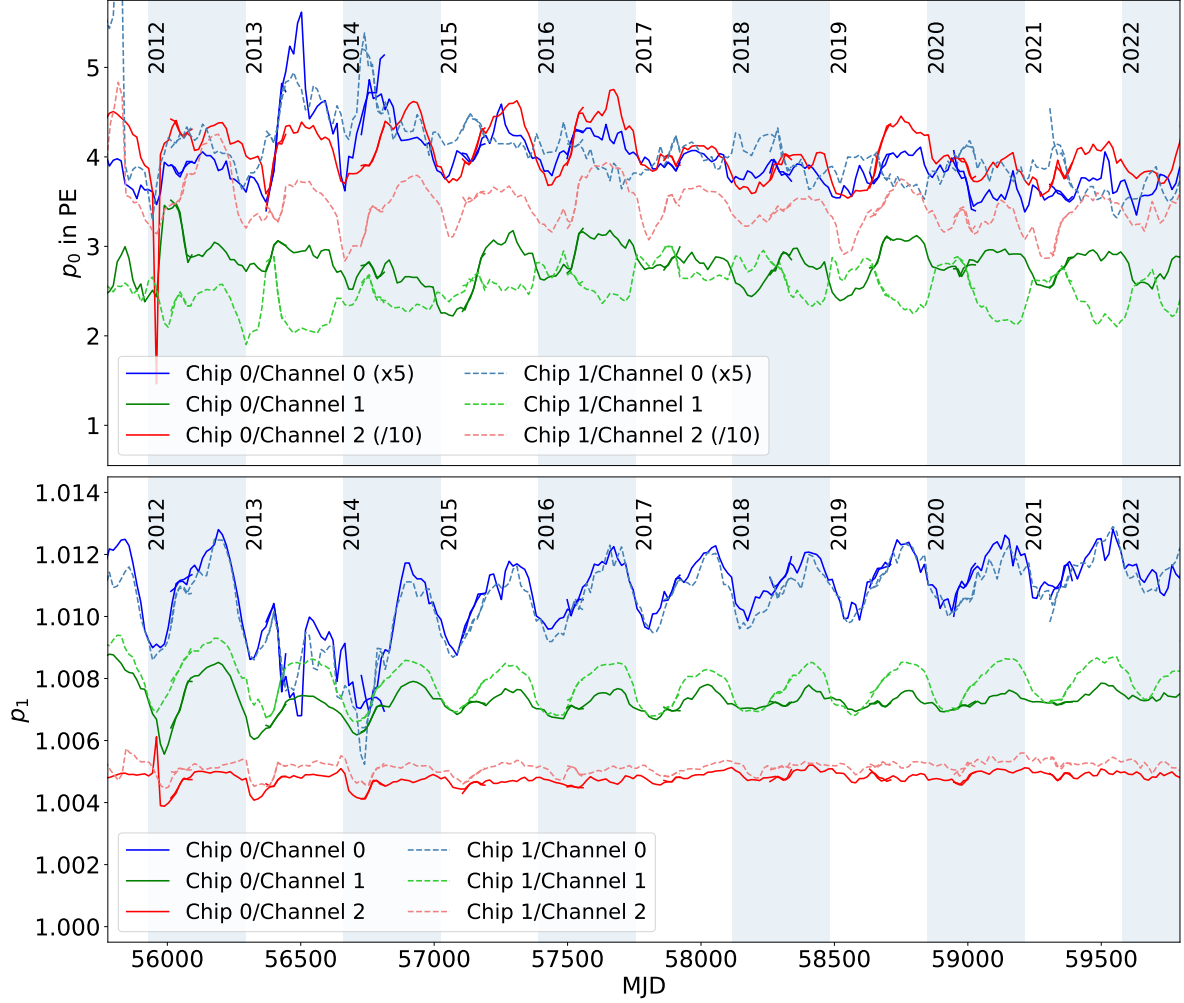
**Figure 6.5:** Linear relation between  $Q_{\text{SLC}}$  and  $Q_{\text{cal}}$  values for HLC pulses, recorded over 8 h in June 2023 by DOM (1-61) and digitized on channel 0 of ATWD chip 0. The linear fit (Eq. (6.3)) is indicated by the red line, which deviates from the “perfect” identity ( $p_0 = 0$  PE and  $p_1 = 1$ , gray dashed line).

When the chip and channel of the ATWD are known for a given SLC pulse, Eq. (6.3) can be applied with the respective set of  $p_0$  and  $p_1$  values in order to acquire calibrated SLC charges. When the chip and channel information, however, are not available — which is especially the case for small air-shower events — the ATWD channel needs to be guessed, based on the uncalibrated SLC charge. Reconstructing which chip the SLC pulse was recorded on is not possible but also not of major importance as both chips on each DOM generally behave quite similarly.

### 6.3.3 ATWD Crossover Points

As described earlier, the pulse charge acquisition uses the highest-gain ATWD channel that is not considered saturated. Hence, a strong correlation between the charge stamp value and the ATWD channel can be observed. Low-charge pulses are captured by channel 0, medium charges by channel 1 and only very high pulse charges fall into ATWD2 (→ Fig. 6.7).

Even though the distributions overlap, charge boundaries can be determined within which a specific channel dominates the waveform digitization. These *crossover points* (COPs) are defined as the charge value at which the *kernel density estimation* (KDE) lines of adjacent ATWD channels intersect. Thus, there are two COPs, one for the transition from ATWD0 to ATWD1 and another one between ATWD1 and ATWD2. Provided with a catalog of COP values for all IceTop DOMs, every SLC pulse can now be calibrated with an accuracy comparable to HLCs. An alternative description of the updated SLC calibration with more mathematical details can be found in chapter 5 of Ref. [99].



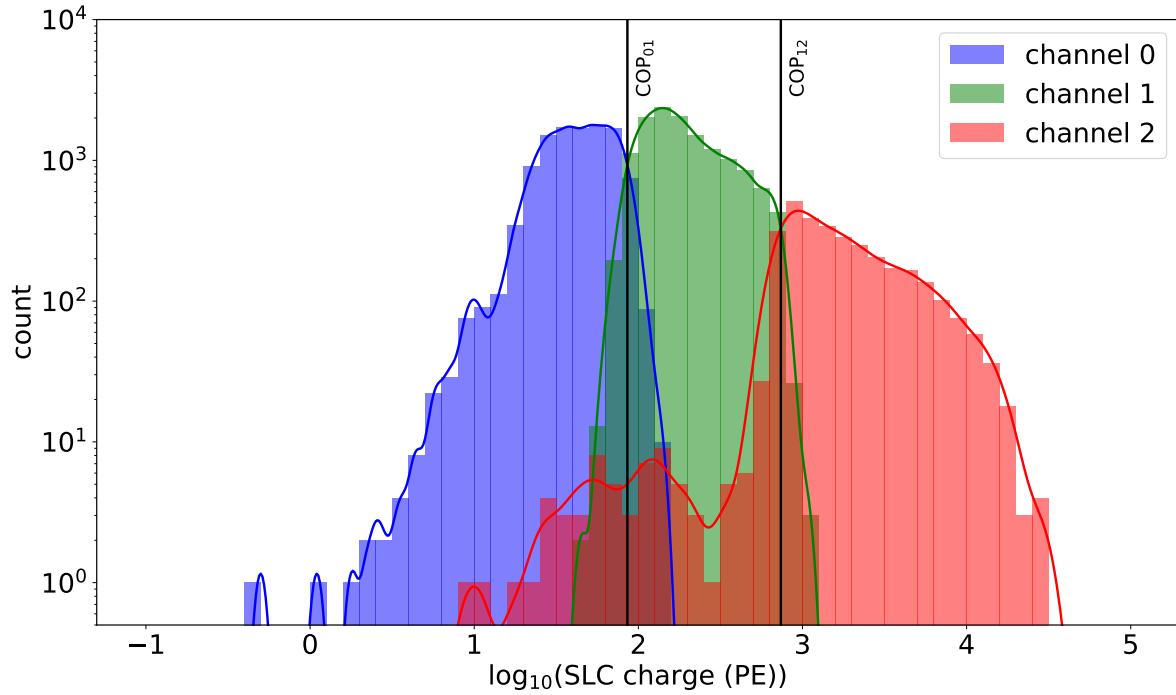
**Figure 6.6:** Seasonal variation of  $p_0$  and  $p_1$  of DOM (1-61) over eleven years of IceTop data-taking. “MJD” refers to the Modified Julian Date format. Courtesy of K. Rawlins.

Similarly to the  $p_0$  and  $p_1$  values, also the COPs move over time. Instead of aggregating SLCs and fitting KDEs to their charge distribution in order to determine the COPs for each data-taking run, the crossover points can also be estimated via a proxy value  $\Theta$  using already available information on the mainboard electronics from the GCD file:

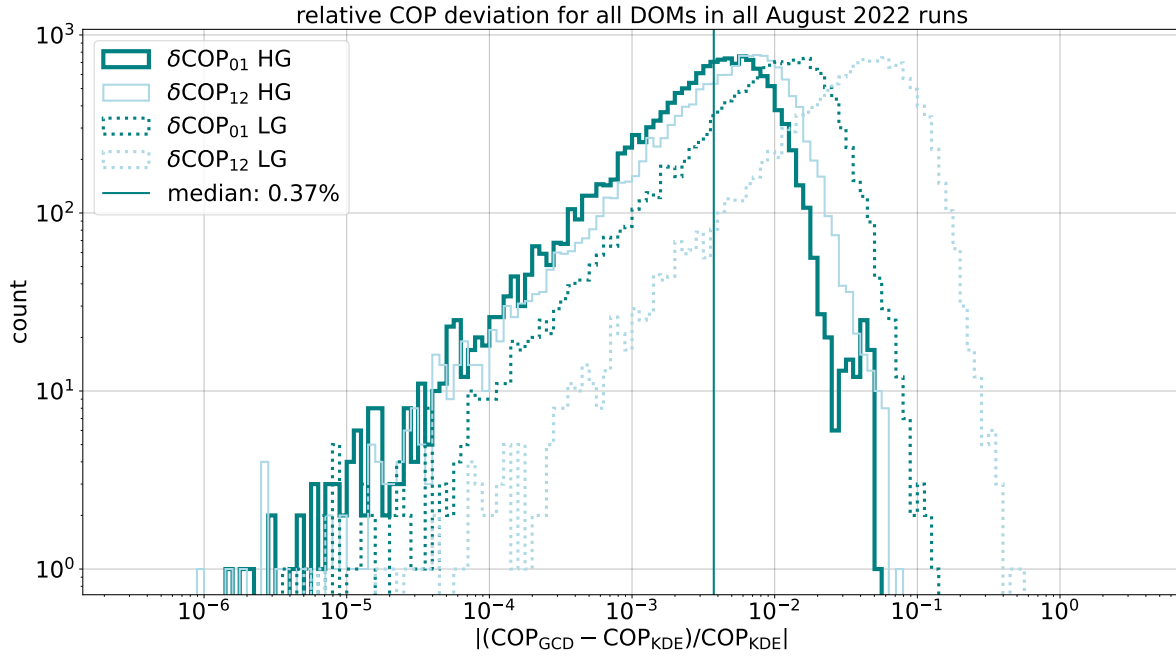
$$\Theta_{01} = \frac{(768 - b_0) \cdot A_0}{G_0 \cdot Z_{FE} \cdot f \cdot G_{PMT} \cdot e}$$

with  $b_0$  the baseline counts of ATWD0,  $A_0$  and  $G_0$  the voltage amplification factor and the gain of channel 0,  $Z_{FE}$  the DOM’s front-end impedance,  $f$  the sampling frequency of the ATWD chips and  $G_{PMT}$  the PMT gain. The direct proportionality between  $\Theta_{01}$  and  $COP_{01}$  eventually is cataloged. In analogy,  $COP_{12}$  can be calculated with the respective quantities for ATWD1.

Figure 6.8 demonstrates the reliability of that approximation method. Deducing the saturation point based on GCD file information works with a median accuracy of 0.37%.



**Figure 6.7:** Histograms of uncalibrated SLC charge stamps from DOM (1-61), both chips, broken down by ATWD channel. KDEs were fitted to all three distributions and the two KDE crossing points are indicated by vertical black lines.



**Figure 6.8:** Accuracy of the estimation of COPs using the GCD proxy method. The distribution of the relative difference of COPs estimated via GCD data from COPs determined via KDE fits is plotted for the four combinations of PMT gain and COP type.

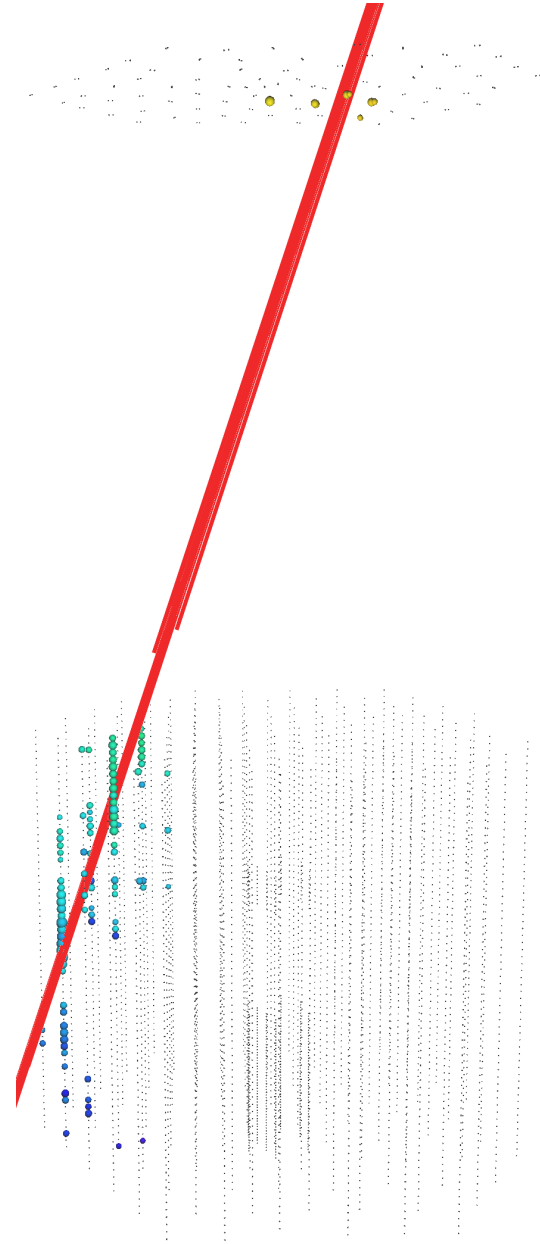
## 6.4 Level 4

IceCube is an exceptional cosmic-ray detector as it combines a square-kilometer Cherenkov tank array with a cubic-kilometer detector sensible to high-energetic muons. This hybrid design provides valuable information for air showers that are detected both at the surface by IceTop tanks and coincidentally deep in the ice below (Fig. 6.9). The coincidence requirement is essential for composition studies, as it facilitates disentangling the energy degeneracy of mass-sensitive observables. A quantity like muon multiplicity can only help reveal the primary mass in combination with the respective primary energy. The former can be estimated using the in-ice signal caused by the high-energy muon bundle, while the latter is reconstructed with information taken from IceTop. The concept of utilizing surface/in-ice coincidences has been pioneered by the precursor of IceCube. The mean logarithmic mass  $\langle \ln(A) \rangle$  of the cosmic radiation has been reconstructed using events that are seen electromagnetically by the scintillator array *SPASE* and muonically by the *AMANDA* in-ice detector [100–102]. With the commissioning of IceCube, those studies kept going on and saw further improvements [70, 98, 103–107].

Based on the Level 3 processing, which has been described in [section 6.2](#), a fourth level of event selection, pulse cleaning and air-shower reconstruction is added. The following paragraphs present the steps that constitute the newly developed *Level 4 Coincidence* processing [108].

### 6.4.1 Coincident Event Selection and Pulse Cleaning

With an input *i3* file read into memory, the first module discards all P-frames of events that neither pass the coincidence filter ( $\rightarrow$  [subsection 4.2.2](#)) nor have *CoincPulses* ascribed to them. This step incidentally acts as a cut on the zenith angle  $\theta$ . From geometric considerations it becomes clear that coincident events can hardly exceed an inclination of about  $30^\circ$ .



**Figure 6.9:** A simulated coincident cosmic-ray event (proton primary, 1.2 PeV) in IceCube. The electromagnetic component of the air shower (not shown) triggers pulses in IceTop. The bundle of high-energetic muons (red) propagates through the ice and emits light that is recorded by DOMs.

Events with a larger zenith angle will either not be contained by IceTop or the accompanying muon bundle will miss the in-ice volume.

Large IceTop events that trigger 25 stations or more sometimes contain a sufficient amount of PMT afterpulses to be identified by the splitter as a separate event which might even pass event selection criteria. The characteristic feature of those fake events is their time delay of approximately  $7\mu\text{s}$  after the larger main event. By iterating through the P-frames inside a frame packet, afterpulse events can be identified by their relative time shift in the CleanedHLCTankPulses with respect to bright events in the same DAQ window and eventually are removed.

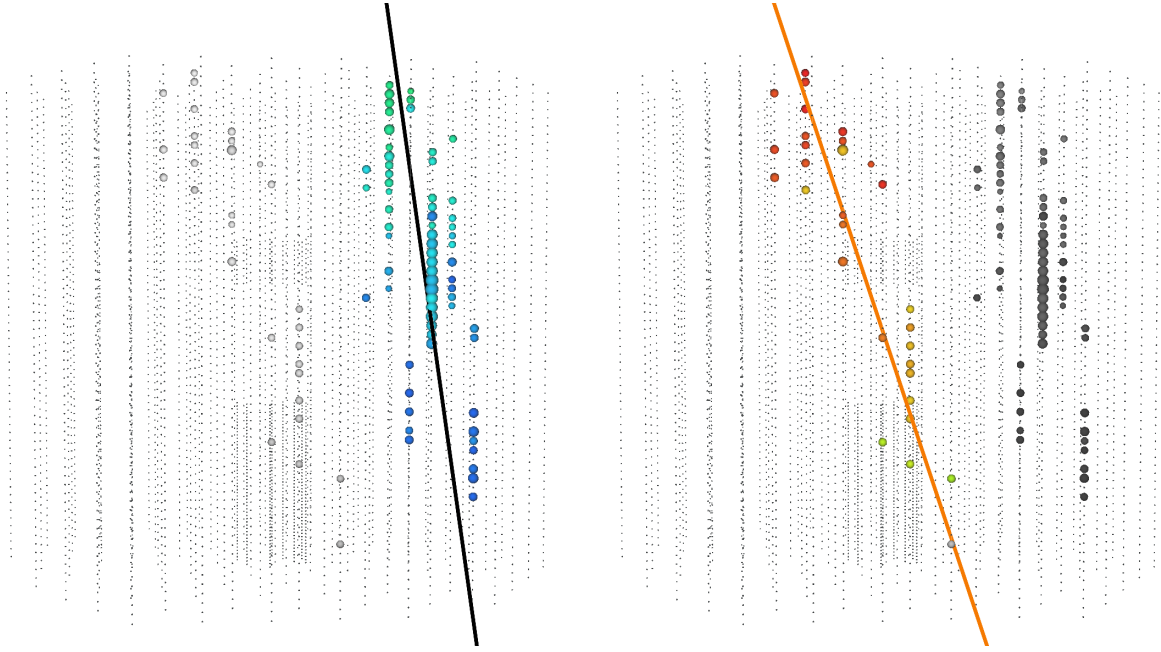
Even after the elimination of afterpulse P-frames, there still are frame packets left with more than one physics frame. Those might originate from independent air showers that are seen by IceTop within the same readout window. From a physics perspective these events are not flawed and could be included in the data analysis, but technical aspects of the further processing inhibit their usage. Some of the next steps (→ [subsection 6.4.2](#) and [subsection 6.4.5](#)) require an unambiguous IceTop pulse series in each frame packet, thus frame packets with more than one P-frame are removed. This leads only to a small deficit of events of ca. 0.02% in data.

Level 3 processing includes a time window cleaning of in-ice pulses that lie in a coincidence time window after the IceTop HLC pulses. This is, however, only performed when IceTopHLC-SeedRTPulses are present in the frame, thus excluding small events from low-energy showers, particularly all 2-station events. In order to extend the range of the coincidence analysis to as low energies as possible, TWC is run once more in Level 4 on those overlooked events, with the distinction that the presence of CleanedHLCTankPulses is sufficient. This is again followed by seeded RT cleaning, resulting in SRTCoincPulses for those low-energy events that lacked time-window-cleaned and cluster-cleaned coincident in-ice pulses in Level 3.

### 6.4.2 Hive splitting

Even though at this point all cases of multiple P-frames in one frame packet have been eliminated, in-ice pulses from coincident high-energy muons that are not related to the air shower detected in IceTop can still be present alongside the muon signals from the main event — even in the SRTCoincPulses. For events with sufficiently high primary energy, the in-ice pulses can be cleaned around the extended Laputop track, which roughly aligns with the center of the bundle of high-energy muons. For lower energies, performing an IceTop-only reconstruction is not possible or not precise enough to allow for a reliable background cleaning of in-ice pulses. A solution is provided by two IceTray modules called IceHive and CoincSuite [109, 110].

IceHive is an advanced event splitter that uses an agglomerative hierarchical clustering algorithm. Any two in-ice pulses that are close-by and causally connected in space-time are considered to belong to the same cluster. Since light propagation is impeded inside the dust layer region of the in-ice detector, pulse series caused by low-inclination muon bundles passing through the detector can be erroneously split into separate sub-events — one



(a) *HiveSplitPulses* (colored) of sub-event 1 with the Monte Carlo track in black.

(b) *HiveSplitPulses* of sub-event 2 with the background track in orange.

**Figure 6.10:** *SRTCoincPulses* of a simulated 1 PeV proton event with muons from an injected 16 TeV background proton shower after Hive splitting.

above and another one below the dust. Immediately following the Hive splitter, CoincSuite runs tests on all pairs of previously split sub-events to check whether the hypothesis of a combined pulse series is valid. In case of a successful test, both splits are re-combined into one P-frame. In this analysis, the *SRTCoincPulses* are taken as input to the event splitter and the split and partially re-combined pulses are labeled *HiveSplitPulses* in order to distinguish them from the original pulse series. Temporarily, individual sub-events are kept in separate *IceHiveSplit* P-frames alongside the *IceTopSplit* P-frames they were created from. Ideally, signal in-ice pulses that stem from the *IceTop* event are now stored in one frame, while potential background pulses were put in different frames. In Figure 6.10 both sub-events that were created by splitting a single set of *SRTCoincPulses* are shown side by side.

### 6.4.3 In-Ice Track Reconstruction

The background rejection process starts with a sequence of track reconstructions on the in-ice pulses. Six parameters are required to define a track: a vertex position in 3-dimensional space, the direction expressed with  $\theta$  and  $\varphi$  as well as time. First, a simple line fit to the *HiveSplitPulses* of each sub-event is performed. This first guess serves as seed for a more advanced reconstruction, which uses a convoluted Pandel probability density function (PDF) [111, 112]. This method also properly takes into account the emission and propagation of Cherenkov light in the ice. The chosen likelihood model is called *SPE* (single photoelectron) and only considers the first pulse of each triggered DOM for the reconstruction, as those are most likely caused by photons that directly hit the sensor without being scattered on their way. The implementation of Pandel track fits in the *IceTray* framework has been motivated

by reconstructing single muons which can be improved by also adding the total number of hits in each DOM, i.e., using a *multi photoelectron* (MPE) likelihood model. The accuracy of muon bundle reconstructions, however, is best with the SPE likelihood. The direction of the resulting SPEfit is varied to create eight different quasi-random<sup>1</sup> seeds for a second iteration of Pandel reconstructions. The one with the highest likelihood is saved in the frame.

#### 6.4.4 Background Tagging

During detector simulation, no information about the origin of each PMT pulse is stored; i.e., there is no tag indicating whether the detected light was emitted by the primary air shower (MCPrimary), a potentially injected background muon bundle or just noise. However, it is possible to quantify the likelihood  $\mathcal{L}$  of a given pulse series to be caused by a specified primary particle.

For every IceHiveSplit frame, the negative log-likelihood  $-\ln \mathcal{L}$  (or nLLH) for the combination of HiveSplitPulses with the MCPrimary track as well as each individual background primary is calculated. For the identification of the signal sub-events and background P-frames, a few cases need to be distinguished:

1. when there is only a single IceHiveSplit P-frame:
  - a) if no background has been injected ( $k = 0$  in Eq. (5.3)) or if the nLLH with respect to the MCPrimary is smaller than the nLLH for all injected background particles<sup>2</sup>, the P-frame is tagged as *signal*
  - b) otherwise the frame is marked as *background*
2. when there are several IceHiveSplit P-frames:
  - a) if the nLLH value for the MCPrimary is below the lowest background value in one or more P-frames, the frame with the lowest  $-\ln \mathcal{L}_{\text{MC}}$  is labeled as *signal*, all other frames receive a *background* tag
  - b) if no P-frame has a  $-\ln \mathcal{L}_{\text{MC}}$  below the background values, all sub-events are considered background

Now, each IceHiveSplit frame contains a tag object indicating whether its cleaned and split HiveSplitPulses are likely to originate from the MCPrimary, i.e., the IceTop air shower, or rather from injected coincident background. This feature will come in handy later, when cuts will be developed for the rejection of frames containing background-triggered in-ice pulses. Therefore, this background tagging module is only activated during the development phase using Monte Carlo events.

<sup>1</sup>Variations of  $\theta$  and  $\varphi$  are generated from a Sobol' sequence [113] which covers the parameter space uniformly.

<sup>2</sup>A minimal  $-\ln \mathcal{L}$  is equivalent to a maximal likelihood  $\mathcal{L}$ .

### 6.4.5 Connecting COGs

As mentioned in [section 6.2](#), events with just two HLC stations not only used to lack cleaned in-ice pulses and a Laputop reconstruction, they also missed out on the estimation of the shower core position  $\vec{r}_{\text{COG}}$  ( $\rightarrow$  Eq. (6.1)). Similarly to the re-run of TWC, ShowerCOGs are computed using CleanedHLCTankPulses for small events. Following that, also the in-ice COG of the HiveSplitPulses in every sub-event frame is determined. In order to generate a first rough estimate of a hypothetical shower axis, the ShowerCOG at the surface and the point on the in-ice-only track, which is closest to the in-ice COG are connected with a line, dubbed “COG2COG”.

Since both vertices, the surface COG and the one deep in the ice, have a time assigned to them in addition to their Cartesian coordinates, it is straightforward to calculate the speed  $\beta_{\text{COG2COG}}$  required for particles to travel from ShowerCOG at the surface to the in-ice COG. Furthermore, the angle  $\Psi$  between the hypothetical COG2COG track and the direction of the reconstructed in-ice-only fit is saved for later use.

### 6.4.6 RockBottom

Despite the conventional Laputop reconstruction not being applicable to small events due to their low number of pulses, a new reconstruction framework, which has been developed by the IceCube [CR](#) group, is capable of reconstructing even low-energy showers. Whereas Laputop only considers IceTop HLC pulses, the new *RockBottom* reconstruction can be configured to take signals from several sub-detectors simultaneously into account [[114](#), [115](#)]. In the future, this flexible design will allow the incorporation of [IACT](#), scintillator and radio signals into a hybrid reconstruction model. In the presented Level 4 coincidence processing, the label “RockBottom” refers to the configuration with IceTop HLCs and SLCs as well as in-ice pulses. The considered pulse series are the cluster-cleaned HLC tank pulses<sup>3</sup> ( $\rightarrow$  [section 6.2](#)), the calibrated SLC tank pulses ( $\rightarrow$  [section 6.3](#)) and the HiveSplitPulses ( $\rightarrow$  [subsection 6.4.2](#)). SLCs outside a  $[-200 \text{ ns}, +800 \text{ ns}]$  time window around the COG2COG core time are omitted as they are likely not related to the main air-shower event.

Similar to Laputop, the likelihood to be maximized has a LDF component  $\mathcal{L}_{\text{LDF}}$  that takes triggered detectors, silent stations and saturated DOMs into account. As a reference distance ( $\rightarrow$  Eq. (6.2)) a value of 125 m is selected and the signal strength expectation is modified by the [RADE](#) snow model. Unlike in Laputop, however, SLCs with a charge of at least 0.7 VEM contribute to the surface likelihood as well, which besides  $\mathcal{L}_{\text{LDF}}$  also includes the time likelihood  $\mathcal{L}_{\text{time}}$ . The shower front curvature is modeled as a Gaussian-modified parabola

$$\Delta t(r) = p_0 \cdot \left( 1 - \exp \left( \frac{-r^2}{p_1^2} \right) \right) + p_2 \cdot r^2,$$

where  $\Delta t$  quantifies the time difference between the curved shower front and a simple plane shower front. The coefficients  $p_0 = 19.41 \text{ ns}$  and  $p_1 = 118.1 \text{ m}$  are fixed, whereas  $p_2$  is a free fitting parameter. The total likelihood is complemented by an in-ice likelihood  $\mathcal{L}_{\text{inice}}$  using

<sup>3</sup>IceTopHLCSeedRTPulses are used if they exist, otherwise CleanedHLCTankPulses. In the latter case, the event is too small for SRT cleaning and the CleanedHLCTankPulses cannot be cleaned any further.

**Table 6.1:** Steps of the RockBottom reconstruction. The seed is the first guess for the minimizer. For each parameter it is indicated whether it is free within an absolute or relative interval, completely free or fixed. Since the contribution of  $\mathcal{L}_{\text{time}}$  is set to zero in the first step, the shower front curvature is not a parameter of the minimization.

step	seed	$S_{125}$	$\beta$	$x$ & $y$	$\theta$ & $\varphi$	$t_0$	$p_2$
1	COG2COG	(-3.0, 8.0)	(2.9, 3.1)	$\pm 200$ m	fixed	fixed	–
2	step 1	(-3.0, 8.0)	(2.0, 4.0)	$\pm 15$ m	free	$\pm 200$ ns	free
3	step 2	(-3.0, 8.0)	(0.0, 10.0)	$\pm 45$ m	free	$\pm 200$ ns	free

the SPE model for the PDF. The combined nLLH

$$-\ln \mathcal{L} = -\ln \mathcal{L}_{\text{LDF}} - \ln \mathcal{L}_{\text{time}} - \ln \mathcal{L}_{\text{inice}}$$

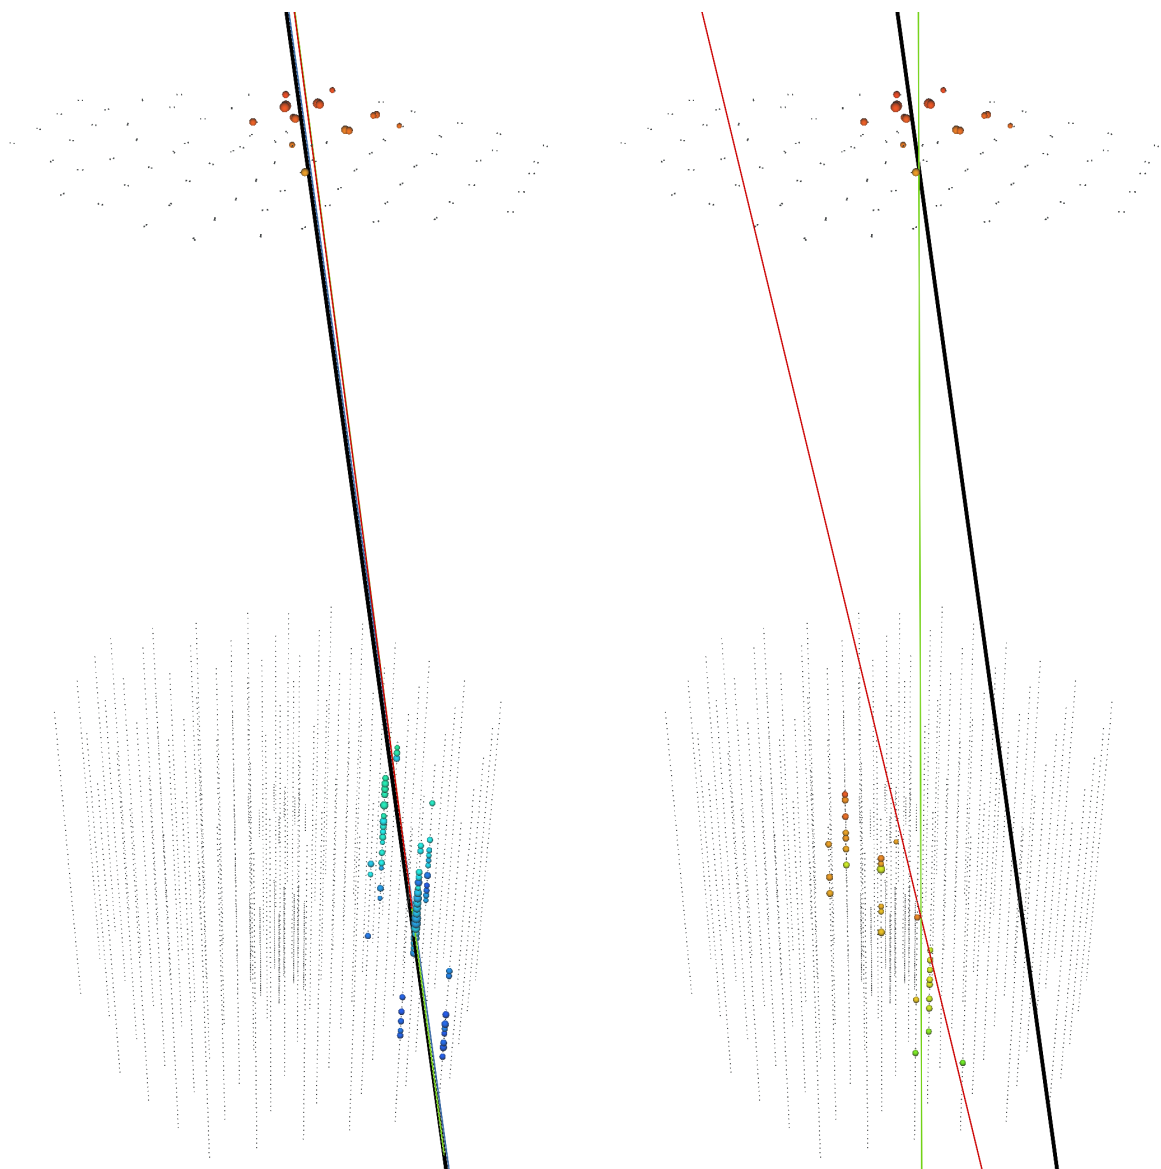
is minimized in three iterations using an implementation of the MINUIT algorithm [116] in IceTray. In contrast to the previously discussed in-ice-only track reconstruction, which only considered  $\mathcal{L}_{\text{inice}}$ , the RockBottom module needs to find the best fit values of eight parameters: the 2-dimensional position and the time of the shower core at the surface, the direction ( $\theta$  and  $\varphi$ ), the shower size  $S_{125}$  and the LDF slope  $\beta$  as well as the  $p_2$  coefficient in the shower front parameterization. The selection of nLLH components and the limits on the free parameters vary in order to encourage stability of the fit. An overview is provided in Table 6.1.

The reconstructed RockBottom track of a simulated event with two IceHiveSplit frames can be seen in Figure 6.11 along with the in-ice-only SPE track and the COG2COG seed.

#### 6.4.7 Background Rejection

At this intermediate point within the Level 4 processing, frame packets consist of a Q-frame, accompanied by one IceTopSplit P-frame. Additionally, one or more IceHiveSplit P-frames may be present, which can contain signal or background in-ice pulses together with their corresponding in-ice fit, the COG2COG track and a RockBottom reconstruction. It is desirable to achieve the same simple frame structure as in Level 3 for the output i3 files, that is, pairs of Q-frames and IceTopSplit P-frames. Frame objects that have been added to the IceHiveSplit frames therefore need to move to the original IceTopSplit frame. The decision whether an IceHiveSplit frame is treated as background — and thus rejected — or expected to contain signal pulses with the corresponding reconstructions is made based on the quantities calculated in subsection 6.4.5. Figure 6.12 shows the distributions of angle  $\Psi$  and speed  $\beta_{\text{COG2COG}}$  for background-tagged sub-events and those that received a label indicating they are likely signal frames.

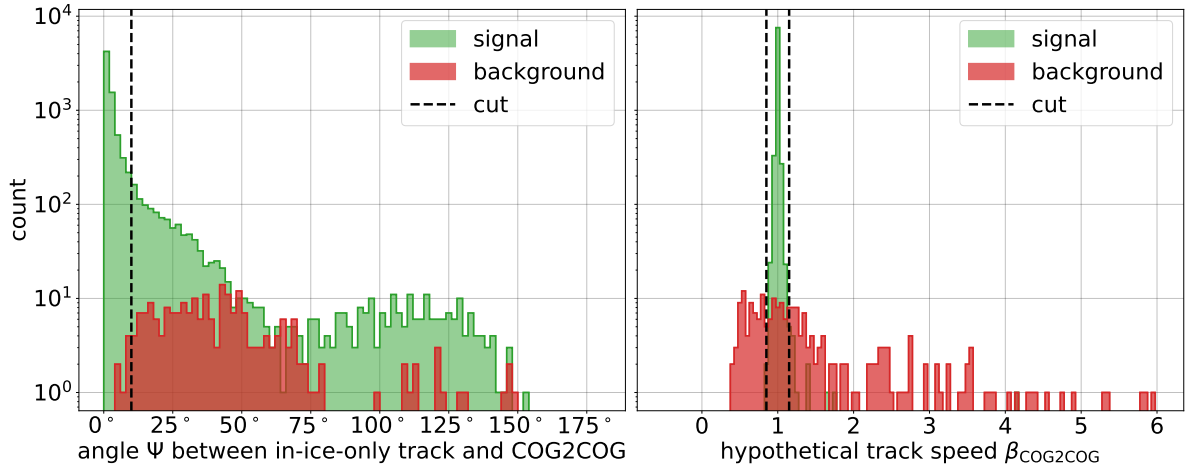
When the in-ice-only track reconstruction SPEFit points back to the ShowerCOG at the surface, that is, the angle  $\Psi$  between SPEFit and COG2COG is small, this is a clear indication that the HiveSplitPulses stored in that frame originate from the air shower that has been detected by IceTop. A cut at  $10^\circ$  keeps most of the signal frames while rejecting the majority of background-tagged sub-events.



**(a)** The signal sub-event with the Monte Carlo track in black, the in-ice SPE reconstruction in red, the hypothetical COG2COG seed track in green and the RockBottom reconstruction in blue.

**(b)** The background sub-event with in-ice (red) and COG2COG track (green). A RockBottom reconstruction is not shown because the minimizer did not converge in this case.

**Figure 6.11:** Track reconstructions in two IceHiveSplit frames from the same DAQ event as in Figure 6.10, including the shared IceTop HLCs and SLCs as well as the respective HiveSplitPulses.



**Figure 6.12:** Distributions of the background rejection quantities for signal-tagged (green) and background-tagged sub-events (red). The simulated primary CRs are protons, helium, oxygen and iron nuclei with about 1 PeV. **Left:** The angle between in-ice SPEFit and the COG connection track. The cut is passed by sub-events with  $\Psi < 10^\circ$ . **Right:** The required speed of the hypothetical COG2COG track. A symmetric cut is applied at  $0.85c$  and  $1.15c$ .

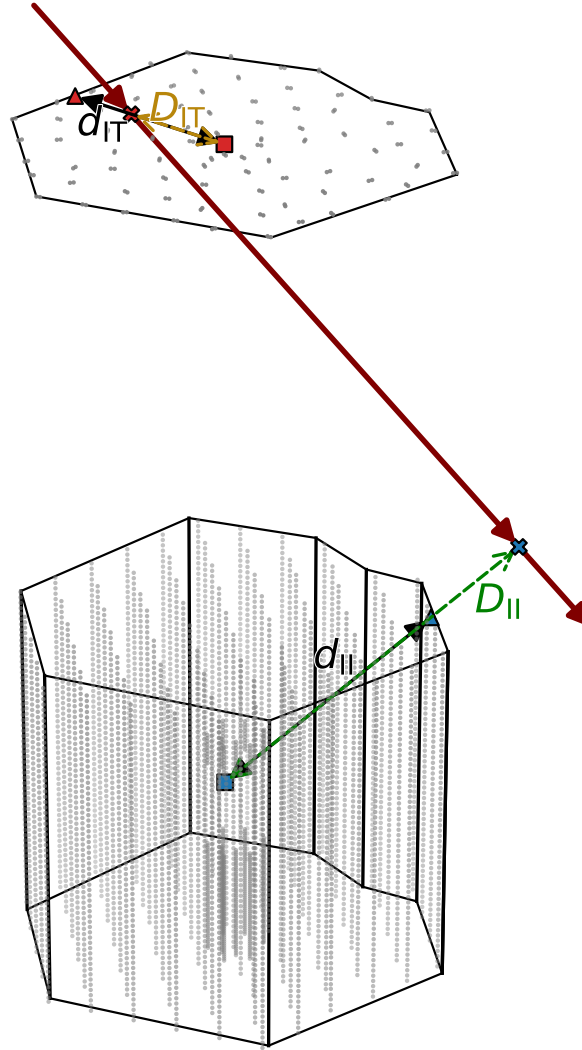
The high-energy muons of an EAS that penetrate the ice and trigger DOMs are highly relativistic. This should be reflected in the required speed  $\beta_{\text{COG2COG}}$  of the hypothetical track connecting surface and in-ice COG. A value close to the speed of light ( $v/c = 1$ ) indicates agreement with the hypothesis of IceTop and in-ice pulses being caused by the same air shower. A hypothetical speed significantly below 1 is caused by HiveSplitPulses appearing too late after the IceTop event for them to be causally related. Unphysically high (or negative) values of  $\beta_{\text{COG2COG}}$  on the other hand are the result of background pulses being triggered almost simultaneously with or even before the IceTop event. IceHiveSplit frames with a  $\beta_{\text{COG2COG}}$  within  $\pm 15\%$  of the speed of light are further considered signal sub-events.

For the final decision on which sub-event is selected to be kept, all IceHiveSplit frames of an event that pass both the speed cut and the angle cut are collected:

- If no frame passes the cuts, nothing is moved to the IceTopSplit frame and the entire event is omitted in the further analysis.
- If one frame passes the cuts, it is considered a signal sub-event. HiveSplitPulses, the RockBottom reconstruction track and other quantities unique to the sub-event are moved to the IceTopSplit frame.
- If multiple frames pass the cuts but one frame has the lowest values in  $\Psi$  as well as  $|\beta_{\text{COG2COG}} - 1|$ , its pulses and reconstructions are taken to the IceTopSplit frame; otherwise none will be used and the entire event is omitted in the further analysis.

#### 6.4.8 Containment

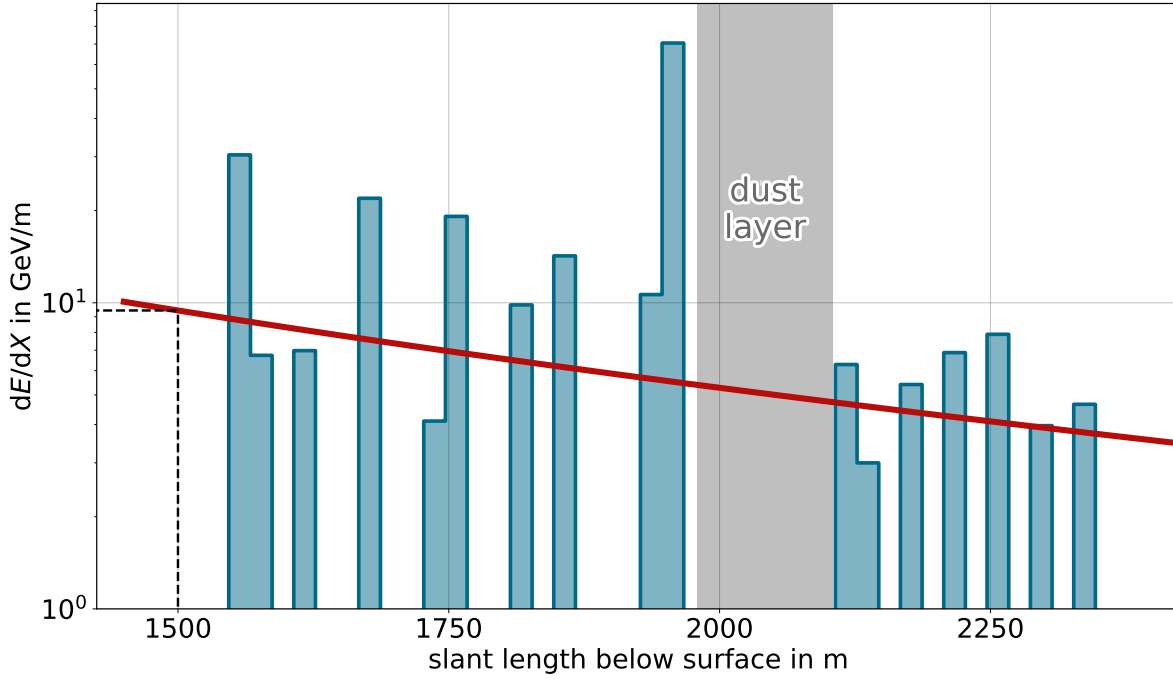
One important tool for the suppression of poorly reconstructed air showers in the final dataset is the removal of events that are not or just hardly contained within the detector. The degree of containment can be quantified with a fractional containment score, which is the



**Figure 6.13:** Definition of fractional containment in IceCube. In this example, the (reconstructed) shower axis (dark red) intersects the IceTop array (red cross) but not the instrumented in-ice volume (blue cross outside the detector). Squares are located at the centers of both sub-detectors while triangles mark the points on the detector hull or border closest to the track. The ratio of the lengths  $D/d$  is the containment value, which is below 1 in IceTop ('IT') and above 1 for the in-ice detector ('II') in this example. Graphic inspired by Fig. 4.5 of Ref. [103].

factor by which the detector would need to grow or contract in order to have a specified track being tangent to the detector. A value below 1 thus indicates containment within the detector area or volume, whereas values greater than 1 correspond to tracks outside the detector. For the quality cuts employed later ( $\rightarrow$  [section 7.4](#)), containment of the reconstructed RockBottom track inside the IceTop array as well as the in-ice detector is calculated (Fig. 6.13).

In preparation for the RockBottom reconstruction ( $\rightarrow$  [subsection 6.4.6](#)), the SLC tank pulses have been cleaned around the hypothetical COG2COG connection track. Now, with the more accurate RockBottom track, this can be re-done and the RockBottom-cleaned SLCs are stored in the output i3 file.



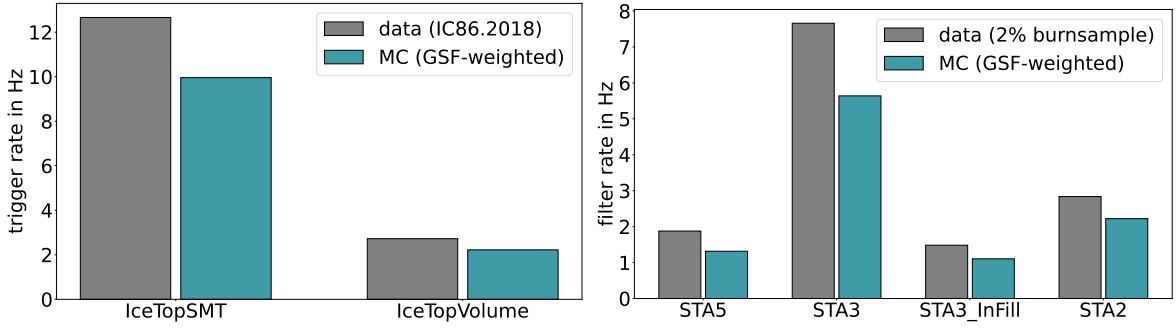
**Figure 6.14:** The profile of in-ice energy loss from muons in an almost vertical coincident CR event, reconstructed from the detected charge deposit. For this event, the dust layer is located at slant lengths 1979–2105 m. The best fit of Eq. (6.4) is shown in red and yields a  $dE/dX_{1500}$  of 9.44 GeV/m.

#### 6.4.9 In-Ice Energy Loss

Level 4 processing is concluded with the reconstruction of the energy loss profile of muons in the ice. By considering the amount of recorded light as a function of slant length  $\varsigma$  (in units of meters), the rate of energy loss can be described by

$$\frac{dE}{d\varsigma}(\varsigma) = \left( \frac{\kappa_m A}{\cos(\theta)} \right) \cdot e^{-b\varsigma} \cdot \gamma_\mu \cdot \left( \frac{E_0}{A} \right)^{\gamma_\mu - 1} \cdot \left[ \left( -\frac{E_0}{A} \right)^{-\gamma_\mu} \left( \frac{a}{\gamma_\mu} - \frac{b}{1 - \gamma_\mu} \cdot \frac{E_0}{A} \right) + E_{\min}^{-\gamma_\mu} \left( \frac{a}{\gamma_\mu} - \frac{b}{1 - \gamma_\mu} \cdot E_{\min} \right) \right], \quad (6.4)$$

which can be derived from the formula for the number of high energy muons (Eq. (2.5)). The constants  $a = 0.23881$  GeV/m for the ionization energy loss and  $b = 3.2852 \cdot 10^{-4}$  /m for radiative losses [117] are also parameters of  $E_{\min}(\varsigma) = (e^{b\varsigma} - 1) \cdot a/b$ . Furthermore, from Eq. (2.5) the parameters  $\kappa_m = 14.5$  GeV and  $\gamma_\mu = 1.757$  are set. For the sake of a stable fit, the mass number  $A$  is fixed at 16, making the best-fit primary energy value  $E_0$  unsuitable for a proper energy estimate [98]. In the literature,  $dE/d\varsigma$  is commonly referred to as  $dE/dX$ . Although this opens the possibility of confusion with the atmospheric slant depth  $X$  (in units of  $\text{g}/\text{cm}^2$ ), in the context of in-ice muon energy loss, the expression  $dE/dX$  will be used in this work as well. Figure 6.14 shows an example of the energy loss profile of a recorded coincident air-shower event with the fit function defined in Eq. (6.4). The interpolated rate of energy loss at a slant length of 1500 m,  $dE/dX_{1500}$ , is a good proxy for the number of TeV muons seen by IceCube and has been used in previous analyses [70].



**Figure 6.15:** *Left:* rates for the two IceTop triggers during the 2018 season and the weighted Monte Carlo dataset (Level 1), *right:* rates in the 2018 burnsample and the weighted Monte Carlo dataset for the four IceTop filters (Level 2).

## 6.5 Monte Carlo Verification

The production of [CORSIKA](#) air showers ([section 5.1](#)) and the simulation of the IceCube detector response ([section 5.2](#)) are carried out by a sequence of programs and scripts with a large number of adjustable parameters, such as the chosen time window in the parameterization of the background injection or the assumed detection efficiency of DOMs. The CORSIKA team and the IceCube collaboration have improved every part of the simulation chain and tweaked the arguments to achieve a more realistic output. Nonetheless, it is inevitable to check the i3 files resulting from the simulation procedure so that they can be used for the analysis. In order to do so, the distributions of low-level (trigger rates, HLC pulse charges, etc.) and high-level quantities (like shower core positions,  $S_{125}$  or cleaned pulse distributions) of the MC dataset are compared with the corresponding distributions of a 2% sub-sample of the [IC86.2018](#) data, called the *burnsample*. In this work, the term *data* refers to experimentally obtained data in contrast to simulated MC data, unless stated otherwise. This section is about low-level comparisons up to Level 2, whereas high-level checks of the observables entering the analysis pipeline will be covered in [subsection 8.1.1](#).

The CORSIKA datasets have been produced following an  $E^{-1}$  spectrum and with an equal number of showers for every primary. For a sensible data/MC comparison, the simulated events therefore need to be weighted to a more realistic energy spectrum. For the verification plots in this chapter, a weighting according to the [GSF](#) composition model has been applied. This model has been fit to a 4-component CR flux consisting of proton, helium, oxygen and iron primaries where the last two represent the mass groups Li–F and Ne–Fe, respectively.

### 6.5.1 Event Rates

As every data analysis starts with the data acquisition, the first part of the low-level check is the comparison of IceTop trigger rates at Level 1. The rate of events having the SMT6 or the 2-station trigger condition fulfilled ( $\rightarrow$  [subsection 4.2.1](#)) within their trigger window is averaged over the entire data-taking season 2018, as well as over the weighted Monte Carlo dataset. [Figure 6.15](#) (left) reveals a clear difference in the trigger rates. For [SMT](#) the data/MC ratio is about 1.27 and for the volume trigger it is close to 1.23. Checks of previous IceCube datasets also show a deficit in the MC rates on a similar level in the data/MC ratio, which is

an artifact of the GSF composition model predicting a slightly lower cosmic-ray flux than seen by IceCube. Furthermore, the anticorrelation of pressure and rate leads to a lower rate in MC, which assumes a higher atmospheric pressure than the average pressure during the data-taking season ( $\rightarrow$  [section 8.4](#)).

