

# Studies of TeV muons in cosmic-ray air showers detected with IceTop and IceCube



**Stef Verpoest**

Supervisor: Prof. dr. Dirk Ryckbosch

This dissertation is submitted for the degree of  
*Doctor of Science: Physics*

September 2022



# Studies of TeV muons in cosmic-ray air showers detected with IceTop and IceCube

Stef Verpoest

Proefschrift ingediend tot het verkrijgen van de academische graad van  
Doctor in de Wetenschappen: Fysica

Promotor: Prof. dr. Dirk Ryckbosch

Co-promotor: Prof. dr. Didar Dobur

Leden van de examenjury:

Prof. dr. Maarten Baes, voorzitter (Universiteit Gent)

Prof. dr. Toon Verstraelen (Universiteit Gent)

Prof. dr. Steven Caluwaerts (Universiteit Gent)

Prof. dr. Archisman Ghosh (Universiteit Gent)

Prof. dr. Frank Schröder (University of Delaware, USA)

Prof. dr. Simona Toscano (Université libre de Bruxelles)

Vakgroep Fysica en Sterrenkunde

Faculteit Wetenschappen

Universiteit Gent

Academiejaar 2022-2023





## Acknowledgements

The completion of this thesis took a lot of time and energy, and I would not have gotten to this point without the many individuals who have supported me in one way or another along the way. I would like to express my sincere gratitude to these people below.

In the first place, I would like to thank my supervisor Prof. dr. Dirk Ryckbosch for the many opportunities and the support he has given me. Through his inspiring lectures, the opportunity to do an internship abroad, my master's thesis in his group, and the past four years as a PhD student, he has undoubtedly had a tremendous impact on my path.

I also owe a lot to my IceCube collaborators, especially the Cosmic Ray Working Group, whose regular feedback was indispensable for my work. In particular, I would like to thank several people at the University of Delaware who invited me to visit their group, which was an important, enriching experience for me: Serap Tilav, Dennis Soldin, David Seckel, Alan Coleman. I also thank the late Thomas Gaisser, working with whom was a wonderful experience and an honor.

I would like to thank my colleagues at INW for the pleasant working environment; especially Liam, Gianni, and Tom, who were so kind to provide me with the necessary distraction and were always available for discussions about work and far beyond that. Furthermore, I'd like to thank some of my predecessors working on IceCube in Ghent. Ward's guidance made my master's thesis a positive experience, which contributed to my choice of continuing in physics. Sam's sporadic presence in the office was very valuable at the start of my PhD, providing me with a clear direction to pursue at a time when I needed it.

In addition, I am grateful to many of my dear friends whose contribution to this work is indirect, but still of vital importance. Alex, Bram, Michiel, Nelis, Victor, Vince, Viktor; thank you for the many fun evenings, dinners, concerts, travels, and so forth throughout the years, and for being supportive and understanding of my frequent absence in the final stages of writing this thesis. Part of the same group, I am deeply grateful to my dear friend Simon, who unfortunately is not with us anymore, but who lives on through the influence he has had on me as a person. I am also thankful to my "physics friends" Seppe, Simon, Thomas, and Tom, whose kindness and curiosity are uplifting

and inspiring. A special thank you to Seppe and Tom for providing me with the best environment imaginable to start a PhD in. To Eline, thank you for the crucial support during challenging times and for being an all-round wonderful person. To Carlos, thank you for being a mentor and a loving friend. Seba and Bien, thank you for the relaxing evenings in times of stress.

Finally, I want to express my gratitude to my family. I thank in particular my parents, Katty and Peter, for encouraging me to follow my own interests and for actively supporting me in doing so throughout my life. I furthermore thank my mother and my sister Lies for making my life easier in a variety of ways during the final stages of my PhD.

# Table of contents

<b>Nederlandstalige samenvatting</b>	<b>xi</b>
<b>Introduction</b>	<b>xv</b>
<b>1 Introduction to cosmic rays</b>	<b>1</b>
1.1 What are cosmic rays? . . . . .	1
1.2 Cosmic-ray measurements . . . . .	2
1.2.1 Type of experiment . . . . .	3
1.2