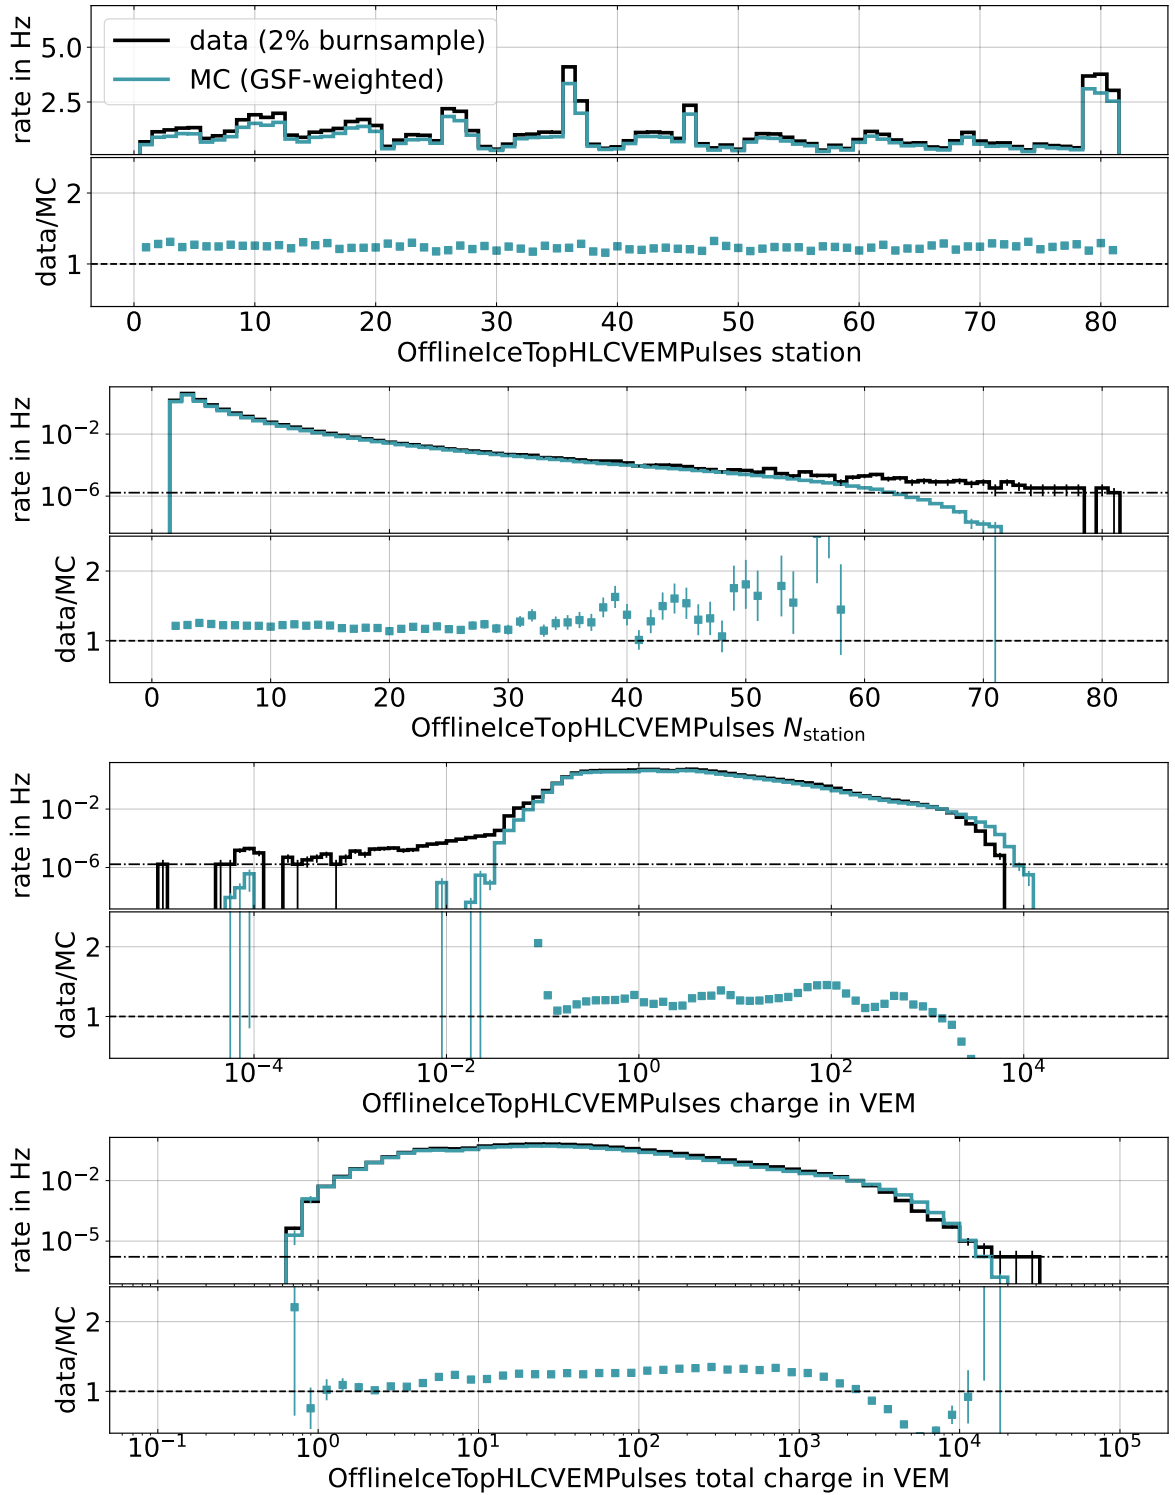
Filter information is added to the i3 files in Level 2. For the comparison of filter rates, the number of events passing one of the IceTop filters in [Table 4.2](#) is determined for the 2018 burnsample with a total duration of about 7 days as well as the for weighted Monte Carlo dataset. The right plot in [Figure 6.15](#) also shows a discrepancy in the filter rates, which is a direct consequence of the similar effect observed in the trigger rates. The ratio values range from about 1.43 (5-station filter) to just below 1.28 (2-station filter).

### 6.5.2 Pulse Series

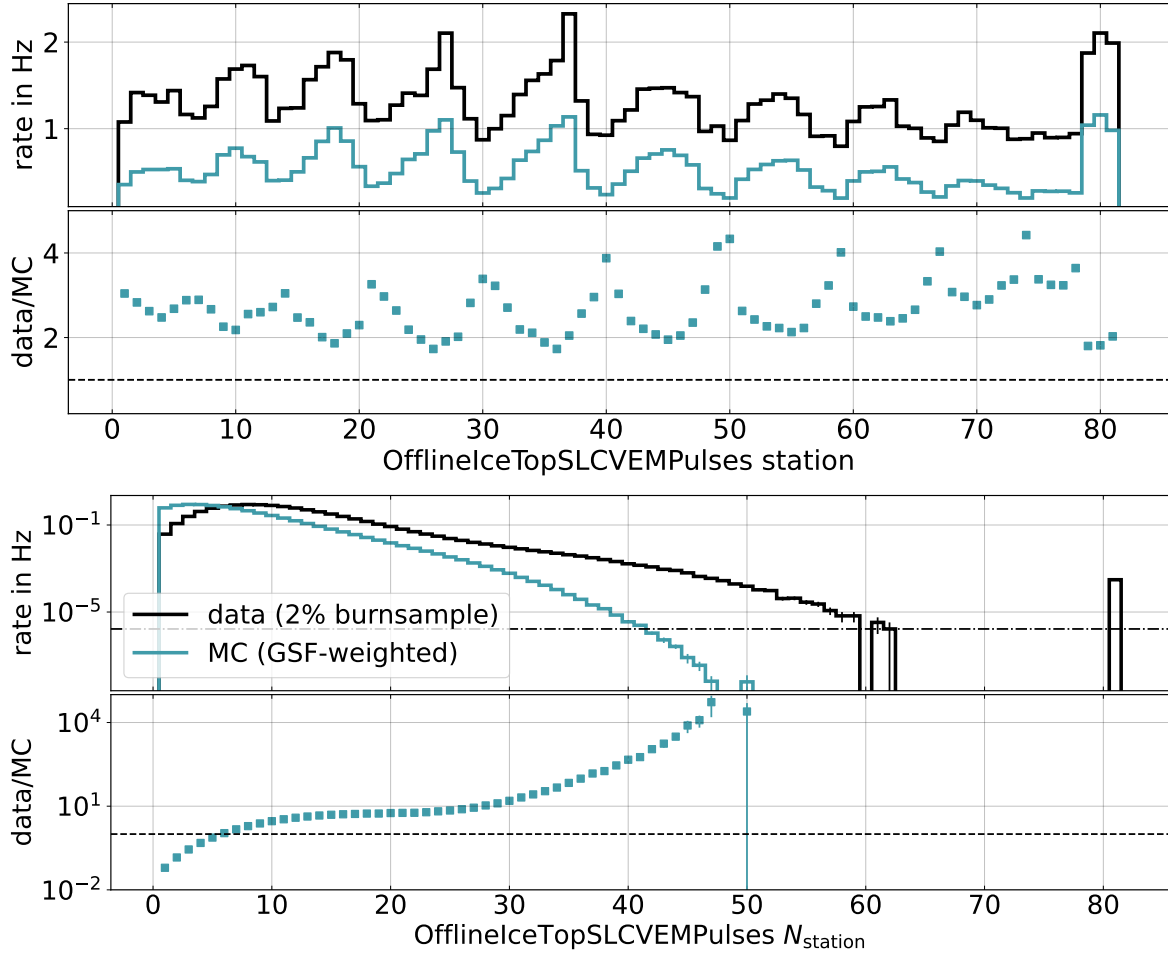
The three classes of pulses that play a role in the CR analysis presented are IceTop HLCs and SLCs as well as in-ice pulses. Their most basic versions at Level 2 are `OfflineIceTopHLCVEMPulses`, `OfflineIceTopSLCVEMPulses` and `InIcePulses`. Before these pulses undergo cleaning and further (re-)calibration steps in Level 3 and 4, their data-MC agreement is checked at this early processing stage. For events to participate in the following MC verification, they need to pass at least one of the four IceTop filters.

[Figure 6.16](#) is a compilation of the HLC distributions in the burnsample data and the weighted simulation sets at Level 2. The top plot presents the occupancy of all 81 IceTop stations. There is a general agreement between the 2018 burnsample and the MC, the constant offset in the data/MC ratio is the result of flux weighting and the pressure discrepancy artifact and will no longer be addressed in this section. However, the HLC rate in all stations is not uniformly distributed. This is for two reasons: due to the uneven snow coverage ( $\rightarrow$  [Fig. 5.7](#)) those tanks that are buried under a deeper layer of snow have a lower chance of detecting electromagnetic shower particles, which results in a lower HLC rate; furthermore, events with HLCs in stations inside the denser InFill region (26, 27, 36, 37, 46, 79, 80 and 81) are more likely to pass an IceTop filter, which manifests itself in a high rate in those stations. The second distribution is of the number of IceTop stations that had HLC pulses in an event. Both histograms overlap mostly until, beyond a station multiplicity of about 30, the MC distribution runs out. This is because of the lack of high-energy events ( $> 100$  PeV) in the simulation sets. The lower two plots show the HLC pulse charges, which demonstrate agreement between data and MC over almost the entire range. However, measured data contains single pulses with very low charges between  $10^{-5}$  and 0.02 VEM, which are barely present in simulations. The reason for this is not yet understood, although it is expected that the distribution of aggregated pulse charges in each DOM does not feature such a tail.

Since SLC charges are calibrated in Level 3, these low-level checks only cover the SLC count rates. [Section 7.7](#) provides high-level verification plots of SLC charges after further processing. The top plot in [Figure 6.17](#) gives the SLC rate of each IceTop station. The offset between data and MC is considerably larger than in HLCs and lies in the order of 0.8 Hz, which results in a non-uniform ratio. The reason for the higher rate in the burnsample is the presence



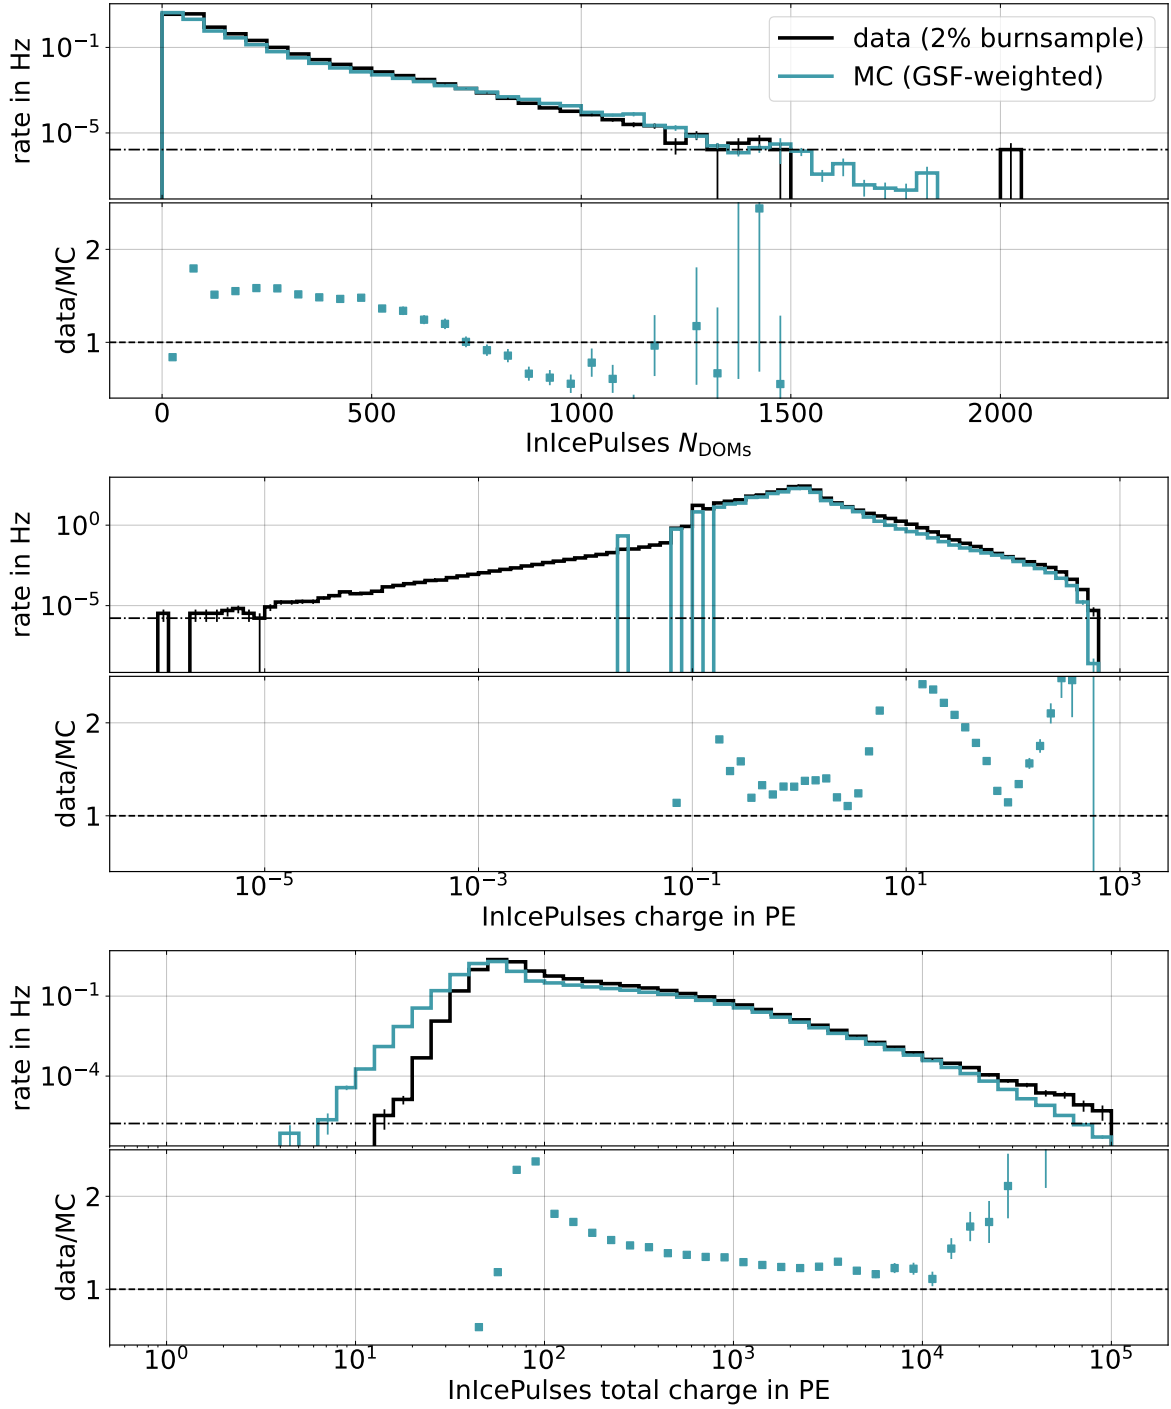
**Figure 6.16:** From top to bottom: distributions of IceTop station index, number of triggered stations, individual pulse charge and total event charge for the `OfflineIceTopHLCVEMPulses` (data and Monte Carlo). The dash-dotted horizontal line in the upper panel of some plots marks the minimum possible rate in the burnsample, that is, a single occurrence in total which corresponds to about  $1.68 \cdot 10^{-6}$  Hz.



**Figure 6.17:** Data / Monte Carlo comparison of *OfflineIceTopSLCVEMPulses*. **Top:** Station occupancy. **Bottom:** Number of hit stations; note the logarithmic vertical axis in the *data/MC* panel. The dash-dotted horizontal line in the upper panel marks the minimum possible rate in the burnsample.

of coincident noise hits — a phenomenon which HLCs are mostly protected from due to their local coincidence condition. The addition of such noise SLCs is not included in the simulation, thus leading to the lower rate compared to data. This also has the effect of MC events ending up with a smaller number of SLC stations than real events measured with IceTop (bottom plot in Fig. 6.17). The peak at 81 stations in the burnsample distribution originates from a specific kind of events that have been triggered by the so-called fixed-rate trigger (FRT). The number of SLC stations in FRT events saturates at 81 because the readout time window is much longer than the standard hit-based triggers. Therefore, the content of a FRT frame does not originate from a single CR event, which is why in a later processing stage such frames will be excluded from the analysis.

In coincident events, the in-ice pulses are predominantly caused by bundles of high-energy muons. This makes them highly relevant for a study of cosmic-ray mass, hence they need to be included in the data-MC verification (Fig. 6.18). The number of in-ice pulse DOMs per event in MC agrees well with data. Events that were triggered by FRT or the “slow particle trigger” tend to predominantly contain noise from a majority of in-ice DOMs and are thus excluded. The center and bottom plots show the distributions of pulse charge and total event



**Figure 6.18:** From top to bottom: distributions of the number of triggered DOMs, individual pulse charge and total event charge for the *InIcePulses* (data and Monte Carlo). The dash-dotted horizontal line in the upper panels marks the minimum possible rate in the burnsample.

charge, respectively. The MC distribution does not reach the same level of high total charge as a consequence of the limited simulated energy range.

### **Key Messages**

*Recorded events receive filtering, cleaning, reconstruction and many more crucial treatments as part of the processing Levels 2 and 3, which now also includes calibration of single tank hits (SLCs) for low-energy events. A new Level 4 processing dedicated to IceTop/in-ice coincidences has been developed. The comparison of low-level Monte Carlo distributions does not expose any severe mismatches with data.*

# CHAPTER 7

## QUALITY CUTS

*“Bauz! da geht die Türe auf,  
Und herein in schnellem Lauf  
Springt der Schneider in die Stub’  
Zu dem Daumen-Lutscher-Bub.  
Weh! Jetzt geht es klipp und klapp  
Mit der Scher’ die Daumen ab,  
Mit der großen scharfen Scher’!  
Hei! da schreit der Konrad sehr.”*

— Heinrich Hoffmann, «Der Struwwelpeter»

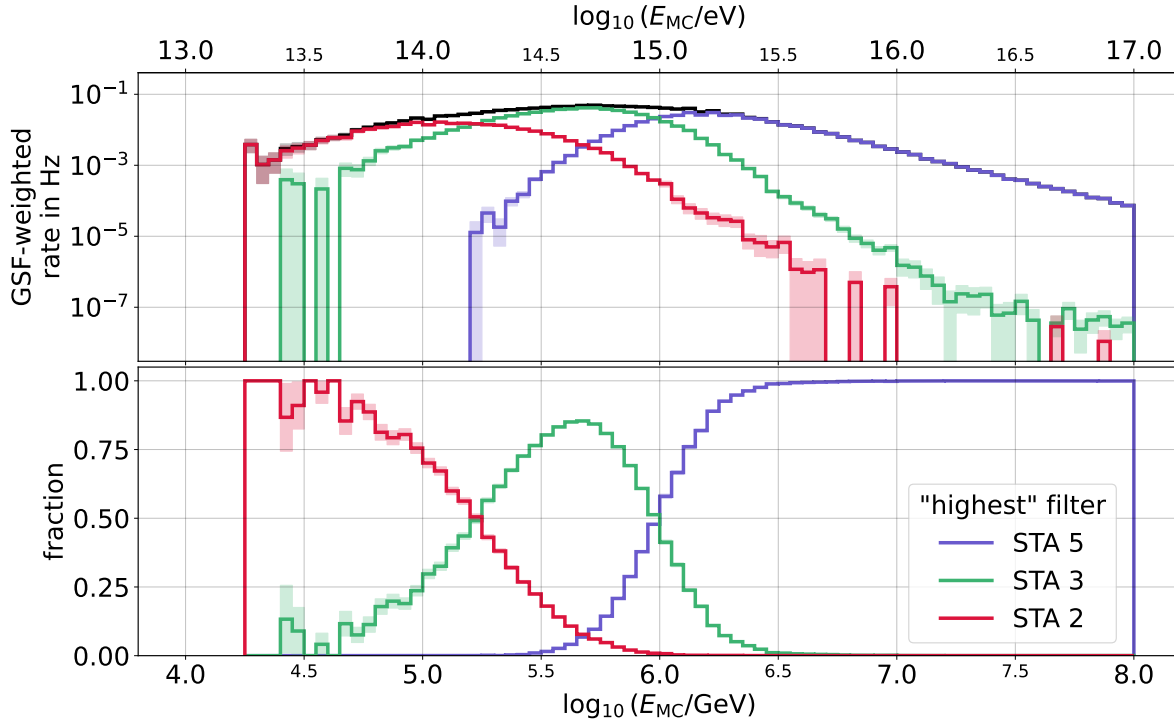
The quality of an analysis strongly depends on the quality of the data that is processed. Ensuring the integrity of the input data is crucial to obtain a valid and credible result. Corrupted data and severely misreconstructed events impede the optimization of analysis methods or can lead to unstable results. Furthermore, the data used for the development of the analysis generally covers only that part of the phase space which is relevant for the work. For example, the energy range of the [MC](#) dataset used to set up this analysis is limited to the scope of interest, whereas measured data extends beyond those points.

A set of selection cuts has been identified, which, when applied to data, remove events with poor reconstruction quality and limit the phase space to coincident air-shower events in the PeV region. Crucially, since the same cuts are used for simulated and measured data, the distributions of the [NN](#) features and target observables must largely agree between both datasets. The following sections provide an overview of the chosen quality cut definitions, their motivation and their effect. In the end, the distributions of pulse information and reconstructed quantities are studied after the application of all quality cuts to ensure comparability of simulations and measured data.

### 7.1 Passed IceTop Filters

The first step of Level 4 processing removed events that did not pass the coincidence filter (→ [subsection 6.4.1](#)). However, this leaves events untouched that only fulfill the minimum filter condition, that is, an InIceSMT trigger and one IceTop [HLC](#) hit (→ [subsection 4.2.2](#)). In order to reduce the background of random single-station HLC coincidences with the in-ice detector and to provide better stability of the track reconstruction chain, events are required to pass at least one IceTop filter (→ [Tab. 4.2](#)).

At low primary energies of 200 TeV and below, the majority of events pass only the 2-station filter ([Fig. 7.1](#)). With increasing energy more events also pass the 3-station or 3-station InFill filter (exclusively or in addition to the 2-station filter), while 2-station events still contribute



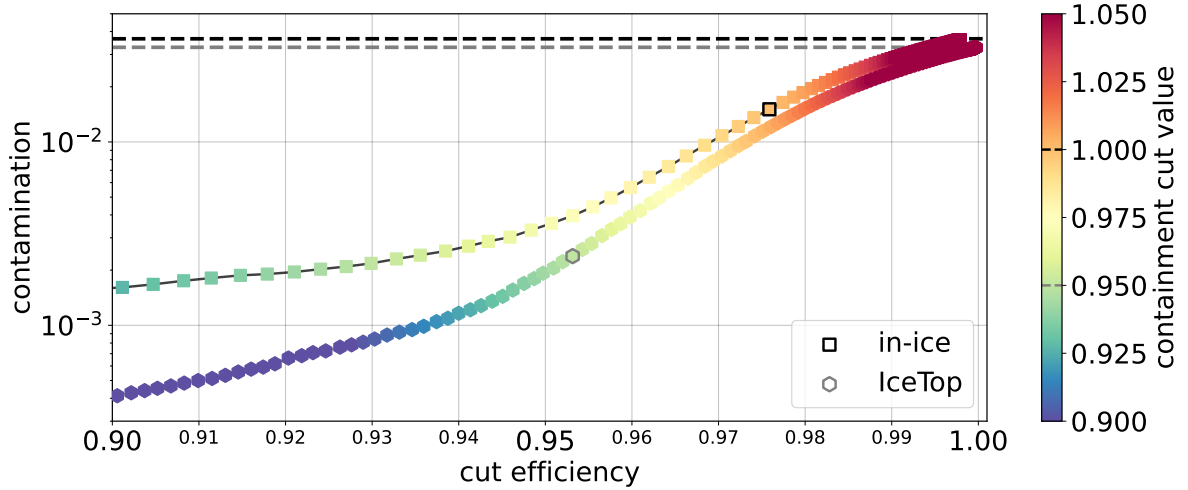
**Figure 7.1:** Distribution of the “highest” filter in Monte Carlo. The “highest” filter is the filter with the highest number of required stations, e.g. if an event passes the 2-station and the 3-station filter, it is considered a STA 3 event in this context. Here, the 3-station and 3-station InFill filter are combined. The shaded bands represent the statistical uncertainty. **Top:** GSF-weighted rates for the three filter groups and all events (black) as a function of energy. **Bottom:** The fraction of events for every highest filter in each energy bin.

significantly to the expected rate. Above 1 PeV, events with five or more triggered IceTop stations dominate the data.

## 7.2 In-Ice Charge

The amount of Cherenkov light seen by the in-ice detector is a vital event property for a composition analysis. An increased multiplicity of high-energy muons in the ice leads to more light emission and thus to a larger number of triggered DOMs recording higher charges. However, the spatial distribution of in-ice charges also gives important insights into the event topology. This is taken into account by the reconstruction of the energy loss profile but can also be used for a selection criterion that rejects events that are not suitable for the analysis.

As a first requirement, the presence of HiveSplitPulses ( $\rightarrow$  subsection 6.4.2) is checked. Only events with a set of these thoroughly cleaned in-ice pulses are considered further. Furthermore, coincident events are expected to have in-ice pulses in the upper part of the detector volume, caused by the down-going bundle of muons. Only events with a nonzero cumulative charge of HiveSplitPulses in the top 32 layers of DOMs ( $Q_{32}^{\text{in-ice}}$ ) will participate in the analysis. This cut removes inclined high-energy events that manage to trigger IceTop stations on the edge of the array and have their muon bundle enter the in-ice detector in the lower half. It also serves as a second background rejection filter.



**Figure 7.2:** For *in-ice* (square markers) and *IceTop* (hexagons) the contamination by truly uncontained events and the cut survival fraction (both weighted to the GSF flux) are plotted for various choices of the respective maximal containment value. Cuts on *in-ice* and *IceTop* containment are handled separately. Those markers with highlighted outlines correspond to the chosen quality cut values, which are also indicated on the color bar. The horizontal dashed lines represent the respective levels of contamination before any containment cut (3.28% in *IceTop* and 3.65% for *in-ice*).

### 7.3 RockBottom Convergence

A reliable RockBottom (**RB**) reconstruction is essential to determine the energy of the primary cosmic ray and to estimate its mass. The **nLLH** minimization needs to converge; otherwise, the event is omitted. The coincidence requirement implicitly applies a soft cut on the zenith angle at about  $40^\circ$ , already during Level 4 processing ( $\rightarrow$  subsection 6.4.1). No explicit cut in this range of inclination is applied; however, events that were clearly misreconstructed are removed by requiring  $\theta_{\text{RB}} < 90^\circ$ .

### 7.4 Containment

In order to improve the shower core resolution and the angular resolution in the dataset, events that are not contained within both sub-detectors need to be rejected. If the core of an air-shower event lies outside the *IceTop* array or if the muon bundle misses the *in-ice* detector volume, signals can still be generated in nearby DOMs and cause triggers. However, the accuracy of the resulting track reconstruction declines drastically for such uncontained events.

The application of a quality cut on fractional containment ( $\rightarrow$  subsection 6.4.8) requires a trade-off between cut efficiency and level of contamination. The efficiency or survival rate of a cut quantifies the fraction of events that pass the selection criterion. Events with a true containment value outside the detector boundaries ( $> 1$ ), which, however, pass the containment cut, lead to contamination of the dataset. An ideal containment cut provides a high survival rate, while not leading to excessive contamination. The relation between cut efficiency and the level of contamination as a function of the chosen cut values can be inspected using Figure 7.2.

The goodness of an air-shower track reconstruction can be described by its resolution of the shower core position and of the arrival direction. Core resolution is defined as the  $\iota$ -percentile of the distribution of the distance between the true and the reconstructed shower core, where

$$\iota = \frac{100}{\sqrt{2\pi}} \int_{-1}^1 e^{-\frac{x^2}{2}} dx \approx 68.27$$

is the central mass of a Gaussian distribution in percent. Similarly, the angular resolution is defined as the  $\iota$ -percentile of the distribution of the opening angle  $\Psi$  between the true and the reconstructed arrival direction.

The resolution as a function of the RockBottom containment can be seen in Figure 7.3 along with the cuts applied on both containment values. Events with RockBottom reconstructions not contained in IceTop have a shower core resolution of about 35 m or worse, while contained events are significantly better reconstructed with a resolution of about 10–20 m. Therefore, it can be justified to set a maximum IceTop containment of 0.95, which roughly corresponds to a 20 m safety margin in relation to the 400–500 m distance from the center of IceTop to the edge. Events with a RockBottom (and also a true) shower axis close to the edge of the in-ice detector are generally still quite accurately reconstructed. An in-ice containment cut at 1.0 is sufficient to remove air showers with muon bundles missing the deep optical detector but also keep the majority of well-reconstructed events.

Figure 7.4 shows the influence of the containment cut on the core and the angular resolution.<sup>1</sup> Above 2 PeV, both the core and the angular resolution get worse due to an increased fraction of uncontained events (open markers). Removing uncontained events leads to a significant improvement in resolution in that energy range (filled markers). For energies of a few PeV and below, the quality cut has a negligible effect, as the shower tracks are generally contained inside the detector. However, the new RockBottom reconstruction outperforms the standard algorithm Laputop, which only takes IceTop pulses into account and is limited to events with at least three HLC stations<sup>2</sup>.

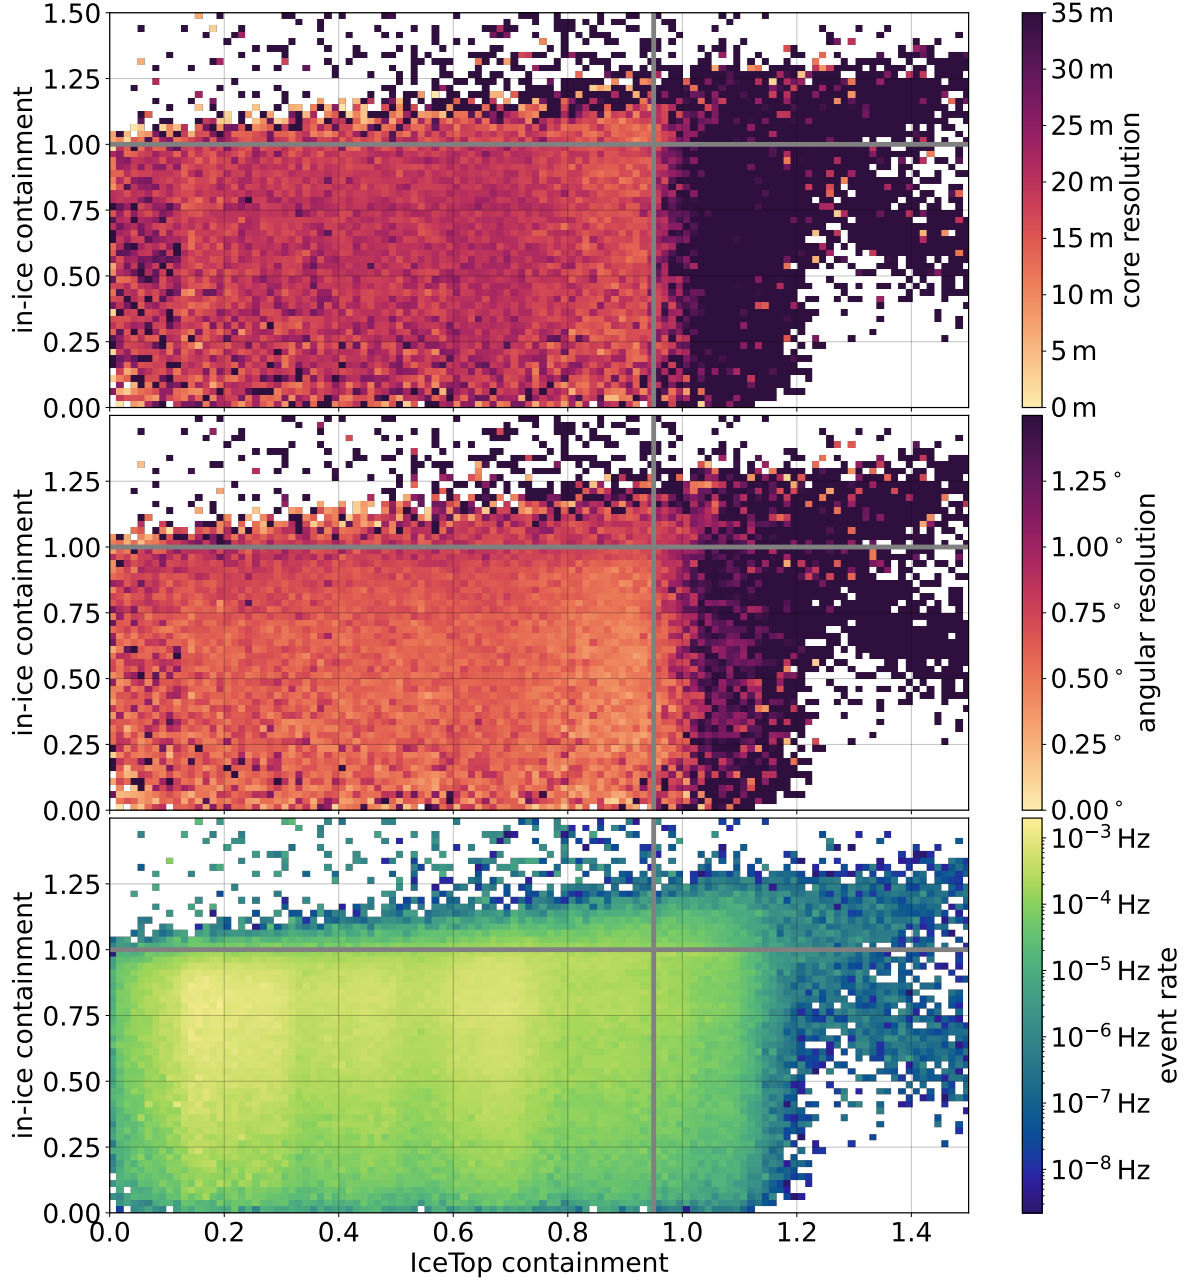
Before applying the containment cut, the level of contamination in IceTop is 3.28% and it is 3.65% for the in-ice detector. The containment cut reduces the contamination to 0.23% and 1.51%, respectively, when weighted according to the GSF flux model.

## 7.5 Shower Size

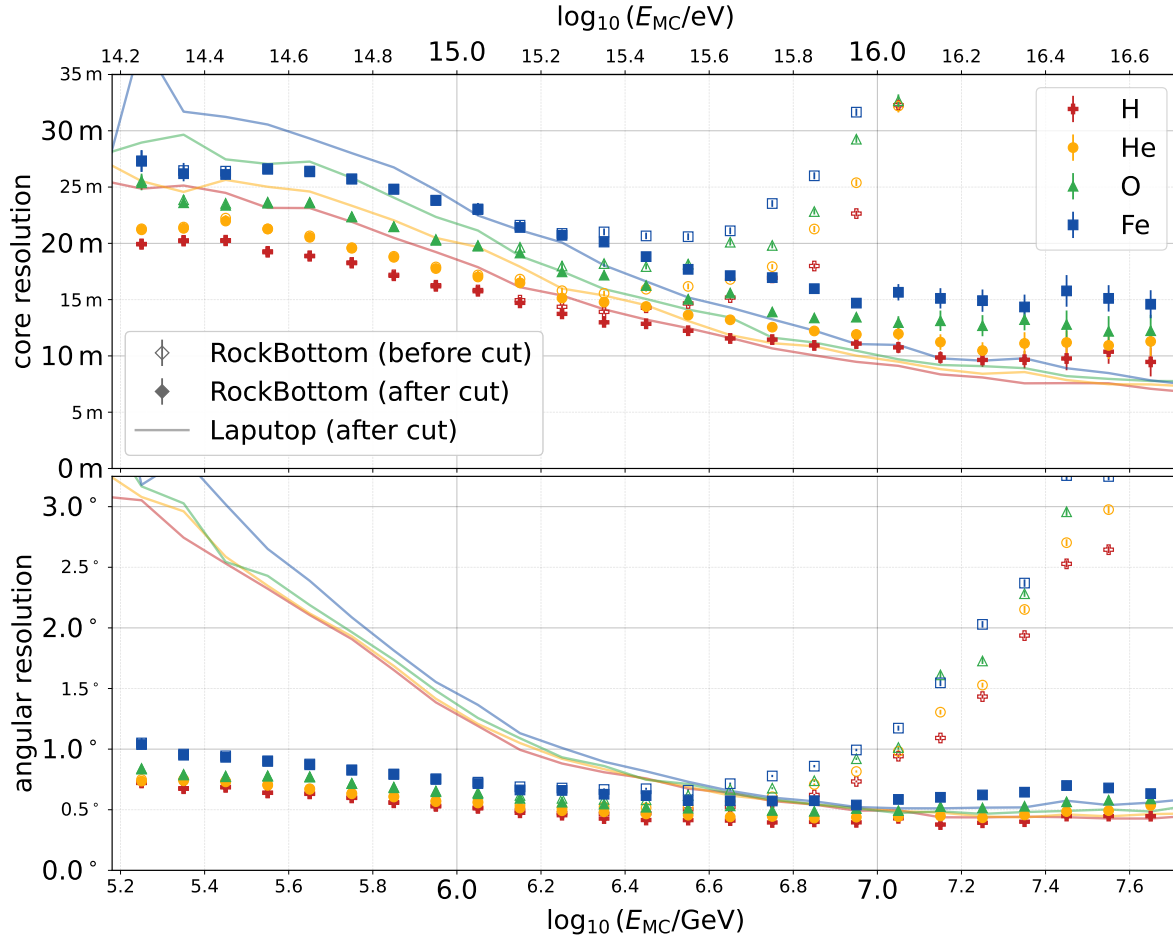
For this analysis, new MC datasets have been produced with primary energies ranging up to 100 PeV ( $\rightarrow$  Tab. 5.1). However, the cosmic-ray spectrum reaches several orders of magnitude higher. In order to set the focus on the region around and below the knee and avoid that measured events with primary energy greater than the maximum simulated

<sup>1</sup>The uncertainty on the  $(100 \times \alpha)$ -percentile  $x_\alpha$  of a quantity  $x$  is estimated as  $\Delta x_\alpha = \frac{1}{f(x_\alpha)} \sqrt{\frac{\alpha(1-\alpha)}{N}}$ , where  $f(x_\alpha)$  denotes the probability density at  $x = x_\alpha$  and  $N$  being the number of samples [118]. For resolution,  $\alpha = \iota/100$ .

<sup>2</sup>Laputop uses a simplified reconstruction algorithm when only three or four HLC stations are present, leading to significantly worse resolution. An improved algorithm for 3-station events has recently been developed in Ref. [99].



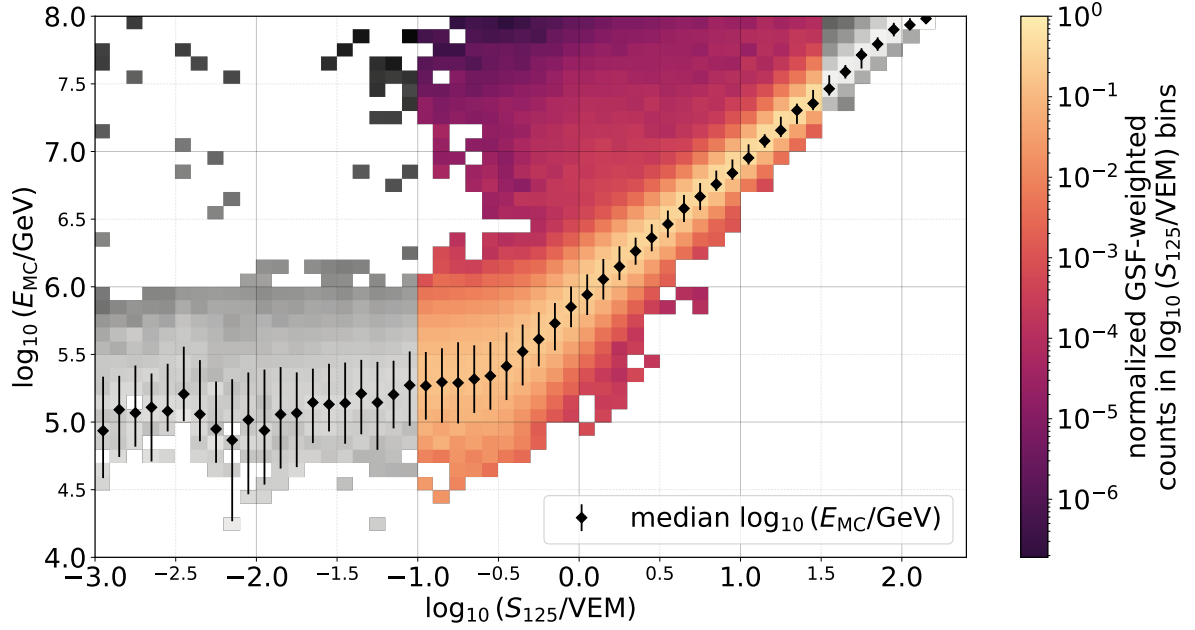
**Figure 7.3:** The GSF-weighted core resolution (*top*) and angular resolution (*center*) of RockBottom in bins of IceTop containment and in-ice containment of the reconstruction. The gray lines represent the quality cut, which removes those events with the worst resolution. **Bottom:** The containment-binned event rate, weighted to the GSF flux model.



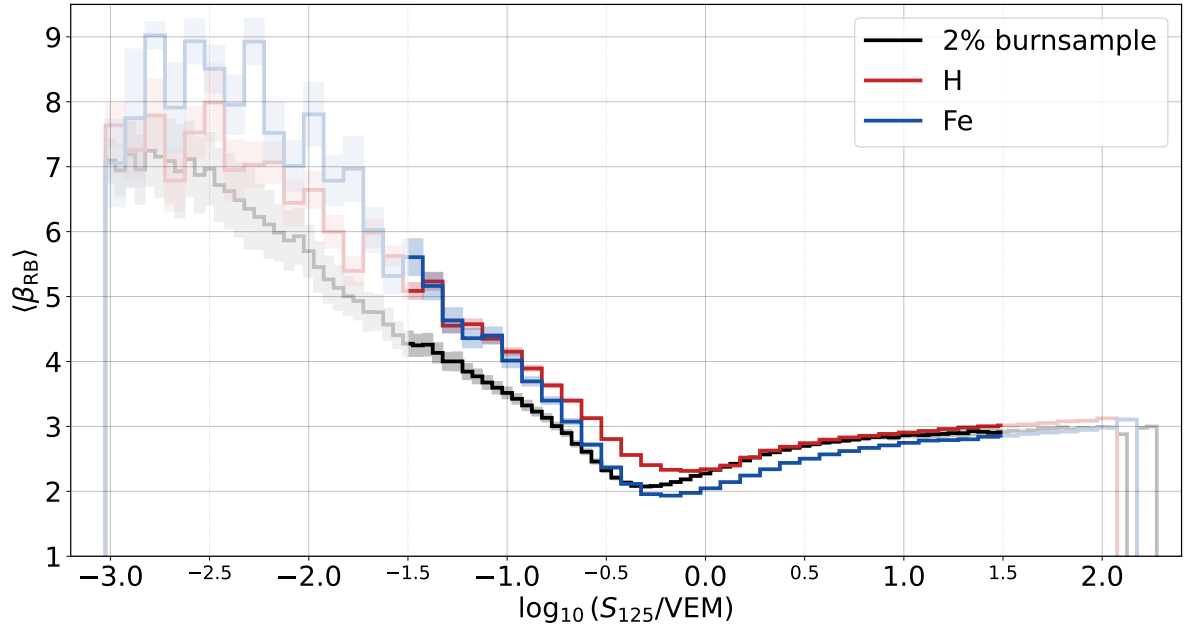
**Figure 7.4:** Effect of the combined IceTop and in-ice containment cut on spatial resolution. Open markers represent the resolution of the RockBottom track after the first three quality cuts but before the containment cut. After the containment cut, the resolution is given by the filled markers. Additionally, the lines demonstrate the resolution of Laputop after identical event selection and quality cuts. **Top:** Shower core resolution. **Bottom:** Angular resolution.

energy find their way into the analysis data, a generous cut on primary energy is required. The RockBottom shower size  $S_{125}$  is a suitable choice for this task, as it correlates strongly with the primary energy ( $\rightarrow$  Fig. 7.5). Thus, an upper cut is applied to remove events with  $\log_{10}(S_{125}/\text{VEM}) > 1.5$  from the dataset. There is a population of high-energy events with a much lower reconstructed  $S_{125}$  than expected. These events have an uncontained true shower core but their reconstructed shower axis intersects the ground inside the 95% containment area, which is why they pass the containment quality cut. RockBottom can be configured to improve the track reconstruction of those events [115] but since their rate is rather low compared to the set of truly contained events, they are not addressed any further.

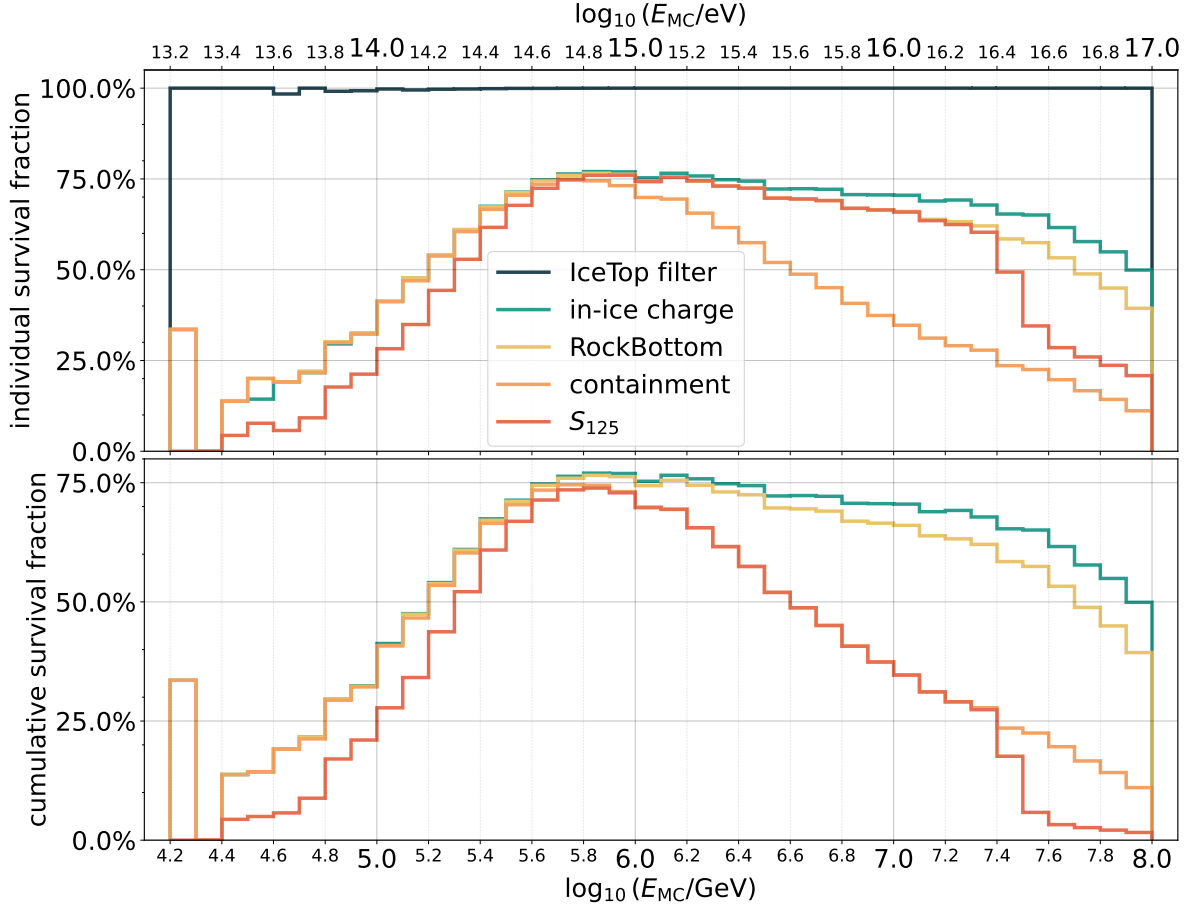
Below a few hundred TeV the linear correlation between shower size and primary energy levels off and the vertical distribution in Figure 7.5 broadens. Furthermore, the dependence of the average LDF slope  $\beta_{RB}$  on  $S_{125}$  features an increasing disagreement between measured data and the simulations, as can be seen in Figure 7.6.



**Figure 7.5:** The correlation between the RockBottom shower size  $S_{125}$  and the true primary energy  $E_{MC}$  with all previously covered cuts (section 7.1–7.4) applied. The  $\log_{10}(S_{125}/\text{VEM})$  range between  $-1$  and  $1.5$  is highlighted. For each column, the GSF-weighted rate is normalized and the median logarithmic energy is shown with a 68% range around it. The linear trend flattens below a  $\log_{10}(S_{125}/\text{VEM})$  of about  $-0.5$  or a few hundred TeV primary energy.



**Figure 7.6:** The mean LDF slope  $\langle\beta_{RB}\rangle$  as a function of the shower size from RockBottom with the previous cuts applied. Since  $\beta_{RB}$  is expected to be smaller on average for heavier primaries, the distribution of measured data (black) is expected to lie between the lines for simulated proton (red) and iron (blue) air showers. The  $\log_{10}(S_{125}/\text{VEM})$  range between  $-1$  and  $1.5$  is highlighted.



**Figure 7.7:** *Top:* The fraction of events surviving each individual quality cut as a function of primary energy. *Bottom:* For the cumulative cut efficiency, each line represents the survival fraction of the respective quality cut, combined with all previous cuts in the sequence.

Since those two unwanted effects challenge a sensible reconstruction of the energy and mass of primary cosmic rays, a second cut on shower size is deployed to remove low-energy events, which are of little use for this analysis. Only events with  $-1 \leq \log_{10}(S_{125}/\text{VEM}) \leq 1.5$  will be used for the analysis.

## 7.6 Survival Rates

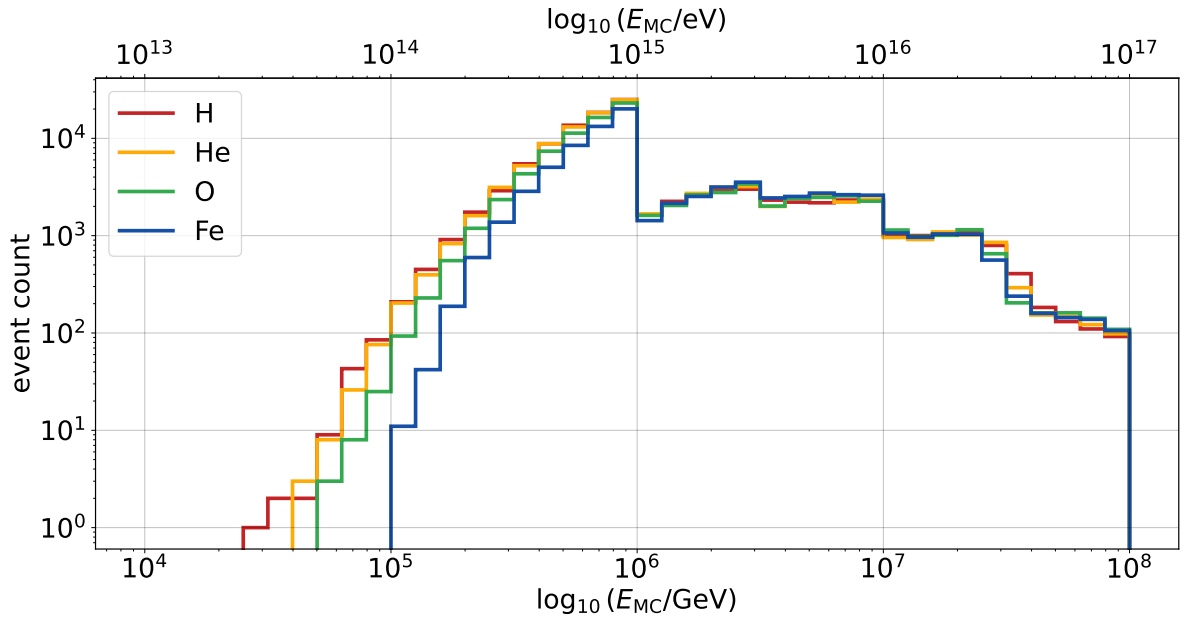
As noted at the beginning of this chapter, the purpose of quality cuts is to clean the dataset by removing incomplete, misreconstructed or otherwise unsuitable samples, which are not supposed to enter the present analysis while still maintaining a decent amount of data in order to limit statistical fluctuations. The previous sections focused on the positive influence provided by the individual event selections. The loss of events in the MC datasets as a function of primary energy is shown in Figure 7.7. The overall reduction in the event rate can be estimated by weighting to a realistic flux model such as GSF ( $\rightarrow$  Tab. 7.1).

For the sake of completeness, one more quality cut should be mentioned. If an event has NaN charges ( $\rightarrow$  section 6.1) in both HLCs of a triggered station, the event is excluded. However, this is extremely rare to happen<sup>3</sup>. The effect of this cut is thus negligible and will not be

<sup>3</sup>Out of over  $8.2 \cdot 10^5$  simulated events at Level 4 and below 100 PeV only three events do not pass this cut.

**Table 7.1:** The GSF-weighted survival fraction of each quality cut individually and for the combination of the first  $n$  cuts.

quality cut reference link	filter passed <a href="#">7.1</a>	in-ice charge <a href="#">7.2</a>	RockBottom <a href="#">7.3</a>	containment <a href="#">7.4</a>	$S_{125}$ <a href="#">7.5</a>
individually	99.86%	65.62%	65.00%	60.71%	60.93%
cumulatively	99.86%	65.60%	64.85%	60.56%	56.57%

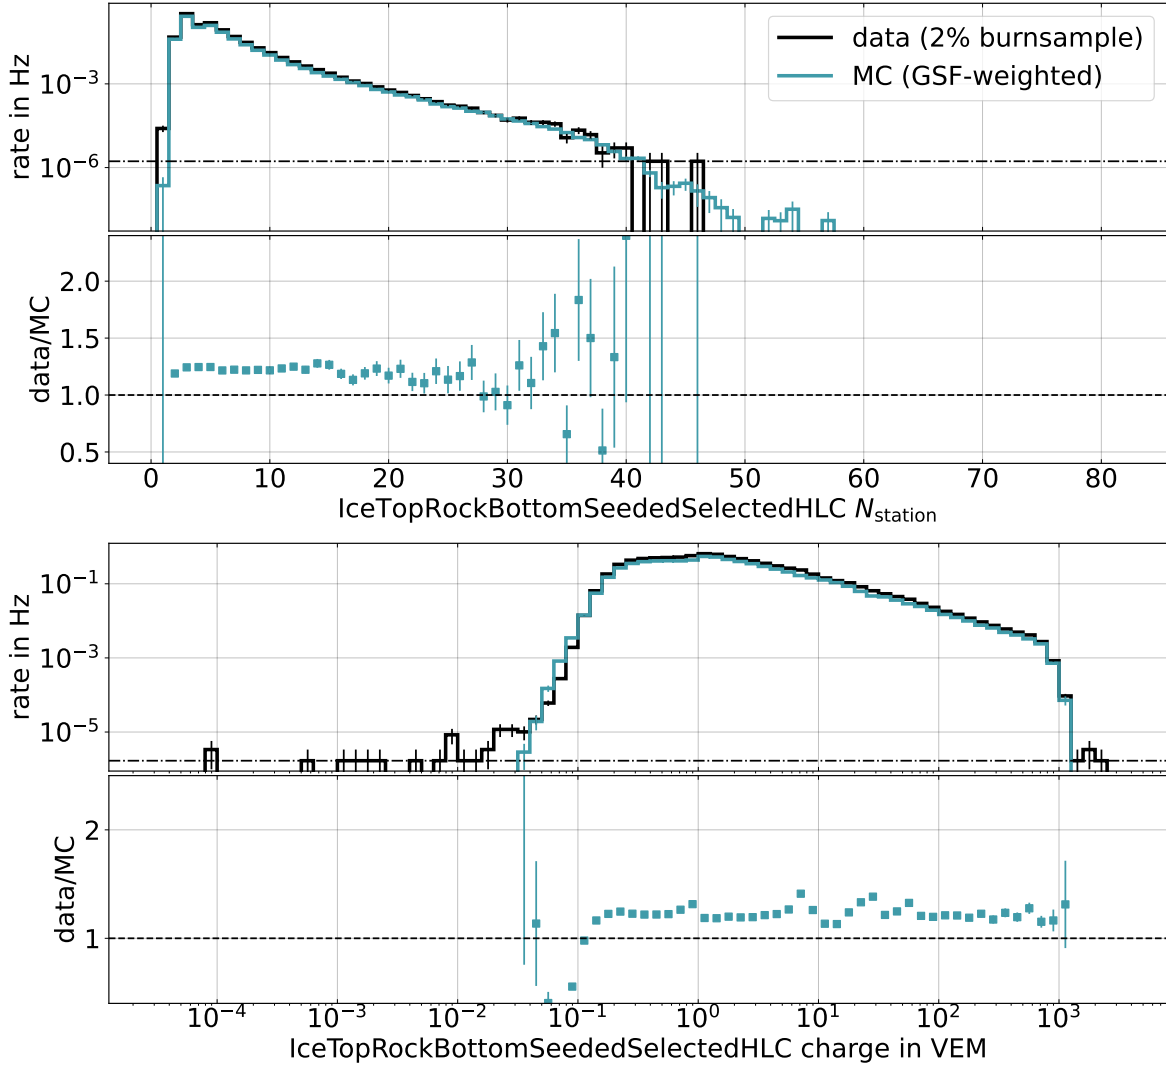
**Figure 7.8:** Distribution of true energy for the four primary types after quality cuts.

considered further. The energy distribution of the simulation set after the application of all quality cuts is shown in Figure 7.8.

Furthermore, there is one last cut that applies only to recorded data. Each data taking run is monitored shortly after its end in order to check for potential problems ( $\rightarrow$  Fig. 4.4). Only runs marked as “good” are considered for the reconstruction of the light spectrum, which, among other things, requires stable trigger and filter rates, as well as a minimum run duration of 10 min. Because of the important role of the in-ice detector in this analysis, it is also necessary that all of the 78 non-DeepCore strings are active.

## 7.7 Data / Monte Carlo Comparison

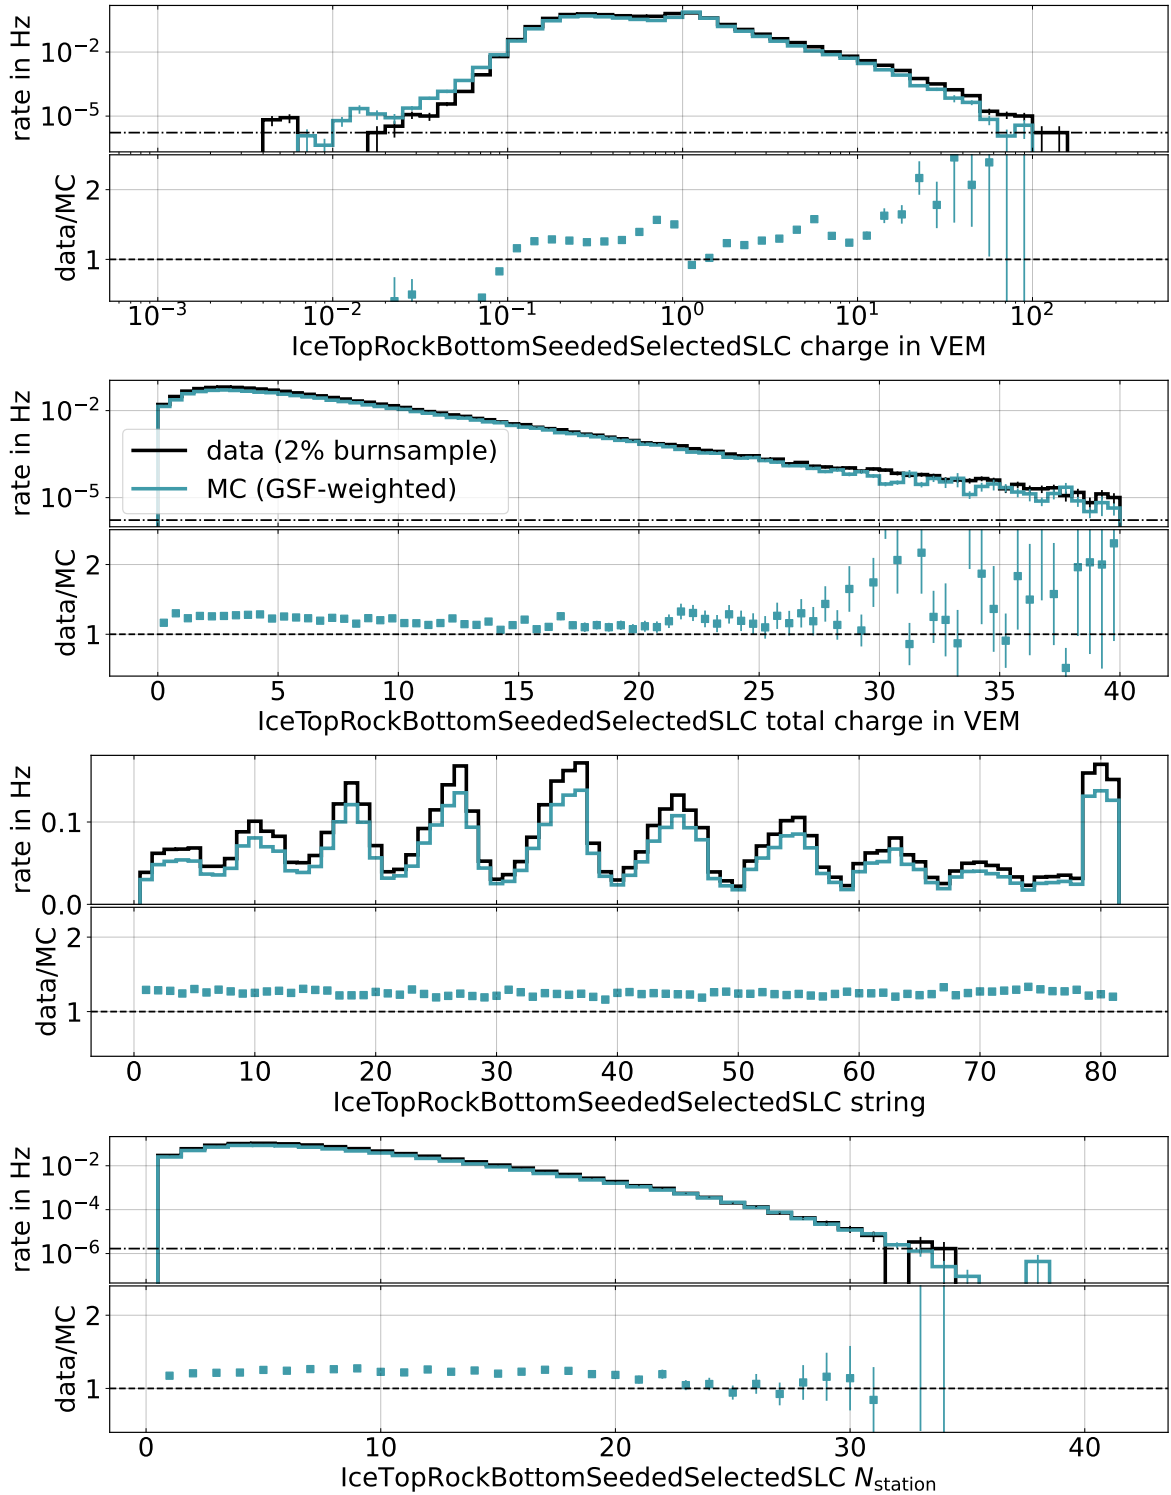
The first iteration of MC verifications has been presented in section 6.5 for Level 1 and 2 distributions. The processing Levels 3 and 4 as well as the application of quality cuts have reshaped the distributions of the HLC and SLC pulse series. The removal of events with a reconstructed  $\log_{10}(S_{125}/\text{VEM}) > 1.5$  leads to a much better agreement in the number of HLC stations between the burnsample data and the Monte Carlo dataset, which has been simulated up to  $10^8$  GeV ( $\rightarrow$  Fig. 7.9). Moreover, the high-level processing and the quality cuts lead to a suppression of very high single charges beyond 1000 VEM as well as extremely low charges  $< 10^{-1.5}$  VEM, resulting in more similar distributions in simulations and data.



**Figure 7.9:** Data / Monte Carlo comparison of *IceTopRockBottomSeededSelectedHLC* pulses. **Top:** Number of hit stations. **Bottom:** Pulse charges. The dash-dotted horizontal line in the upper panels marks the minimum possible rate in the burnsample, that is, a single occurrence in total.

The improvements in data-MC agreements achieved by the high-level processing and quality cuts are especially visible in the SLCs. Figure 7.10 contrasts the charge distribution of Level 4 SLC pulses of data and Monte Carlo. Despite a jump in the rate ratio at around 1 VEM in the individual pulses, this effect seems to be averaged out in the summed SLC charge for each event. Comparing with Figure 6.17, the two advanced distributions of time window cleaned SLC stations match considerably better owing to the reduction of noise SLCs, which significantly contaminated the burnsample data at Level 2. This cleaning of the single-tank hits also leads to a much better agreement in the multiplicity of SLC stations. Despite the overall improvement of the data-MC agreement, the systematic offset in the ratio of about 20% remains. This is due to the flux-model weighting and a mismatch between average real pressure and the simulated atmosphere, as discussed in subsection 6.5.1.

Appendix A contains further MC verification plots at Level 4 after quality cuts for RockBottom core position and direction as well as the background rejection features  $\beta_{\text{COG2COG}}$  and  $\Psi$ .



**Figure 7.10:** From top to bottom: charge distribution, total event charge, station occupancy and number of hit stations for the IceTopRockBottomSeededSelectedSLCs (data and simulation). The dash-dotted horizontal line in the upper panel of some plots marks the minimum possible rate in the burnsample.

## CHAPTER 7. QUALITY CUTS

### Key Messages

*After passing through the full processing, events need to also fulfill the conditions of a set of quality cuts. Those selection criteria remove events that do not pass any IceTop filter, did not deposit charge in the upper part of the in-ice detector, have a failing or obviously inaccurate RockBottom reconstruction, are assigned a  $S_{125}$  either too low to be reliable or too high to participate in the analysis, or have a reconstructed track that is not contained in either IceTop or the in-ice detector. Eventually, the collection of those cuts remove a substantial amount of events. However, the remaining dataset reveals significantly improved agreement between simulated and measured low-level distributions of pulse information and reconstruction parameters.*

# CHAPTER 8

## LIGHT SPECTRUM RECONSTRUCTION

*“Allting är mycket osäkert, och det är just det som lugnar mig.”*

— Tove Jansson, *Trollvinter*

This chapter covers the analysis steps needed to process a collection of cosmic-ray events in order to determine the fraction of the light flux component, which eventually is transformed into a combined spectrum of protons and helium nuclei. Beginning with setting up a neural network and running the inference on it (→ [section 8.1](#)), the model output is investigated to quantify the bias and resolution of the primary energy estimation and to develop a calibration method for the light fraction in a set of cosmic-ray events, which can be done in two different ways (a conservative one with neutral assumptions, the other one seeded with a composition model; → [section 8.2](#)). Systematic uncertainties (→ [section 8.3](#) and [8.5](#)) and correction factors for the reconstructed light flux (→ [section 8.4](#)) are discussed. Recovery tests in [section 8.6](#) demonstrate the ability to reconstruct an injected spectrum of protons and helium.

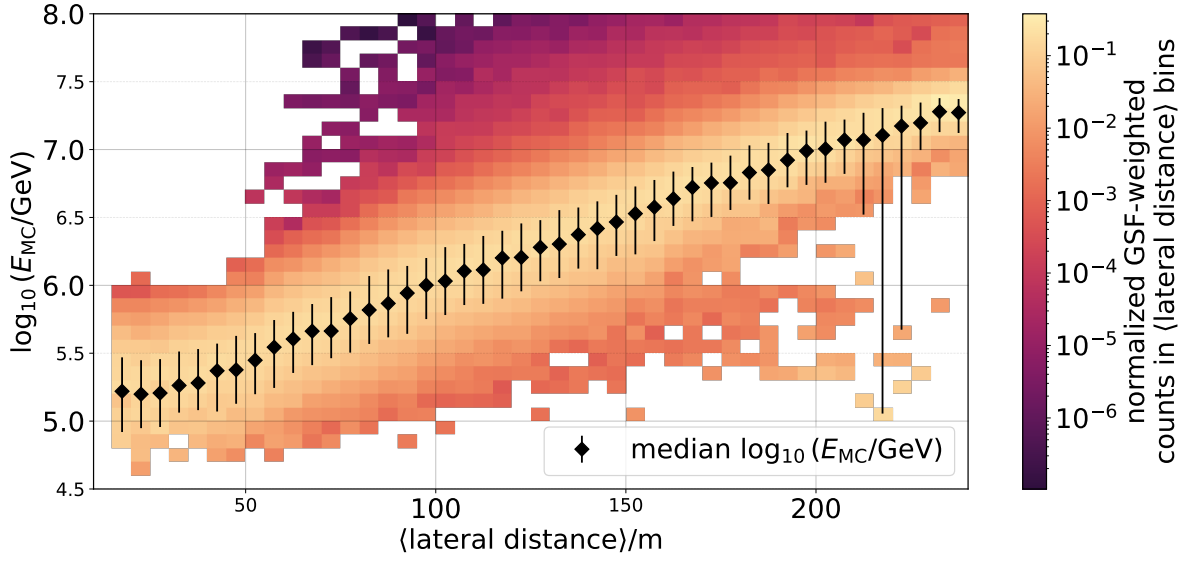
Since the analysis method is based on Monte Carlo simulations, all fits, calibrations and comparisons performed in this chapter depend on the hadronic interaction model that has been used in the [CORSIKA](#) simulation. The analysis is developed on Sibyll 2.3d simulations and the main final results will be based on this. However, the spectrum reconstruction can be repeated with simulations of alternative hadronic interaction models. A comparison of the resulting fluxes obtained using Sibyll, [QGSJet](#) and [EPOS](#) will be provided in the next chapter as well as the individual spectra in the appendix.

### 8.1 The Neural Network Model

The measurement of the proton and helium spectrum of cosmic rays requires the reconstruction of two essential attributes of the detected events: their primary energy and their primary mass. This task is performed by a neural network (→ [section 3.2](#)), which is fed with information on every single event that passes the quality cuts described in the previous chapter. Those features include raw quantities that are taken directly from the recorded pulses as well as outputs from the advanced RockBottom air-shower reconstruction.

#### 8.1.1 Model Inputs

The [NN](#) model in use takes a total of 15 input features. Those are categorized into energy features and mass features. The former could be used for a standalone energy regression, whereas the reconstruction of primary mass requires the combination of feature types to deconvolve the energy-mass degeneracy.



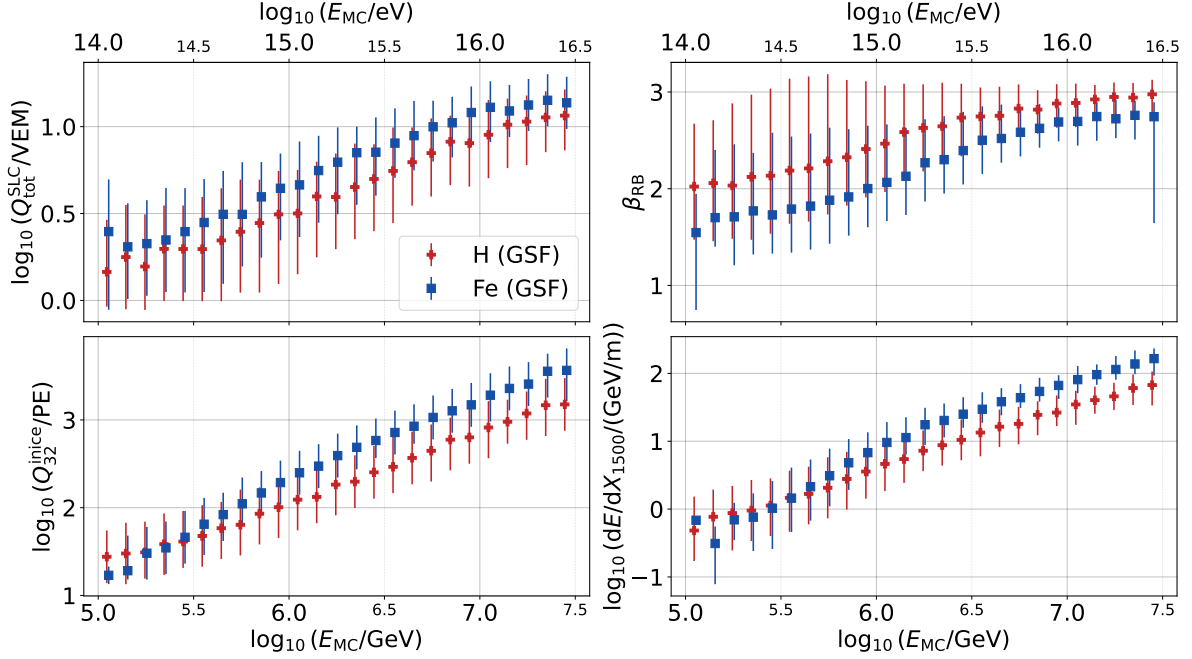
**Figure 8.1:** The correlation between the average lateral distance of HLC pulses from the shower core and the true primary energy. The GSF-weighted event rate is normalized in each bin of  $\langle \text{lateral distance} \rangle$  and the median MC energy is indicated, which follows a linear trend.

### Energy Features

As the primary energy-sensitive feature for events at full trigger efficiency, the shower size  $S_{125}$  from RockBottom’s best fit [LDF](#) is taken. Since the relation between primary energy and  $S_{125}$  changes with the inclination of the event [70], the zenith angle  $\theta_{\text{RB}}$  of the reconstructed RockBottom track is added as an input feature as well.

With decreasing energy, the linear relation with the shower size flattens gradually ( $\rightarrow$  Fig. 7.5). In order to provide the NN with an approximation of the lateral signal distribution of smaller events, a set of low-level information of the shower footprint is supplied in the form of six more inputs. For each station with [HLC](#) hits, the pulse charges are summed up and sorted in descending order. The three largest values  $Q_1$ – $Q_3$  are given as input to the network together with the horizontal distances  $r_1$ – $r_3$  from the RockBottom shower core to the respective stations where the location of a station is defined as the charge-weighted position of the triggered [DOMs](#). In case that only two IceTop stations have HLC pulses, the slot for the third highest station charge and its distance to the shower core are filled with a dummy value. The set of low-level energy features is completed with the total HLC charge summed over all active stations ( $Q_{\text{tot}}^{\text{HLC}}$ ) and the averaged lateral distance of HLC pulses to the reconstructed shower core. Figure 8.1 shows how the mean lateral distance of an event has a linear relation with the primary energy that extends to lower energies than  $S_{125}$  ( $\rightarrow$  Fig. 7.5).

The model is expected to use primarily the shower size for energy reconstruction of large events and relies on the low-level HLC inputs for smaller events. In order to support the network in making this distinction and learn that the dummy values in 2-station events are no physical quantities, but rather should be ignored, a categorical feature is added. If an event passes the 5-station filter ( $\rightarrow$  [subsection 4.2.2](#)), this feature holds the number 5. In case of a smaller event passing only the 3-station or 3-station InFill filter, the “highest filter” input is 3.



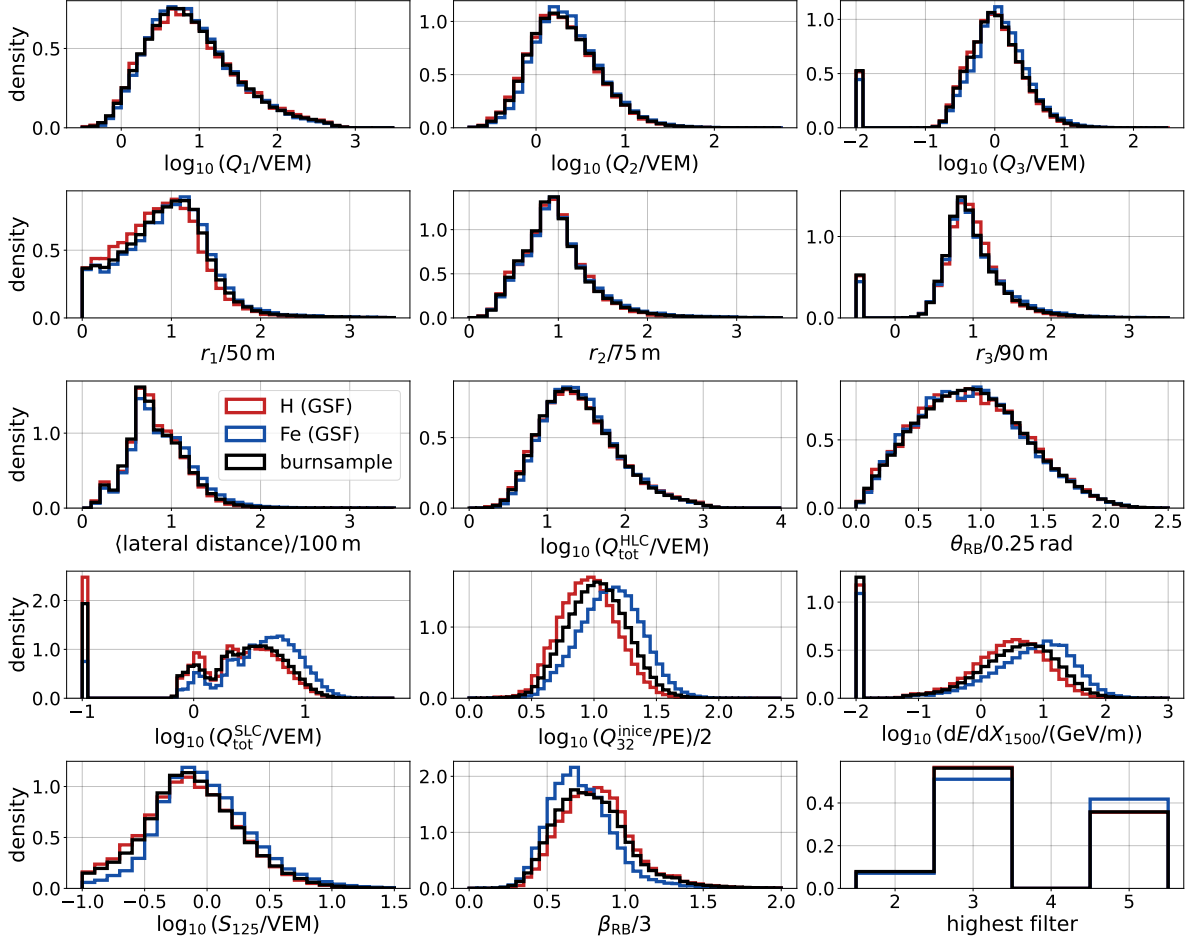
**Figure 8.2:** Energy dependence of the four mass features separately for proton and iron events. Markers represent the binned median while the bars indicate the 68% around the median. Red and blue markers are slightly shifted relative to each other for better visual separation.

The remaining events have only two HLC stations, which is reflected in the filter feature accordingly. The reconstructed zenith angle and the filter information are considered indirect energy features, as they help the network interpret the main energy features.

### Mass Features

While the energy predictors all probe the electromagnetic shower component at the surface, the mass-sensitive features are gathered from both IceTop and the in-ice array. The slope  $\beta_{RB}$  of the RockBottom LDF is influenced by the primary mass and can therefore help distinguish heavy from light cosmic rays. HLCs are essential for the energy reconstruction as they are created from electromagnetic shower particles, **SLCs** on the other hand are generally triggered by single GeV muons. Adding up their charge per event provides a proxy for the number of low-energy muons, which in turn is an indication for the primary type. Only SLC charges between 0.7 **VEM** and 5 VEM are considered as those are most likely to be of muonic origin [119]. If an event has no SLC pulse in this range, a dummy value is used for  $Q_{tot}^{SLC}$ .

A similar logic can be applied to the TeV muons that are detected deep in the ice. However, as described by Eq. (2.5), with primary energies below 1 PeV the difference in muon multiplicity starts to vanish (and flip) for muons exceeding about 1 TeV ( $\rightarrow$  Fig. 2.5). For that reason, not only the established mass-sensitive features  $dE/dX_{1500}$  [70] is employed for this analysis but also the charges of HiveSplitPulses in the top 32 layers of DOMs are summed up to create the last network input  $Q_{32}^{inice}$ . The restriction to the top half of the in-ice DOMs allows to reduce the influence from the highest-energy muons, which are able to propagate through the entire array to the bottom without decaying. DeepCore DOMs are excluded from the in-ice charge

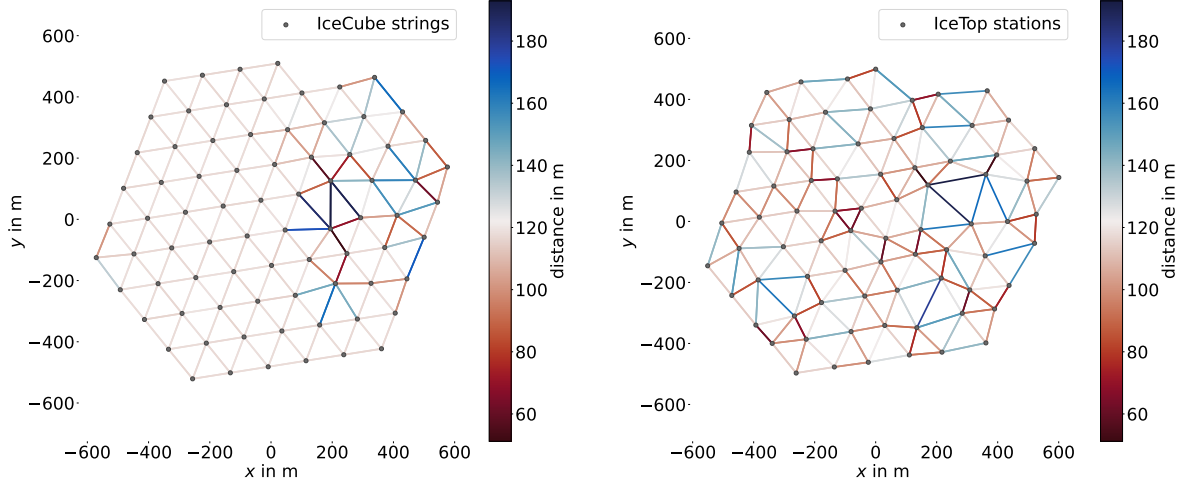


**Figure 8.3:** The distributions of the 15 NN input features. For simulated events, proton and iron distributions are plotted separately and weighted according to the GSF flux model. In order to facilitate the comparison of the shapes among proton and iron as well as with the burnsample distributions, all histograms are normalized.

sum as they provide a higher density of optical sensors, which could lead to a bias towards a higher total charge in the detector center.

An overview of the four mass-sensitive network inputs is given in Figure 8.2. There, the degeneracy of primary energy and mass becomes apparent. The distinction of primary types is generally only possible in combination with information on the primary energy. The discriminating power of the mass features diminishes towards decreasing primary energy. However, the point where the difference between proton and iron events vanishes is reached at distinct energies, which is why the combination of all features is expected to enhance the ability of the NN to separate light from heavy cosmic rays.

Figure 8.3 gives an overview of all 15 input features of the neural network. As part of the preprocessing ( $\rightarrow$  section 3.1), most features are transformed by taking the logarithm, scaling by some constant factor or both. GSF-weighted proton and iron events are compared with the 2% burnsample dataset. The distributions of energy features are virtually mass-independent, whereas systematic shifts can be observed in the four mass inputs with the recorded data between the two extreme primary types.



(a) Inter-string spacing of the IceCube optical array. (b) Distance between the centers of IceTop stations.

**Figure 8.4:** Spacing between neighboring in-ice strings and IceTop stations, indicated by the color of the connecting lines. The additional DeepCore strings and InFill stations are excluded.

### 8.1.2 Network Architecture

The chosen type of machine learning model needs to be suitable for the structure of both the input features and the targets. Although tree-based methods excel in separating dummy values in the inputs, as they are present in four of the identified features, they can generally only be trained for a single target. However, because of the energy-mass degeneracy in the air-shower observables, it is favorable to reconstruct primary energy and mass simultaneously. For this work, the entire ML infrastructure uses the Python library *PyTorch* [120], which provides all tools for building, training and comparing NN models.

The regular arrangement of the in-ice strings suggests the application of a CNN as it has been used in previous IceCube analyses [121, 122]. However, DOMs are almost equidistantly spaced only in most parts of the deep in-ice array (Fig. 8.4a). The offset of IceTop stations from the string location causes a rather irregular grid of detectors at the surface as can be seen in Figure 8.4b. In addition, the InFill stations 79, 80 and 81 introduce even more irregularities in the surface grid, which would require a separate treatment of those three stations. This, however, undermines the idea of “neighborhood”, which is part of the core concept of convolutional neural networks. An alternative approach using graph neural networks has been explored in Ref. [107].

Therefore, instead of a CNN, this work builds on a multilayer perceptron. This brings the additional advantage that high-level features, such as the reconstructed shower core and LDF parameters from RockBottom, can be easily fed into the model.

The two network targets are

- the logarithmic primary energy  $\log_{10}(E/\text{GeV})$  and
- the *proton score*  $\omega_p$ , which is 1 for proton-induced events and 0 for all other events.

Even though this target definition is designed for the selection of protons, the network's  $\omega_p$  output works just as well for the separation of proton and helium events from heavier primaries. Despite its name, the proton score will be used for reconstructing the fraction of light primary nuclei in the cosmic-ray flux. In this work, the term “light” is used synonymously to the combination of proton and helium primaries.

The network architecture is determined through hyperparameter optimization with the ML tracking tool *weights & biases* [123]. The set of free hyperparameters consists of the number of hidden layers (the *network depth*), their individual *width* (that is the number of nodes), the activation function and the initial learning rate. A Bayesian search algorithm narrows down the space of hyperparameter combinations that lead to a minimal total loss. The optimized network configuration has the shape 15-150-40-25-2; that is, the 15 input features are propagated through three hidden layers with narrowing width culminating in the two output nodes for energy and mass estimation. Each hidden layer is followed by a batch normalization layer. As an activation function the “sigmoid linear unit”

$$\text{SiLU}(x) = x \cdot \sigma_{\log}(x)$$

is used. The initial learning rate  $\eta$  is set to  $10^{-2}$ .

### 8.1.3 Training

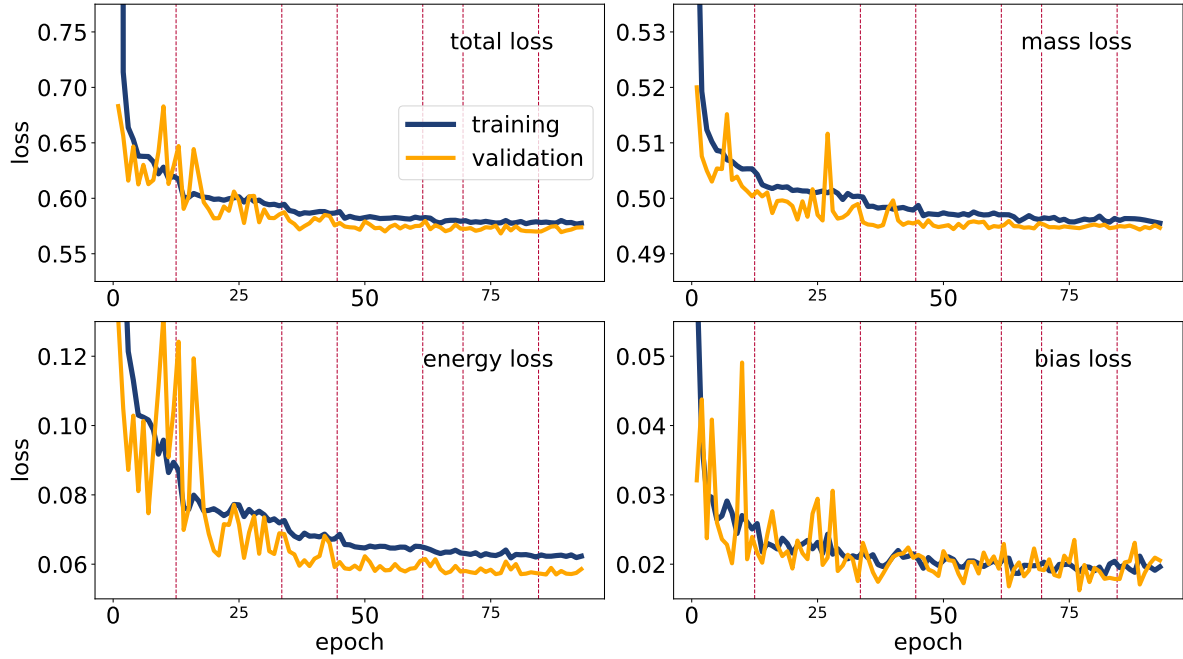
The MC dataset is shuffled and divided into three sets at a ratio of 59.5 : 10.5 : 30 for training, validation and testing. The NN reads in the input tensors in batches of 256 events at a time. Every training epoch consists of a learning step and a monitoring step.

On each batch of training data, the model returns its  $256 \times 2$  outputs, which are forwarded to the loss calculation. The total loss has three contributions:

- mass loss: BCE (Eq. (3.4)) scores the classification of events into the two classes “proton” and “not proton” using the  $\omega_p$  target node
- energy loss: the regression of logarithmic primary energy is assessed using the FVU (Eq. (3.3)).
- bias loss: inspired by Ref. [124], one more loss term is introduced, aiming to restrain the difference in energy bias  $\sum_{i=1}^n (\log_{10}(\hat{E}_i) - \log_{10}(E_i))/n$  among the mass groups<sup>1</sup>. Since the largest disagreement is expected between proton and iron events, the loss is defined as  $|\text{bias}_H - \text{bias}_{Fe}|$

The contribution to the loss calculation by each individual event is weighted by  $E^{-2}$ . This is a compromise between the omission of loss weighting, which would favor high-energy events which are generally easier to reconstruct their energy and mass, and weighting with a more realistic spectral index such as  $\gamma = 2.7$ , which increases the risk of neglecting high energies

<sup>1</sup>The term  $\log_{10}(\hat{E})$  denotes the predicted logarithmic energy from the energy target node, not the logarithm of predicted energy.



**Figure 8.5:** Training (blue) and validation loss (orange) curves for the total cost function and the three loss components. Vertical red lines indicate drops in the learning rate.

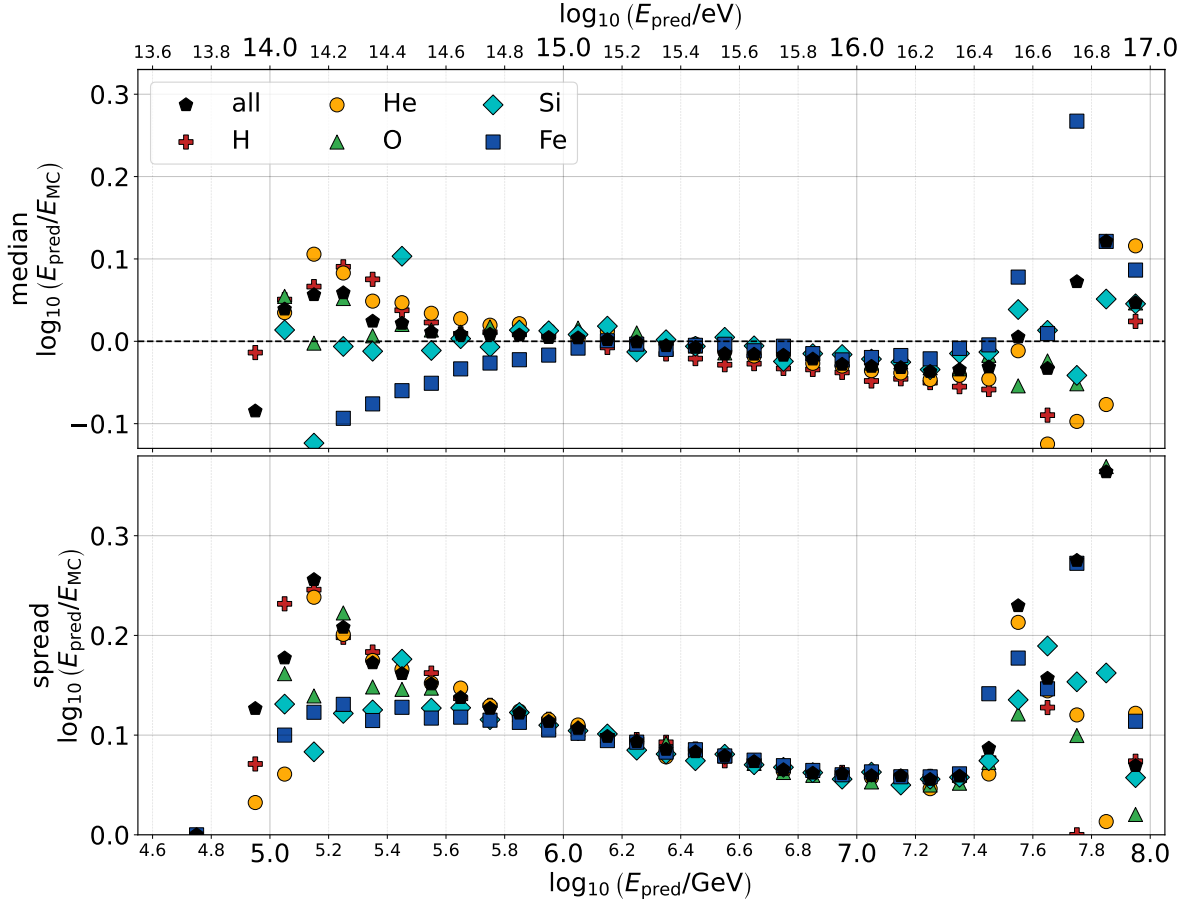
at the expense of optimizing for sub-PeV energies. The NN's cost function  $C(f_{\theta}(x), y)$  is the sum of mass loss, energy loss and bias loss. To complete the training iteration on the batch, the gradients  $\nabla_{\theta} C$  are determined and the backpropagation step is performed, which updates the network parameters before the next batch is processed.

In order to monitor the training process in real time and identify overtraining, after every training epoch, the model is temporarily set to inference mode; that is, gradients are not tracked and the weights and biases are frozen. Similarly to the training step, the validation set is loaded in batches and the  $E^{-2}$ -weighted total loss is calculated on the network outputs.

Finding a minimum of the cost function in a high-dimensional parameter space can be challenging. Convergence of the loss can be facilitated by starting the training with a rather high learning rate  $\eta$  and reducing it every few epochs. Here, a learning rate scheduler has been implemented, which divides the learning rate by  $\sqrt[3]{10}$  as soon as the total loss has not reached a new minimum for seven consecutive epochs. The training comes to an end when  $\eta$  falls below  $5 \cdot 10^{-5}$ . Figure 8.5 shows the total cost as a function of training epoch as well as the loss curves of its three contributions for both the training and the validation set. The initially large loss falls rapidly during the first few epochs. Subsequent reductions of the learning rate lead to drops of the losses but in the end phase no further improvement is achievable and  $\eta$  keeps shrinking until it reaches the stopping criterion, thus concluding the training.

#### 8.1.4 Model Evaluation

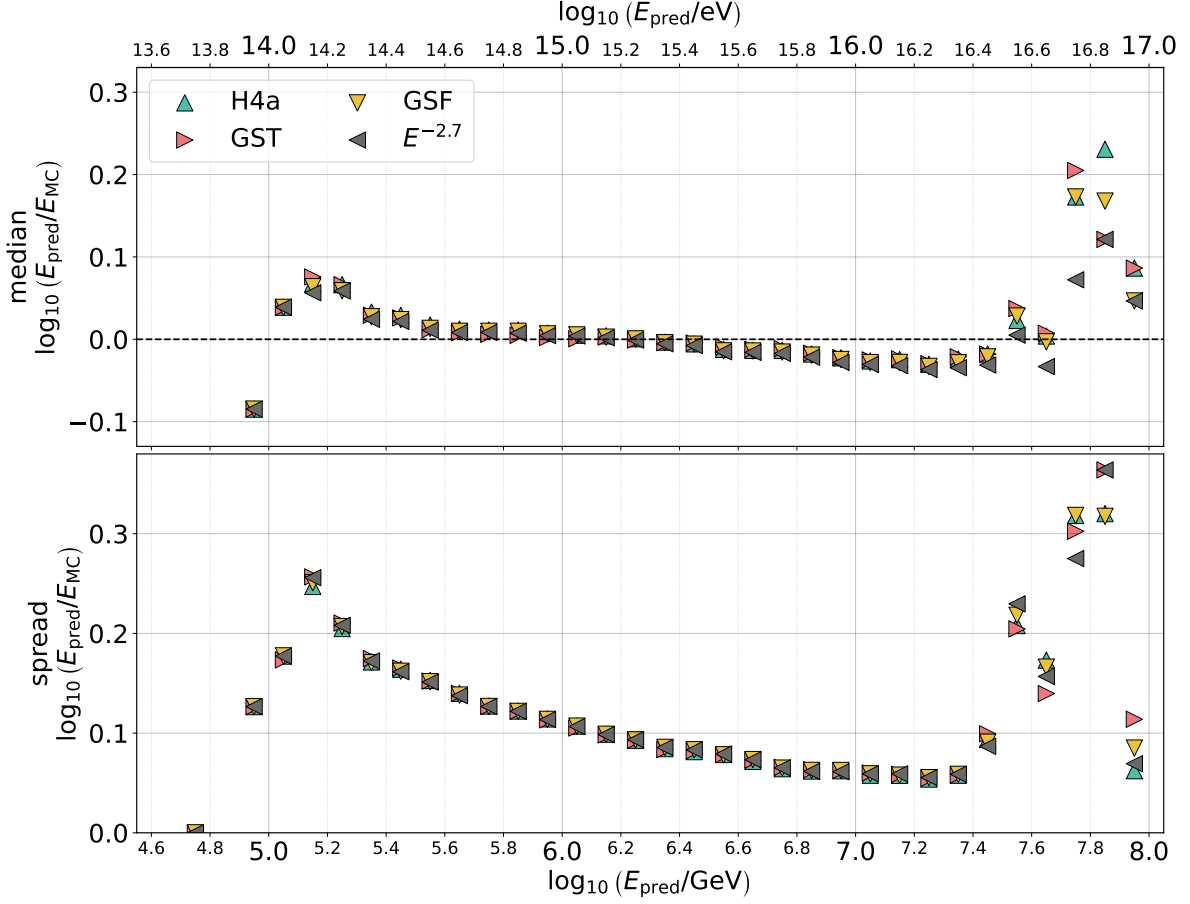
With the NN being trained, the performance of its capability in predicting the primary energy and type of a cosmic-ray event needs to be investigated.



**Figure 8.6: Top:** Energy bias, represented by the median of  $\Delta \log E$  for five different primary types and the merged set of all primaries combined. **Bottom:** Energy resolution. The calculation of the spread requires the correction of the energy bias, which for the individual primaries is the respective bias in the panel above. While this gives an accurate impression of the width of each  $\Delta \log E$  distribution, the all-particle bias will be used for bias correction in this analysis. Bias and resolution are calculated with an  $E^{-2.7}$  weighting and binned in  $\log_{10}(E_{\text{pred}})$ . Only test data and the silicon simulations are used.

### Energy Regression

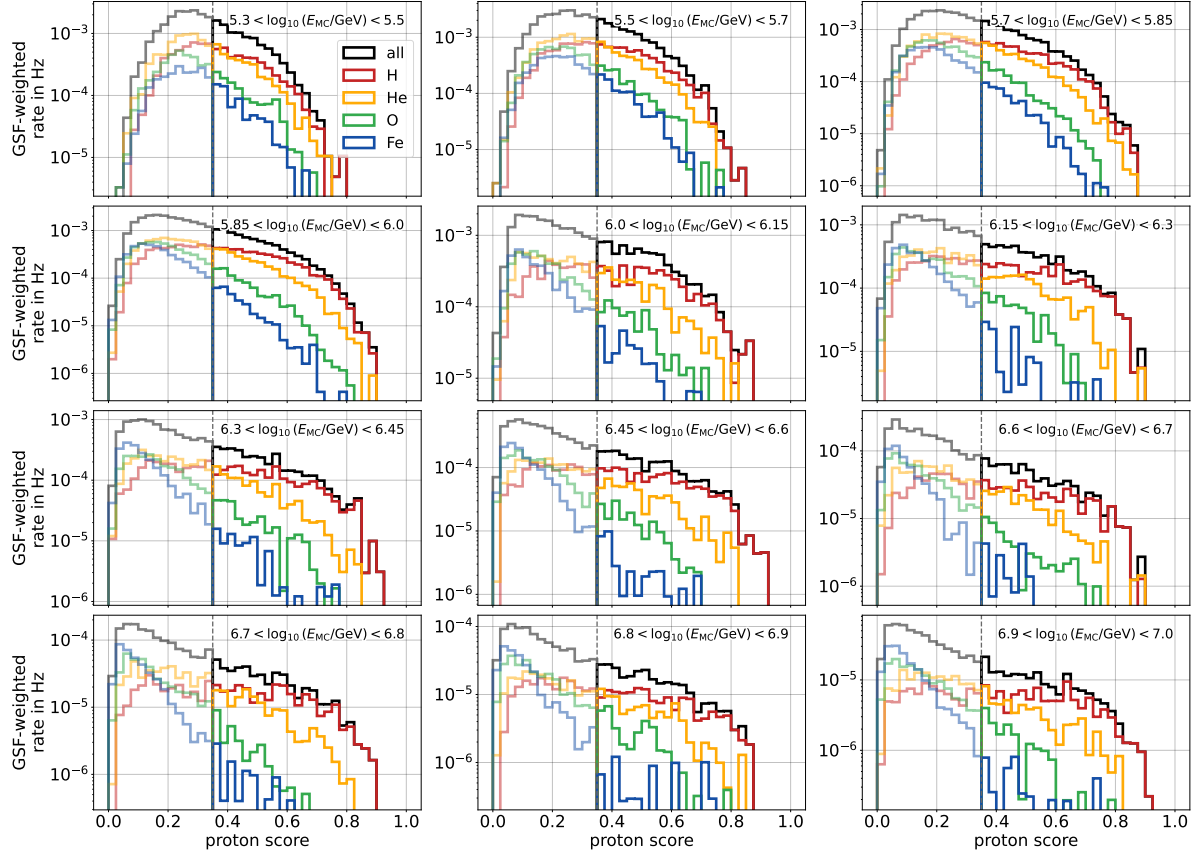
Even though energy loss and bias loss have been reduced as much as possible by selecting an appropriate model architecture and training of the NN, systematic shifts of the predicted logarithmic energy with respect to the true energy can still occur. The top panel of Figure 8.6 gives an impression of the energy bias by plotting the median  $\Delta \log E = \log_{10}(E_{\text{pred}}) - \log_{10}(E_{\text{MC}})$  in bins of  $\log_{10}(E_{\text{pred}})$ . At a few PeV primary energy, the bias almost vanishes and shows virtually no mass dependence. Beyond this energy range, the bias deviates from zero and differs among primary types. This effect is especially noticeable for iron-induced showers below 1 PeV. Figure 8.6 also includes an all-particle bias, which is weighted with the assumption of an  $E^{-2.7}$  spectrum with 20% fraction of each primary. Such an all-particle energy bias estimation will be used to reduce the errors of the NN's energy prediction. The expression  $\log_{10}(E_{\text{pred}})$  is used for the energy output of the network, whereas  $\log_{10}(E_{\text{reco}})$  refers to the bias-corrected quantity.



**Figure 8.7:** Energy bias (top) and resolution (bottom) for the primaries H, He, O and Fe combined and weighted according to composition models as well as the scenario of equal flux fractions with an  $E^{-2.7}$  spectrum.

Energy resolution describes the width of the  $\Delta \log E$  distribution. Still, after the bias correction, this spread can cause misallocation of events into wrong energy bins. It is therefore necessary to find a suitable balance between a sufficiently large bin size to limit bin migration on one hand and maintaining a meaningful segmentation of the energy spectrum needed for an energy-binned analysis on the other hand. Figure 8.6 shows in the bottom panel the energy-binned spread of the distribution of  $\Delta \log E$ , which is defined as the  $\iota$ -percentile of  $|\Delta \log E - \text{bias}|$ . The resolution follows a general trend of improvement with increasing energy but suffers and fluctuates above 30 PeV due to reduced statistics, partly caused by the  $S_{125}$  quality cut. The low number of training samples is also responsible for the worse energy resolution below approximately 400 TeV (recall Fig. 7.8). The fading linearity between primary energy and the e.m.-sensitive features contributes further to the decline in energy resolution at low energy.

In the further course of this analysis, both the weighting with the GSF flux model and the  $E^{-2.7}$  spectrum with equal contributions of the four main mass groups will play a substantial role. Figure 8.7 compares the all-particle bias and resolution under those two flux assumptions as well as weighted to the H4a and the GST composition models. Up to 30 PeV, the metrics for all four scenarios agree well. Based on this study, the investigated energy range is segmented



**Figure 8.8:** Distributions of the proton score NN output for the test dataset, weighted to GSF, binned in true energy. The four distributions of individual mass groups overlap largely but they are offset with respect to each other. The vertical line represents the classification threshold at 0.35 for estimating the light fraction.

into logarithmic energy bins of size 0.2 from  $10^{5.3}$  GeV to  $10^{5.7}$  GeV, 0.15 in the range  $10^{5.7}$  GeV to  $10^{6.6}$  GeV and 0.1 between  $10^{6.6}$  GeV and  $10^{7.0}$  GeV.

### Mass Classification

The second part of the examination of the network output covers the proton score target. As per design, it is 1 for proton-induced events and 0 otherwise. However, the distribution of the proton score response to the test dataset is not bimodal but rather resembles an exponential distribution at high energy that transitions towards a Gaussian-like shape at low energies ( $\rightarrow$  Fig. 8.8). Broken down by primary type, however, shifts in the  $\omega_p$  distributions become visible despite the overlap, which is a legacy of the overlapping mass features ( $\rightarrow$  Fig. 8.2). Selecting events with a score exceeding a certain threshold yields a subset that is contaminated by showers initiated by heavy CRs but is dominated by light primaries. Due to the energy-dependent discrimination power of the mass-sensitive network inputs, the overlap is largest at low energies and shrinks with increasing primary energy. The next section describes a method that exploits this systematic shift of the proton score in order to estimate the fraction of light cosmic rays in any given set of events.

## 8.2 Light Fraction Calibration

The proportion of light primaries, that is, protons and helium nuclei, in a dataset cannot be determined directly from a cut on  $\omega_p$ . Classifying all events with an assigned proton score above a certain threshold  $c_\omega$  as “light” leads to a considerable amount of misclassifications in both directions, that is, false negatives and false positives. For a given mixed set of events, one value of  $c_\omega$  can be identified, which yields  $FP = FN$  and thus results in the correct fraction of light events as the misclassifications cancel out. However, because of the complex shape of the proton score distributions ( $\rightarrow$  Fig. 8.8), such a threshold will generally not work for any other set of events. Choosing a threshold sufficiently high to achieve a pure subset without false positives does also not represent a solution since the small amount of those very proton-like events is subject to relatively large statistical fluctuations. When scaling up the light fraction to account for the overwhelming amount of false negatives, those uncertainties grow accordingly.

Instead, an energy-independent threshold value of  $c_\omega = 0.35$  is set<sup>2</sup> and the relation between the fraction of events with a score  $> c_\omega$ , or the *predicted fraction*, and the *true fraction* of light events is studied in detail. This relation can then be used to develop a calibration to convert “raw” predicted fractions into corrected fractions, which approximate the true light fraction. For this calibration method, two variants are developed that differ at several points. They are referred to as the *agnostic* and the *GSF-informed* versions. The former is based on neutral assumptions with a conservative definition for the limits of what is physically possible, whereas the latter is oriented towards a realistic CR composition model. The concept of fraction calibration will be described in detail for the agnostic scenario first. Differences in the alternative GSF case will be pointed out afterwards.

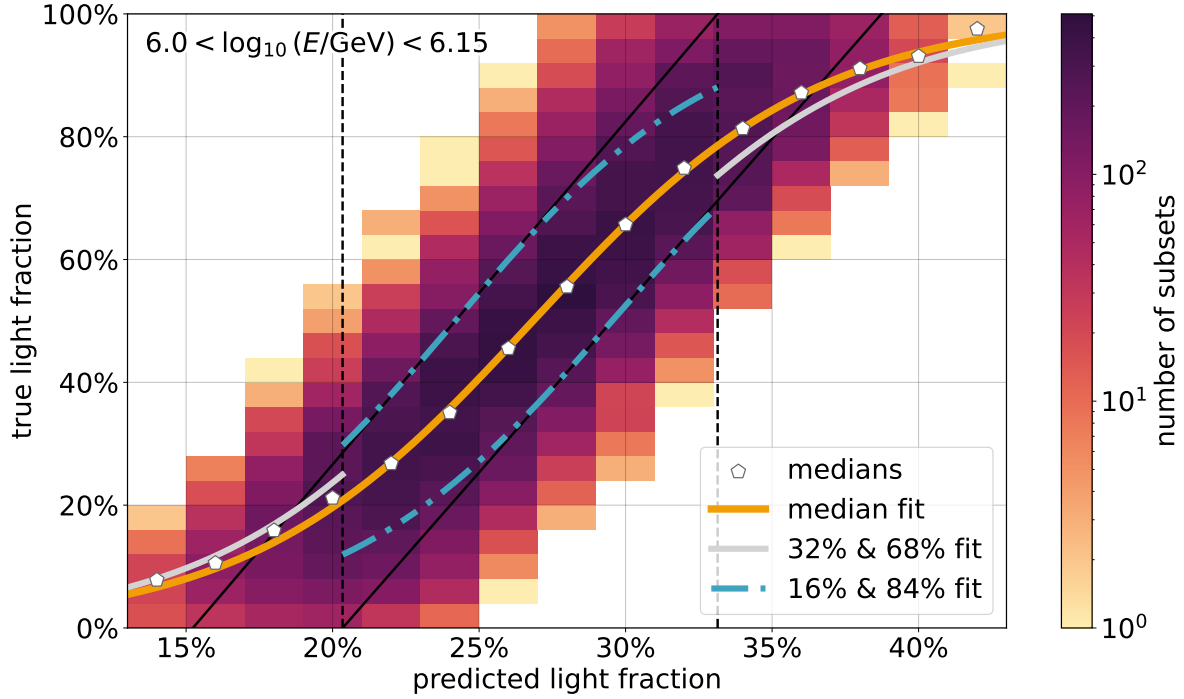
### 8.2.1 Method 1: agnostic

In a first step, the test dataset (30% of the full simulated dataset) is passed through the neural network. The energy predictions are bias-corrected using the binned all-particle energy bias and all events that are assigned a  $\omega_p > c_\omega$  are labeled as “predicted light” events. For every bin of true primary energy, a subset of random size and with a random true fraction of light events is created. Those subsets are merged and subsequently binned in  $\log_{10}(E_{\text{reco}})$ . For every bin of reconstructed energy, the flux-weighted fraction of predicted light events within the sample is determined. Now, for each energy bin, the predicted and weighted light fraction (binned in reconstructed energy) can be compared with the true weighted light fraction (binned in true energy).

This process is repeated 50 000 times in order to explore the phase space of subset size and light fraction. A similarity check is performed for each subset to avoid the usage of multiple event selections that resemble each other. For every new candidate subset  $M_i$ , the Jaccard similarity score [125]

$$S_J(M_i, M_j) = \frac{|M_i \cap M_j|}{|M_i \cup M_j|}$$

<sup>2</sup>The value of 0.35 will be motivated in the course of the following paragraphs.



**Figure 8.9:** For the agnostic approach and the logarithmic energy bin 6.0–6.15, this plot presents the relation between predicted and true light fraction. Every entry of the histogram is taken from one of 25 330 subsets. The median true fraction in each slice of predicted fraction can be described by a logistic sigmoid function (orange). The sigmoid fits to the 16-, 32-, 68- and 84-percentiles are shown in the ranges where they are needed. The solid black lines are tangents to the blue sigmoid functions at their steepest point and the dashed lines mark the point where they intersect the vertical 0% and 100% edges. Figure B.1 includes the histograms for all energy bins along with the sigmoid fits.

for all 500 previously accepted subsets<sup>3</sup> is calculated.  $M_j$  is added to the set of subsets only if all those  $S_j$  values are below a similarity threshold, which starts at 1 (i.e., every set is accepted) and decreases linearly towards 0 (only entirely disjoint sets are accepted) at the last iteration. Due to the omission of too similar subsets, in the end, the set of subsets is substantially smaller than the total of 50 000 created subsets.

In the final step, all these sets that passed the similarity condition are stored and their collection of true and predicted fractions can be compared to construct the calibration ( $\rightarrow$  Fig. 8.9). The extent of predicted fractions is considerably more compressed and constrained to the range of about 10–50% in contrast to the 0–100% spanned by the true fractions. This is the consequence of the rather high FN rate that comes with the chosen  $c_\omega$ . In order to unfold the relation between true and predicted light fraction, the median true fraction is calculated in bins of predicted fraction. Those can be fit with a sigmoid function

$$\sigma_{\log}(x) = \frac{1}{1 + e^{-p_0(x-p_1)}} \quad (8.1)$$

with the slope  $p_0$  and the vertical shift  $p_1$  as free parameters. It is this sigmoid function that is used for the calibration of the light fraction in each energy bin. The best-fit parameters are

<sup>3</sup>All  $n$  accepted subsets are used for comparison as long as  $n < 500$ .

listed in Table B.1 in the appendix along with the 2-dimensional histograms for all energy bins in Figure B.1.

As a result of the broad distribution in the vertical dimension, the calibrated light fraction comes with an uncertainty. Since its magnitude is tied to the amount of simulated events that are available to create subsets, but is not influenced by the statistics of the data the calibration is tested on, it is considered systematic and referred to as “sampling uncertainty”.

The sampling uncertainty can be quantified by repeating the sigmoid fit procedure with the binned  $(50 \pm \iota/2)$ -percentiles<sup>4</sup>. Between the resulting two lines, approximately 68% of subsets are contained for a given predicted light fraction. However, using the 16- and 84-percentile lines only works well in the central range of light fractions. The distributions of true fraction become increasingly asymmetric as they approach 0% or 100%, respectively, and the central 68% around the median cannot be consulted as a sensible confidence interval of the calibrated fraction any longer. Instead, for low fractions an upper limit can be defined by fitting a sigmoid function to the 68-percentiles of subsets in each slice of predicted fraction. Vice versa, the 32-percentiles are taken for lower limits on high fractions. In order to determine the transition points from the “nominal” fraction range to the domain of limits, tangents are constructed at the steepest point of the 16- and 84-percentile fits. The functional form of such a tangent is

$$f(x) = \frac{(p_0x - p_0p_1 + 2)}{4}.$$

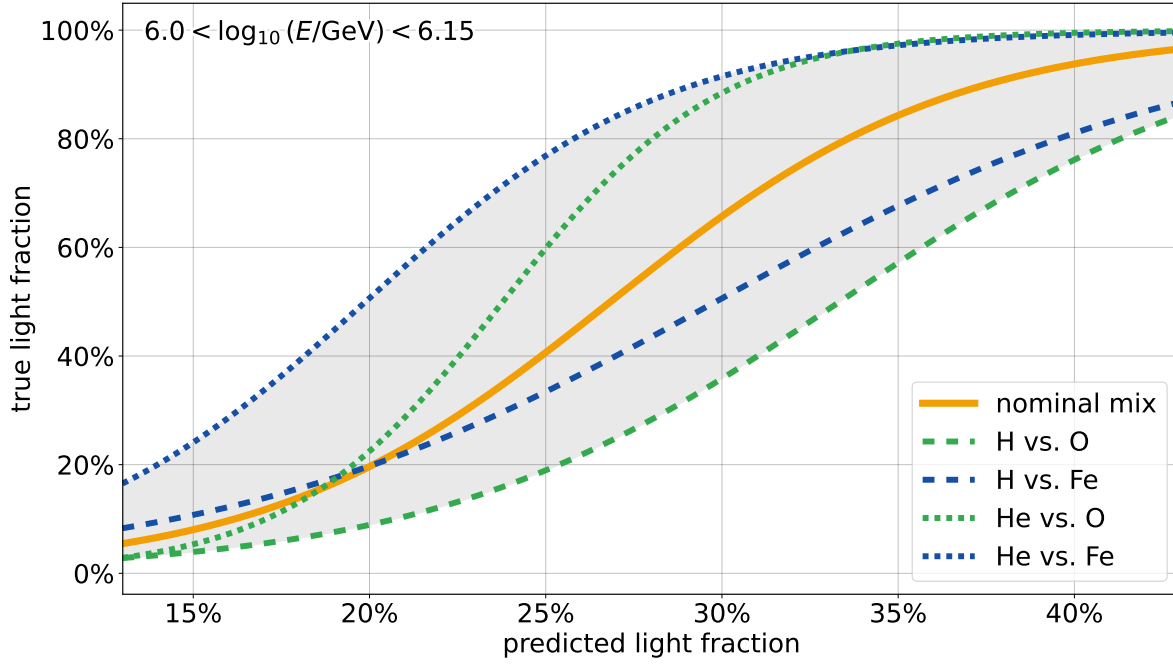
For predicted fractions  $x$  that have  $f_{16}(x) < 0$ , upper limits are calculated, whereas lower limits are necessary for predicted fractions where  $f_{84}(x) > 1$ .

It is preferred to have as little vertical spread for any given predicted fraction as this results in a small sampling uncertainty. The possible range of the classification threshold  $c_\omega$  has been scanned from 0 to 1 to find a value that minimizes the vertical spread. This study did not yield any clear optimal threshold value. However, with very high or very low  $c_\omega$ , the probability of obtaining a predicted fraction close to 0% or 100%, respectively, increases. The resulting pile-up of subsets near those edges perturbs the relation between predicted and true light fraction. Therefore, the value 0.35 is selected from the range of possible thresholds.

The weighting of true and predicted fractions in the agnostic version does not differentiate between true primary mass. The underlying flux model assumes equal fractions of all four mass groups (proton, helium, oxygen and iron) with an unbroken  $E^{-2.7}$  spectrum. This simple assumption serves as the nominal model in the agnostic case. However, the described method of classifying events with  $\omega_p > c_\omega$  as light events is sensitive to the relative composition of cosmic rays. In order to conservatively cover the full range of possible compositions, the process of generating subsets of the MC test set and the subsequent fitting of sigmoid functions to the median true fractions is repeated four times with the restriction that subsets contain only events initiated by

---

<sup>4</sup> $(50 - \iota/2) \approx 15.9$  and  $(50 + \iota/2) \approx 84.1$ , thus the percentiles are also referred to as 16- and 84-percentiles. Accordingly for the  $\iota$ - and  $(100 - \iota)$ -percentiles.



**Figure 8.10:** The definition of composition uncertainty for the agnostic approach using the logarithmic energy bin 6.0–6.15 as an example. The nominal sigmoid fit from Figure 8.9 is shown together with the four most extreme composition assumptions which are used to determine the systematic uncertainty due to composition within the light and the heavy flux component, which is indicated by the gray area.

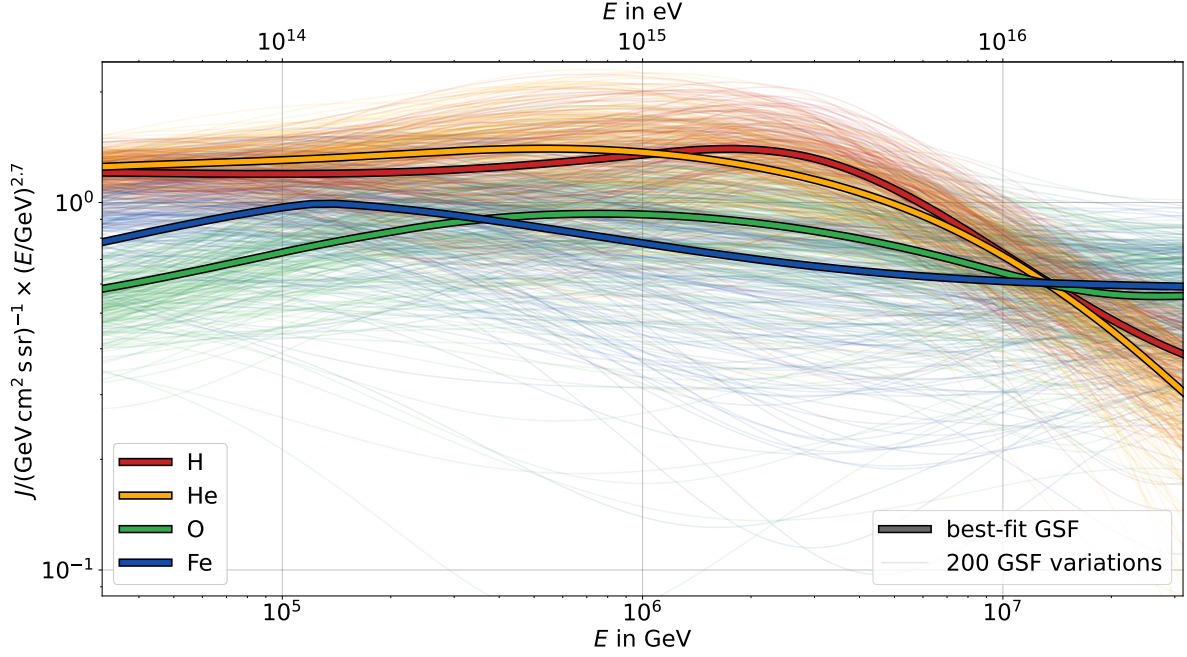
- protons and oxygen nuclei,
- helium and oxygen nuclei,
- protons and iron nuclei and
- helium and iron nuclei,

respectively. Those are the most extreme possible compositions that include both a light and a heavy component. Figure 8.10 shows these four fit lines along with the nominal sigmoid, which represents an equal mixture within the light and heavy components, respectively. The range of true light fraction enclosed by those fit lines is interpreted as the systematic “composition uncertainty”.

### 8.2.2 Method 2: GSF-informed

The alternative GSF-informed approach differs from the agnostic one in several points. First, the correction of the energy bias is done using the GSF-weighted<sup>5</sup> bias instead of the  $E^{-2.7}$ -weighted values ( $\rightarrow$  Fig. 8.7). The GSF-weights are furthermore used in the calculation of the weighted light fractions, where they replace the  $E^{-2.7}$ -weights of the agnostic scenario. Histograms of the predicted and true light fraction can be found in Appendix B in Figure B.2. Furthermore, the estimation of the composition uncertainty is designed in a fundamentally

<sup>5</sup>As of 2025, two versions of GSF exist: the initial release in Ref. [17] and a recent update published in 2024 [18], which is not yet available to the cosmic-ray community. Therefore, in this work, a modification of the original version is used.

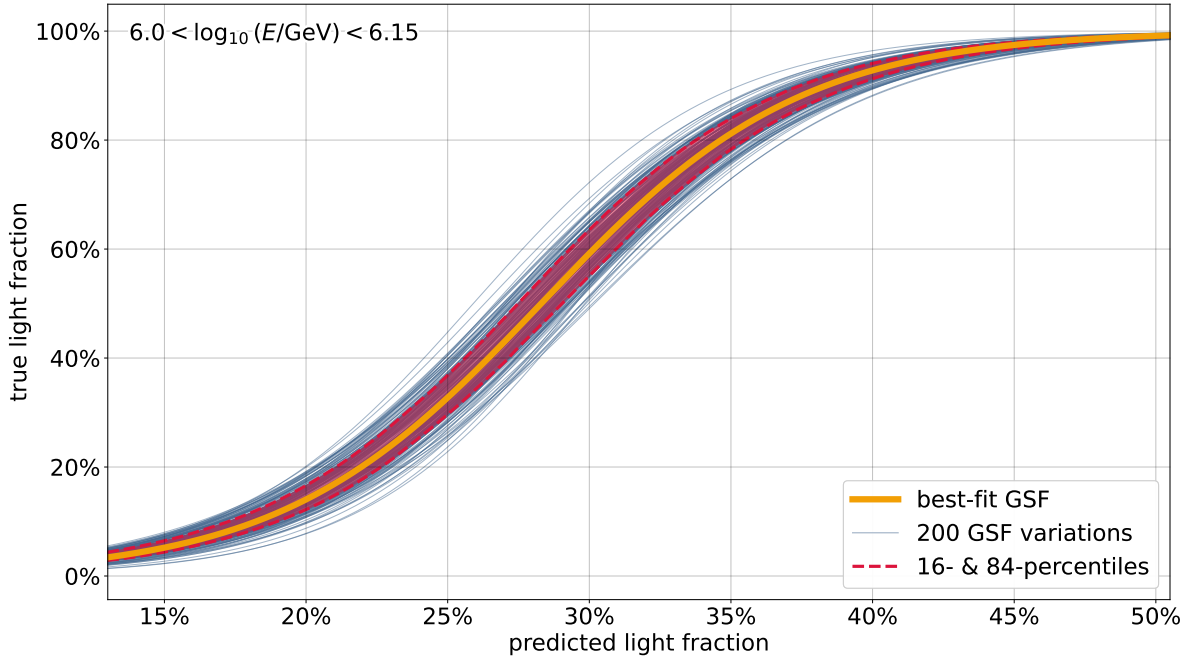


**Figure 8.11:** The best-fit GSF flux model (thick lines) in the energy range of interest together with 200 variations, generated using the GSF covariance (thin lines). The agnostic weighting to  $E^{-2.7}$  would be represented by four identical horizontal lines given the chosen flux scaling with  $E^{2.7}$ .

different way. As the only cosmic-ray composition model, GSF provides uncertainties on the best spline fit. Taking into account the covariance between the four mass groups, random variations of GSF can be generated, which is visualized in Figure 8.11. For each of 200 such GSF alterations, the sampling of subsets is repeated, true and predicted light fractions are weighted according to the respectively modified flux and the sigmoid function fitted to the median true fractions. This culminates in 200 calibration lines for every energy bin ( $\rightarrow$  Fig. 8.12). The composition uncertainty range is defined by the 16-percentile and the 84-percentile at any given predicted light fraction. It is obvious that the set of GSF variations spans a narrower vertical range than the conservative limits assumed in the agnostic case. Consequently, the composition uncertainty is a larger contribution to the total systematic uncertainty in the agnostic case, compared to the GSF approach. The best-fit parameters for the light fraction calibration with the GSF-informed method are listed in Table B.2 in the appendix.

### 8.3 Efficiency Correction

The combined proton and helium spectrum is determined by taking the product of the reconstructed differential flux of all kinds of hadronic cosmic rays, as defined in Eq. (1.1), and the calibrated light fraction. The calculation of the former includes the division by area. For this, it is crucial to recognize that the detection efficiency is not universal. The chance of triggering on an incident air shower changes most notably with its primary energy but also depends on the primary type, inclination and environmental conditions (the latter being discussed in sections 8.4 and 8.5). Low-energy showers contain fewer electromagnetic particles and have a narrower lateral spread at observation level. Therefore, the probability of

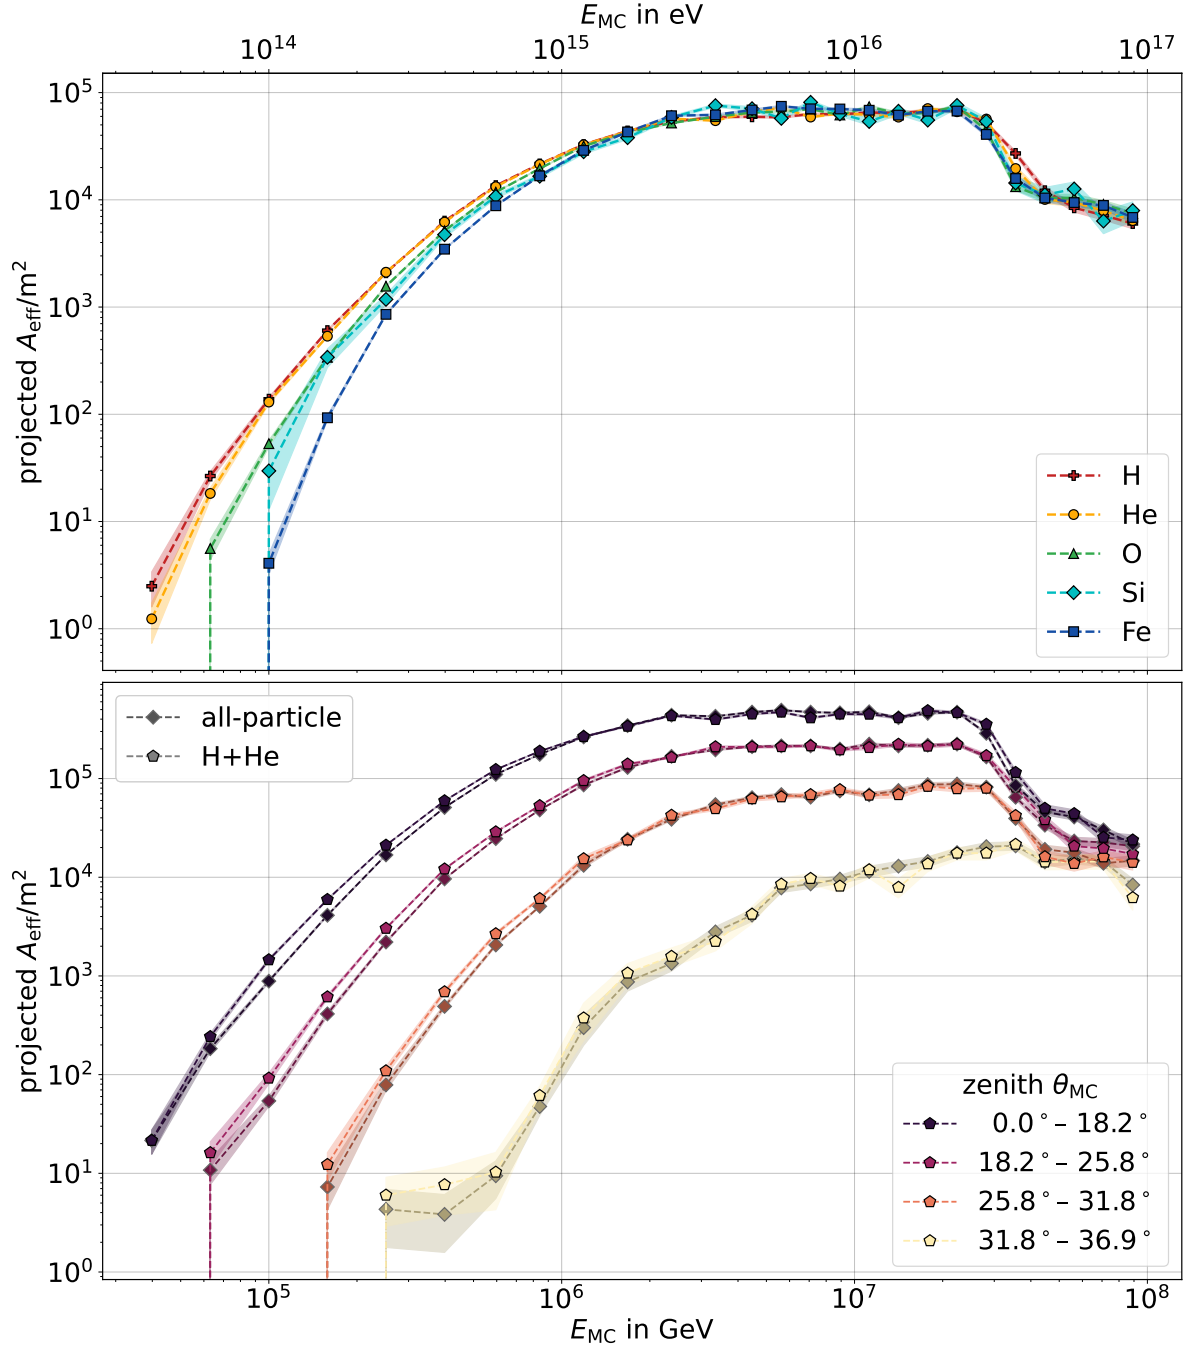


**Figure 8.12:** The nominal sigmoid line fit (orange) for the GSF approach and the logarithmic energy bin 6.0–6.15. The calibration line provided by the best-fit GSF model is surrounded by 200 lines (blue), emerging from random variations of the flux model. The 68% composition uncertainty area is indicated in red. This is the equivalent to Figure 8.10 for the GSF-informed method.

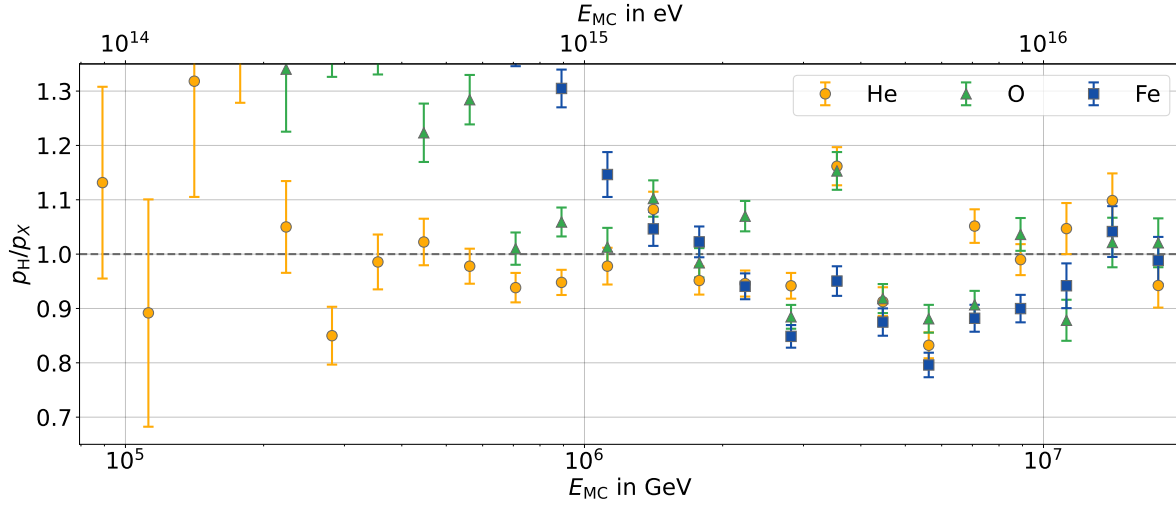
causing a sufficient number of HLCs needed to trigger data acquisition drops rapidly with decreasing primary energy. Furthermore, the multiple event processing steps up to Level 4 combined with the subsequent quality cuts lead to an overall reduction of total efficiency, that is, the fraction of incident events that pass through event selection.

The energy-dependent efficiency can be observed in Figure 8.13, which shows the projected effective area as a function of primary energy. Efficiency saturates at about 3 PeV and drops again above 20 PeV due to the  $S_{125}$  quality cut. Above approximately 1 PeV, the efficiency becomes independent of the primary type. For lower energies, the efficiency of heavy primaries is noticeably lower than that of proton and helium showers. Vertical showers have a higher efficiency than inclined events, mainly due to the in-ice coincidence requirement. At low energies, the increase in slant depth  $X$  with larger zenith angle also contributes to the reduced efficiency, that is, small inclined air showers are increasingly likely to die out before reaching the ground.

It is necessary that the effective areas of proton and helium events do not disagree significantly. Otherwise, the efficiency correction of the combined H+He spectrum would vary with the relative fractions of protons and helium nuclei within the light component. The ratio of proton efficiency  $p_H$  and helium efficiency  $p_{He}$  is depicted in Figure 8.14. The efficiencies of proton and helium events are considered sufficiently similar, allowing for a combined treatment of both primaries. For the detailed mathematical development of the confidence interval for the ratio of efficiencies, consult [Appendix D](#).



**Figure 8.13:** Projected effective area as a function of primary energy and weighted to an  $E^{-2.7}$  flux. **Top:** Effective area broken down by primary type. **Bottom:** The all-particle (diamonds) and the combined H+He effective area (pentagons), binned in true zenith angle.



**Figure 8.14:** Efficiency ratio between proton and helium, oxygen and iron events.

The efficiency correction in zenith bins implies that also the solid angle needs to be calculated for each individual interval of  $\theta$ . For every zenith bin, the covered solid angle is

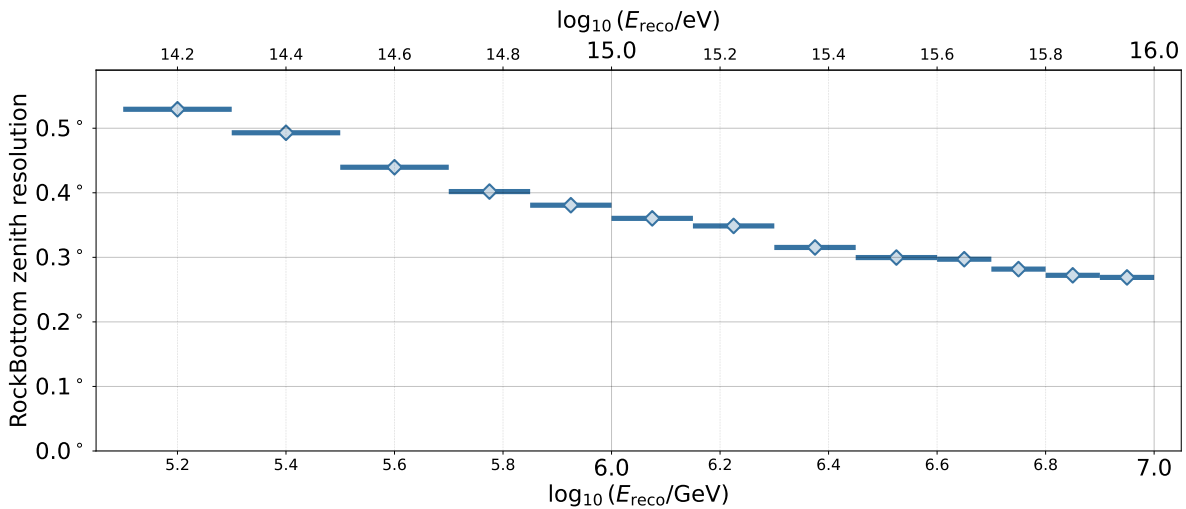
$$\Omega = \pi \left( \sin^2(\theta_{\max}) - \sin^2(\theta_{\min}) \right).$$

The zenith resolution of RockBottom after all quality cuts<sup>6</sup> is shown in Figure 8.15. The uncertainty on the solid angle can be estimated with

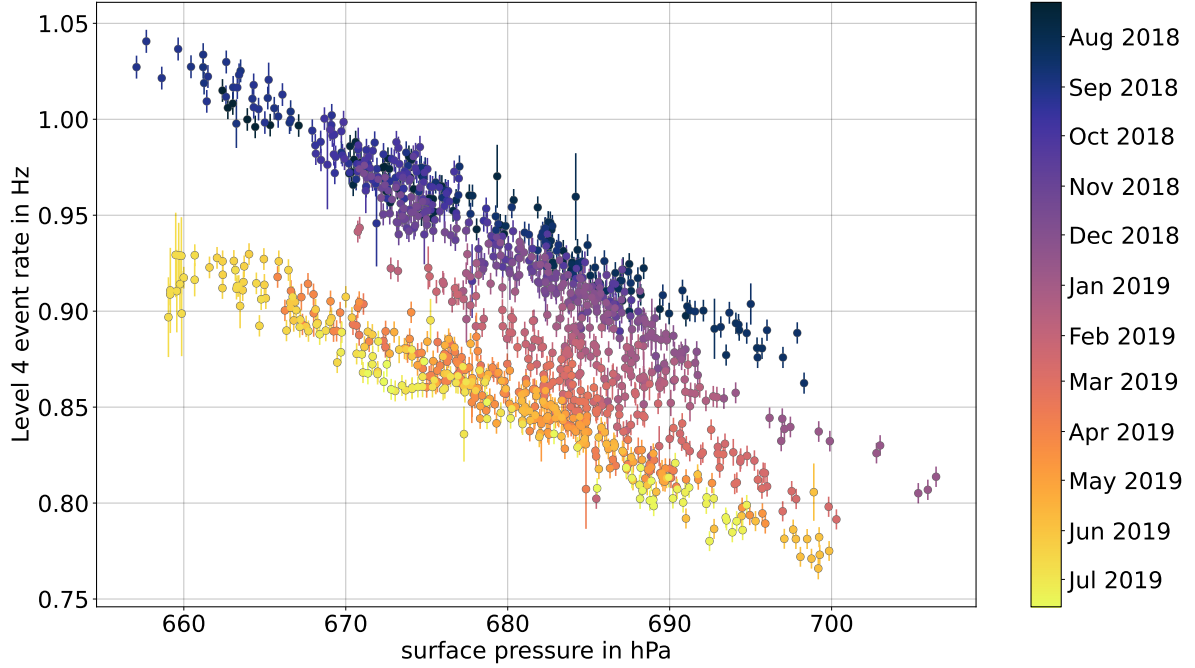
$$\Delta\Omega = \pi \cdot \Delta\theta \sqrt{\sin^2(2\theta_{\min}) + \sin^2(2\theta_{\max})},$$

where  $\Delta\theta$  is the zenith resolution in the respective energy bin.

<sup>6</sup>Those values are systematically smaller than the angular resolution in Figure 7.4, which also includes azimuthal deviations of the reconstructed direction and does not benefit from the  $S_{125}$  cut.



**Figure 8.15:** Zenith resolution of the RockBottom reconstruction, binned in bias-corrected reconstructed logarithmic energy.



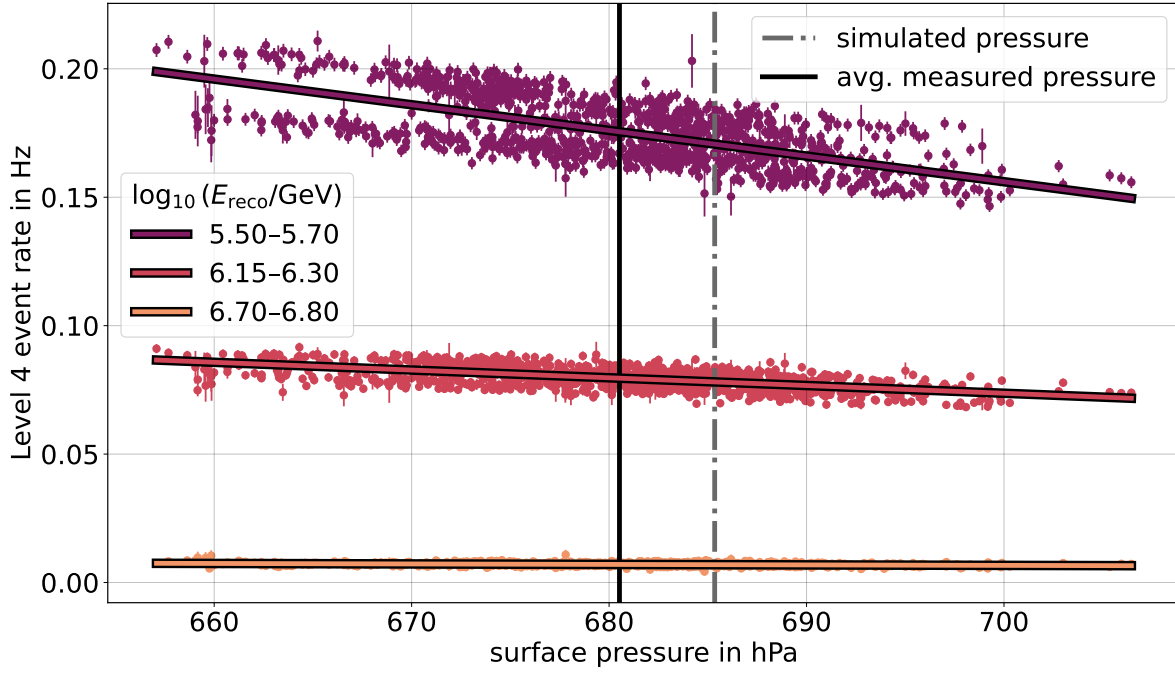
**Figure 8.16:** The rate of measured events at Level 4 versus the atmospheric pressure. Each point represents one data-taking run between July 2018 and July 2019.

The uncertainty on the projected effective area — which is a consequence of limited MC statistics — and the uncertainty on solid angle are added in quadrature to yield the combined “acceptance uncertainty”.

#### 8.4 Pressure Correction

Evaluating the modeled atmospheric profile used in the CORSIKA simulations (→ Fig. 5.5) at the average altitude of IceTop tanks of 2830 m above sea level gives a mass overburden  $\Xi(2830 \text{ m})$  of  $698.62 \text{ g/cm}^2$ . With a gravitational acceleration of  $9.81 \text{ m/s}^2$  this corresponds to a surface air pressure of 685.35 hPa. This simulated pressure represents the average atmospheric conditions in April at the South Pole. The actual average measured pressure during the data-taking time period from July 2018 to July 2019 was almost one percent lower, at 680.53 hPa. The surface pressure is continuously measured at the *South Pole Atmospheric Research Observatory*, which is operated by NOAA. As mentioned in subsection 5.4.1, both the trigger rate at IceTop and deep in the ice correlate with the air pressure. The barometric effect on the event rate after quality cuts is visible in Figure 8.16 where the total all-particle rate is plotted versus the measured air pressure for each run of the IC86.2018 detection season. The clear negative correlation between event rate and pressure is superimposed with a continuous decrease of rate over time. The latter is caused by the steady accumulation of snow on top of the Cherenkov tanks, leading to increased attenuation of the electromagnetic shower component and thus a higher trigger threshold.

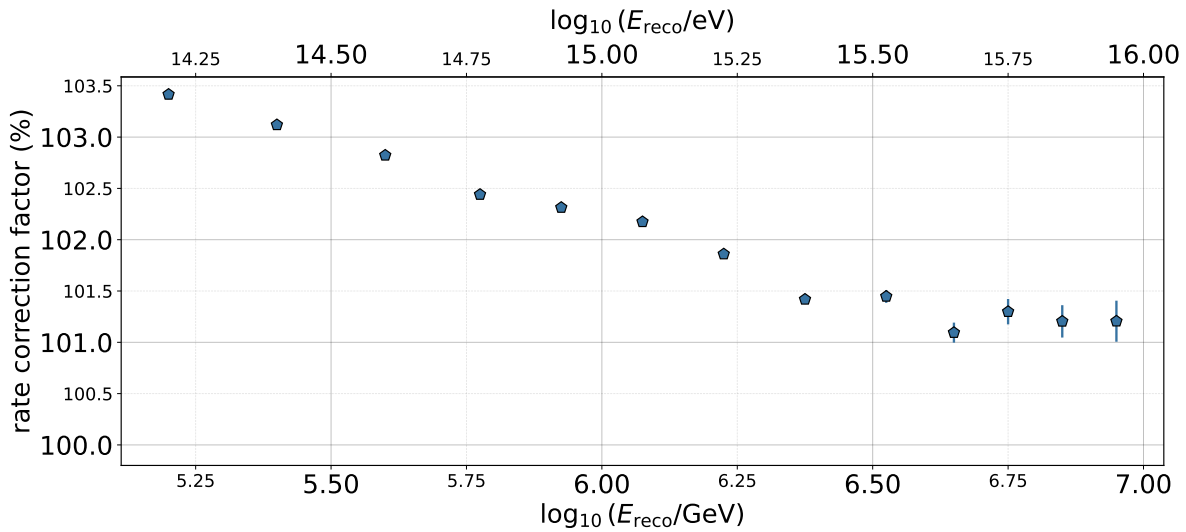
The effect of air pressure on the rate of recorded events is not universal over the full energy range. The relative increase of rate with decreasing air pressure is strongest in bins of low reconstructed primary energy. In Figure 8.17 event rate in three different bins of reconstructed



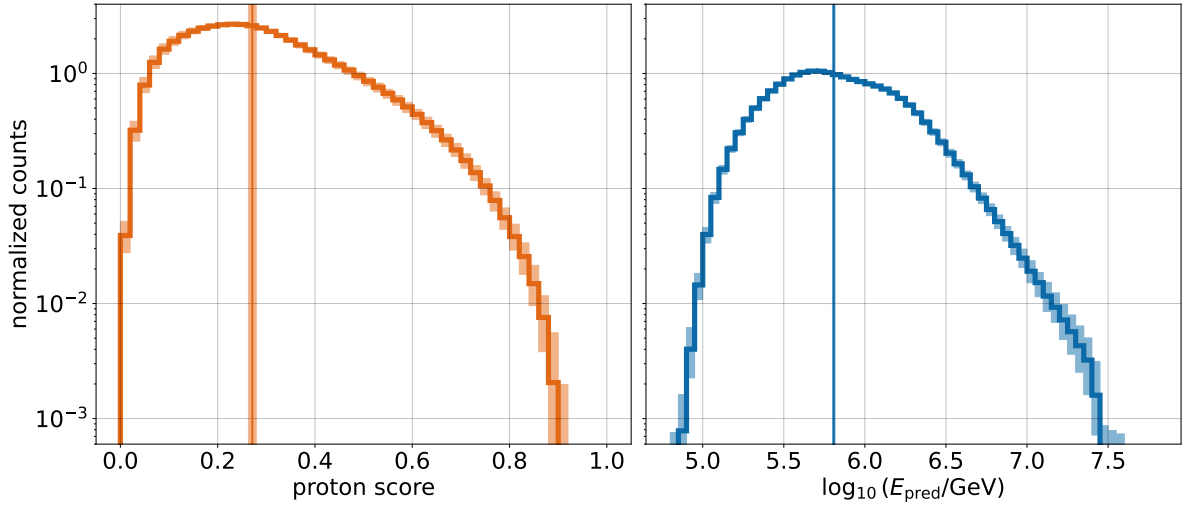
**Figure 8.17:** Event rate versus atmospheric pressure in three selected bins of reconstructed energy. Lines are fit to the respective data points. Their slope is taken to correct the rate for the difference in simulated and average measured pressure.

energy (corrected for the agnostic energy bias) is plotted over the corresponding air pressure together with line fits, which are necessary to determine correction factors to account for this barometric effect.

Taking the line fit slopes for all energy bins, the average 4.82 hPa pressure “deficit” in the measured data with respect to the simulation translates to the correction factors shown in Figure 8.18, which compensate for the slight mismatch in trigger efficiency in the MC. The uncertainty on these factors contributes to the total systematic uncertainty budget.



**Figure 8.18:** Rate correction factors to account for the difference in simulated and average measured air pressure.



**Figure 8.19:** Run-to-run variation of the NN outputs proton score (left) and  $\log_{10}(E/\text{GeV})$  (right). All distributions are normalized, the bold histograms represent the respective median distribution with the shaded area around being the central 68% range. The average of all median proton score and logarithmic energy distributions are indicated by vertical lines together with the standard deviation.

## 8.5 Environmental Variation

In addition to the flux correction due to the difference in the *mean* surface pressure between data and simulation, the fluctuations of the reconstructed proton and helium flux on a short time scale also need to be taken into account. Not only does the total event rate fluctuate due to changes in the environmental conditions, but also the distributions of the NN output may be affected by variations of the atmosphere.

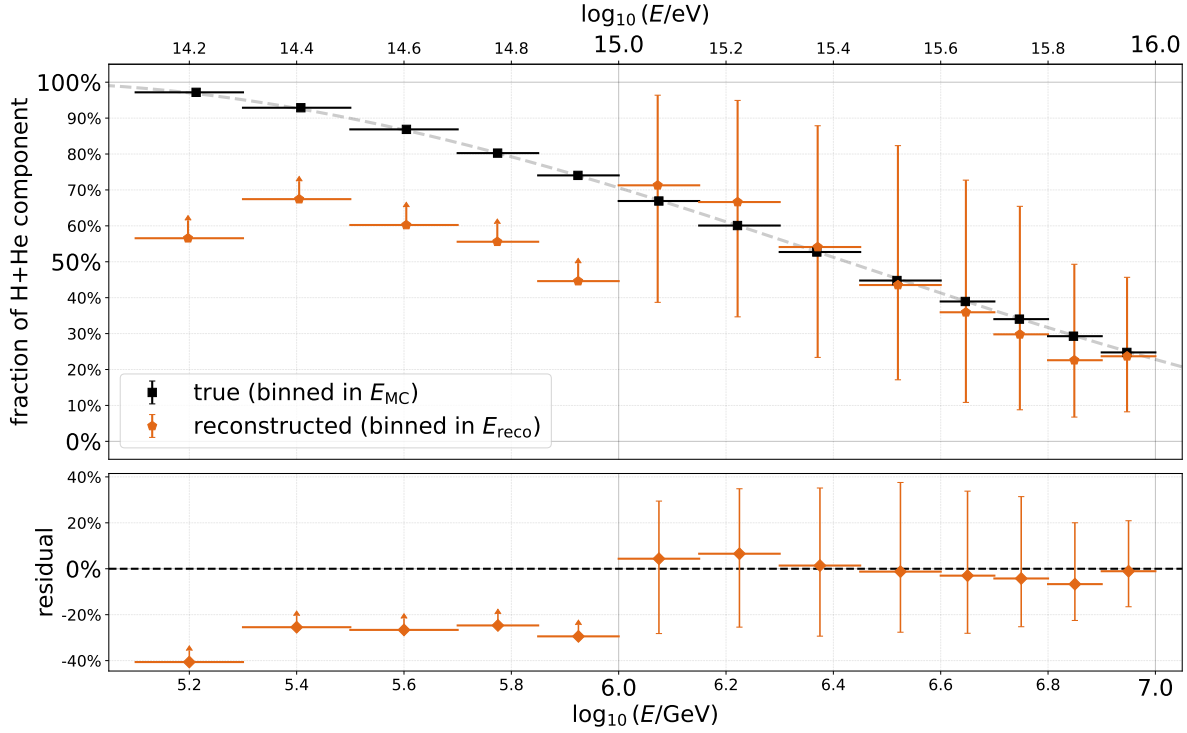
The run-to-run variation of both the proton score and the predicted logarithmic energy are shown in Figure 8.19. The median proton score is on average  $0.271 \pm 0.009$  but the mean uncalibrated light fraction is  $(31.8 \pm 1.8)\%$ , ranging from 28.1% to 35.4% in the most extreme cases. The reconstruction of primary energy, on the other hand, is hardly affected by seasonal variations with the averaged median  $\log_{10}(E_{\text{pred}}/\text{GeV})$  being  $5.808 \pm 0.009$ .

The fluctuations of the mass and the energy output of the NN propagate through the entire calculation of the combined proton and helium flux. In order to estimate the systematic uncertainty due to environmental changes<sup>7</sup>, the calibrated light fraction is multiplied by the all-particle event rate for each run. This yields a quantity that is proportional to the light flux and can therefore be used to estimate the relative uncertainty on the light flux caused by external influences — the “environmental uncertainty”.

## 8.6 Recovery Tests

Before applying the fraction calibration to the data recorded by IceCube over the course of one year, tests are necessary to validate the method. First, the performance of reconstructing the light fraction of an injected distribution is evaluated. For this, again, subsets of the MC test dataset are generated obeying the true light fraction specified by some smooth function

<sup>7</sup>This includes variations of the air temperature and pressure as well as the increase in snow height. Changes of the properties of the deep glacial ice on a time scale of one year are negligible.

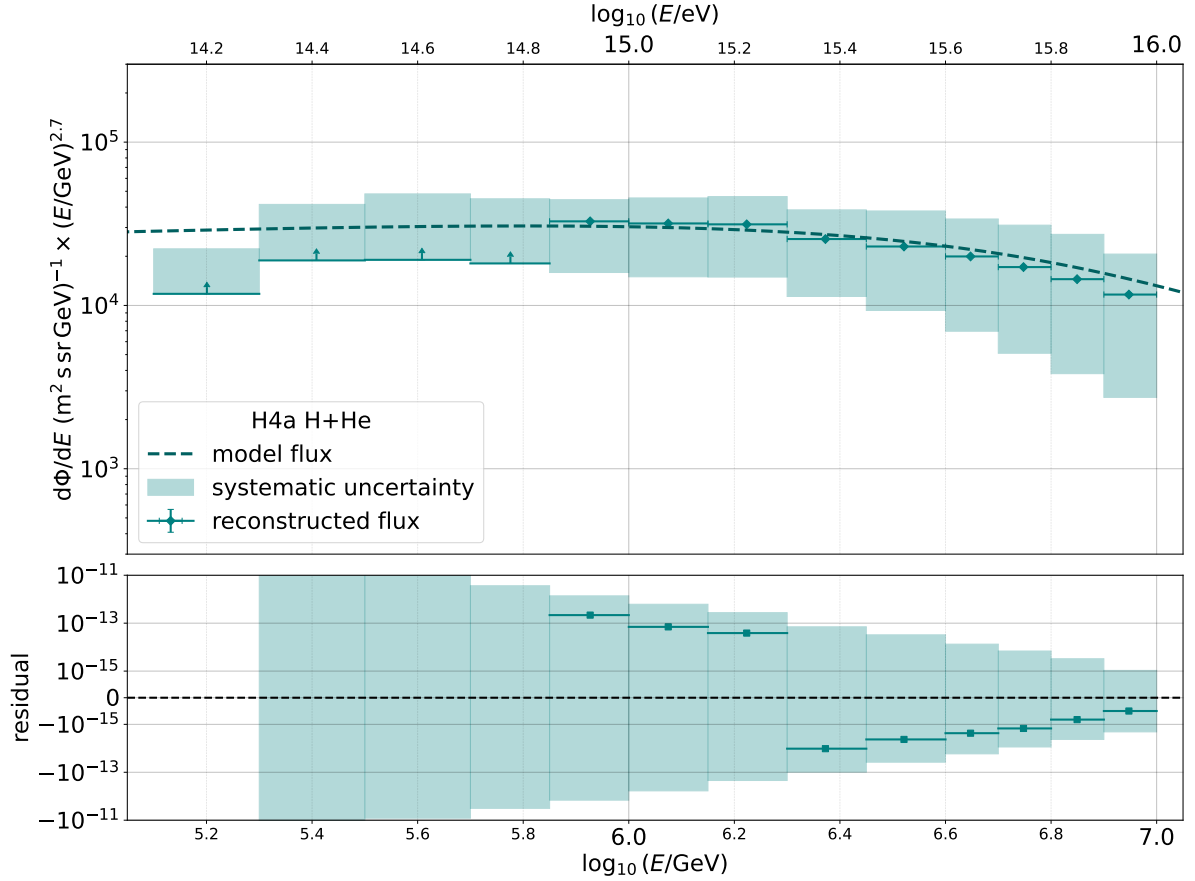


**Figure 8.20:** Light fraction recovery test for the agnostic approach. The true H+He fraction of the subset of test simulations (black markers in the top panel) follows  $0.5 \cdot (\sin(3 + \log_{10}(E/\text{GeV})) + 1)$  (gray dashed line). The orange markers represent the calibrated light fraction with systematic uncertainty (sampling and composition contribution). The bottom panel has the fraction residual.

that is bound between 0 and 1. Like for the development of the calibration, the sets are then binned in  $\log_{10}(E_{\text{reco}})$  to determine the predicted light fraction and the ratios H : He and O : Fe are close to 50 : 50. Finally, the calibration is applied to the fractions. An example for this is presented in Figure 8.20 where the agnostic fraction calibration is applied to a set of events that follows a broad sine to describe the light fraction as a function of  $\log_{10}(E)$ . Below 1 PeV, only lower limits are provided, while above, the central values follow the true quantities closely. Over the entire energy range, the conservative uncertainties obtained from the agnostic approach contain the true fractions.

In a similar manner, the reconstruction of a combined proton and helium flux can be checked. For this test, the true fractions as well as the ratios between the four mass groups (proton, helium, oxygen and iron) follow the H4a composition model. The calibrated fractions are multiplied by the sum of H4a-weights in each energy bin and divided by the effective area as well as the bin width to yield the differential light flux. Since the recovery tests are performed on simulated data, no pressure correction is necessary. The result is shown in Figure 8.21. Lower limits on the reconstructed flux are available up to 700 TeV. In addition to the conservative systematic uncertainty of the agnostic method, this is caused by the relatively high light fraction of H4a compared to competing models of almost 70% in the sub-PeV regime. The reconstructed H+He flux agrees well with the injected spectrum.

Even though a fraction calibration has been created for the  $5.1 < \log_{10}(E/\text{GeV}) < 5.3$  bin, it is not included in the analysis. At such low primary energies, the energy bias becomes



**Figure 8.21:** Light flux recovery test for the agnostic approach. The combined flux of protons and helium nuclei of the H4a composition model (dashed line in the top panel) is reconstructed by the calibration method described in the text. The bottom panel shows the residuals on a so-called “symmetric log-scale”.

increasingly mass-dependent and the resolution of both energy and zenith angle is worst ( $\rightarrow$  Figs. 8.6, 8.7 and 8.15). Furthermore, because of the very low survival rate (including trigger efficiency, processing, reconstruction and quality cuts) below 200 TeV, the efficiency correction is considered unreliable.

### Key Messages

The energy and mass outputs of the neural network are used to develop an energy-binned calibration for the estimation of the light fraction in a set of cosmic-ray events. Since air showers induced by protons and helium nuclei have almost identical efficiency, which is also higher than that of heavier primaries, the combined spectrum of protons and helium is targeted. Two distinct methods can be applied for this, a conservative “agnostic” approach with neutral assumptions and one seeded with the GSF composition model, which accordingly yields a smaller spread in the reconstructed light fraction. The systematic uncertainties comprise contributions from 1. sampling statistics in the fraction calibration, 2. the relative composition within the light and heavy flux components, 3. acceptance, that is, effective area and reconstructed zenith, 4. a correction factor to account for changes in the event rate due to pressure differences and 5. environmental variation caused by changing meteorological conditions and accumulating snow.



# CHAPTER 9

## RESULTS

*“I need to connect these two lines, but my first idea didn’t have enough math.”*

— Randall Munroe, *xkcd* #2048: ‘Curve-Fitting’

The reconstruction method elaborated in the previous chapter is now applied to one year worth of [CR](#) data, measured by IceCube. Since the 2-station trigger started operating in 2016, the Monte Carlo datasets have been produced with simulated snow heights as measured in November 2018<sup>1</sup>. The data of the [IC86.2018](#) season is divided into 1582 runs and spans from July 10, 2018 to July 17, 2019<sup>2</sup>. Since runs have to be marked as “good”, that is, sufficiently long and with stable rates, with all non-DeepCore strings active to be considered for the evaluation, the number of usable runs reduces to

1146. The total duration of the selected runs is  $3.14 \cdot 10^7$  s which corresponds to 363.3 days. During that time, 28.1 million events have been recorded that pass the full processing sequence and all quality cuts. This corresponds to an average rate of usable events of 0.895 Hz. The [IC86.2018](#) season started with 77 inoperative [DOMs](#) but in December 2018 two more in-ice modules had a short circuit.

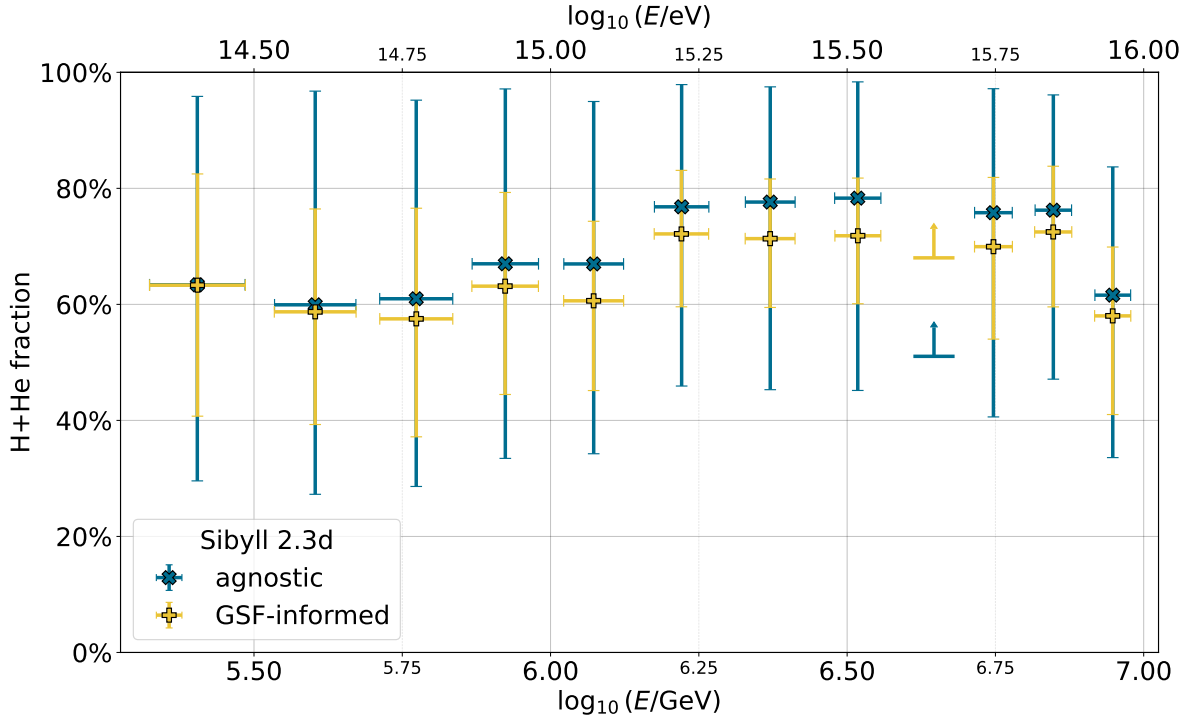
First, the fraction and the differential flux of light cosmic rays are reconstructed in [section 9.1](#) and [section 9.2](#), respectively. This is followed by a comparison of the light fraction obtained with different hadronic interaction models in [section 9.3](#). Eventually, the reconstructed flux is put into context by comparing it to composition models (→ [section 9.4](#)) and measurements performed by other experiments (→ [section 9.5](#)).

**Table 9.1:** *Calibrated fraction of the light cosmic-ray flux for the agnostic method (left) and the GSF-informed approach (right) of Figure 9.1. Lower limits are indicated by arrows. The hadronic interaction model used for the reconstruction is Sibyll 2.3d.*

$\log_{10}(E/\text{GeV})$	agnostic	GSF-informed
5.30–5.50	$63^{+33}_{-34} \%$	$63^{+20}_{-23} \%$
5.50–5.70	$60^{+37}_{-33} \%$	$59^{+18}_{-20} \%$
5.70–5.85	$61^{+35}_{-33} \%$	$58^{+20}_{-21} \%$
5.85–6.00	$67^{+31}_{-34} \%$	$63^{+17}_{-19} \%$
6.00–6.15	$67^{+29}_{-33} \%$	$61^{+14}_{-16} \%$
6.15–6.30	$77^{+22}_{-31} \%$	$72^{+11}_{-13} \%$
6.30–6.45	$78^{+20}_{-33} \%$	$71^{+11}_{-12} \%$
6.45–6.60	$78^{+21}_{-34} \%$	$72^{+10}_{-12} \%$
6.60–6.70	$\uparrow 51 \%$	$\uparrow 68 \%$
6.70–6.80	$76^{+22}_{-36} \%$	$70^{+12}_{-16} \%$
6.80–6.90	$76^{+20}_{-30} \%$	$72^{+12}_{-13} \%$
6.90–7.00	$62^{+23}_{-29} \%$	$58^{+12}_{-18} \%$

<sup>1</sup>Since the simulation of detector response is very time consuming and several TB of disk space are required for each simulated year, for the general purpose datasets Monte Carlo is usually produced for only every third year. As datasets for the snow heights in 2012 and 2015 already existed, 2018 was chosen to ensure the existence of 2-station triggers and, at the same time, have as shallow snow as possible.

<sup>2</sup>These just over 372 days correspond to significantly fewer intervals of eight hours, which is the typical length of a full data-taking run. In case of occurring issues with hardware or software or when changes are made to the active detector configuration, the current run is stopped and a new one is initiated. This leads to a considerably shorter average run duration.



**Figure 9.1:** Reconstructed fraction of the light component using the agnostic calibration (blue) and the GSF-informed variant (yellow). The error bars represent the systematic uncertainty (sampling and composition) in the vertical direction and the estimated energy resolution horizontally.

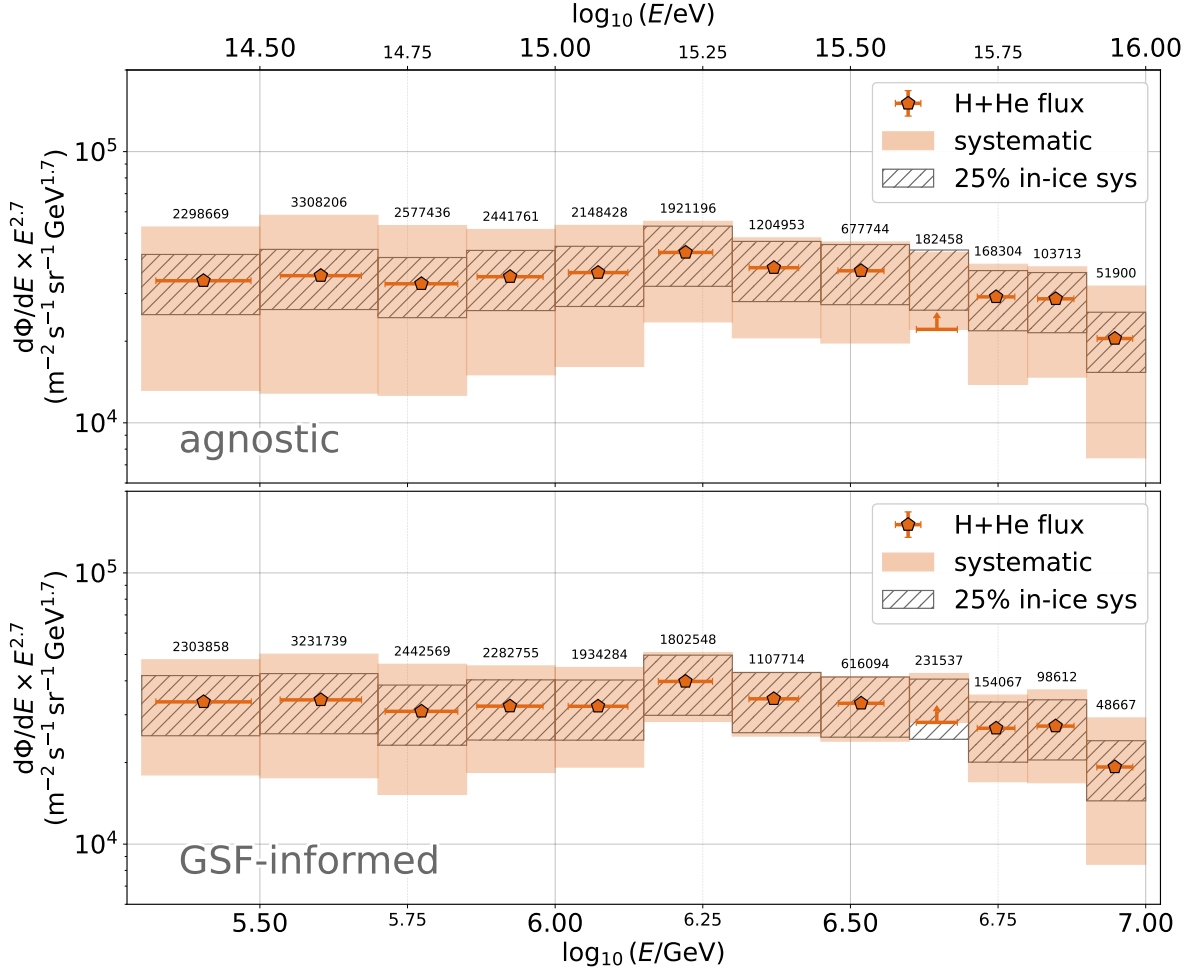
## 9.1 Light Cosmic-Ray Fraction

The result of applying the calibrations of the light fraction of proton and helium primaries that have been developed on the Sibyll 2.3d simulation set are presented in Figure 9.1 and Table 9.1. The calibrated fraction of the combined proton and helium flux component in the agnostic model increases from 60% at around 400 TeV primary energy up to 78% at approximately 3 PeV. The agnostic reconstruction yields systematically higher fractions compared to the values obtained with the GSF method, which follows a similar trend with the light fraction ranging between 58% and 72%. The agnostic version of the light fraction calibration is more conservative in estimating the composition uncertainty, leading, in total, to significantly larger systematic uncertainties. The light fraction in the energy bin  $6.6 < \log_{10}(E/\text{GeV}) < 6.7$  can only be restricted by a lower limit for both approaches. The reason for this is that the predicted light fraction from the NN output is above the highest value that allows for a central fraction value ( $\rightarrow$  “6.6–6.7” panels in Figures B.1 and B.2). It can thus be stated that in the 200 TeV–10 PeV energy range cosmic rays are not purely light but have a mixed composition with at least a few percent contribution from the oxygen mass group or heavier primaries.

## 9.2 Differential Light Cosmic-Ray Flux

Having estimated the fraction of light cosmic rays, their differential flux  $J = d\Phi/dE$  can be calculated, taking into account the pressure correction factors ( $\rightarrow$  section 8.4). The energy-binned differential light flux is determined by multiplying the calibrated light fraction with the total number of events that fall into the respective bins and dividing by the effective area

## 9.2. DIFFERENTIAL LIGHT COSMIC-RAY FLUX

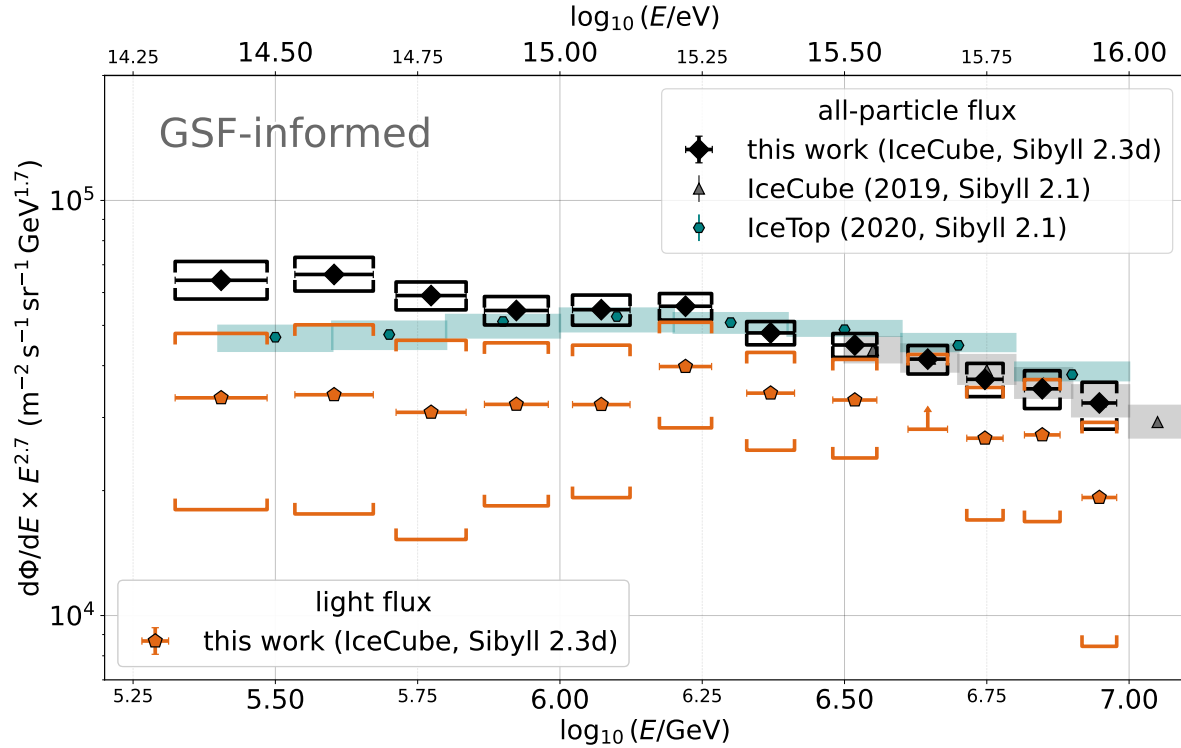


**Figure 9.2:** The reconstructed light cosmic-ray flux (protons and helium) around the knee region of the spectrum, scaled with  $E^{2.7}$  for the agnostic (top) and the GSF-informed method (bottom). Error bars indicate energy resolution horizontally and statistical uncertainty in the vertical direction, which is too small to be visible. Orange shaded boxes represent the total systematic uncertainty. The estimated systematic uncertainty due to the ice model is displayed separately by gray hatched boxes. The number of reconstructed light events, respectively their limit, for each energy bin is provided in the plot. The method is based on Sibyll 2.3d simulations.

(the projected  $A_{\text{eff}}$  corrected with  $1/\cos(\theta_{\text{RB}})$ ), the solid angle element  $\Omega$  that has been taken into account in the calculation of effective area, the total duration of selected runs and the width of the energy bin in linear space.

In addition to the composition and sampling uncertainties, which are propagated forward from the light fraction calibration, further contributions to the systematic uncertainty of the combined proton and helium flux emerge from the uncertainty on acceptance (via  $\Delta A_{\text{eff}}$  and  $\Delta\theta$ ), on flux variations due to environmental fluctuations and on the pressure correction factors.

Figure 9.2 shows the reconstructed light cosmic-ray flux between 200 TeV and 10 PeV primary energy, respectively lower limits on it. Due to the high flux of cosmic rays around the knee, the statistical uncertainty is too small to be visible. In contrast, systematic uncertainties are of



**Figure 9.3:** The reconstructed light cosmic-ray flux (orange) together with the “all-particle” spectrum (black), both with the Sibyll 2.3d hadronic interaction model and the GSF-informed calibration method. A cross-check with all-particle spectra from Refs. [70, 126] (gray and teal) indicates agreement above  $\sim 2$  PeV where efficiency is independent of the primary mass.

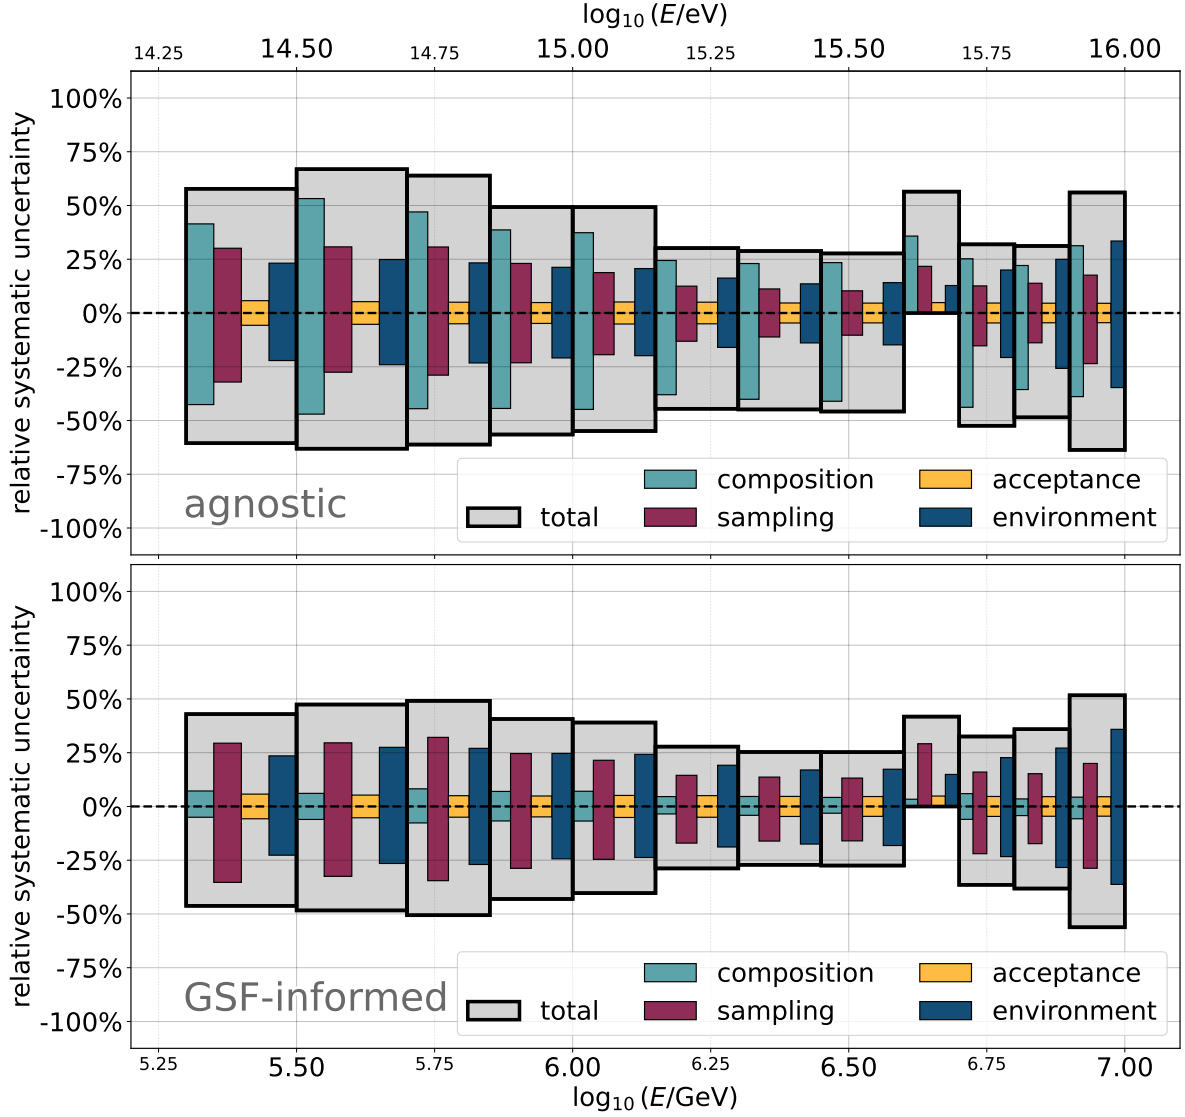
significant magnitude. As it was not possible to investigate the systematic uncertainty arising from the ice model used in the simulation ( $\rightarrow$  subsection 5.4.3), a conservative estimate of  $\pm 25\%$  (inspired by Ref. [70]) is assumed and presented separately from the systematic uncertainties that have been studied before.

For nominal flux values, the systematic uncertainties are added in quadrature for the upward and downward component, respectively. In the case of lower limits, the flux corresponding to the median-sigmoid fit is chosen as the origin to calculate the downward uncertainty<sup>3</sup>, vice versa for upper limits and the upward uncertainty. In addition, the approximated in-ice systematic in Figure 9.2 is centered around this median flux for limits.

Lower limits are accompanied by an upper limit, which corresponds to the entire CR flux being composed of protons and helium nuclei, that is, a light fraction of 100%. Even when the calibrated light fraction is quantified by central values with uncertainties above and below, including the relative uncertainty on acceptance and environmental effects can cause the uncertainty interval on the fraction to slightly exceed 100%. However, it is not necessary to convert those measurements into lower limit as the excess is well covered by the systematic uncertainty (acceptance, environmental and pressure correction) on the “all-particle” flux as demonstrated in Figure 9.3 ( $\log_{10}(E/\text{GeV})$  bins 6.6–6.7 and 6.8–6.9). This spectrum is obtained

<sup>3</sup>Taking 100% as starting point would overestimate the contribution of the sampling uncertainty to the total systematics, choosing instead the limit value itself as origin would lead to a too conservative, namely linear, addition of systematic uncertainties.

## 9.2. DIFFERENTIAL LIGHT COSMIC-RAY FLUX



**Figure 9.4:** The systematic uncertainty for the agnostic (top) and the GSF-informed method (bottom), also broken down by its contributions. The relative uncertainty on the pressure correction factors is too small to be visible and thus omitted here. For lower limits, only the uncertainty upward is considered. The underlying hadronic interaction model for the analysis method is Sibyll 2.3d.

by calculating the differential flux via the GSF- respectively the  $E^{-2.7}$ -weighted effective area and omitting the multiplication with the light fraction. Below approximately 2 PeV, where efficiency becomes mass-dependent, this “all-particle” flux is not considered sensible due to its strong dependence on the composition. However, in the range of full efficiency it can be used for a cross-check with previous analyses of the cosmic-ray all-particle spectrum with IceCube and IceTop [70, 126], which does not reveal any tension, although it should be noted that this work presents the first IceCube result of this kind using Sibyll 2.3d as hadronic interaction model in the CORSIKA simulation.

The magnitude of the total relative systematic uncertainties on the light flux and their individual contributions are presented in Figure 9.4. For the light flux obtained with the agnostic method, the relative sampling uncertainties range roughly between 11% and 33%.

**Table 9.2:** Reconstructed differential flux of light cosmic rays from the agnostic method, developed on Sibyll 2.3d Monte Carlo. Statistical and total systematic uncertainties are provided as well. Lower limits are indicated by arrows. The second column contains the mean logarithmic energy in each bin.

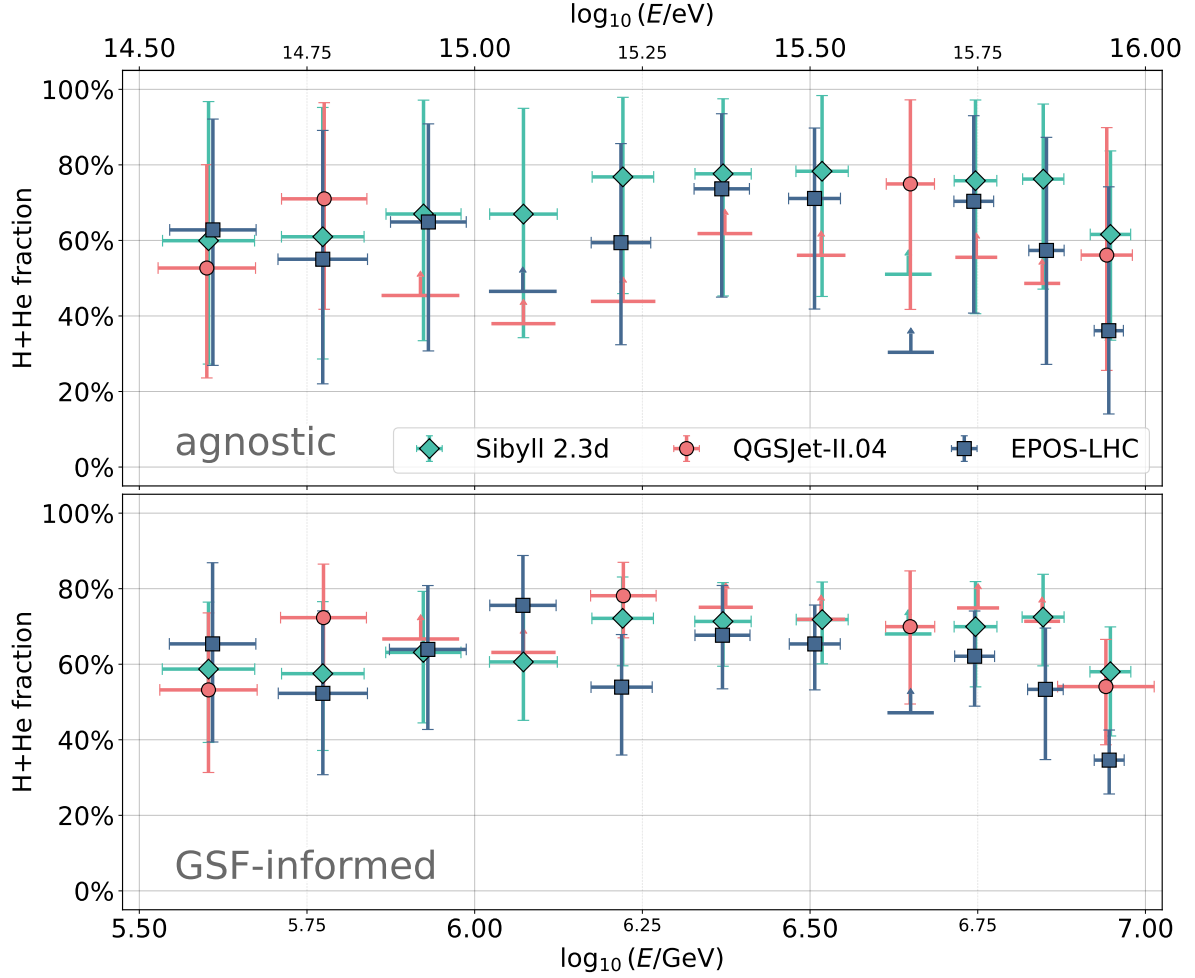
$\log_{10}(E/\text{GeV})$	$\langle \log_{10}(E/\text{GeV}) \rangle$	Flux $\pm$ stat. $\pm$ sys. ( $\text{m}^2 \text{ sr GeV}^{-1}$ )
5.30–5.50	$5.40 \pm 0.09$	$(8.551 \pm 0.009^{+4.941}_{-5.179}) \cdot 10^{-11}$
5.50–5.70	$5.60 \pm 0.07$	$(2.5940 \pm 0.0024^{+1.7347}_{-1.6372}) \cdot 10^{-11}$
5.70–5.85	$5.77 \pm 0.07$	$(8.399 \pm 0.009^{+5.368}_{-5.139}) \cdot 10^{-12}$
5.85–6.00	$5.92 \pm 0.06$	$(3.506 \pm 0.004^{+1.728}_{-1.982}) \cdot 10^{-12}$
6.00–6.15	$6.07 \pm 0.06$	$(1.4377 \pm 0.0015^{+0.7080}_{-0.7894}) \cdot 10^{-12}$
6.15–6.30	$6.22 \pm 0.05$	$(6.786 \pm 0.007^{+2.049}_{-3.023}) \cdot 10^{-13}$
6.30–6.45	$6.37 \pm 0.05$	$(2.3551 \pm 0.0028^{+0.6786}_{-1.0566}) \cdot 10^{-13}$
6.45–6.60	$6.52 \pm 0.04$	$(9.167 \pm 0.015^{+2.540}_{-4.201}) \cdot 10^{-14}$
6.60–6.70	$6.65 \pm 0.04$	$\uparrow (2.507 \pm 0.012^{+0.892}_{-0.930}) \cdot 10^{-14}$
6.70–6.80	$6.75 \pm 0.04$	$(1.772 \pm 0.006^{+0.566}_{-0.930}) \cdot 10^{-14}$
6.80–6.90	$6.85 \pm 0.04$	$(9.30 \pm 0.04^{+2.90}_{-4.52}) \cdot 10^{-15}$
6.90–7.00	$6.95 \pm 0.04$	$(3.562 \pm 0.026^{+1.997}_{-2.269}) \cdot 10^{-15}$

**Table 9.3:** Reconstructed differential flux of light cosmic rays (GSF-informed method, Sibyll 2.3d). Statistical and total systematic uncertainties are provided as well. Lower limits are indicated by arrows. The second column contains the mean logarithmic energy in each bin.

$\log_{10}(E/\text{GeV})$	$\langle \log_{10}(E/\text{GeV}) \rangle$	Flux $\pm$ stat. $\pm$ sys. ( $\text{m}^2 \text{ sr GeV}^{-1}$ )
5.30–5.50	$5.40 \pm 0.09$	$(8.547 \pm 0.009^{+3.670}_{-3.951}) \cdot 10^{-11}$
5.50–5.70	$5.60 \pm 0.07$	$(2.5311 \pm 0.0024^{+1.1999}_{-1.2244}) \cdot 10^{-11}$
5.70–5.85	$5.77 \pm 0.07$	$(7.958 \pm 0.009^{+3.909}_{-4.028}) \cdot 10^{-12}$
5.85–6.00	$5.92 \pm 0.06$	$(3.277 \pm 0.004^{+1.333}_{-1.411}) \cdot 10^{-12}$
6.00–6.15	$6.07 \pm 0.06$	$(1.2945 \pm 0.0016^{+0.5053}_{-0.5220}) \cdot 10^{-12}$
6.15–6.30	$6.22 \pm 0.05$	$(6.370 \pm 0.007^{+1.773}_{-1.837}) \cdot 10^{-13}$
6.30–6.45	$6.37 \pm 0.05$	$(2.1693 \pm 0.0029^{+0.5496}_{-0.5886}) \cdot 10^{-13}$
6.45–6.60	$6.52 \pm 0.04$	$(8.325 \pm 0.015^{+2.104}_{-2.289}) \cdot 10^{-14}$
6.60–6.70	$6.65 \pm 0.04$	$\uparrow (3.191 \pm 0.011^{+1.146}_{-0.593}) \cdot 10^{-14}$
6.70–6.80	$6.75 \pm 0.04$	$(1.626 \pm 0.006^{+0.529}_{-0.593}) \cdot 10^{-14}$
6.80–6.90	$6.85 \pm 0.04$	$(8.85 \pm 0.04^{+3.18}_{-3.39}) \cdot 10^{-15}$
6.90–7.00	$6.95 \pm 0.04$	$(3.353 \pm 0.027^{+1.734}_{-1.885}) \cdot 10^{-15}$

The conservative composition uncertainty is of the order of 23–54%. The influence of environmental variations from run to run is reflected in the systematic budget with an energy-dependent contribution between 13% and 35%. The effective area and solid angle lead to a comparably minor fraction of the overall uncertainty, with 4.6–5.8%.

The most notable difference in the set of systematic uncertainties of the GSF-informed method with respect to the agnostic approach is the reduced contribution of the composition uncertainty (3.2–8.2%), which is on a similar level as the acceptance uncertainty. In both scenarios,

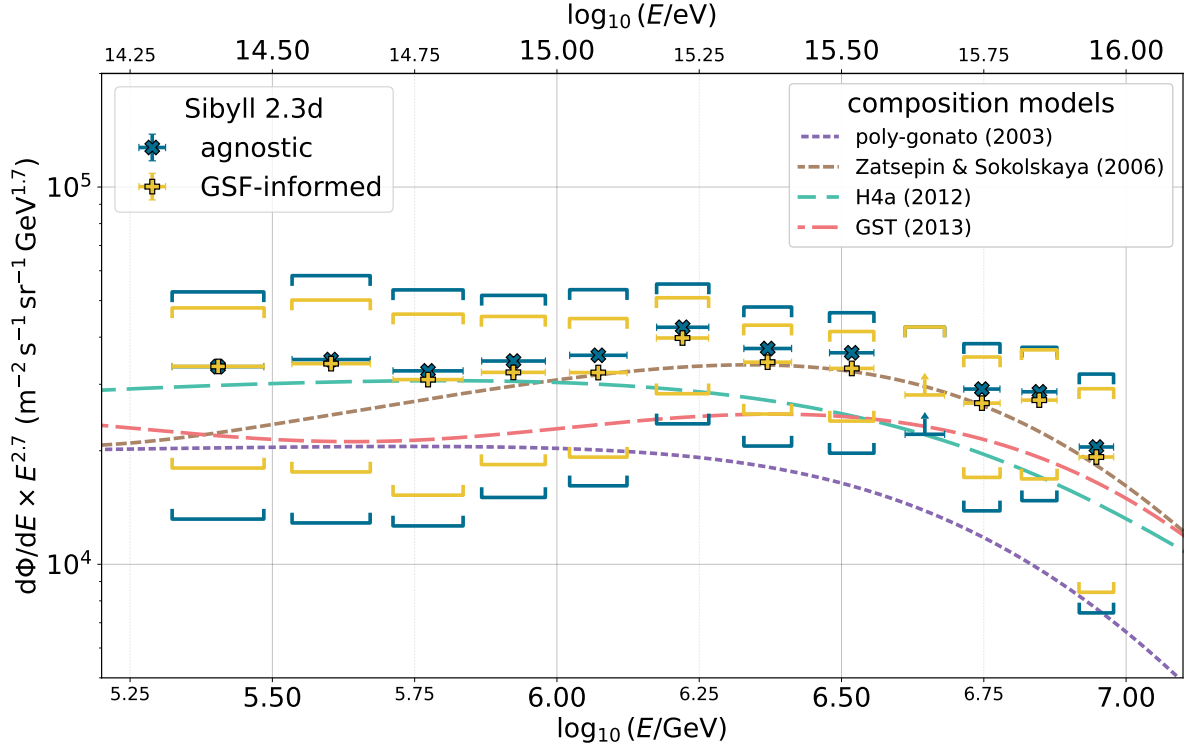


**Figure 9.5:** The reconstructed light fraction based on the Sibyll 2.3d hadronic interaction model (green), compared with the results obtained from the same analysis procedure using QGSJet-II.04 (red) and EPOS-LHC (blue) for the Monte Carlo production, respectively. The top panel shows the agnostic results, whereas the bottom panel compares the fractions obtained with the GSF-informed method. Markers are located horizontally at the mean logarithmic energy in the respective bin.

the total uncertainty is minimal around  $10^{6.5} \text{ GeV}$ . As primary energy increases, statistics decline in both data (due to the steep spectrum, resulting in larger run-to-run fluctuations) and **MC** (leading to fewer possible subsets and thus a larger sampling uncertainty). The uncertainty on the pressure correction factors is included in the total systematic as well, but since its relative scale is on the order of permille it is not shown in Figure 9.4. The reconstructed light flux values are condensed in Table 9.2 for the agnostic method and in Table 9.3 for the GSF result.

### 9.3 Comparison of Hadronic Interaction Models

All described analysis steps can be performed with alternative simulation datasets that use QGSJet-II.04 and EPOS-LHC as hadronic interaction models, respectively. While the neural network model remains unchanged, the energy bias correction, the calibration of the predicted light fraction and the quantification of acceptance and environmental uncertainty have to be conducted again for both models. The histograms and sigmoid fits that have



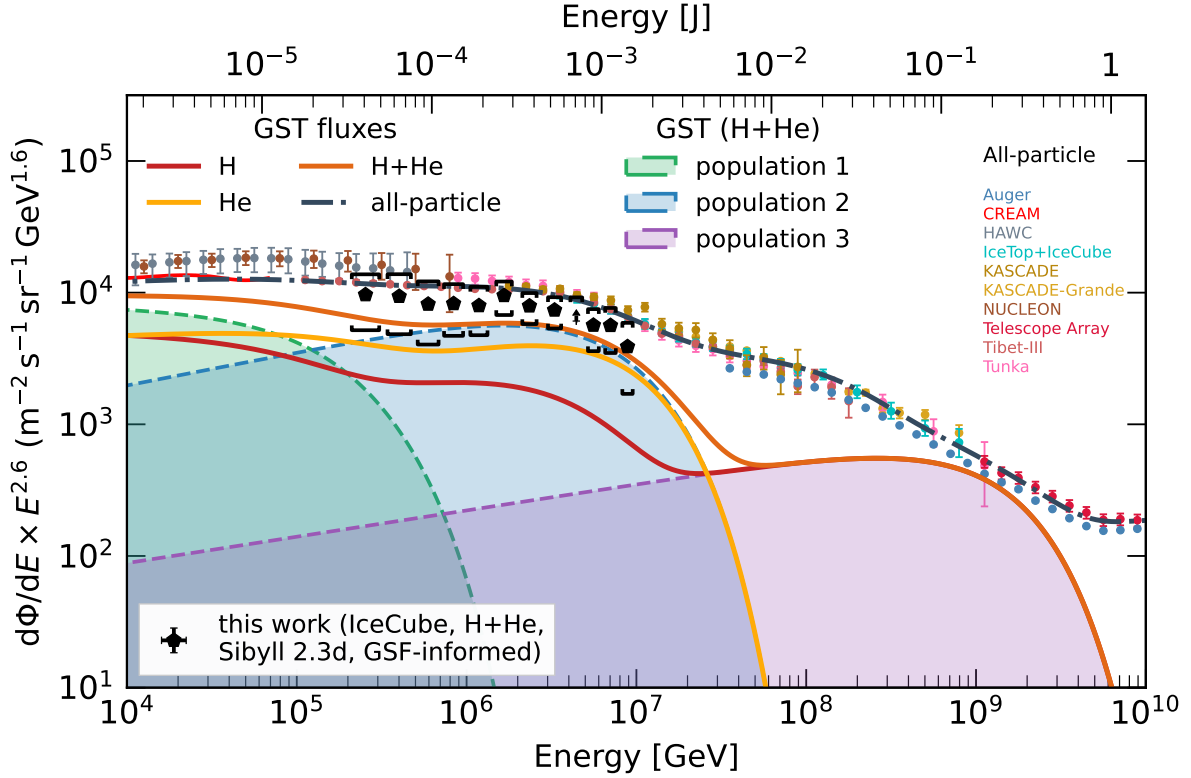
**Figure 9.6:** Compilation of the reconstructed light flux using the agnostic and the GSF-informed method. Systematic uncertainties are represented by brackets. The combined proton and helium flux of four different composition models is included for comparison.

been determined for the light fraction calibration with QGSJet and EPOS Monte Carlo can be found in Figures B.3–B.6 in Appendix B. The calibration in the energy bin  $10^{5.3}$ – $10^{5.5}$  GeV suffers from the reduced statistics at low energy in the alternative datasets. Therefore, the energy range in plots including results that are based on QGSJet or EPOS simulations only start at  $10^{5.5}$  GeV.

The resulting distributions of light fractions are mostly in agreement within their combined sampling and composition uncertainty ( $\rightarrow$  Fig. 9.5). The heaviest composition is predicted by the analysis based on EPOS-LHC. QGSJet-II.04 on the other hand generally leads to a higher fraction of light cosmic-ray primaries compared to EPOS and Sibyll. Even with the GSF-informed method, more than half of the fractions obtained with QGSJet are defined by lower limits — some even higher than the nominal fractions obtained from the other two hadronic interaction models. Refer to Appendix C for the flux values and plots as well as systematic uncertainties of the results based on QGSJet and EPOS.

#### 9.4 Comparison with Composition Models

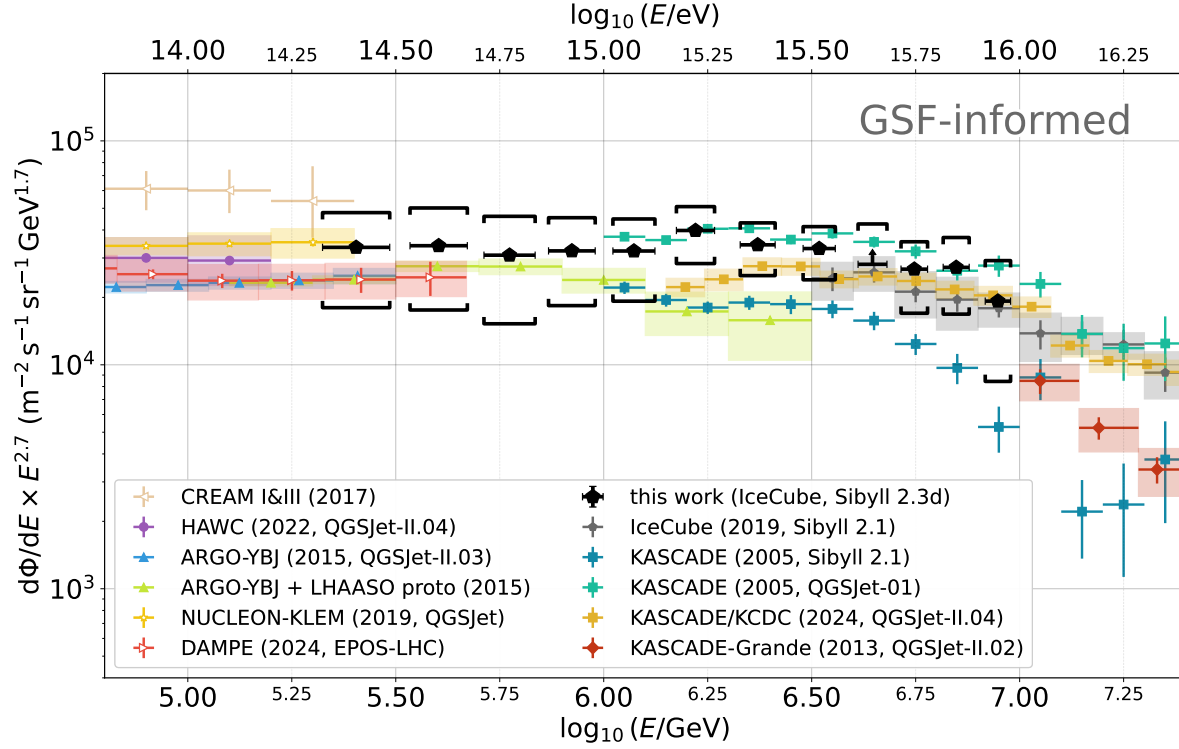
The reconstructed light flux emerging from the GSF-informed approach and its total uncertainty are contained entirely within the systematic uncertainty interval of the conservative agnostic result, as shown in Figure 9.6. In comparison of the result with flux models of the past 25 years, none stands out as a perfect match. The light H4a flux [15] reveals good agreement up to about 1.5 PeV ( $\approx 10^{6.2}$  GeV) where it begins to fall faster than the reconstructed



**Figure 9.7:** Adaptation of Figure 1.1, which is adopted from Ref. [11]. The energy is zoomed in on 10 TeV–10 EeV, only all-particle measurements in this range are shown and the scaling has been changed to  $E^{2.6}$ . The reconstructed light flux (black markers) is shown together with the H+He spectrum (orange) of GST, both broken down by primary type (red and yellow) as well as source population (green, blue and purple).

H+He flux, which rather implies a continuation of the power-law spectrum with  $\gamma \approx 2.7$  until it softens at about 5 PeV. From 2 PeV ( $\approx 10^{6.3}$  GeV) on, the flux appears to be best described by the model of Zatsepin and Sokolskaya [127]. The obtained light flux only barely contains the poly-gonato model [14] within the systematic uncertainty at low energy. The reconstructed flux above 2 PeV even excludes poly-gonato with both calibration methods at the 68% level.

The reconstructed flux of cosmic-ray protons and helium nuclei is systematically higher than that of all four composition models it is compared to in Figure 9.6. It is worth pointing out that all of them [14–16, 127] have been developed before the release of Sibyll 2.3d [81]. Figure 9.7 contrasts the light cosmic-ray flux obtained with the GSF-informed method with a collection of all-particle measurements in the TeV to EeV energy range. The contributions to the H+He flux from populations 1, 2 (both Galactic origin) and 3 (extragalactic) are included as well as the total proton and helium fluxes according to the GST model [16]. Between 1 PeV and 10 PeV the measured light flux closely follows the shape prescribed by the second population of the GST H+He spectrum. Below 700 TeV, however, the increasing dominance of the first population of Galactic cosmic rays becomes apparent.



**Figure 9.8:** The light cosmic-ray spectrum measured with one year of IceCube data (black). The reconstruction is based on simulations using the Sibyll 2.3d hadronic interaction model and the fraction calibration is informed by the GSF model. Error bars represent energy resolution and statistical uncertainty in the horizontal and vertical direction, respectively. The total systematic uncertainty is indicated by brackets. Those do not include ice-model systematics. The combined proton and helium flux as measured in the past by IceCube and other experiments is shown in various colors. Boxes indicate energy bin size and systematic uncertainty. Open markers represent results obtained via direct detection. For references, see the text.

### 9.5 Comparison with Other Experiments

In order to set the reconstructed light spectrum into context, it is compared with results from various other experiments as well as a previously published proton and helium flux by IceCube [70]<sup>4</sup>. Besides the new IceCube result obtained in this work, Figure 9.8 shows the combined proton and helium fluxes from CREAM [128]<sup>5</sup>, HAWC [129], ARGO-YBJ and a LHAASO prototype [130, 131], NUCLEON-KLEM [132]<sup>5</sup>, DAMPE [133], KASCADE [134, 135]<sup>5</sup> and KASCADE-Grande [136]<sup>5</sup>. The low-energy light spectrum agrees well with direct observations and ground-based air-shower experiments that are sensitive below 300 TeV, such as ARGO-YBJ and HAWC. On the high-energy side, the challenges that come with the indirect observation of cosmic rays lead to larger variations. Differences in the detector setup and type, the observation altitude and the hadronic interaction model used as a base for each analysis can affect the eventually reconstructed light flux. The new light flux result is in agreement with the QGSjet flux from KASCADE and the combined proton and helium flux measured by IceCube in the PeV to EeV energy range. A knee just below 1 PeV as declared by

<sup>4</sup>The proton and helium spectra are added and correlation-corrected for this comparison.

<sup>5</sup>Individual spectra for protons and helium nuclei have been published and are added here.

ARGO-YBJ cannot be confirmed, instead a tension with the softer ARGO-YBJ spectrum can be observed starting at about 1.5 PeV. A similar disagreement is present with the Sibyll 2.1 result of KASCADE but it follows the same spectral shape as the light flux obtained in this work based on Sibyll 2.3d. Post-LHC hadronic interaction models generally produce higher muon multiplicities than older models such as Sibyll 2.1 [81]. Even though a direct comparison of the light flux obtained with Sibyll 2.1 and a variant of Sibyll 2.3 in the PeV energy range does not currently exist, previous studies [135, 137] imply a lighter composition as a consequence of using post-LHC models for the interpretation of cosmic-ray data.

The best-fit proton+helium spectrum obtained via the GSF-informed approach suggests a slight deviation from a pure  $E^{-2.7}$  power law with a marginally steeper slope below 500 TeV, whereas between 1 PeV and 2 PeV the spectrum appears to soften. This is the same trend as predicted by the transition from population 1 to population 2 in the GST flux model. An indication of a light knee is observed around 8 PeV.

### Key Messages

*The analysis developed in the previous chapters is applied to one year of IceCube data. The calibrated light fraction confirms that cosmic rays in the PeV energy range are mostly protons and helium nuclei, between 60% and 80% of the total flux; that is, the light component dominates in this energy range. The differential H+He spectrum follows an  $E^{-2.7}$  power law in first approximation, however, a dip can be observed at about 500 TeV, similar to the prediction of the GST model. A light knee appears at an energy of slightly above 8 PeV. Significant tensions are neither observed among hadronic interaction models nor with composition models. This is the first measurement of the H+He flux in the energy range from 300 TeV to 1.4 PeV with post-LHC models, closing the gap between direct and indirect measurements.*



# CHAPTER 10

## CONCLUSION

*“Kids, you tried your best and you failed miserably. The lesson is... never try.”*

— Homer Simpson, *The Simpsons*

In this thesis, the spectrum of light cosmic rays, that is, protons and helium nuclei, between 200 TeV and 10 PeV has been reconstructed. Data of air-shower events taken over the course of one year by the IceCube Neutrino Observatory has been analyzed to identify light cosmic rays and determine their primary energy.

Since this is the first composition study at IceCube that includes low-energy events with only a few hundred TeV primary energy, a new large database of event simulations had to be created. It includes three of the latest hadronic interaction models (Sibyll 2.3d, QGSJet-II.04 and EPOS-LHC) and five different primary particle types (protons, helium, oxygen, silicon and iron). Both the shower development and the detector response have been simulated. For the first time, the in-ice muon background caused by unrelated coincident air showers has been included in the simulation at IceCube. The validity of the new Monte Carlo datasets has been verified on low-level pulse information at several data processing stages as well as on reconstructed high-level event parameters.

A specific kind of muon-rich pulses in the surface array IceTop is a valuable observable for investigating the composition of cosmic rays. The charge calibration of those SLC pulses had to undergo a revision to make them reliably usable in the sub-PeV energy range.

The standard processing of cosmic-ray events had to be extended for the specialized event selection of coincidences between IceTop and the in-ice array. Among other features, this new Level 4 processing includes the rejection of in-ice background and a hybrid air-shower reconstruction that simultaneously fits to the signal recorded in both sub-detectors.

A neural network has been constructed to process a collection of 15 input features taken from IceTop, the in-ice array and the hybrid shower reconstruction. The model has been trained and applied to events that have passed a set of newly developed quality cuts. Those selection criteria ensure good agreement between the simulated training data and the measured data taken at the South Pole. The resolution of shower core position and incident direction obtained from the hybrid reconstruction improves significantly compared to the standard algorithm.

The two network outputs are estimators for energy and mass of the primaries and needed to be further calibrated to obtain the energy-binned differential flux of light cosmic rays. For estimating the fraction of light primaries, two distinct methods have been developed. One is based on an equal mix of primaries in the light and heavy component, respectively, with an  $E^{-2.7}$  spectrum (the agnostic method) and defines very conservative limits on the composition of light and heavy cosmic rays. The second approach is seeded with the

## CHAPTER 10. CONCLUSION

GSF model of cosmic-ray composition and has its uncertainties oriented along the confidence of that model. Recovery tests have been performed to demonstrate the reliability of the light fraction calibration and the reconstructed H+He flux.

The analysis result confirms a dominant light component in the cosmic-ray flux below and around the knee. The combined fraction of protons and helium ranges from 60% (300–700 TeV) to almost 80% (1.5–4 PeV). Between 4 PeV and 5 PeV it is only possible to determine a lower limit on the light fraction. The agnostic approach systematically predicts a slightly lighter composition compared to the GSF-informed method.

In first order, the reconstructed light spectrum agrees with an  $E^{-2.7}$  power law. However, after scaling with  $E^{2.7}$  indications for a concave spectral feature can be observed around 500 TeV, followed by a soft knee-like structure above 1 PeV. This is in accordance with the transition from the first to the second Galactic population in the GST model. Due to the high flux of PeV cosmic rays and the large acceptance of IceCube, statistical uncertainties on the resulting spectrum are negligible. On the other hand, large systematic uncertainties restrain clear evidences of the apparent spectral features. The agnostic method is dominated by the conservative limits on the assumed composition, whereas systematic effects arising from the calibration method and changing environmental conditions are the largest uncertainties of the GSF-informed result. No tensions are observed between the results obtained with different hadronic interaction models, although QGSJet-II.04 tends to predict a lighter composition than Sibyll 2.3d, whereas EPOS-LHC yields systematically lower fractions of the light component, even dropping below 40% beyond 8 PeV. The reconstructed light flux with Sibyll 2.3d is higher than predicted by established composition models, but most of them agree with the new result within systematic uncertainties.

This is the first composition study in IceCube that is based on a post-LHC version of Sibyll. Moreover, this work closes the gap in the light cosmic-ray spectrum between direct and indirect measurements. At lower energies, around 200 TeV, the reconstructed flux agrees well with observations, for example, by DAMPE and NUCLEON. On the high-energy end, the overlap with previous results by IceCube and various KASCADE spectra shows agreement, in general, as well. Merely, the light spectrum by KASCADE based on the Sibyll 2.1 model creates tension with the presented result. However, the KASCADE measurement has the same spectral shape but only differs notably in the absolute flux. Furthermore, the sub-PeV light knee declared by ARGO-YBJ is not confirmed by this analysis.

A possible continuation of this work in the future would need to reduce the systematic uncertainties in order to make more significant statements on features in the spectrum of protons and helium nuclei. Template fits might be a promising alternative to the classification of light events based on the proton score with subsequent fraction calibration. There is also a potential of reducing the environmental uncertainty on the differential flux, which is a major contribution in this analysis, by developing a pressure correction not only on the average atmosphere but on short time scales, like on the basis of data-taking runs or on even shorter intervals. On that note, despite the increasing need for computing resources and storage space, simulating more extreme, that is, summer and winter, atmospheres in addition to the

average April model, will be of great benefit for all sorts of cosmic-ray analyses that use IceTop data. In this analysis, the magnitude of the systematic uncertainty due to the ice model could not be investigated. Such a study would provide valuable insight into how much uncertainties on the most recent ice model propagate to the final result. The deployment of the surface array enhancement in the following years will help lower the energy threshold. This will improve the sensitivity to low-energy proton and helium events but will also allow the analysis of heavier primaries in the PeV energy range. Lastly, during the work on this thesis, the hadronic interaction models have received updates. The effect of those changes could be studied by producing new Monte Carlo datasets based on QGSJet-III.01 [138], EPOS-LHC-R [139] and Sibyll★ [140].

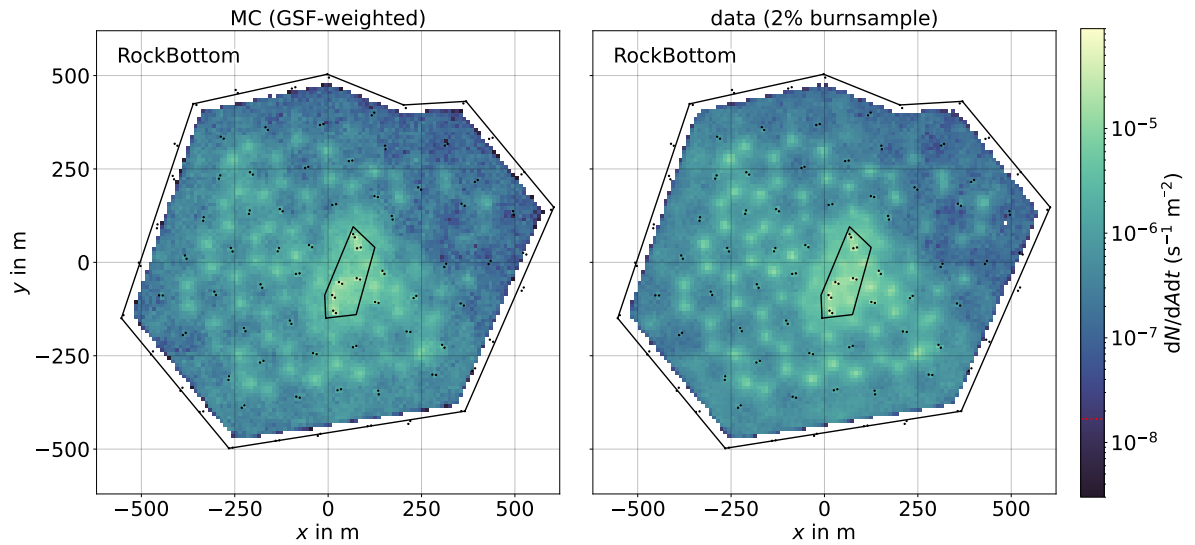
This analysis has demonstrated that it is possible to measure the energy and study the composition of cosmic rays below 1 PeV with IceCube. The result provides a test of the validity of hadronic interaction models and contributes to the understanding of Galactic cosmic radiation.



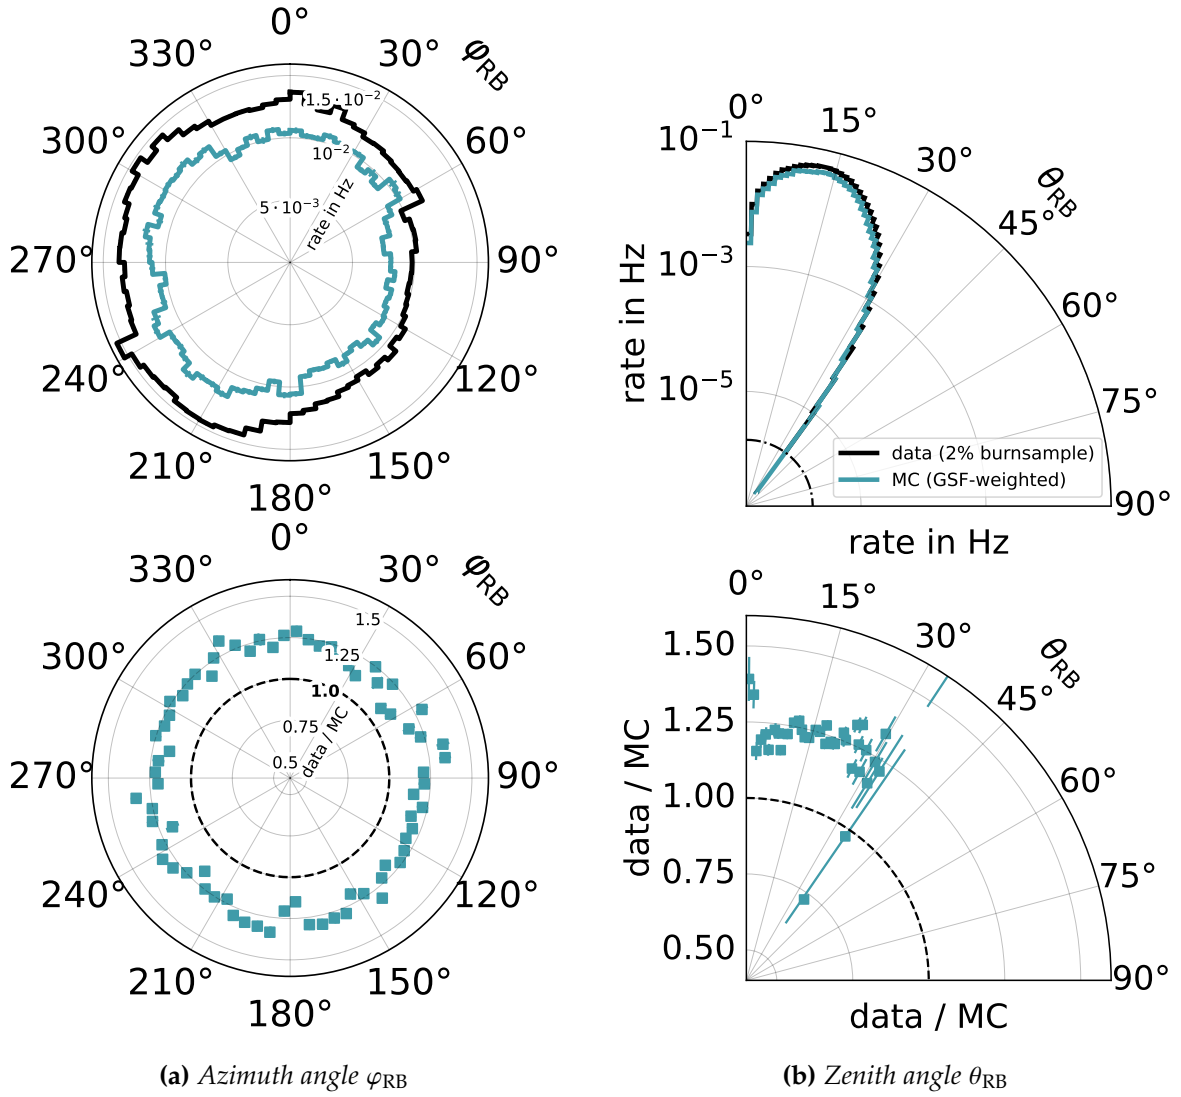
# APPENDIX A

## ADDENDUM MONTE CARLO VERIFICATION

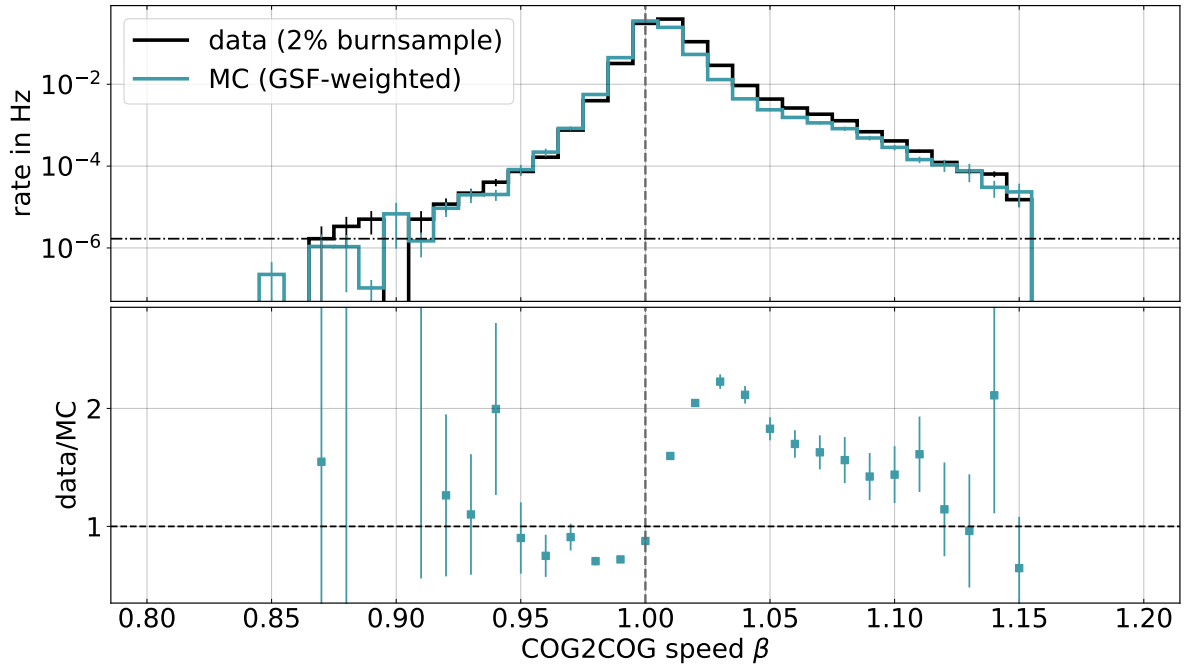
In this appendix, additional high-level data/MC verification plots of the Sibyll 2.3d dataset are collected. Quality cuts are applied. The spatial distributions of core positions as reconstructed by RockBottom are shown in Figure A.1 for weighted simulations and burnsample data. Data-MC agreement in the distribution of reconstructed RB arrival direction are presented in Figure A.2, broken down by zenith and azimuth angle. A slight shift in the distributions of the COG2COG speed are visible in Figure A.3. Even though this is not understood yet, it is not considered concerning since the scale of this shift is well contained within the allowed regime of the background rejection cut ( $\rightarrow$  Fig. 6.12). Good agreement can be observed in Figure A.4 for the angle  $\Psi$  between the hypothetical COG2COG track and the pure in-ice SPE reconstruction.



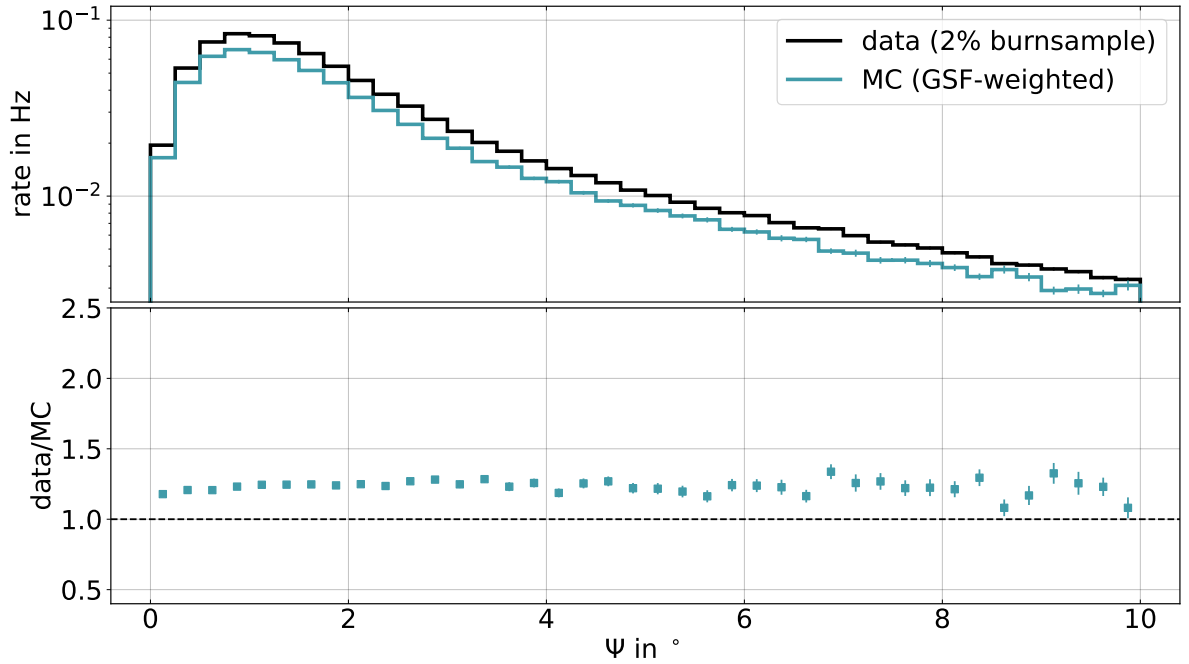
**Figure A.1:** Reconstructed RockBottom core positions after quality cuts for the GSF-weighted simulation (left) and the 2% burnsample data (right). Due to the limited exposure time of the burnsample, rates below  $1.68 \cdot 10^{-8} \text{ Hz/m}^2$  can only be achieved by weighted MC but not in data (red dotted line in the colorbar). The outline of the detector array and the InFill are indicated by black lines.



**Figure A.2:** Reconstructed RockBottom direction components with GSF-weighted MC (teal) and burnsample data (black) in the top panels and the data/MC ratio at the bottom.



**Figure A.3:** The speed of the hypothetical COG2COG connection track in data (black) and weighted simulations (teal). The bottom panel shows the data/MC ratio.



**Figure A.4:** The angle  $\Psi$  between the COG2COG track and the in-ice only SPE reconstruction track. The top panel shows data and GSF-weighted MC, the bottom panel has the data/MC ratio.



# APPENDIX B

## LIGHT FRACTION CALIBRATION

Tables B.1 and B.2 contain the fit parameters of the sigmoid function (8.1) to the 16-, 32-, 50-, 68- and 84-percentiles of the 2-dimensional fraction calibration histograms ( $\rightarrow$  Figs. B.1 and B.2) for the agnostic and the GSF-informed method, applied to Sibyll 2.3d Monte Carlo. Furthermore, the histograms for QGSJet-II.04 (Figs. B.3 and B.4) and EPOS-LHC (Figs. B.5 and B.6) are provided.

**Table B.1:** Sigmoid best-fit parameters  $p_0$  and  $p_1$  for Eq. (8.1) with the agnostic version and the Sibyll 2.3d hadronic interaction model.

$\log_{10}(E/\text{GeV})$	percentile	$p_0$	$p_1$
5.10–5.30	16	$20.48 \pm 0.22$	$0.3832 \pm 0.0005$
	32	$19.88 \pm 0.20$	$0.3572 \pm 0.0004$
	50	$20.29 \pm 0.21$	$0.3319 \pm 0.0005$
	68	$21.67 \pm 0.24$	$0.3079 \pm 0.0006$
	84	$24.7 \pm 0.4$	$0.2834 \pm 0.0009$
5.30–5.50	16	$29.27 \pm 0.22$	$0.36490 \pm 0.00026$
	32	$27.53 \pm 0.18$	$0.34958 \pm 0.00023$
	50	$27.36 \pm 0.18$	$0.33526 \pm 0.00023$
	68	$29.01 \pm 0.19$	$0.32135 \pm 0.00024$
	84	$32.77 \pm 0.24$	$0.30806 \pm 0.00028$
5.50–5.70	16	$30.21 \pm 0.19$	$0.33300 \pm 0.00020$
	32	$28.03 \pm 0.15$	$0.32081 \pm 0.00019$
	50	$27.29 \pm 0.14$	$0.30949 \pm 0.00019$
	68	$28.41 \pm 0.16$	$0.29658 \pm 0.00022$
	84	$30.98 \pm 0.21$	$0.28272 \pm 0.00026$
5.70–5.85	16	$26.97 \pm 0.18$	$0.30775 \pm 0.00024$
	32	$24.91 \pm 0.16$	$0.29350 \pm 0.00023$
	50	$23.97 \pm 0.15$	$0.27921 \pm 0.00024$
	68	$24.83 \pm 0.17$	$0.26479 \pm 0.00026$
	84	$27.46 \pm 0.21$	$0.2480 \pm 0.0004$

*Continued on next page*

**Table B.1** – *continued from previous page*

$\log_{10}(E/\text{GeV})$	percentile	$p_0$	$p_1$
5.85–6.00	16	$24.88 \pm 0.16$	$0.29014 \pm 0.00026$
	32	$23.15 \pm 0.14$	$0.27498 \pm 0.00024$
	50	$22.50 \pm 0.13$	$0.26106 \pm 0.00024$
	68	$23.03 \pm 0.14$	$0.24760 \pm 0.00025$
	84	$25.38 \pm 0.18$	$0.23161 \pm 0.00029$
6.00–6.15	16	$21.75 \pm 0.13$	$0.29537 \pm 0.00025$
	32	$20.74 \pm 0.12$	$0.28169 \pm 0.00024$
	50	$20.60 \pm 0.12$	$0.26839 \pm 0.00023$
	68	$21.13 \pm 0.12$	$0.25543 \pm 0.00023$
	84	$22.33 \pm 0.13$	$0.24187 \pm 0.00025$
6.15–6.30	16	$21.88 \pm 0.13$	$0.27213 \pm 0.00026$
	32	$20.77 \pm 0.12$	$0.25782 \pm 0.00023$
	50	$20.42 \pm 0.12$	$0.24525 \pm 0.00023$
	68	$20.70 \pm 0.12$	$0.23288 \pm 0.00023$
	84	$21.88 \pm 0.13$	$0.21941 \pm 0.00025$
6.30–6.45	16	$20.02 \pm 0.11$	$0.26221 \pm 0.00025$
	32	$19.01 \pm 0.10$	$0.24825 \pm 0.00023$
	50	$18.79 \pm 0.10$	$0.23587 \pm 0.00022$
	68	$19.13 \pm 0.10$	$0.22308 \pm 0.00023$
	84	$20.00 \pm 0.11$	$0.21002 \pm 0.00024$
6.45–6.60	16	$20.36 \pm 0.11$	$0.26142 \pm 0.00025$
	32	$19.33 \pm 0.10$	$0.24784 \pm 0.00022$
	50	$18.86 \pm 0.09$	$0.23546 \pm 0.00021$
	68	$18.95 \pm 0.10$	$0.22307 \pm 0.00022$
	84	$19.78 \pm 0.10$	$0.21012 \pm 0.00023$
6.60–6.70	16	$20.97 \pm 0.16$	$0.2473 \pm 0.0004$
	32	$19.60 \pm 0.13$	$0.22856 \pm 0.00030$
	50	$18.86 \pm 0.13$	$0.21194 \pm 0.00029$
	68	$18.95 \pm 0.13$	$0.19552 \pm 0.00029$
	84	$19.92 \pm 0.13$	$0.1791 \pm 0.0004$

*Continued on next page*

**Table B.1** – *continued from previous page*

$\log_{10} (E/\text{GeV})$	percentile	$p_0$	$p_1$
6.70–6.80	16	$18.50 \pm 0.12$	$0.2722 \pm 0.0004$
	32	$17.33 \pm 0.10$	$0.2533 \pm 0.0004$
	50	$16.82 \pm 0.10$	$0.23594 \pm 0.00029$
	68	$16.95 \pm 0.10$	$0.21999 \pm 0.00029$
	84	$17.56 \pm 0.11$	$0.20347 \pm 0.00030$
6.80–6.90	16	$19.62 \pm 0.12$	$0.2711 \pm 0.0004$
	32	$18.27 \pm 0.10$	$0.25423 \pm 0.00028$
	50	$17.81 \pm 0.10$	$0.23873 \pm 0.00027$
	68	$18.14 \pm 0.10$	$0.22364 \pm 0.00027$
	84	$19.45 \pm 0.11$	$0.20743 \pm 0.00028$
6.90–7.00	16	$15.83 \pm 0.09$	$0.3047 \pm 0.0004$
	32	$14.92 \pm 0.08$	$0.2854 \pm 0.0004$
	50	$14.46 \pm 0.07$	$0.2647 \pm 0.0004$
	68	$14.38 \pm 0.07$	$0.24763 \pm 0.00029$
	84	$14.91 \pm 0.07$	$0.23260 \pm 0.00027$

**Table B.2:** *Sigmoid best-fit parameters  $p_0$  and  $p_1$  for Eq. (8.1) with the GSF-informed version and the Sibyll 2.3d hadronic interaction model.*

$\log_{10} (E/\text{GeV})$	percentile	$p_0$	$p_1$
5.10–5.30	16	$19.29 \pm 0.21$	$0.3750 \pm 0.0005$
	32	$19.00 \pm 0.20$	$0.3467 \pm 0.0005$
	50	$19.56 \pm 0.20$	$0.3208 \pm 0.0006$
	68	$21.06 \pm 0.24$	$0.2970 \pm 0.0008$
	84	$23.8 \pm 0.4$	$0.2736 \pm 0.0011$

*Continued on next page*

## APPENDIX B. LIGHT FRACTION CALIBRATION

**Table B.2** – *continued from previous page*

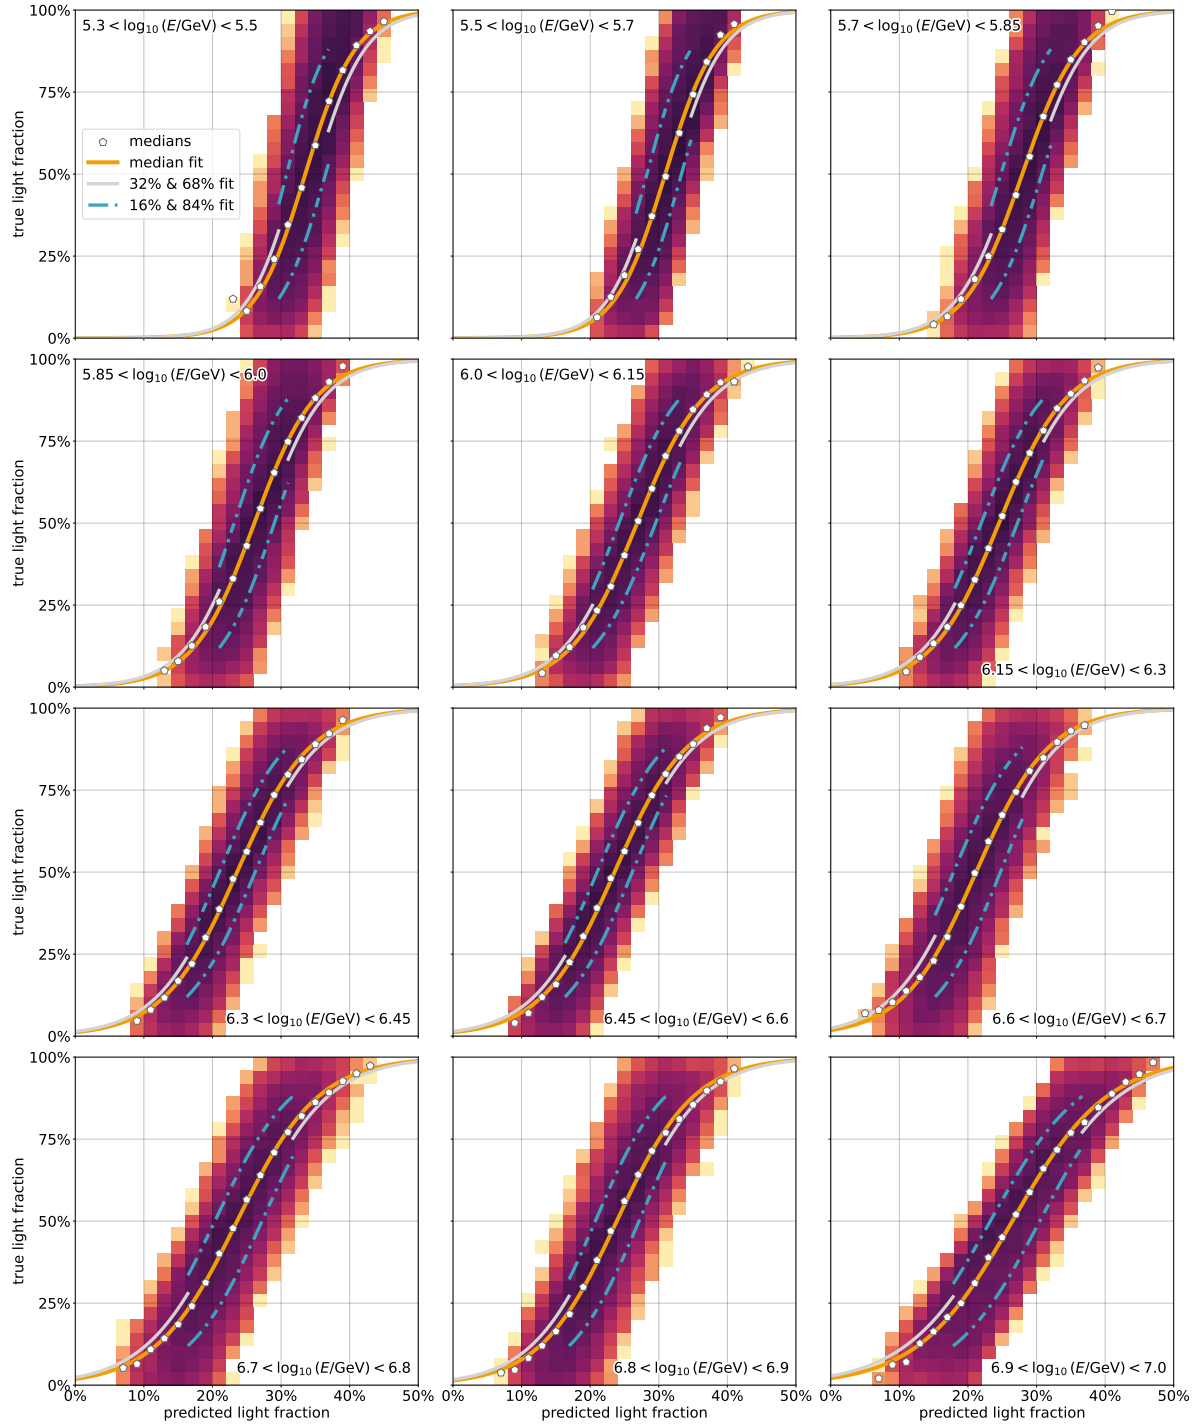
$\log_{10}(E/\text{GeV})$	percentile	$p_0$	$p_1$
5.30–5.50	16	$29.39 \pm 0.25$	$0.36752 \pm 0.00028$
	32	$27.96 \pm 0.20$	$0.35051 \pm 0.00026$
	50	$27.82 \pm 0.19$	$0.33546 \pm 0.00027$
	68	$29.64 \pm 0.20$	$0.32208 \pm 0.00029$
	84	$33.49 \pm 0.26$	$0.3099 \pm 0.0004$
5.50–5.70	16	$33.58 \pm 0.22$	$0.33643 \pm 0.00021$
	32	$30.29 \pm 0.17$	$0.32327 \pm 0.00020$
	50	$29.62 \pm 0.16$	$0.31199 \pm 0.00021$
	68	$30.73 \pm 0.17$	$0.30088 \pm 0.00023$
	84	$32.77 \pm 0.21$	$0.28856 \pm 0.00027$
5.70–5.85	16	$30.64 \pm 0.21$	$0.31409 \pm 0.00023$
	32	$27.58 \pm 0.18$	$0.29998 \pm 0.00025$
	50	$26.35 \pm 0.17$	$0.28615 \pm 0.00026$
	68	$26.72 \pm 0.18$	$0.27283 \pm 0.00029$
	84	$28.81 \pm 0.21$	$0.2577 \pm 0.0004$
5.85–6.00	16	$28.14 \pm 0.19$	$0.29985 \pm 0.00026$
	32	$25.48 \pm 0.16$	$0.28426 \pm 0.00026$
	50	$24.54 \pm 0.14$	$0.27073 \pm 0.00026$
	68	$24.66 \pm 0.15$	$0.25799 \pm 0.00027$
	84	$26.56 \pm 0.18$	$0.2435 \pm 0.0004$
6.00–6.15	16	$23.90 \pm 0.15$	$0.30993 \pm 0.00026$
	32	$22.42 \pm 0.13$	$0.29600 \pm 0.00025$
	50	$21.83 \pm 0.12$	$0.28297 \pm 0.00026$
	68	$22.06 \pm 0.12$	$0.27040 \pm 0.00027$
	84	$23.13 \pm 0.13$	$0.25832 \pm 0.00030$
6.15–6.30	16	$23.87 \pm 0.16$	$0.28721 \pm 0.00027$
	32	$22.10 \pm 0.13$	$0.27247 \pm 0.00027$
	50	$21.38 \pm 0.12$	$0.25940 \pm 0.00027$
	68	$21.40 \pm 0.12$	$0.24699 \pm 0.00027$
	84	$22.26 \pm 0.13$	$0.2340 \pm 0.0004$

*Continued on next page*

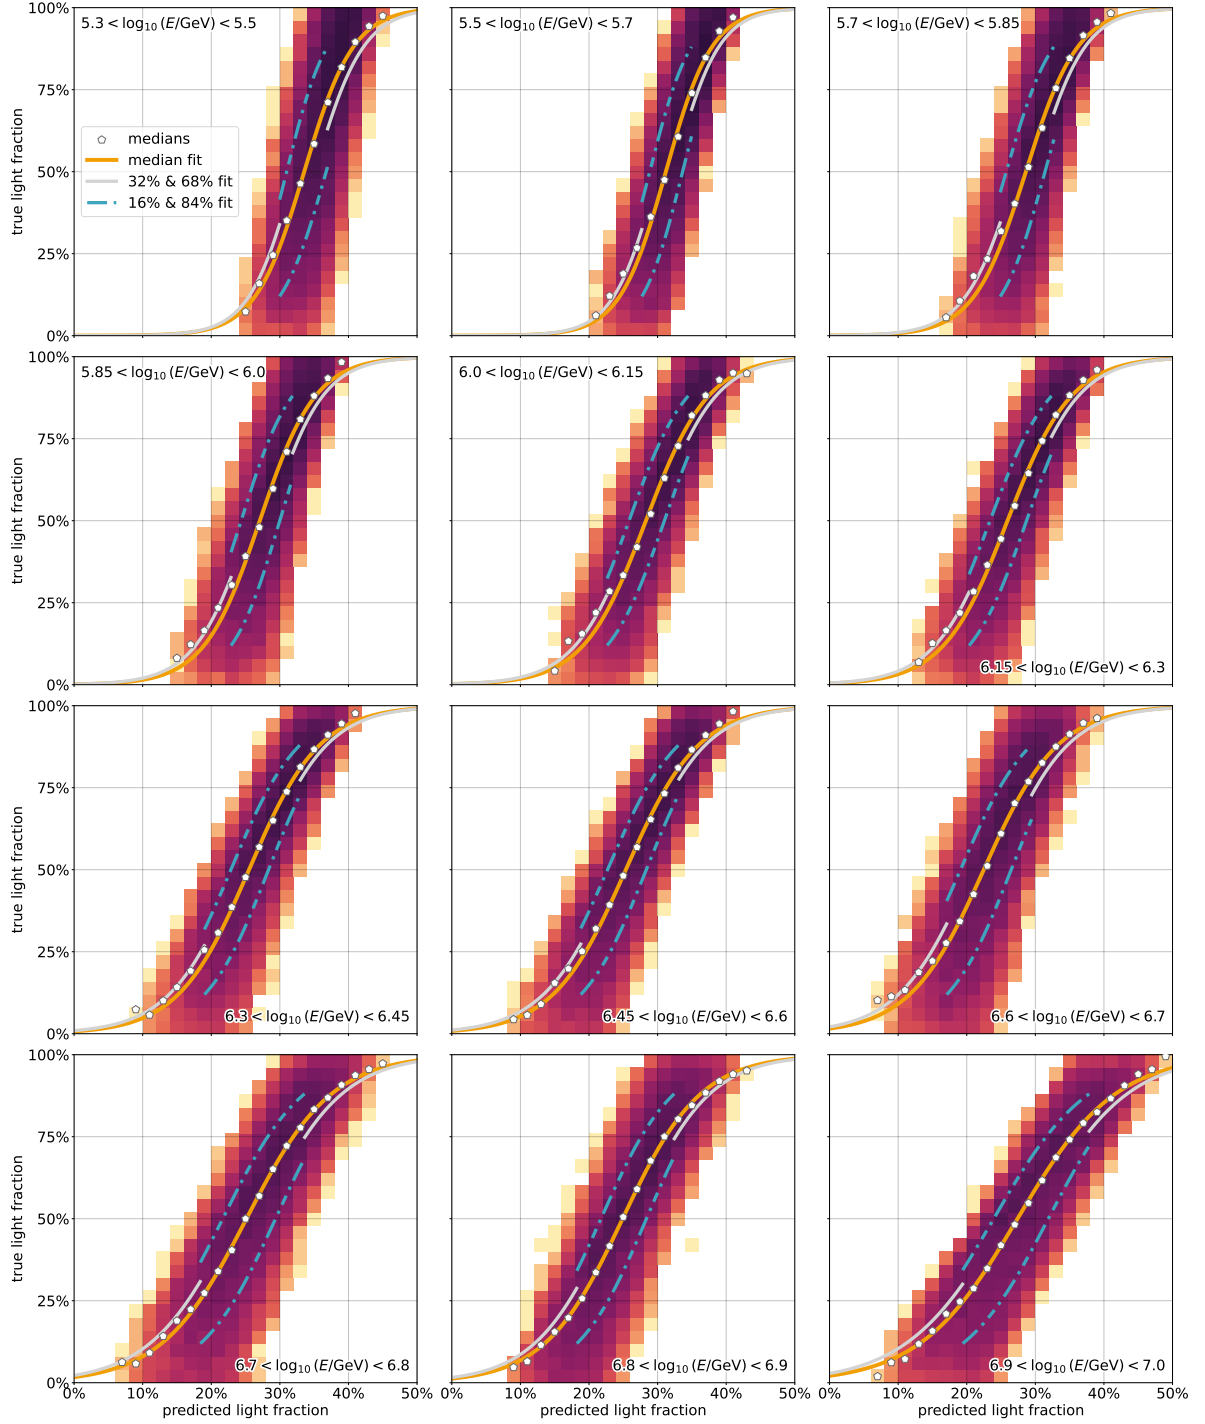
**Table B.2** – *continued from previous page*

$\log_{10}(E/\text{GeV})$	percentile	$p_0$	$p_1$
6.30–6.45	16	$21.38 \pm 0.13$	$0.28348 \pm 0.00029$
	32	$19.96 \pm 0.11$	$0.26831 \pm 0.00027$
	50	$19.33 \pm 0.10$	$0.25498 \pm 0.00027$
	68	$19.29 \pm 0.10$	$0.24187 \pm 0.00028$
	84	$19.79 \pm 0.10$	$0.22864 \pm 0.00030$
6.45–6.60	16	$20.92 \pm 0.13$	$0.28347 \pm 0.00029$
	32	$19.52 \pm 0.11$	$0.26741 \pm 0.00028$
	50	$18.77 \pm 0.10$	$0.25365 \pm 0.00027$
	68	$18.69 \pm 0.09$	$0.24022 \pm 0.00027$
	84	$19.22 \pm 0.10$	$0.22702 \pm 0.00028$
6.60–6.70	16	$21.28 \pm 0.17$	$0.2648 \pm 0.0004$
	32	$19.30 \pm 0.14$	$0.2438 \pm 0.0004$
	50	$18.44 \pm 0.13$	$0.2248 \pm 0.0004$
	68	$18.59 \pm 0.13$	$0.2076 \pm 0.0004$
	84	$19.25 \pm 0.13$	$0.1909 \pm 0.0004$
6.70–6.80	16	$18.23 \pm 0.14$	$0.2937 \pm 0.0004$
	32	$16.76 \pm 0.11$	$0.2713 \pm 0.0004$
	50	$16.11 \pm 0.10$	$0.2513 \pm 0.0004$
	68	$16.16 \pm 0.10$	$0.2338 \pm 0.0004$
	84	$16.75 \pm 0.10$	$0.2167 \pm 0.0004$
6.80–6.90	16	$19.69 \pm 0.13$	$0.2837 \pm 0.0004$
	32	$18.07 \pm 0.11$	$0.26541 \pm 0.00029$
	50	$17.54 \pm 0.10$	$0.24896 \pm 0.00029$
	68	$17.70 \pm 0.10$	$0.23328 \pm 0.00029$
	84	$18.71 \pm 0.11$	$0.2174 \pm 0.0004$
6.90–7.00	16	$16.00 \pm 0.09$	$0.3189 \pm 0.0004$
	32	$14.61 \pm 0.09$	$0.2971 \pm 0.0004$
	50	$14.00 \pm 0.08$	$0.2739 \pm 0.0004$
	68	$13.91 \pm 0.07$	$0.2547 \pm 0.0004$
	84	$14.35 \pm 0.07$	$0.23927 \pm 0.00030$

## APPENDIX B. LIGHT FRACTION CALIBRATION

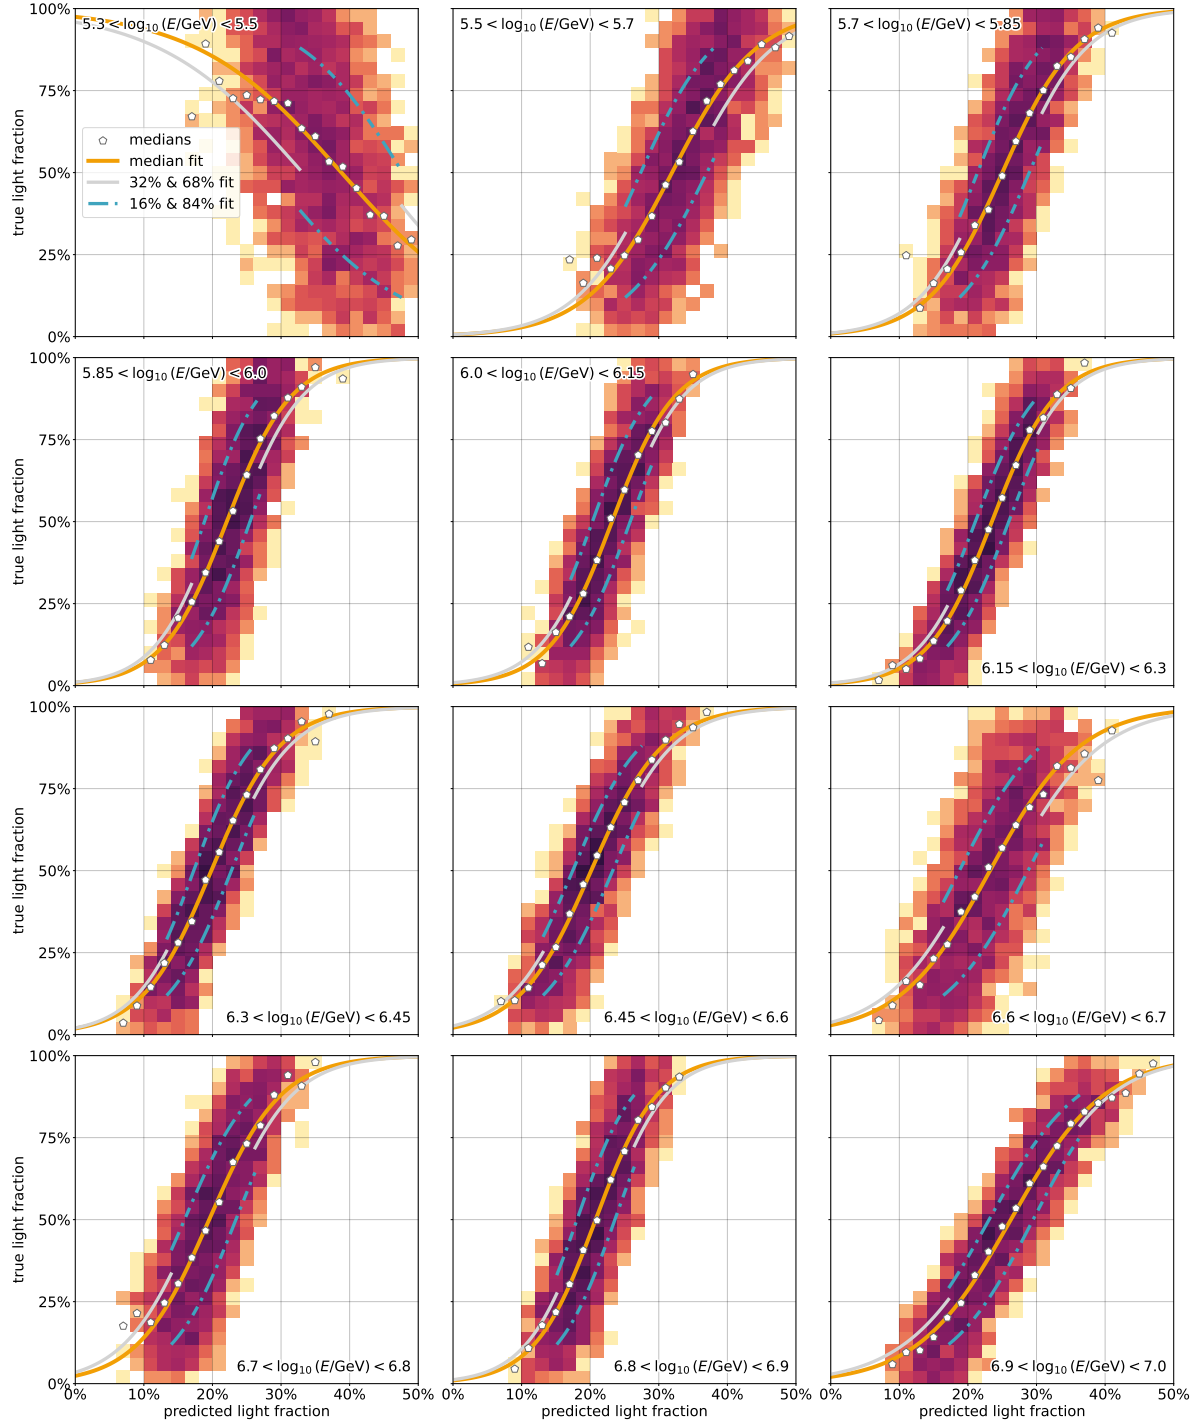


**Figure B.1:** The subset histograms with their sigmoid fits in all energy bins for the agnostic case and Sibyll 2.3d Monte Carlo simulations.

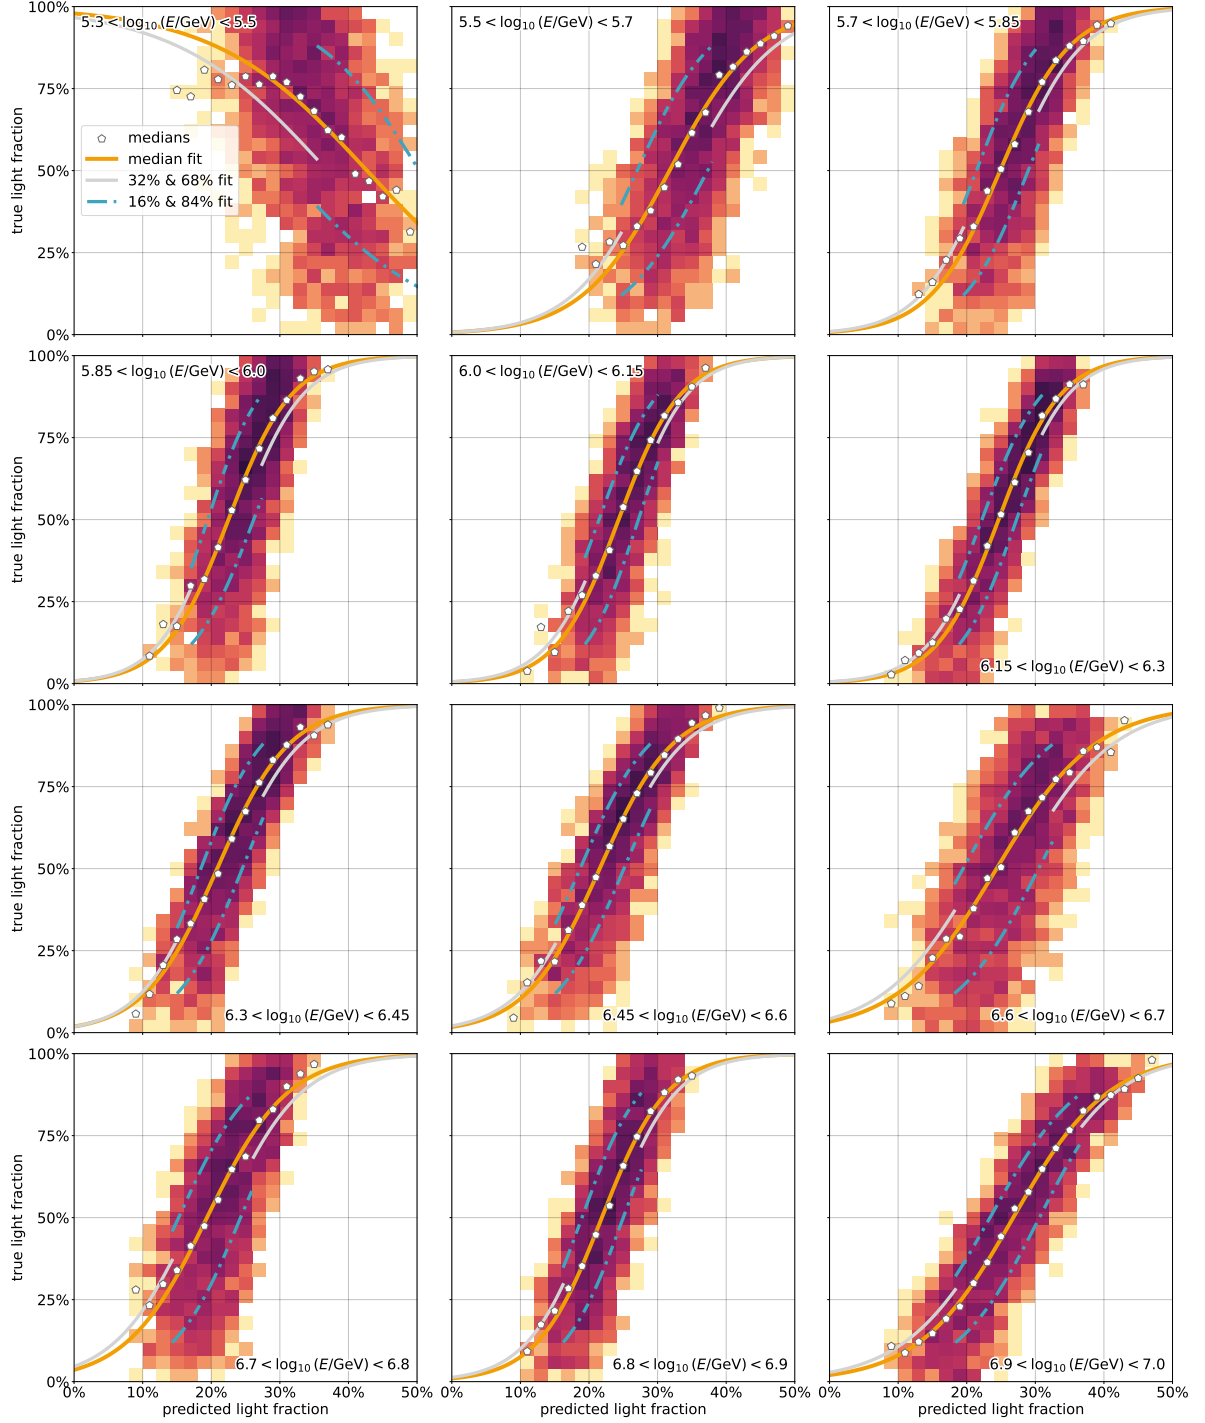


**Figure B.2:** The subset histograms with their sigmoid fits in all energy bins for the GSF-informed case and Sibyll 2.3d Monte Carlo simulations.

## APPENDIX B. LIGHT FRACTION CALIBRATION

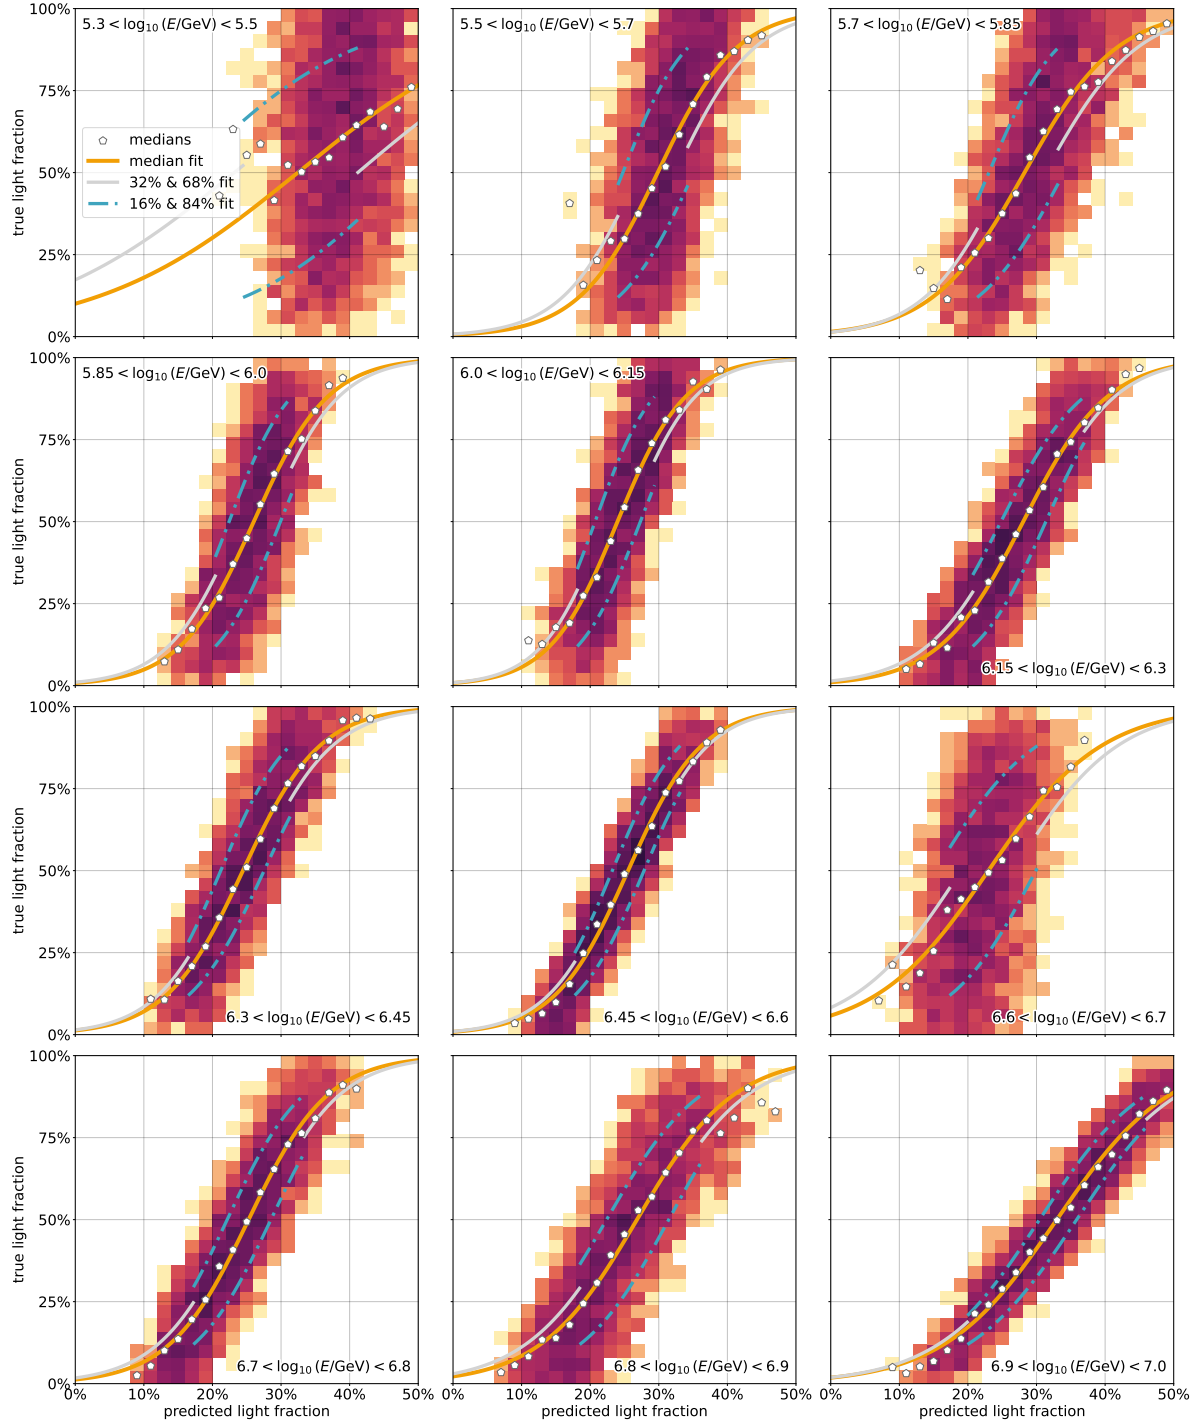


**Figure B.3:** The subset histograms with their sigmoid fits in all energy bins for the agnostic case and QGSJet-II.04 Monte Carlo simulations. The energy bin  $10^{5.3}-10^{5.5}$  GeV will be excluded in the plots of fraction and flux of the proton+helium component.

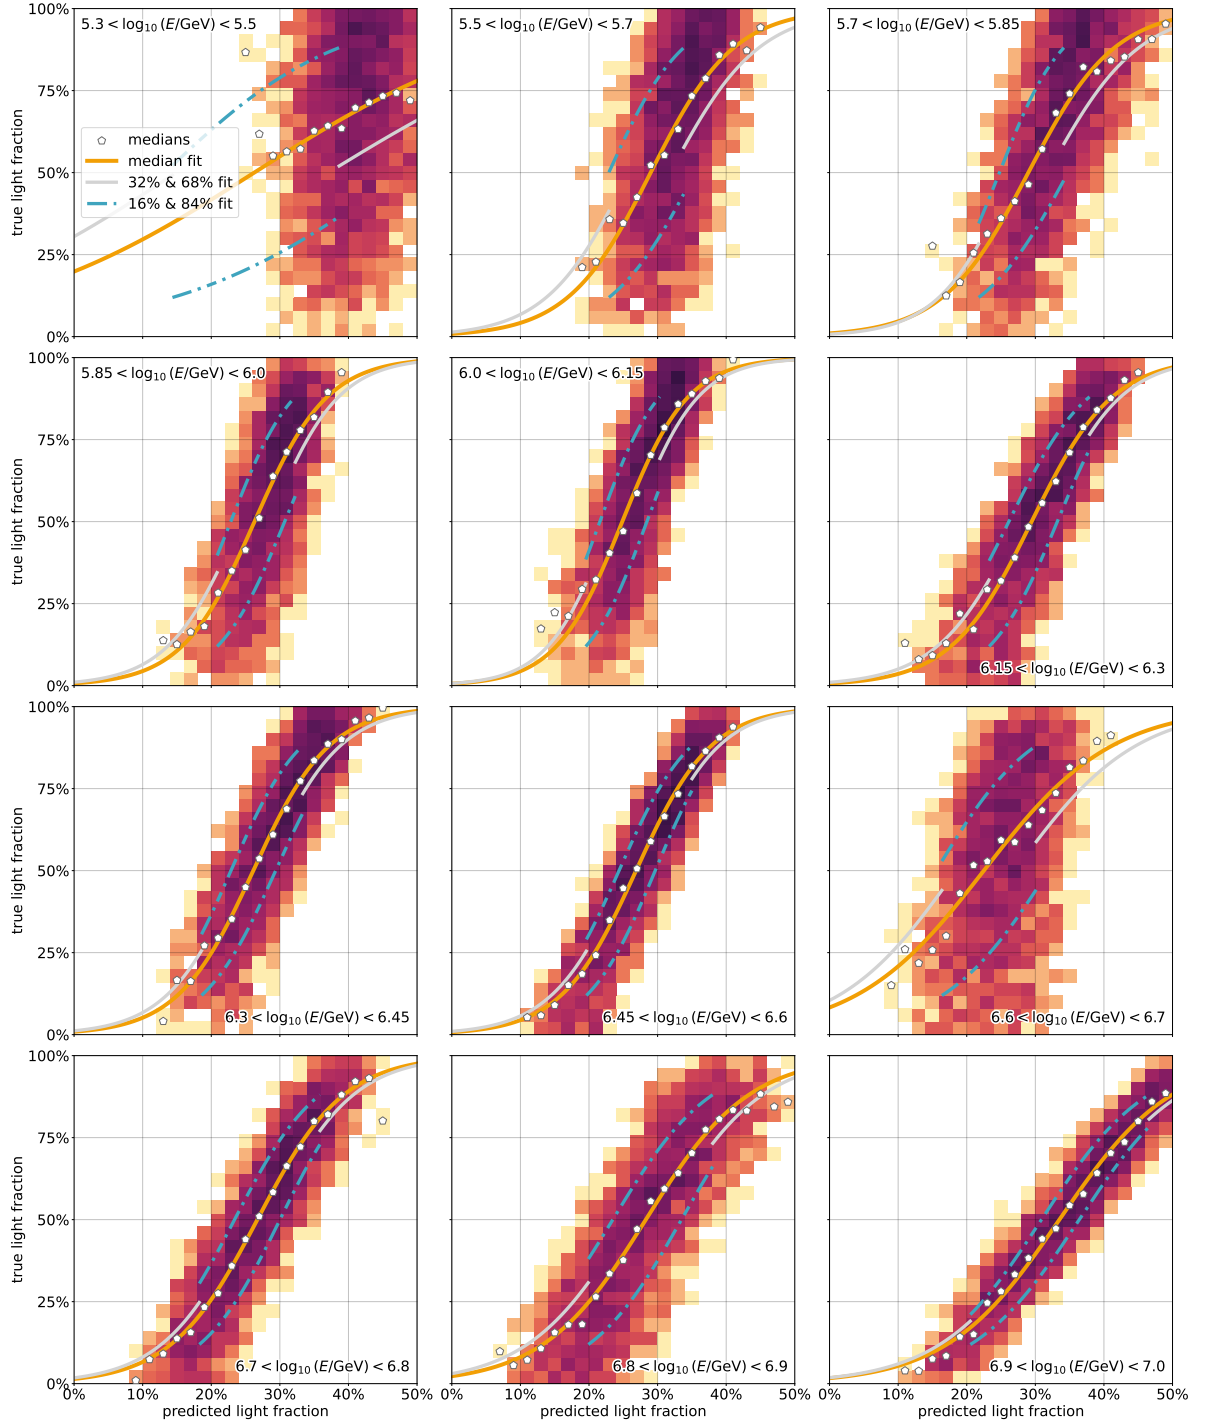


**Figure B.4:** The subset histograms with their sigmoid fits in all energy bins for the GSF-informed case and QGSjet-II.04 Monte Carlo simulations. The energy bin  $10^{5.3}$ – $10^{5.5}$  GeV will be excluded in the plots of fraction and flux of the proton+helium component.

## APPENDIX B. LIGHT FRACTION CALIBRATION



**Figure B.5:** The subset histograms with their sigmoid fits in all energy bins for the agnostic case and EPOS-LHC Monte Carlo simulations. The energy bin  $10^{5.3}$ – $10^{5.5}$  GeV will be excluded in the plots of fraction and flux of the proton+helium component.



**Figure B.6:** The subset histograms with their sigmoid fits in all energy bins for the GSF-informed case and EPOS-LHC Monte Carlo simulations. The energy bin  $10^{5.3}$ – $10^{5.5}$  GeV will be excluded in the plots of fraction and flux of the proton+helium component.



# APPENDIX C

## LIGHT FLUX WITH QGSJET-II.04 AND EPOS-LHC

The reconstructed differential H+He flux is tabulated in Tables C.1 (agnostic) and C.2 (GSF) for the analysis based on QGSJet-II.04 and in Tables C.3 (agnostic) and C.4 (GSF) for EPOS-LHC. The corresponding plots can be found in Figures C.1 and C.2. The systematic uncertainties and their contributions in each energy bin are shown in Figures C.3 and C.4. For all results obtained with the alternative hadronic interaction models QGSJet-II.04 and EPOS-LHC the energy range  $10^{5.5}$ – $10^{7.0}$  GeV is considered.

**Table C.1:** Reconstructed differential flux of light cosmic rays from the agnostic method, developed on QGSJet-II.04 Monte Carlo. Statistical and total systematic uncertainties are provided as well. Lower limits are indicated by arrows.

$\log_{10}(E/\text{GeV})$	$\langle \log_{10}(E/\text{GeV}) \rangle$	Flux $\pm$ stat. $\pm$ sys. ( $\text{m}^2 \text{ s sr GeV}^{-1}$ )
5.50–5.70	$5.60 \pm 0.08$	$(2.5441 \pm 0.0029^{+1.5231}_{-1.5765}) \cdot 10^{-11}$
5.70–5.85	$5.78 \pm 0.07$	$(1.2785 \pm 0.0010^{+0.5220}_{-0.5870}) \cdot 10^{-11}$
5.85–6.00	$5.92 \pm 0.06$	$\uparrow (2.386 \pm 0.004^{+1.179}) \cdot 10^{-12}$
6.00–6.15	$6.07 \pm 0.05$	$\uparrow (8.974 \pm 0.021^{+4.953}) \cdot 10^{-13}$
6.15–6.30	$6.22 \pm 0.05$	$\uparrow (4.319 \pm 0.009^{+2.131}) \cdot 10^{-13}$
6.30–6.45	$6.37 \pm 0.05$	$\uparrow (1.8287 \pm 0.0030^{+0.4220}) \cdot 10^{-13}$
6.45–6.60	$6.52 \pm 0.04$	$\uparrow (5.867 \pm 0.015^{+1.627}) \cdot 10^{-14}$
6.60–6.70	$6.65 \pm 0.04$	$(4.338 \pm 0.012^{+1.450}_{-2.141}) \cdot 10^{-14}$
6.70–6.80	$6.75 \pm 0.04$	$\uparrow (1.120 \pm 0.007^{+0.306}) \cdot 10^{-14}$
6.80–6.90	$6.846 \pm 0.027$	$\uparrow (6.61 \pm 0.05^{+2.29}) \cdot 10^{-15}$
6.90–7.00	$6.94 \pm 0.04$	$(7.28 \pm 0.05^{+5.34}_{-5.04}) \cdot 10^{-15}$

**Table C.2:** Reconstructed differential flux of light cosmic rays from the GSF-informed method, developed on QGSJet-II.04 Monte Carlo. Statistical and total systematic uncertainties are provided as well. Lower limits are indicated by arrows.

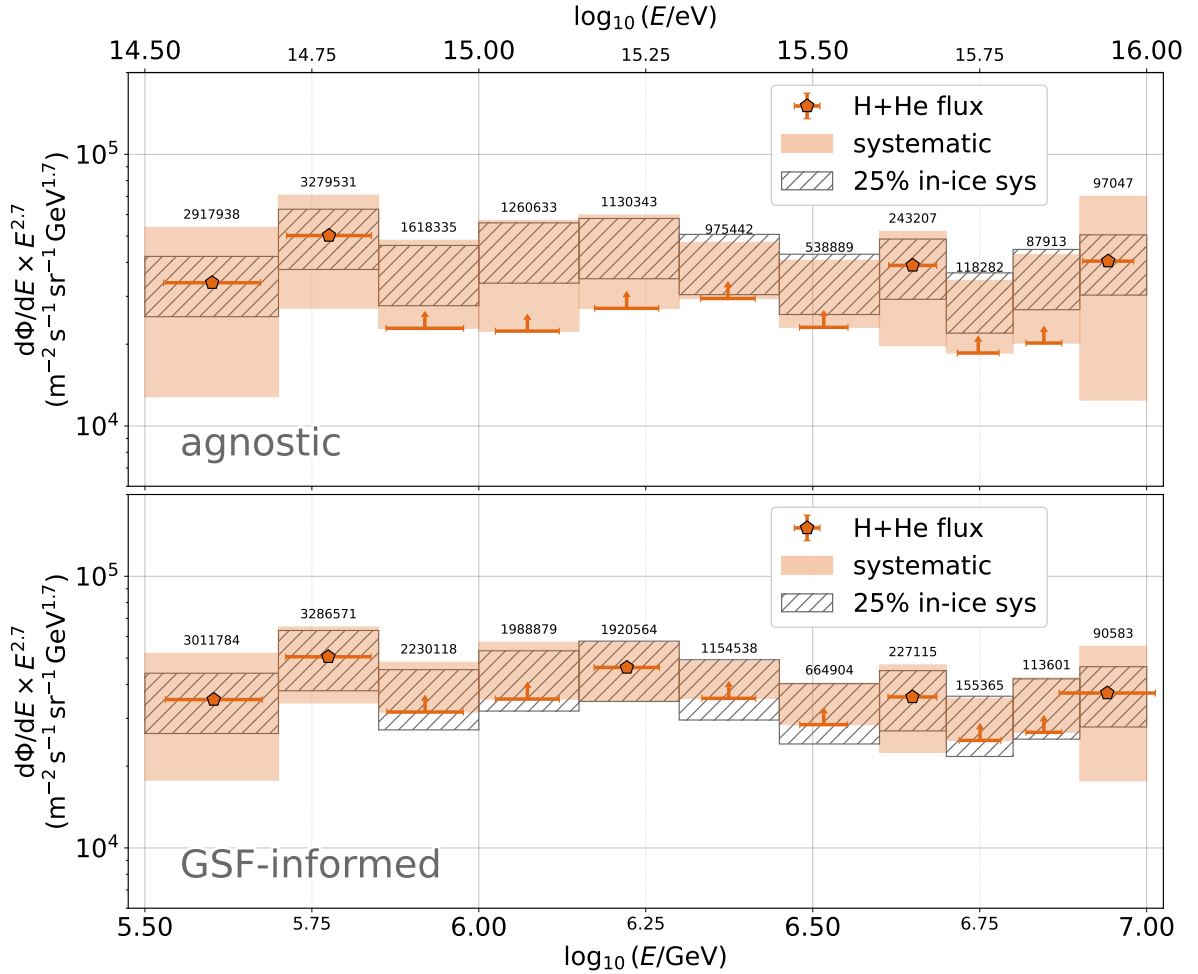
$\log_{10}(E/\text{GeV})$	$\langle \log_{10}(E/\text{GeV}) \rangle$	Flux $\pm$ stat. $\pm$ sys. ( $\text{m}^2 \text{ s sr GeV}^{-1}$ )
5.50–5.70	$5.60 \pm 0.08$	$(2.6170 \pm 0.0029^{+1.2563}_{-1.2952}) \cdot 10^{-11}$
5.70–5.85	$5.77 \pm 0.07$	$(1.2949 \pm 0.0010^{+0.3722}_{-0.4189}) \cdot 10^{-11}$
5.85–6.00	$5.92 \pm 0.06$	$\uparrow (3.293 \pm 0.004^{+1.250}) \cdot 10^{-12}$
6.00–6.15	$6.07 \pm 0.05$	$\uparrow (1.4178 \pm 0.0017^{+0.5859}) \cdot 10^{-12}$
6.15–6.30	$6.22 \pm 0.05$	$(7.342 \pm 0.007^{+1.719}_{-1.821}) \cdot 10^{-13}$
6.30–6.45	$6.37 \pm 0.05$	$\uparrow (2.1873 \pm 0.0028^{+0.5650}) \cdot 10^{-13}$
6.45–6.60	$6.52 \pm 0.04$	$\uparrow (7.235 \pm 0.013^{+2.099}) \cdot 10^{-14}$
6.60–6.70	$6.65 \pm 0.04$	$(4.007 \pm 0.013^{+1.252}_{-1.501}) \cdot 10^{-14}$
6.70–6.80	$6.75 \pm 0.04$	$\uparrow (1.477 \pm 0.006^{+0.390}) \cdot 10^{-14}$
6.80–6.90	$6.846 \pm 0.027$	$\uparrow (8.73 \pm 0.05^{+2.85}) \cdot 10^{-15}$
6.90–7.00	$6.94 \pm 0.08$	$(6.75 \pm 0.05^{+3.27}_{-3.55}) \cdot 10^{-15}$

**Table C.3:** Reconstructed differential flux of light cosmic rays from the agnostic method, developed on EPOS-LHC Monte Carlo. Statistical and total systematic uncertainties are provided as well. Lower limits are indicated by arrows.

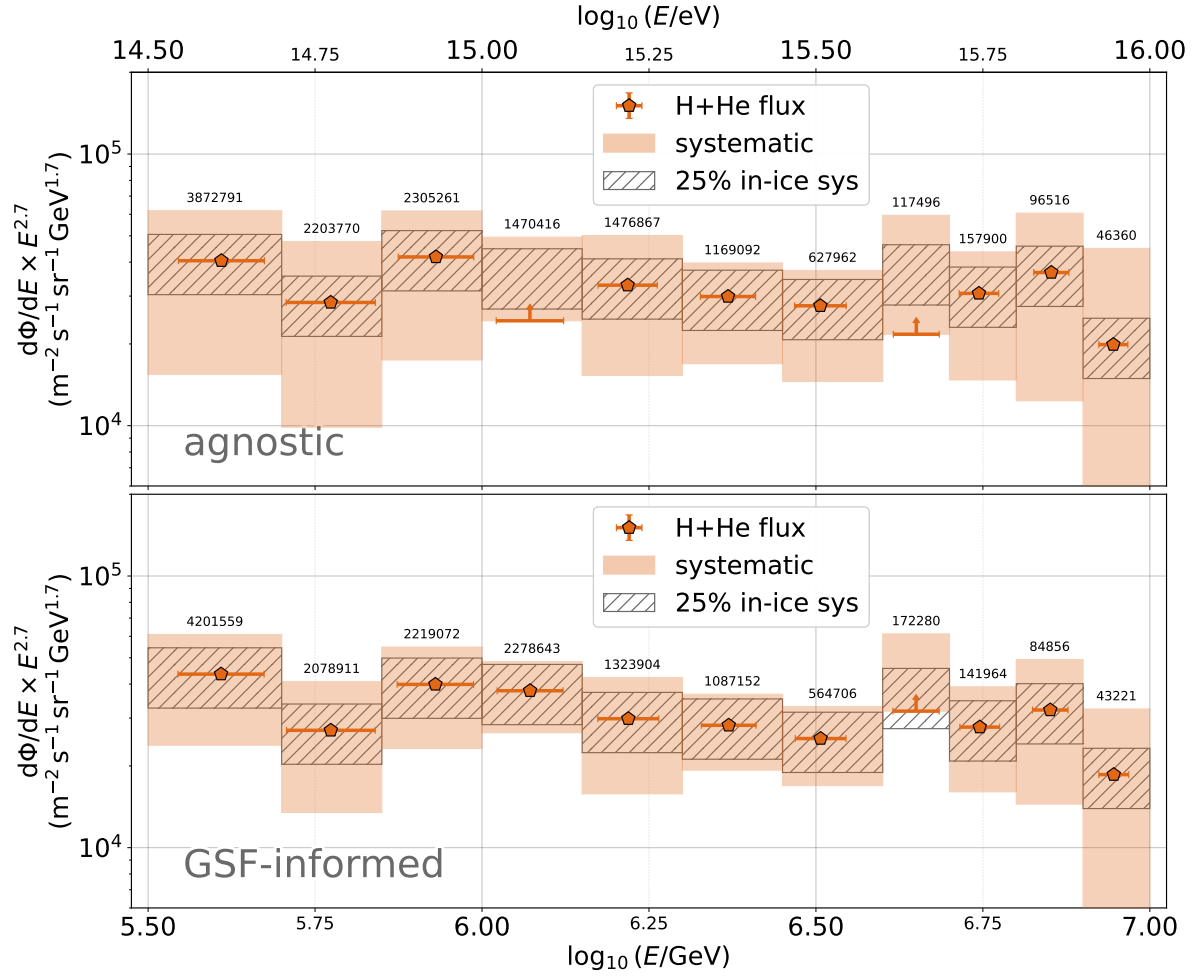
$\log_{10}(E/\text{GeV})$	$\langle \log_{10}(E/\text{GeV}) \rangle$	Flux $\pm$ stat. $\pm$ sys. ( $\text{m}^2 \text{ s sr GeV}^{-1}$ )
5.50–5.70	$5.61 \pm 0.07$	$(2.8911 \pm 0.0024^{+1.5270}_{-1.7874}) \cdot 10^{-11}$
5.70–5.85	$5.77 \pm 0.07$	$(7.334 \pm 0.009^{+4.952}_{-4.787}) \cdot 10^{-12}$
5.85–6.00	$5.93 \pm 0.06$	$(4.051 \pm 0.005^{+1.930}_{-2.360}) \cdot 10^{-12}$
6.00–6.15	$6.07 \pm 0.06$	$\uparrow (9.829 \pm 0.017^{+5.492}) \cdot 10^{-13}$
6.15–6.30	$6.22 \pm 0.05$	$(5.351 \pm 0.008^{+2.800}_{-2.862}) \cdot 10^{-13}$
6.30–6.45	$6.37 \pm 0.05$	$(1.9099 \pm 0.0024^{+0.6280}_{-0.8289}) \cdot 10^{-13}$
6.45–6.60	$6.51 \pm 0.04$	$(7.472 \pm 0.014^{+2.606}_{-3.537}) \cdot 10^{-14}$
6.60–6.70	$6.65 \pm 0.04$	$\uparrow (2.403 \pm 0.020^{+2.478}) \cdot 10^{-14}$
6.70–6.80	$6.744 \pm 0.030$	$(1.897 \pm 0.007^{+0.800}_{-0.987}) \cdot 10^{-14}$
6.80–6.90	$6.852 \pm 0.027$	$(1.154 \pm 0.007^{+0.755}_{-0.766}) \cdot 10^{-14}$
6.90–7.00	$6.945 \pm 0.022$	$(3.52 \pm 0.05^{+4.43}_{-3.20}) \cdot 10^{-15}$

**Table C.4:** Reconstructed differential flux of light cosmic rays from the GSF-informed method, developed on EPOS-LHC Monte Carlo. Statistical and total systematic uncertainties are provided as well. Lower limits are indicated by arrows.

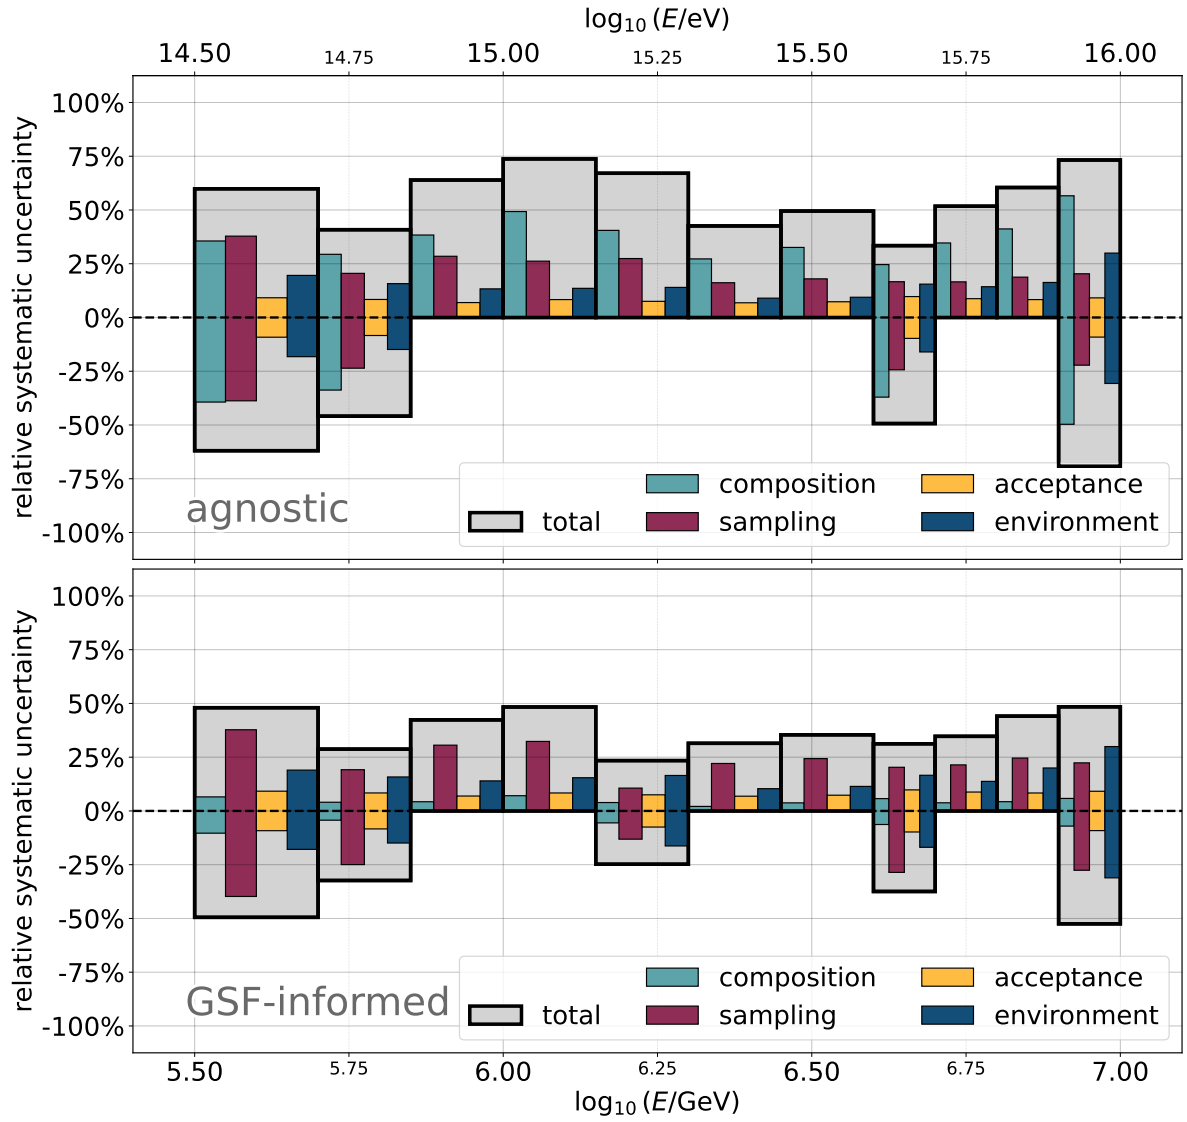
$\log_{10}(E/\text{GeV})$	$\langle \log_{10}(E/\text{GeV}) \rangle$	Flux $\pm$ stat. $\pm$ sys. ( $\text{m}^2 \text{s sr GeV}^{-1}$ )
5.50–5.70	$5.61 \pm 0.07$	$(3.1175 \pm 0.0024^{+1.2430}_{-1.4109}) \cdot 10^{-11}$
5.70–5.85	$5.77 \pm 0.07$	$(6.968 \pm 0.010^{+3.562}_{-3.494}) \cdot 10^{-12}$
5.85–6.00	$5.93 \pm 0.06$	$(3.887 \pm 0.005^{+1.448}_{-1.630}) \cdot 10^{-12}$
6.00–6.15	$6.07 \pm 0.05$	$(1.5256 \pm 0.0014^{+0.4281}_{-0.4568}) \cdot 10^{-12}$
6.15–6.30	$6.22 \pm 0.05$	$(4.822 \pm 0.008^{+2.012}_{-2.272}) \cdot 10^{-13}$
6.30–6.45	$6.37 \pm 0.05$	$(1.7910 \pm 0.0026^{+0.5396}_{-0.5664}) \cdot 10^{-13}$
6.45–6.60	$6.51 \pm 0.04$	$(6.803 \pm 0.014^{+2.114}_{-2.235}) \cdot 10^{-14}$
6.60–6.70	$6.65 \pm 0.04$	$\uparrow (3.529 \pm 0.017^{+2.735}) \cdot 10^{-14}$
6.70–6.80	$6.745 \pm 0.030$	$(1.704 \pm 0.008^{+0.694}_{-0.719}) \cdot 10^{-14}$
6.80–6.90	$6.851 \pm 0.027$	$(1.022 \pm 0.007^{+0.545}_{-0.562}) \cdot 10^{-14}$
6.90–7.00	$6.946 \pm 0.023$	$(3.27 \pm 0.05^{+2.44}_{-2.48}) \cdot 10^{-15}$



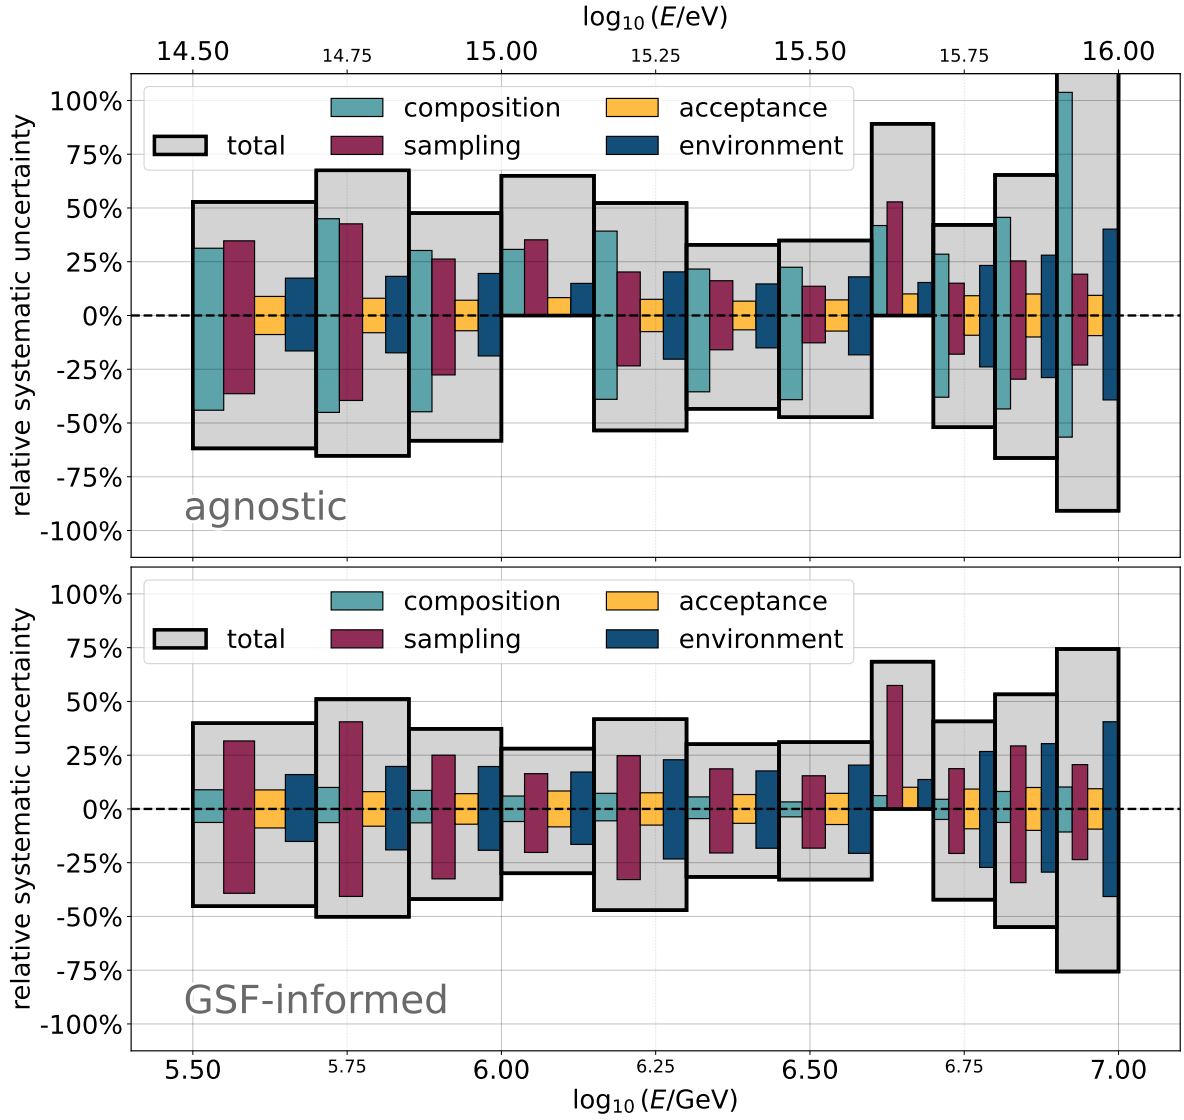
**Figure C.1:** The reconstructed light cosmic-ray flux (protons and helium) around the knee region of the spectrum, scaled with  $E^{2.7}$  for the agnostic (top) and the GSF-informed method (bottom), based on QGSJet-II.04 simulations. For details, refer to Figure 9.2.



**Figure C.2:** The reconstructed light cosmic-ray flux (protons and helium) around the knee region of the spectrum, scaled with  $E^{2.7}$  for the agnostic (top) and the GSF-informed method (bottom), based on EPOS-LHC simulations. For details, refer to Figure 9.2.



**Figure C.3:** The systematic uncertainty for the agnostic (top) and the GSF-informed method (bottom), also broken down by its contributions, for the QGSJet-II.04 hadronic interaction model.



**Figure C.4:** The systematic uncertainty for the agnostic (top) and the GSF-informed method (bottom), also broken down by its contributions, for the EPOS-LHC hadronic interaction model.

# APPENDIX D

## CONFIDENCE INTERVALS FOR EFFICIENCY RATIOS

Determining detection efficiencies comes down to counting the number of successfully detected events  $k$  and comparing it with the total number of incident events  $n$ . The detection probability, that is, the efficiency, can then be estimated as

$$\hat{p} = \frac{k}{n}. \quad (\text{D.1})$$

The number of successes  $k$  in a Bernoulli process of length  $n$  is commonly described by the discrete probability mass function of the binomial distribution

$$f_{\text{binom}}(k|n, p) = \binom{n}{k} p^k (1-p)^{n-k}, \quad (\text{D.2})$$

where  $k, n \in \mathbb{N}$  and  $p \in [0, 1]$  is the probability of success, which is independent of the individual trial. Alternatively, one can consider the continuous PDF of  $p$  when interpreting  $f_{\text{binom}}$  as a beta distribution

$$f_{\text{beta}}(x|\alpha, \beta) = \frac{x^{\alpha-1} (1-x)^{\beta-1}}{B(\alpha, \beta)} \quad (\text{D.3})$$

with the beta function  $B(\alpha, \beta)$ , which can be expressed by gamma functions:

$$B(\alpha, \beta) = \frac{\Gamma(\alpha) \Gamma(\beta)}{\Gamma(\alpha + \beta)}$$

Comparing (D.2) with (D.3) reveals that the PDF of  $p$  in a Bernoulli process can be described by a beta distribution with the substitutions

$$\begin{aligned} x &= p \\ \alpha &= k + 1 \\ \beta &= n - k + 1. \end{aligned}$$

The expectation value of  $p$  therefore is

$$\begin{aligned} E[p] &= \int_0^1 p f_{\text{beta}}(p|k+1, n-k+1) dp \\ &= \int_0^1 \frac{p^{k+1} (1-p)^{n-k}}{B(k+1, n-k+1)} dp \\ &= \int_0^1 \frac{p^{k+1} (1-p)^{n-k}}{B(k+2, n-k+1)} \frac{k+1}{n+2} dp \\ &= \underbrace{\int_0^1 f_{\text{beta}}(p|k+2, n-k+1) dp}_{=1} \frac{k+1}{n+2} = \frac{k+1}{n+2}, \end{aligned} \quad (\text{D.4})$$

which converges towards the point estimate (D.1) for large sample sizes. For the transition from line 2 to line 3 in (D.4) the identity  $B(\alpha, \beta) = B(\alpha+1, \beta) \cdot (\alpha+\beta)/\alpha$  has been applied.

## APPENDIX D. CONFIDENCE INTERVALS FOR EFFICIENCY RATIOS

While there are several approaches to estimate the two-sided confidence interval around the mean of a beta distribution, here, an approximation of the method proposed by Clopper and Pearson [141] was chosen. The  $(50 \pm \iota/2)$ -percentiles of the probability density function are set as the lower and upper boundaries of the  $1\sigma$  confidence interval, which can be identified from the cumulative distribution function (CDF)<sup>1</sup>.

For the ratio of two random variables that follow beta distributions (D.3)  $f_{\text{beta}}(p_1|\alpha_1, \beta_1)$  and  $f_{\text{beta}}(p_2|\alpha_2, \beta_2)$ , respectively, the PDF of

$$w = \frac{p_1}{p_2}$$

has been formulated by Pham-Gia [142]. It is defined piecewise as

$$f_{\text{ratio}}(w) = \frac{B(\alpha_1 + \alpha_2, \beta_2) w^{\alpha_1 - 1} {}_2F_1(\alpha_1 + \alpha_2, 1 - \beta_1; \alpha_1 + \alpha_2 + \beta_2; w)}{B(\alpha_1, \beta_1) B(\alpha_2, \beta_2)} \quad \text{for } 0 \leq w < 1 \quad (\text{D.5})$$

and

$$f_{\text{ratio}}(w) = \frac{B(\alpha_1 + \alpha_2, \beta_1) w^{-(\alpha_2 + 1)} {}_2F_1(\alpha_1 + \alpha_2, 1 - \beta_2; \alpha_1 + \alpha_2 + \beta_1; w^{-1})}{B(\alpha_1, \beta_1) B(\alpha_2, \beta_2)} \quad \text{for } w \geq 1. \quad (\text{D.6})$$

${}_2F_1(a_1, a_2; b_1; z)$  is the most common member of the family of generalized hypergeometric functions

$${}_pF_q(a_1, \dots, a_p; b_1, \dots, b_q; z) = \sum_{n=0}^{\infty} \frac{(a_1)_n \dots (a_p)_n}{(b_1)_n \dots (b_q)_n} \frac{z^n}{n!},$$

where  $(q)_n$  denotes the rising Pochhammer symbol<sup>2</sup>

$$\begin{aligned} (q)_n &= \prod_{k=0}^{n-1} (q + k) \\ &= \frac{\Gamma(q + n)}{\Gamma(q)}. \end{aligned}$$

---

<sup>1</sup>The CDF is given by the regularized incomplete beta function  $I_x(\alpha, \beta) = \int_0^x t^{\alpha-1} (1-t)^{\beta-1} dt / B(\alpha, \beta)$ .

<sup>2</sup>In the context of hypergeometric functions this notation is common, whereas generally this denotes the *falling* Pochhammer symbol  $\Gamma(q+1)/\Gamma(q-n+1)$ .

Since Ref. [142] provides only the PDFs, for the application in this work the cumulative distribution function for (D.5) and (D.6) had to be determined via integration over  $w'$ :

$$\begin{aligned}
F_{\text{ratio}}(w) &= \int_0^w f_{\text{ratio}}(w') dw' && \text{for } 0 \leq w < 1 \\
&= \Theta \int_0^w w'^{\alpha_1-1} {}_2F_1(\alpha_1 + \alpha_2, 1 - \beta_1; \alpha_1 + \alpha_2 + \beta_2; w') dw' \\
&= \Theta \int_0^w (w')^{\alpha_1-1} \sum_{n=0}^{\infty} \frac{(\alpha_1 + \alpha_2)_n (1 - \beta_1)_n}{(\alpha_1 + \alpha_2 + \beta_2)_n} \frac{(w')^n}{n!} dw' \\
&= \Theta \sum_{n=0}^{\infty} \frac{(\alpha_1 + \alpha_2)_n (1 - \beta_1)_n}{(\alpha_1 + \alpha_2 + \beta_2)_n} \frac{w^{n+\alpha_1}}{(n + \alpha_1) n!} \\
&= \Theta w^{\alpha_1} \sum_{n=0}^{\infty} \frac{(\alpha_1 + \alpha_2)_n (1 - \beta_1)_n}{(\alpha_1 + \alpha_2 + \beta_2)_n} \frac{w^n}{n!} \frac{\alpha_1}{(n + \alpha_1) \alpha_1} \\
&= \Theta \frac{w^{\alpha_1}}{\alpha_1} \sum_{n=0}^{\infty} \frac{(\alpha_1 + \alpha_2)_n (1 - \beta_1)_n}{(\alpha_1 + \alpha_2 + \beta_2)_n} \frac{(\alpha_1)_n}{(\alpha_1 + 1)_n} \frac{w^n}{n!} \\
&= B(\alpha_1 + \alpha_2, \beta_2) \frac{w^{\alpha_1}}{\alpha_1} \frac{{}_3F_2(\alpha_1, \alpha_1 + \alpha_2, 1 - \beta_1; \alpha_1 + 1, \alpha_1 + \alpha_2 + \beta_2; w)}{B(\alpha_1, \beta_1) B(\alpha_2, \beta_2)}, \tag{D.7}
\end{aligned}$$

where in the intermediate steps the substitution  $\Theta = \frac{B(\alpha_1 + \alpha_2, \beta_2)}{B(\alpha_1, \beta_1) B(\alpha_2, \beta_2)}$  is used and  $\frac{(\alpha_1)_n}{(\alpha_1 + 1)_n} = \frac{\alpha_1}{n + \alpha_1}$  is needed to express the power series as a higher-order hypergeometric function in the final representation again. Similarly:

$$\begin{aligned}
F_{\text{ratio}}(w) &= 1 - \int_w^{\infty} f_{\text{ratio}}(w') dw' && \text{for } w \geq 1 \\
&= 1 - B(\alpha_1 + \alpha_2, \beta_1) \frac{w^{-\alpha_2}}{\alpha_2} \frac{{}_3F_2(\alpha_2, \alpha_1 + \alpha_2, 1 - \beta_2; \alpha_2 + 1, \alpha_1 + \alpha_2 + \beta_1; w^{-1})}{B(\alpha_1, \beta_1) B(\alpha_2, \beta_2)} \tag{D.8}
\end{aligned}$$

Like the beta distribution (D.3), the ratio distribution (D.5, D.6) is asymmetric, too. The two-sided confidence interval around the expectation value<sup>3</sup>

$$E[w] = E\left[\frac{p_1}{p_2}\right] = E[p_1] \cdot E[p_2^{-1}] = \frac{\alpha_1}{\alpha_1 + \beta_1} \frac{\alpha_2 + \beta_2 - 1}{\alpha_2 - 1}$$

can be obtained in a Clopper–Pearson-like way as well by finding the  $(50 \pm \iota/2)$ -percentiles of the PDF. This requires finding the roots of the vertically shifted CDF (D.7, D.8), which can only be done numerically.

---

<sup>3</sup>For the calculation of the expectation value of  $p_2^{-1}$  a similar trick as in (D.4) can be used to shift the argument of the beta function by  $-1$ .



## BIBLIOGRAPHY

- [1] V. F. Hess. “Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten”. In: *Phys. Z.* 13 (1912), pp. 1084–1091.
- [2] M. G. Aartsen et al. “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. In: *Science* 361.6398 (2018), eaat1378. doi: [10.1126/science.aat1378](https://doi.org/10.1126/science.aat1378). arXiv: [1807.08816](https://arxiv.org/abs/1807.08816) [astro-ph.HE].
- [3] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Phys. Rev. Lett.* 119.16 (2017), p. 161101. doi: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). arXiv: [1710.05832](https://arxiv.org/abs/1710.05832) [gr-qc].
- [4] T. K. Gaisser, R. Engel, and E. Resconi. *Cosmic Rays and Particle Physics*. 2nd ed. Cambridge University Press, June 2016. doi: [10.1017/CB09781139192194](https://doi.org/10.1017/CB09781139192194).
- [5] D. J. Bird et al. “Detection of a Cosmic Ray with Measured Energy Well beyond the Expected Spectral Cutoff due to Cosmic Microwave Radiation”. In: *Astrophys. J.* 441 (1995), pp. 144–150. doi: [10.1086/175344](https://doi.org/10.1086/175344). arXiv: [astro-ph/9410067](https://arxiv.org/abs/astro-ph/9410067).
- [6] R. U. Abbasi et al. “An extremely energetic cosmic ray observed by a surface detector array”. In: *Science* 382.6673 (2023), abo5095. doi: [10.1126/science.abo5095](https://doi.org/10.1126/science.abo5095). arXiv: [2311.14231](https://arxiv.org/abs/2311.14231) [astro-ph.HE].
- [7] C. D. Anderson. “The Positive Electron”. In: *Phys. Rev.* 43 (1933), pp. 491–494. doi: [10.1103/PhysRev.43.491](https://doi.org/10.1103/PhysRev.43.491).
- [8] S. H. Neddermeyer and C. D. Anderson. “Note on the Nature of Cosmic Ray Particles”. In: *Phys. Rev.* 51 (1937), pp. 884–886. doi: [10.1103/PhysRev.51.884](https://doi.org/10.1103/PhysRev.51.884).
- [9] C. M. G. Lattes et al. “Processes Involving Charged Mesons”. In: *Nature* 159 (1947), pp. 694–697. doi: [10.1038/159694a0](https://doi.org/10.1038/159694a0).
- [10] G. D. Rochester and C. C. Butler. “Evidence for the Existence of New Unstable Elementary Particles”. In: *Nature* 160 (1947), pp. 855–857. doi: [10.1038/160855a0](https://doi.org/10.1038/160855a0).
- [11] C. Evoli. *The Cosmic-Ray Energy Spectrum*. Dec. 2020. doi: [10.5281/zenodo.4396125](https://doi.org/10.5281/zenodo.4396125).
- [12] H. Moraal. “Cosmic-Ray Modulation Equations”. In: *Space Sci. Rev.* 176 (2013), pp. 299–319. doi: [10.1007/s11214-011-9819-3](https://doi.org/10.1007/s11214-011-9819-3).
- [13] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. doi: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [14] J. R. Hörandel. “On the knee in the energy spectrum of cosmic rays”. In: *Astropart. Phys.* 19 (2003), pp. 193–220. doi: [10.1016/S0927-6505\(02\)00198-6](https://doi.org/10.1016/S0927-6505(02)00198-6). arXiv: [astro-ph/0210453](https://arxiv.org/abs/astro-ph/0210453).
- [15] T. K. Gaisser. “Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio”. In: *Astropart. Phys.* 35 (2012), pp. 801–806. doi: [10.1016/j.astropartphys.2012.02.010](https://doi.org/10.1016/j.astropartphys.2012.02.010). arXiv: [1111.6675](https://arxiv.org/abs/1111.6675) [astro-ph.HE].
- [16] T. K. Gaisser, T. Stanev, and S. Tilav. “Cosmic Ray Energy Spectrum from Measurements of Air Showers”. In: *Front. Phys. (Beijing)* 8 (2013), pp. 748–758. doi: [10.1007/s11467-013-0319-7](https://doi.org/10.1007/s11467-013-0319-7). arXiv: [1303.3565](https://arxiv.org/abs/1303.3565) [astro-ph.HE].
- [17] H. P. Dembinski et al. “Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to  $10^{11}$  GeV”. In: *PoS ICRC2017*. 2018, p. 533. doi: [10.22323/1.301.0533](https://doi.org/10.22323/1.301.0533). arXiv: [1711.11432](https://arxiv.org/abs/1711.11432) [astro-ph.HE].

## BIBLIOGRAPHY

- [18] K. Fujisue et al. “Global Spline Fit (GSF) 2024”. In: *PoS UHECR2024*. 2025, p. 087. doi: [10.22323/1.484.0087](https://doi.org/10.22323/1.484.0087).
- [19] C. C. Finlay et al. “International Geomagnetic Reference Field: the eleventh generation”. In: *Geophys. J. Int.* 183.3 (2010), pp. 1216–1230. doi: [10.1111/j.1365-246X.2010.04804.x](https://doi.org/10.1111/j.1365-246X.2010.04804.x).
- [20] R. Durrer and A. Neronov. “Cosmological Magnetic Fields: Their Generation, Evolution and Observation”. In: *Astron. Astrophys. Rev.* 21 (2013), p. 62. doi: [10.1007/s00159-013-0062-7](https://doi.org/10.1007/s00159-013-0062-7). arXiv: [1303.7121](https://arxiv.org/abs/1303.7121) [[astro-ph.CO](#)].
- [21] J. G. Learned and K. Mannheim. “High-energy neutrino astrophysics”. In: *Ann. Rev. Nucl. Part. Sci.* 50 (2000), pp. 679–749. doi: [10.1146/annurev.nucl.50.1.679](https://doi.org/10.1146/annurev.nucl.50.1.679).
- [22] W. Baade and F. Zwicky. “Cosmic Rays from Super-Novae”. In: *Proc. Nat. Acad. Sci.* 20.5 (1934), p. 259. doi: [10.1073/pnas.20.5.259](https://doi.org/10.1073/pnas.20.5.259).
- [23] E. Fermi. “On the Origin of the Cosmic Radiation”. In: *Phys. Rev.* 75 (1949), pp. 1169–1174. doi: [10.1103/PhysRev.75.1169](https://doi.org/10.1103/PhysRev.75.1169).
- [24] L. O’C. Drury. “Origin of Cosmic Rays”. In: *Astropart. Phys.* 39-40 (2012), pp. 52–60. doi: [10.1016/j.astropartphys.2012.02.006](https://doi.org/10.1016/j.astropartphys.2012.02.006). arXiv: [1203.3681](https://arxiv.org/abs/1203.3681) [[astro-ph.HE](#)].
- [25] A. M. Hillas. “The Origin of Ultra-High-Energy Cosmic Rays”. In: *Ann. Rev. Astron. Astrophys.* 22 (1984), pp. 425–444. doi: [10.1146/annurev.aa.22.090184.002233](https://doi.org/10.1146/annurev.aa.22.090184.002233).
- [26] B. Peters. “Primary cosmic radiation and extensive air showers”. In: *Nuovo Cim.* 22.4 (1961), pp. 800–819. doi: [10.1007/bf02783106](https://doi.org/10.1007/bf02783106).
- [27] R. Alves Batista et al. “Open Questions in Cosmic-Ray Research at Ultrahigh Energies”. In: *Front. Astron. Space Sci.* 6 (2019), p. 23. doi: [10.3389/fspas.2019.00023](https://doi.org/10.3389/fspas.2019.00023). arXiv: [1903.06714](https://arxiv.org/abs/1903.06714) [[astro-ph.HE](#)].
- [28] A. Coleman et al. “Ultra high energy cosmic rays – The intersection of the Cosmic and Energy Frontiers”. In: *Astropart. Phys.* 149 (2023), p. 102819. doi: [10.1016/j.astropartphys.2023.102819](https://doi.org/10.1016/j.astropartphys.2023.102819). arXiv: [2205.05845](https://arxiv.org/abs/2205.05845) [[astro-ph.HE](#)].
- [29] A. Aab et al. “Large-scale cosmic-ray anisotropies above 4 EeV measured by the Pierre Auger Observatory”. In: *Astrophys. J.* 868.1 (2018), p. 4. doi: [10.3847/1538-4357/aae689](https://doi.org/10.3847/1538-4357/aae689). arXiv: [1808.03579](https://arxiv.org/abs/1808.03579) [[astro-ph.HE](#)].
- [30] W. D. Apel et al. “Ankle-like Feature in the Energy Spectrum of Light Elements of Cosmic Rays Observed with KASCADE-Grande”. In: *Phys. Rev. D* 87 (2013), p. 081101. doi: [10.1103/PhysRevD.87.081101](https://doi.org/10.1103/PhysRevD.87.081101). arXiv: [1304.7114](https://arxiv.org/abs/1304.7114) [[astro-ph.HE](#)].
- [31] K. Greisen. “End to the cosmic ray spectrum?” In: *Phys. Rev. Lett.* 16 (1966), pp. 748–750. doi: [10.1103/PhysRevLett.16.748](https://doi.org/10.1103/PhysRevLett.16.748).
- [32] G. T. Zatsepin and V. A. Kuzmin. “Upper limit of the spectrum of cosmic rays”. In: *JETP Lett.* 4 (1966), pp. 78–80.
- [33] W. Bothe and W. Kolhörster. “Das Wesen der Höhenstrahlung”. In: *Z. Physik* 56.11 (1929), pp. 751–777. doi: [10.1007/BF01340137](https://doi.org/10.1007/BF01340137).
- [34] B. Rossi. “Über die Eigenschaften der durchdringenden Korpuskularstrahlung im Meeresniveau”. In: *Z. Physik* 82.3 (1933), pp. 151–178. doi: [10.1007/BF01341486](https://doi.org/10.1007/BF01341486).
- [35] P. Auger et al. “Extensive Cosmic-Ray Showers”. In: *Rev. Mod. Phys.* 11 (1939), pp. 288–291. doi: [10.1103/RevModPhys.11.288](https://doi.org/10.1103/RevModPhys.11.288).
- [36] W. Heitler. *The Quantum Theory of Radiation*. Clarendon Press, 1944.

- [37] J. Matthews. “A Heitler model of extensive air showers”. In: *Astropart. Phys.* 22 (2005), pp. 387–397. doi: [10.1016/j.astropartphys.2004.09.003](https://doi.org/10.1016/j.astropartphys.2004.09.003).
- [38] J. W. Elbert. “Multiple Muons Produced by Cosmic Ray Interactions”. In: *DUMAND Summer Workshop, Volume 2-UHE Interactions, Neutrino Astronomy*. Vol. 2. 1978, pp. 101–121.
- [39] M. Ackermann et al. “Optical properties of deep glacial ice at the South Pole”. In: *J. Geophys. Res.* 111.D13 (2006), p. D13203. doi: [10.1029/2005JD006687](https://doi.org/10.1029/2005JD006687).
- [40] R. Turcotte-Tardif. “Radio Measurements of Cosmic Rays at the South Pole”. PhD thesis. Karlsruher Institut für Technologie, 2022. doi: [10.5445/IR/1000160782](https://doi.org/10.5445/IR/1000160782).
- [41] J. Abraham et al. “The Fluorescence Detector of the Pierre Auger Observatory”. In: *Nucl. Instrum. Meth. A* 620 (2010), pp. 227–251. doi: [10.1016/j.nima.2010.04.023](https://doi.org/10.1016/j.nima.2010.04.023). arXiv: [0907.4282](https://arxiv.org/abs/0907.4282) [astro-ph.IM].
- [42] H. He. “Design of the LHAASO detectors”. In: *Radiat. Detect. Technol. Methods* 2.1 (2018), p. 7. doi: [10.1007/s41605-018-0037-3](https://doi.org/10.1007/s41605-018-0037-3).
- [43] T. Antoni et al. “The cosmic-ray experiment KASCADE”. In: *Nucl. Instrum. Meth. A* 513 (2003), pp. 490–510. doi: [10.1016/S0168-9002\(03\)02076-X](https://doi.org/10.1016/S0168-9002(03)02076-X).
- [44] W. D. Apel et al. “The KASCADE-Grande experiment”. In: *Nucl. Instrum. Meth. A* 620 (2010), pp. 202–216. doi: [10.1016/j.nima.2010.03.147](https://doi.org/10.1016/j.nima.2010.03.147).
- [45] A. U. Abeysekara et al. “The High-Altitude Water Cherenkov (HAWC) observatory in México: The primary detector”. In: *Nucl. Instrum. Meth. A* 1052 (2023), p. 168253. doi: [10.1016/j.nima.2023.168253](https://doi.org/10.1016/j.nima.2023.168253). arXiv: [2304.00730](https://arxiv.org/abs/2304.00730) [astro-ph.HE].
- [46] G. Aielli et al. “Layout and performance of RPCs used in the Argo-YSB experiment”. In: *Nucl. Instrum. Meth. A* 562 (2006), pp. 92–96. doi: [10.1016/j.nima.2006.02.136](https://doi.org/10.1016/j.nima.2006.02.136).
- [47] A. Aab et al. “The Pierre Auger Cosmic Ray Observatory”. In: *Nucl. Instrum. Meth. A* 798 (2015), pp. 172–213. doi: [10.1016/j.nima.2015.06.058](https://doi.org/10.1016/j.nima.2015.06.058). arXiv: [1502.01323](https://arxiv.org/abs/1502.01323) [astro-ph.IM].
- [48] H. Kawai et al. “Telescope Array Experiment”. In: *Nucl. Phys. B Proc. Suppl.* 175-176 (2008), pp. 221–226. doi: [10.1016/j.nuclphysbps.2007.11.002](https://doi.org/10.1016/j.nuclphysbps.2007.11.002).
- [49] G. James et al. *An Introduction to Statistical Learning. with Applications in R*. 2nd ed. Springer, 2013. doi: [10.1007/978-3-031-38747-0](https://doi.org/10.1007/978-3-031-38747-0).
- [50] C. H. Achen. “What does “explained variance “explain?: Reply”. In: *Political Analysis* 2 (1990), pp. 173–184. doi: [10.1093/pan/2.1.173](https://doi.org/10.1093/pan/2.1.173).
- [51] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization”. In: *ICLR 2015*. 2015. arXiv: [1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG].
- [52] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. “Learning representations by back-propagating errors”. In: *Nature* 323.6088 (1986), pp. 533–536. doi: [10.1038/323533a0](https://doi.org/10.1038/323533a0).
- [53] M. A. Nielsen. *Neural Networks and Deep Learning*. Determination Press, 2015.
- [54] S. Ioffe and C. Szegedy. “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift”. In: *PMLR* 37. 2015, pp. 448–456. arXiv: [1502.03167](https://arxiv.org/abs/1502.03167) [cs.LG].
- [55] J. S. Avva et al. “Particle Physics with the Cosmic Microwave Background with SPT-3G”. In: *J. Phys. Conf. Ser.* 1468.1 (2020). Ed. by Masayuki Nakahata, p. 012008. doi: [10.1088/1742-6596/1468/1/012008](https://doi.org/10.1088/1742-6596/1468/1/012008). arXiv: [1911.08047](https://arxiv.org/abs/1911.08047) [astro-ph.CO].

## BIBLIOGRAPHY

- [56] K. Akiyama et al. “First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole”. In: *Astrophys. J. Lett.* 875 (2019), p. L1. doi: [10.3847/2041-8213/ab0ec7](https://doi.org/10.3847/2041-8213/ab0ec7). arXiv: [1906.11238](https://arxiv.org/abs/1906.11238) [astro-ph.GA].
- [57] P. W. Gorham et al. “Constraints on the ultrahigh-energy cosmic neutrino flux from the fourth flight of ANITA”. In: *Phys. Rev. D* 99.12 (2019), p. 122001. doi: [10.1103/PhysRevD.99.122001](https://doi.org/10.1103/PhysRevD.99.122001). arXiv: [1902.04005](https://arxiv.org/abs/1902.04005) [astro-ph.HE].
- [58] E. Andres et al. “Observation of high-energy neutrinos using Cherenkov detectors embedded deep in Antarctic ice”. In: *Nature* 410 (2001), pp. 441–443. doi: [10.1038/35068509](https://doi.org/10.1038/35068509).
- [59] L. Heuermann, R. Abbasi et al. “Three-year performance of the IceAct telescopes at the IceCube Neutrino Observatory”. In: *PoS ICRC2023*. 2023, p. 367. doi: [10.22323/1.444.0367](https://doi.org/10.22323/1.444.0367). arXiv: [2307.13969](https://arxiv.org/abs/2307.13969) [astro-ph.HE].
- [60] M. G. Aartsen et al. “IceCube-Gen2: the window to the extreme Universe”. In: *J. Phys. G* 48.6 (2021), p. 060501. doi: [10.1088/1361-6471/abbd48](https://doi.org/10.1088/1361-6471/abbd48). arXiv: [2008.04323](https://arxiv.org/abs/2008.04323) [astro-ph.HE].
- [61] A. Ishihara, M. G. Aartsen et al. “The IceCube Upgrade – Design and Science Goals”. In: *PoS ICRC2019*. 2020, p. 1031. doi: [10.22323/1.358.1031](https://doi.org/10.22323/1.358.1031). arXiv: [1908.09441](https://arxiv.org/abs/1908.09441) [astro-ph.HE].
- [62] M. G. Aartsen et al. “The IceCube Neutrino Observatory: Instrumentation and Online Systems”. In: *JINST* 12.03 (2017), P03012. doi: [10.1088/1748-0221/12/03/P03012](https://doi.org/10.1088/1748-0221/12/03/P03012). arXiv: [1612.05093](https://arxiv.org/abs/1612.05093) [astro-ph.IM].
- [63] R. Abbasi et al. “The Design and Performance of IceCube DeepCore”. In: *Astropart. Phys.* 35 (2012), pp. 615–624. doi: [10.1016/j.astropartphys.2012.01.004](https://doi.org/10.1016/j.astropartphys.2012.01.004). arXiv: [1109.6096](https://arxiv.org/abs/1109.6096) [astro-ph.IM].
- [64] S. Euler. “Observation of oscillations of atmospheric neutrinos with the IceCube Neutrino Observatory”. PhD thesis. RWTH Aachen University, 2014.
- [65] R. Abbasi et al. “Calibration and Characterization of the IceCube Photomultiplier Tube”. In: *Nucl. Instrum. Meth. A* 618 (2010), pp. 139–152. doi: [10.1016/j.nima.2010.03.102](https://doi.org/10.1016/j.nima.2010.03.102). arXiv: [1002.2442](https://arxiv.org/abs/1002.2442) [astro-ph.IM].
- [66] R. Abbasi et al. “The IceCube data acquisition system: Signal capture, digitization, and timestamping”. In: *Nucl. Instrum. Meth. A* 601.3 (2009), pp. 294–316. doi: [10.1016/j.nima.2009.01.001](https://doi.org/10.1016/j.nima.2009.01.001). arXiv: [0810.4930](https://arxiv.org/abs/0810.4930) [physics.ins-det].
- [67] J. E. Dickinson et al. “The new South Pole air shower experiment: SPASE-2”. In: *Nucl. Instrum. Meth. A* 440 (2000), pp. 95–113. doi: [10.1016/S0168-9002\(99\)00788-3](https://doi.org/10.1016/S0168-9002(99)00788-3).
- [68] J. E. Dickinson et al. “A new air-Cherenkov array at the South Pole”. In: *Nucl. Instrum. Meth. A* 440 (2000), pp. 114–123. doi: [10.1016/S0168-9002\(99\)00789-5](https://doi.org/10.1016/S0168-9002(99)00789-5).
- [69] R. Abbasi et al. “IceTop: The surface component of IceCube”. In: *Nucl. Instrum. Meth. A* 700 (2013), pp. 188–220. doi: [10.1016/j.nima.2012.10.067](https://doi.org/10.1016/j.nima.2012.10.067). arXiv: [1207.6326](https://arxiv.org/abs/1207.6326) [astro-ph.IM].
- [70] M. G. Aartsen et al. “Cosmic ray spectrum and composition from PeV to EeV using 3 years of data from IceTop and IceCube”. In: *Phys. Rev. D* 100.8 (2019), p. 082002. doi: [10.1103/PhysRevD.100.082002](https://doi.org/10.1103/PhysRevD.100.082002). arXiv: [1906.04317](https://arxiv.org/abs/1906.04317) [astro-ph.HE].
- [71] R. Koirala. “Extension of the IceTop energy spectrum to 250 TeV and application of the constant intensity cut method to IceTop data”. PhD thesis. University of Delaware, 2019.

- [72] D. Heck et al. *CORSIKA: A Monte Carlo code to simulate extensive air showers*. Tech. rep. Forschungszentrum Karlsruhe, 1998.
- [73] R. Engel et al. “Towards a Next Generation of CORSIKA: A Framework for the Simulation of Particle Cascades in Astroparticle Physics”. In: *Comput. Softw. Big Sci.* 3.1 (2019), p. 2. doi: [10.1007/s41781-018-0013-0](https://doi.org/10.1007/s41781-018-0013-0). arXiv: [1808.08226](https://arxiv.org/abs/1808.08226) [[astro-ph.IM](#)].
- [74] J.-M. Alameddine et al. “Simulating radio emission from particle cascades with CORSIKA 8”. In: *Astropart. Phys.* 166 (2025), p. 103072. doi: [10.1016/j.astropartphys.2024.103072](https://doi.org/10.1016/j.astropartphys.2024.103072). arXiv: [2409.15999](https://arxiv.org/abs/2409.15999) [[astro-ph.HE](#)].
- [75] W. R. Nelson, H. Hirayama, and D. W. O. Rogers. *EGS4 Code System*. Tech. rep. Stanford Linear Accelerator Center, Menlo Park, CA (USA), 1985. doi: [10.2172/1453993](https://doi.org/10.2172/1453993).
- [76] A. Ferrari et al. *FLUKA: A multi-particle transport code*. Tech. rep. Stanford Linear Accelerator Center, Stanford University, Stanford, CA (USA), 2005. doi: [10.2172/877507](https://doi.org/10.2172/877507).
- [77] T. T. Böhlen et al. “The FLUKA Code: Developments and Challenges for High Energy and Medical Applications”. In: *Nucl. Data Sheets* 120 (2014), pp. 211–214. doi: [10.1016/j.nds.2014.07.049](https://doi.org/10.1016/j.nds.2014.07.049).
- [78] K. Werner, F.-M. Liu, and T. Pierog. “Parton ladder splitting and the rapidity dependence of transverse momentum spectra in deuteron-gold collisions at RHIC”. In: *Phys. Rev. C* 74 (2006), p. 044902. doi: [10.1103/PhysRevC.74.044902](https://doi.org/10.1103/PhysRevC.74.044902). arXiv: [hep-ph/0506232](https://arxiv.org/abs/hep-ph/0506232).
- [79] N. N. Kalmykov and S. S. Ostapchenko. “The Nucleus–nucleus interaction, nuclear fragmentation, and fluctuations of extensive air showers”. In: *Phys. Atom. Nucl.* 56 (1993), pp. 346–353.
- [80] R. S. Fletcher et al. “SIBYLL: An Event generator for simulation of high-energy cosmic ray cascades”. In: *Phys. Rev. D* 50 (1994), pp. 5710–5731. doi: [10.1103/PhysRevD.50.5710](https://doi.org/10.1103/PhysRevD.50.5710).
- [81] F. Riehn et al. “Hadronic interaction model Sibyll 2.3d and extensive air showers”. In: *Phys. Rev. D* 102.6 (2020), p. 063002. doi: [10.1103/PhysRevD.102.063002](https://doi.org/10.1103/PhysRevD.102.063002). arXiv: [1912.03300](https://arxiv.org/abs/1912.03300) [[hep-ph](#)].
- [82] S. Ostapchenko. “QGSJET-II: physics, recent improvements, and results for air showers”. In: *EPJ Web Conf.* 52 (2013). Ed. by U. Gensh and M. Walter, p. 02001. doi: [10.1051/epjconf/20125202001](https://doi.org/10.1051/epjconf/20125202001).
- [83] T. Pierog et al. “EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider”. In: *Phys. Rev. C* 92.3 (2015), p. 034906. doi: [10.1103/PhysRevC.92.034906](https://doi.org/10.1103/PhysRevC.92.034906). arXiv: [1306.0121](https://arxiv.org/abs/1306.0121) [[hep-ph](#)].
- [84] J. Albrecht et al. “The Muon Puzzle in cosmic-ray induced air showers and its connection to the Large Hadron Collider”. In: *Astrophys. Space Sci.* 367.3 (2022), p. 27. doi: [10.1007/s10509-022-04054-5](https://doi.org/10.1007/s10509-022-04054-5). arXiv: [2105.06148](https://arxiv.org/abs/2105.06148) [[astro-ph.HE](#)].
- [85] L. D. Landau and I. Pomeranchuk. “The Limits of Applicability of the Theory of Bremsstrahlung by Electrons and of the Creation of Pairs at Large Energies”. In: *Dokl. Akad. Nauk SSSR* 92 (1953), pp. 586–588. doi: [10.1016/b978-0-08-010586-4.50080-8](https://doi.org/10.1016/b978-0-08-010586-4.50080-8).
- [86] A. B. Migdal. “Bremsstrahlung and pair production in condensed media at high-energies”. In: *Phys. Rev.* 103 (1956), pp. 1811–1820. doi: [10.1103/PhysRev.103.1811](https://doi.org/10.1103/PhysRev.103.1811).
- [87] S. Agostinelli et al. “GEANT4—a simulation toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. doi: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).

## BIBLIOGRAPHY

- [88] J. Eisch. “A Multi-Variate Fit to the Chemical Composition of the Cosmic-Ray Spectrum”. PhD thesis. University of Wisconsin–Madison, 2014.
- [89] J.-H. Koehne et al. “PROPOSAL: A tool for propagation of charged leptons”. In: *Comput. Phys. Commun.* 184.9 (2013), pp. 2070–2090. doi: [10.1016/j.cpc.2013.04.001](https://doi.org/10.1016/j.cpc.2013.04.001).
- [90] C. Kopper. *CLSim*. [github.com/claudiok/clsim](https://github.com/claudiok/clsim).
- [91] J. Lundberg et al. “Light tracking for glaciers and oceans: Scattering and absorption in heterogeneous media with Photonics”. In: *Nucl. Instrum. Meth. A* 581 (2007), pp. 619–631. doi: [10.1016/j.nima.2007.07.143](https://doi.org/10.1016/j.nima.2007.07.143). arXiv: [astro-ph/0702108](https://arxiv.org/abs/astro-ph/0702108).
- [92] D. Chirkin, M. Rongen, R. Abbasi et al. “An improved mapping of ice layer undulations for the IceCube Neutrino Observatory”. In: *PoS ICRC2023* (2023), p. 975. doi: [10.22323/1.444.0975](https://doi.org/10.22323/1.444.0975). arXiv: [2307.13951](https://arxiv.org/abs/2307.13951) [[astro-ph.HE](https://arxiv.org/archive/hep)].
- [93] S. De Ridder. “Sensitivity of IceCube Cosmic Ray measurements to the hadronic interaction models”. PhD thesis. Ghent University, 2019.
- [94] S. Tilav et al. “Atmospheric Variations as Observed by IceCube”. In: *PoS ICRC2009*. 2010. arXiv: [1001.0776](https://arxiv.org/abs/1001.0776) [[astro-ph.HE](https://arxiv.org/archive/hep)].
- [95] K. Rawlins, R. Abbasi et al. “Accounting for changing snow over 10 years of IceTop, and its impact on the all-particle cosmic ray spectrum”. In: *PoS ICRC2023*. 2023, p. 377. doi: [10.22323/1.444.0377](https://doi.org/10.22323/1.444.0377).
- [96] R. Abbasi et al. “In situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory”. In: *The Cryosphere* 18.1 (2024), pp. 75–102. doi: [10.5194/tc-18-75-2024](https://doi.org/10.5194/tc-18-75-2024).
- [97] The IceCube Collaboration. *IceTray*. URL: [docs.icecube.aq/icetray/main](https://docs.icecube.aq/icetray/main).
- [98] T. Feusels. “Measurement of cosmic ray composition and energy spectrum between 1 PeV and 1 EeV with IceTop and IceCube”. PhD thesis. Ghent University, 2013.
- [99] F. Bontempo. “Enhanced Gamma-Hadron Separation at the IceCube Neutrino Observatory”. PhD thesis. Karlsruher Institut für Technologie, 2025. doi: [10.5445/IR/1000182568](https://doi.org/10.5445/IR/1000182568).
- [100] K. Rawlins. “Measuring the Composition of Cosmic Rays with the SPASE and AMANDA detectors”. PhD thesis. University of Wisconsin–Madison, 2001.
- [101] J. Ahrens et al. “Measurement of the cosmic ray composition at the knee with the SPASE-2/AMANDA-B10 detectors”. In: *Astropart. Phys.* 21 (2004), pp. 565–581. doi: [10.1016/j.astropartphys.2004.04.007](https://doi.org/10.1016/j.astropartphys.2004.04.007).
- [102] K. G. Andeen, C. Song, and K. Rawlins. “Measuring Cosmic Ray Composition at the Knee with SPASE-2 and AMANDA-II”. In: *PoS ICRC2007*. 2007, pp. 165–170. arXiv: [0711.0353](https://arxiv.org/abs/0711.0353) [[astro-ph](https://arxiv.org/archive/hep)].
- [103] K. Andeen. “First Measurements of Cosmic Ray Composition from 1-50 PeV using New Techniques on Coincident Data from the IceCube Neutrino Observatory”. PhD thesis. University of Wisconsin–Madison, 2011.
- [104] R. Abbasi et al. “Cosmic Ray Composition and Energy Spectrum from 1-30 PeV Using the 40-String Configuration of IceTop and IceCube”. In: *Astropart. Phys.* 42 (2013), pp. 15–32. doi: [10.1016/j.astropartphys.2012.11.003](https://doi.org/10.1016/j.astropartphys.2012.11.003). arXiv: [1207.3455](https://arxiv.org/abs/1207.3455) [[astro-ph.HE](https://arxiv.org/archive/hep)].
- [105] D. Soldin. “Composition from high  $p_T$  muons in IceCube”. In: *EPJ Web Conf.* 99 (2015), p. 06001. doi: [10.1051/epjconf/20159906001](https://doi.org/10.1051/epjconf/20159906001). arXiv: [1411.4448](https://arxiv.org/abs/1411.4448) [[astro-ph.HE](https://arxiv.org/archive/hep)].

- [106] P. Koundal, M. Plum, J. Saffer, R. Abbasi et al. “Study of mass composition of cosmic rays with IceTop and IceCube”. In: *PoS ICRC2021*. 2021, p. 323. doi: [10.22323/1.395.0323](https://doi.org/10.22323/1.395.0323). arXiv: [2107.09626](https://arxiv.org/abs/2107.09626) [astro-ph.HE].
- [107] P. Koundal. “Elemental Composition of Cosmic Rays: Analysis of IceCube data using Graph Neural Networks”. PhD thesis. Karlsruher Institut für Technologie, 2023. doi: [10.5445/IR/1000169558](https://doi.org/10.5445/IR/1000169558).
- [108] J. Saffer, R. Abbasi, et al. “Towards the Composition of sub-PeV Cosmic Rays at IceCube”. In: *PoS TAUP2023*. 2024, p. 137. doi: [10.22323/1.441.0137](https://doi.org/10.22323/1.441.0137).
- [109] M. Zoll. “Preparation for the next solar WIMP Analysis with IceCube. Advances in simulation, filtering, event topology identification and analysis approach”. Licentiate Thesis. Stockholm University, 2014.
- [110] M. Zoll. “A search for solar dark matter with the IceCube neutrino detector”. PhD thesis. Stockholm University, 2016.
- [111] D. Pandel. “Bestimmung von Wasser- und Detektorparametern und Rekonstruktion von Myonen bis 100 TeV mit dem Baikal Neutrinoteleskop NT-72”. Diploma Thesis. Humboldt-Universität zu Berlin, 1996.
- [112] N. van Eijndhoven, O. Fadiran, and G. Japaridze. “Implementation of a Gauss convoluted Pandel PDF for track reconstruction in Neutrino Telescopes”. In: *Astropart. Phys.* 28 (2007), pp. 456–462. doi: [10.1016/j.astropartphys.2007.09.001](https://doi.org/10.1016/j.astropartphys.2007.09.001). arXiv: [0704.1706](https://arxiv.org/abs/0704.1706) [astro-ph].
- [113] I. A. Antonov and V. M. Saleev. “An economic method of computing  $LP\tau$ -sequences”. In: *USSR Comput. Maths. Math. Phys.* 19.1 (1979), pp. 252–256. doi: [10.1016/0041-5553\(79\)90085-5](https://doi.org/10.1016/0041-5553(79)90085-5).
- [114] A. S. Leszczynska. “Potential of the IceTop Enhancement with a Scintillation Detector Array”. PhD thesis. Karlsruher Institut für Technologie, 2020. doi: [10.5445/IR/1000131245](https://doi.org/10.5445/IR/1000131245).
- [115] A. Leszczynska et al. “A multi-detector EAS reconstruction framework for IceCube”. In: *PoS ICRC2023*. 2023, p. 366. doi: [10.22323/1.444.0366](https://doi.org/10.22323/1.444.0366).
- [116] F. James and M. Roos. “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations”. In: *Comput. Phys. Commun.* 10 (1975), pp. 343–367. doi: [10.1016/0010-4655\(75\)90039-9](https://doi.org/10.1016/0010-4655(75)90039-9).
- [117] D. Chirkin and W. Rhode. “Muon Monte Carlo: a new high precision tool for muon propagation through matter”. In: *PoS ICRC2001*. 2001, p. 1017.
- [118] T. Matsuyama. “Estimation of uncertainty of percentile values in particle size distribution analysis as a function of number of particles”. In: *J.APT* 30.11 (2019), pp. 2616–2619. doi: [10.1016/j.appt.2019.08.008](https://doi.org/10.1016/j.appt.2019.08.008).
- [119] R. Abbasi et al. “Density of GeV muons in air showers measured with IceTop”. In: *Phys. Rev. D* 106.3 (2022), p. 032010. doi: [10.1103/PhysRevD.106.032010](https://doi.org/10.1103/PhysRevD.106.032010). arXiv: [2201.12635](https://arxiv.org/abs/2201.12635) [hep-ex].
- [120] J. Ansel et al. “Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation”. In: *ASPLOS 2024*. 2024, pp. 929–947. doi: [10.1145/3620665.3640366](https://doi.org/10.1145/3620665.3640366).
- [121] R. Abbasi et al. “A Convolutional Neural Network based Cascade Reconstruction for the IceCube Neutrino Observatory”. In: *JINST* 16 (2021), P07041. doi: [10.1088/1748-0221/16/07/P07041](https://doi.org/10.1088/1748-0221/16/07/P07041). arXiv: [2101.11589](https://arxiv.org/abs/2101.11589) [hep-ex].

## BIBLIOGRAPHY

- [122] R. Abbasi et al. “Observation of high-energy neutrinos from the Galactic plane”. In: *Science* 380.6652 (2023), adc9818. doi: [10.1126/science.adc9818](https://doi.org/10.1126/science.adc9818). arXiv: [2307.04427](https://arxiv.org/abs/2307.04427) [astro-ph.HE].
- [123] L. Biewald. *Experiment Tracking with Weights and Biases*. 2020. URL: [wandb.com](https://wandb.com).
- [124] S. T. Hahn. “Methods for Estimating Mass-Sensitive Observables of Ultra-High Energy Cosmic Rays using Artificial Neural Networks”. PhD thesis. Karlsruher Institut für Technologie, 2022. doi: [10.5445/IR/1000154770](https://doi.org/10.5445/IR/1000154770).
- [125] P. Jaccard. “Étude comparative de la distribution florale dans une portion des Alpes et des Jura”. In: *Bull. Soc. Vaud. Sc.* 37 (1901), pp. 547–579. doi: [10.5169/seals-266450](https://doi.org/10.5169/seals-266450).
- [126] M. G. Aartsen et al. “Cosmic ray spectrum from 250 TeV to 10 PeV using IceTop”. In: *Phys. Rev. D* 102.12 (2020), p. 122001. doi: [10.1103/PhysRevD.102.122001](https://doi.org/10.1103/PhysRevD.102.122001). arXiv: [2006.05215](https://arxiv.org/abs/2006.05215) [astro-ph.HE].
- [127] V. I. Zatsepin and N. V. Sokolskaya. “Three component model of cosmic ray spectra from 100 GeV to 100 PeV”. In: *Astron. Astrophys.* 458 (2006), pp. 1–5. doi: [10.1051/0004-6361:20065108](https://doi.org/10.1051/0004-6361:20065108). arXiv: [astro-ph/0601475](https://arxiv.org/abs/astro-ph/0601475).
- [128] Y. S. Yoon et al. “Proton and Helium Spectra from the CREAM-III Flight”. In: *Astrophys. J.* 839.1 (2017), p. 5. doi: [10.3847/1538-4357/aa68e4](https://doi.org/10.3847/1538-4357/aa68e4). arXiv: [1704.02512](https://arxiv.org/abs/1704.02512) [astro-ph.HE].
- [129] A. Albert et al. “Cosmic ray spectrum of protons plus helium nuclei between 6 and 158 TeV from HAWC data”. In: *Phys. Rev. D* 105.6 (2022), p. 063021. doi: [10.1103/PhysRevD.105.063021](https://doi.org/10.1103/PhysRevD.105.063021). arXiv: [2204.06662](https://arxiv.org/abs/2204.06662) [astro-ph.HE].
- [130] B. Bartoli et al. “Cosmic ray proton plus helium energy spectrum measured by the ARGO-YBJ experiment in the energy range 3–300 TeV”. In: *Phys. Rev. D* 91.11 (2015), p. 112017. doi: [10.1103/PhysRevD.91.112017](https://doi.org/10.1103/PhysRevD.91.112017). arXiv: [1503.07136](https://arxiv.org/abs/1503.07136) [hep-ex].
- [131] B. Bartoli et al. “Knee of the cosmic hydrogen and helium spectrum below 1 PeV measured by ARGO-YBJ and a Cherenkov telescope of LHAASO”. In: *Phys. Rev. D* 92.9 (2015), p. 092005. doi: [10.1103/PhysRevD.92.092005](https://doi.org/10.1103/PhysRevD.92.092005). arXiv: [1502.03164](https://arxiv.org/abs/1502.03164) [astro-ph.HE].
- [132] N. Gorbunov et al. “Energy spectra of abundant cosmic-ray nuclei in the NUCLEON experiment”. In: *Adv. Space Res.* 64.12 (2019), pp. 2546–2558. doi: [10.1016/j.asr.2019.10.004](https://doi.org/10.1016/j.asr.2019.10.004). arXiv: [1809.05333](https://arxiv.org/abs/1809.05333) [astro-ph.IM].
- [133] F. Alemanno et al. “Measurement of the cosmic p+He energy spectrum from 50 GeV to 0.5 PeV with the DAMPE space mission”. In: *Phys. Rev. D* 109.12 (2024), p. L121101. doi: [10.1103/PhysRevD.109.L121101](https://doi.org/10.1103/PhysRevD.109.L121101). arXiv: [2304.00137](https://arxiv.org/abs/2304.00137) [astro-ph.HE].
- [134] T. Antoni et al. “KASCADE measurements of energy spectra for elemental groups of cosmic rays: Results and open problems”. In: *Astropart. Phys.* 24 (2005), pp. 1–25. doi: [10.1016/j.astropartphys.2005.04.001](https://doi.org/10.1016/j.astropartphys.2005.04.001). arXiv: [astro-ph/0505413](https://arxiv.org/abs/astro-ph/0505413).
- [135] M. Yu. Kuznetsov et al. “Energy spectra of elemental groups of cosmic rays with the KASCADE experiment data and machine learning”. In: *JCAP* 05 (2024), p. 125. doi: [10.1088/1475-7516/2024/05/125](https://doi.org/10.1088/1475-7516/2024/05/125). arXiv: [2312.08279](https://arxiv.org/abs/2312.08279) [astro-ph.HE].
- [136] W. D. Apel et al. “KASCADE-Grande measurements of energy spectra for elemental groups of cosmic rays”. In: *Astropart. Phys.* 47 (2013), pp. 54–66. doi: [10.1016/j.astropartphys.2013.06.004](https://doi.org/10.1016/j.astropartphys.2013.06.004). arXiv: [1306.6283](https://arxiv.org/abs/1306.6283) [astro-ph.HE].
- [137] S. Verpoest. “Studies of TeV muons in cosmic-ray air showers detected with IceTop and IceCube”. PhD thesis. Ghent University, 2022.

- [138] S. Ostapchenko. “QGSJET-III model of high energy hadronic interactions: The formalism”. In: *Phys. Rev. D* 109.3 (2024), p. 034002. DOI: [10.1103/PhysRevD.109.034002](https://doi.org/10.1103/PhysRevD.109.034002). arXiv: [2401.06202](https://arxiv.org/abs/2401.06202) [hep-ph].
- [139] T. Pierog and K. Werner. “EPOS LHC-R : up-to-date hadronic model for EAS simulations”. In: *PoS ICRC2023*. 2023, p. 230. DOI: [10.22323/1.444.0230](https://doi.org/10.22323/1.444.0230).
- [140] F. Riehn, A. Fedynitch, and R. Engel. “Sibyll★”. In: *Astropart. Phys.* 160 (2024), p. 102964. DOI: [10.1016/j.astropartphys.2024.102964](https://doi.org/10.1016/j.astropartphys.2024.102964). arXiv: [2404.02636](https://arxiv.org/abs/2404.02636) [hep-ph].
- [141] C. J. Clopper and E. S. Pearson. “The use of confidence or fiducial limits illustrated in the case of the binomial”. In: *Biometrika* 26.4 (1934), pp. 404–413. DOI: [10.1093/biomet/26.4.404](https://doi.org/10.1093/biomet/26.4.404).
- [142] T. Pham-Gia. “Distributions of the ratios of independent beta variables and applications”. In: *Commun. Statist.-Theory Meth.* 29.12 (2000), pp. 2693–2715. DOI: [10.1080/03610920008832632](https://doi.org/10.1080/03610920008832632).



## ACKNOWLEDGMENTS

*In memory of IceTop DOM 74-61 “Orc”, which died during my monitoring shift on March 11, 2022.*

This thesis is the yield of many years of intense work — not only by myself but also the many people who supported and guided me along the way. I would like to take this opportunity to express my gratitude to them.

First and foremost, I want to thank my referees Prof. Ralph Engel and Prof. Kathrin Valerius for supporting my research and for their expertise.

Next, I would like to express my gratitude to Andreas Haungs, who invited me to join KIT for the PhD in the middle of the turbulent times of 2020. Even though, this topic turned out significantly more challenging than anticipated, it is thanks to his invaluable supervision that this work could come to a good end.

My warmest thanks go to Agnieszka Leszczyńska, who has been there for me since day one. The simulation production would not have been conducted on that level of professionalism without her help. I sincerely thank Donghwa Kang, who is not only an expert on air showers since the KASCADE era, but also knows how to provide moral support, always with a bright smile even in the most stressful times. I deeply appreciate Fahim Varsi’s expertise on the spectrum analysis of light cosmic rays. His input sparked new ideas that eventually lead to the method I developed in this thesis.

I would like to reciprocate gratitude to Paras Koundal, my “IceCube-machine-learning-composition companion”. Our regular meetings — in the office, on our walks to the “KASCADE meadow” or via a transatlantic zoom connection — lead to new ideas for tackling this challenge on both sides. Furthermore, I would like to thank Mark Weyrauch for every single remark, which always made me look at the problem from another perspective, and Federico Bontempo for his critical view and his help with the CORSIKA simulations.

Huge thanks go to the “cosmic-ray composition posse”, especially Matthias Plum and Karen Andeen for bringing the composition call back to life and thereby providing an informal discussion forum for a constructive exchange without feeling dumb. I want to particularly thank Kath Rawlins. It was a pleasure to work together with her on improving the SLC calibration, most notably during that cold dark week of Alaskan winter.

Thank you to the broader cosmic-ray working group, specifically Alan Coleman, Paolo Desiati, Hermann Kolanoski, Frank Schröder, Serap Tilav and Stef Verpoest, for their essential help in various ways and for sometimes keeping the weekly call shorter than two hours. I would not have been able to unblind my analysis if it had not been for Larissa Paul and Segev BenZvi, who reviewed my work thoroughly.

At WIPAC in Madison, I would like to thank Juan Carlos Díaz-Vélez, Kevin Meagher and David Schultz for their technical support regarding the simulation and priority on the cluster. Thanks to Patrick Hatch for introducing me to IceHive and Anatoli Fedynitch, who provided me with the tools to explore GSF and its many faces.

## BIBLIOGRAPHY

Thanks to Ben Flaggs, who showed me around Delaware and D.C. as well as Marc Jacquart for shooting amazing photos at the South Pole for me after I introduced him to IceCube. Special thanks go to Jamie Zvirzdin for her marvelous editing and great writing advice.

The technical support with CORSIKA and FLUKA by Tanguy Pierog and Alfredo Ferrari, respectively, is greatly appreciated. For all kinds of helpful discussions I would like to also thank Vika Tokareva, Steffen Hahn and Nikos Karastathis. My sincere thanks go to Sabine Bucher, Peichen Chen, Anna Friedrich and Doris Wochele, who keep everything running.

At long last, I want to express my deepest gratitude to my parents. Only owing to them, I have been able to pursue this path in academia. Thank you for everything.

The author gratefully acknowledges the computing time provided on the high-performance computer HoreKa by the National High-Performance Computing Center at KIT (NHR@KIT). This center is jointly supported by the Federal Ministry of Education and Research and the Ministry of Science, Research and the Arts of Baden-Württemberg, as part of the National High-Performance Computing (NHR) joint funding program (<https://www.nhr-verein.de/en/our-partners>). HoreKa is partly funded by the German Research Foundation (DFG).

The L<sup>A</sup>T<sub>E</sub>Xtemplate for this thesis is based on <https://www.overleaf.com/latex/templates/caltech-thesis-latex-template-with-logo/kgwqvkbfbkwm>. The template is used and modified under Creative Commons BY 4.0 (<https://creativecommons.org/licenses/by/4.0/>).

# LIST OF ACRONYMS AND INITIALISMS

## A

**ADC** Analog-to-Digital Converter [67](#)

**AGN** Active Galactic Nucleus [11](#)

**AMANDA** Antarctic Muon And Neutrino Detector Array [35](#), [58](#), [71](#)

**AMS-02** Alpha Magnetic Spectrometer (version 2) [12](#)

**ANITA** Antarctic Impulsive Transient Antenna [35](#)

**ARGO-YBJ** Astrophysical Radiation with Ground-based Observatory at Yangbajing [23](#), [132](#)

**ATWD** Analog Transient Waveform Digitizer [39](#), [66](#)

## B

**BCE** Binary Cross Entropy [30](#), [104](#)

**BICEP** Background Imaging of Cosmic Extragalactic Polarization [35](#)

## C

**CDF** Cumulative Density Function [162](#)

**CMB** Cosmic Microwave Background [9](#), [35](#)

**CME** Coronal Mass Ejection [6](#)

**CNN** Convolutional Neural Network [30](#), [103](#)

**COG** Center of Gravity [63](#)

**COP** Crossover Point [68](#)

**CORSIKA** Cosmic Ray Simulations for KASCADE [45](#), [81](#), [99](#), [127](#), [175](#)

**CR** Cosmic Ray [5](#), [13](#), [51](#), [75](#), [108](#), [123](#), [181](#)

**CREAM** Cosmic Ray Energetics And Mass [12](#), [132](#)

## D

**DAMPE** Dark Matter Particle Explorer [132](#), [136](#)

**DAQ** Data Acquisition [35](#), [63](#)

**DOM** Digital Optical Module [36](#), [50](#), [61](#), [88](#), [100](#), [123](#), [175](#), [181](#)

**DSL** Dark Sector Lab [35](#), [57](#)

## E

**EAS** Extensive Air Shower [13](#), [38](#), [46](#), [63](#)

**EPOS** Energy Conserving Quantum Mechanical Multiple Scattering Approach Based on Partons Off-Shell Remnants Splitting of Parton Ladders [46](#), [99](#), [129](#), [135](#), [143](#), [155](#)

## F

**fADC** fast Analog-to-Digital Converter [67](#)

**FLUKA** Fluktuiierende Kaskade, engl: Fluctuating Cascade [46](#), [176](#)

## LIST OF ACRONYMS AND INITIALISMS

FN False Negatives 31, 109

FP False Positives 31, 109

FPGA Field Programmable Gate Array 66

FRT Fixed Rate Trigger 84

FVU Fraction of Variance Unexplained 30, 104

### G

GCD Geometry, Calibration and Detector status 43, 69

GRB Gamma-Ray Burst 11

GSF Global Spline Fit 8, 81, 90, 107, 124, 136, 143, 175

GST Gaisser, Stanev and Tilav 8, 107, 131, 136

GZK Greisen, Zatsepin and Kuzmin 12

### H

H4a Hillas model with four source populations 8, 107, 130

HAWC High Altitude Water Cherenkov Experiment 22, 132

HG High-Gain 39, 61

HLC Hard Local Coincidence 40, 62, 87, 100

### I

IACT Imaging Air-Cherenkov Telescope 21, 35, 75

IC86 IceCube detector configuration with 86 strings 39, 57, 81, 117, 123

ICL IceCube Lab 35, 57

ISM Interstellar Medium 7

### K

KASCADE Karlsruhe Shower Core and Array Detector 18, 45, 132, 136, 175

KDE Kernel Density Estimation 68

KIT Karlsruhe Institute of Technology 22, 175

KLEM Kinematic Lightweight Energy Meter 132

### L

LC Local Coincidence 40, 67

LDF Lateral Distribution Function 64, 92, 100, 181

LG Low-Gain 39, 62

LHAASO Large High Altitude Air Shower Observatory 18, 132

LHC Large Hadron Collider 5, 46, 129, 135, 143, 155

### M

MAE Mean Absolute Error 30

MAPO Martin A. Pomerantz Observatory 35, 57

MC Monte Carlo [45](#), [74](#), [87](#), [104](#), [129](#), [139](#)

MIP Minimum Ionizing Particle [21](#), [62](#)

ML Machine Learning [25](#), [103](#), [181](#)

MLP Multilayer Perceptron [29](#)

MPE Multi Photoelectron [74](#)

MSE Mean Squared Error [30](#)

## N

NaN Not a Number [63](#), [94](#)

NIAC Non-Imaging Air-Cherenkov Telescope [21](#)

nLLH Negative Log-Likelihood [74](#), [89](#)

NN Neural Network [27](#), [87](#), [99](#), [124](#)

NOAA National Oceanic and Atmospheric Administration [55](#), [117](#)

## P

PAMELA Payload for Antimatter Matter Exploration and Light-nuclei Astrophysics [12](#)

PDF Probability Density Function [73](#), [161](#)

PE Photoelectron [43](#), [62](#)

PMT Photomultiplier Tube [37](#), [50](#), [61](#)

PnF Processing and Filtering [35](#), [61](#)

PROPOSAL Propagator with Optimal Precision and Optimized Speed for All Leptons [51](#)

## Q

QGSJet Quark-Gluon String model with Jets [46](#), [99](#), [129](#), [135](#), [143](#), [155](#)

## R

RADE Radius Dependent [58](#), [75](#)

RB RockBottom [89](#), [100](#), [125](#), [139](#)

## S

SGD Stochastic Gradient Descent [31](#)

SLC Soft Local Coincidence [41](#), [61](#), [95](#), [101](#), [135](#), [175](#)

SMT Simple Multiplicity Trigger [41](#), [56](#), [81](#)

SNR Supernova Remnant [9](#)

SPASE South Pole Air Shower Experiment [38](#), [71](#)

SPE Single Photoelectron [73](#)

SPT South Pole Telescope [35](#), [57](#)

SRT Seeded Pulse Cleaning Using Radius and Time [63](#)

## T

TA Telescope Array [23](#)

## LIST OF ACRONYMS AND INITIALISMS

**TP** True Positives [31](#)

**TWC** Time Window Cleaning [64](#)

### U

**UHECR** Ultra-High-Energy Cosmic Ray [9](#), [23](#)

**UV** Ultraviolet [21](#)

### V

**VEM** Vertical Equivalent Muon [62](#), [92](#), [101](#)

### W

**WIPAC** Wisconsin IceCube Particle Astrophysics Center [43](#), [175](#)

## LIST OF SYMBOLS AND CONSTANTS

Symbol	Name
$A$	nucleon number
$A_{\text{(eff)}}$	(effective) area
$a$	neuron activation
$b$	neuron bias
$C$	cost function
$c$	speed of light: $299\,792\,458 \frac{\text{m}}{\text{s}}$
$c_{\omega}$	classification threshold on $\omega_{\text{p}}$
$E$	energy
$e$	elementary charge: $1.602 \cdot 10^{-19} \text{ C}$
$f_{\theta}$	network response function
$h$	height / altitude (above sea level)
$J$	differential CR flux
$L$	number of network layers
$\mathcal{L}$	likelihood
$Q$	DOM charge
$R$	rigidity
$r$	radius
$\vec{r}$	position in IceCube coordinates
$S_{125}$	LDF shower size
$t$	time
$w$	neuron weight
$X_{\text{(max)}}$	slant depth (of shower maximum)
$x_i$	$i^{\text{th}}$ predictor (in ML model)
$y_i$	$i^{\text{th}}$ target (in ML model)

## LIST OF SYMBOLS AND CONSTANTS

Symbol	Name
$Z$	atomic number
$z$	neuron output
$\beta_{\text{(RB)}}$	LDF slope (of RockBottom)
$\beta_{\text{COG2COG}}$	speed of COG2COG track in units of $c$
$\gamma$	spectral index
$\eta$	learning rate
$\theta$	zenith angle
$\vartheta$	free network parameters
$\iota$	Gaussian central mass in percent: 68.27
$\Xi$	vertical atmospheric depth
$\rho_{\text{(air)}}$	(air) density
$\sigma$	activation function
$\sigma_{\text{log}}$	logistic sigmoid
$\varphi$	azimuth angle
$\Psi$	angle between two directions in 3D space
$\Omega$	solid angle
$\omega_{\text{p}}$	proton score

# EIDESSTATTLICHE VERSICHERUNG

Eidesstattliche Versicherung gemäß § 13 Absatz 2 Ziffer 3 der Promotionsordnung des Karlsruher Instituts für Technologie (KIT) für die KIT-Fakultät für Physik:

1. Bei der eingereichten Dissertation zu dem Thema  
SPECTRUM OF LIGHT COSMIC RAYS IN THE PEV ENERGY RANGE MEASURED WITH ICECUBE  
handelt es sich um meine eigenständig erbrachte Leistung.
2. Ich habe nur die angegebenen Quellen und Hilfsmittel benutzt und mich keiner unzulässigen Hilfe Dritter bedient. Insbesondere habe ich wörtlich oder sinngemäß aus anderen Werken übernommene Inhalte als solche kenntlich gemacht.
3. Die Arbeit oder Teile davon habe ich bislang nicht an einer Hochschule des In- oder Auslands als Bestandteil einer Prüfungs- oder Qualifikationsleistung vorgelegt.  
Titel der Arbeit: **Spectrum of Light Cosmic Rays in the PeV Energy Range Measured with IceCube**  
Hochschule und Jahr: **Karlsruher Institut für Technologie, 2025**  
Art der Prüfungs- oder Qualifikationsleistung: **Dissertation**
4. Die Richtigkeit der vorstehenden Erklärungen bestätige ich.
5. Die Bedeutung der eidesstattlichen Versicherung und die strafrechtlichen Folgen einer unrichtigen oder unvollständigen eidesstattlichen Versicherung sind mir bekannt.

Ich versichere an Eides statt, dass ich nach bestem Wissen die reine Wahrheit erklärt und nichts verschwiegen habe.

---

Ort und Datum

---

Unterschrift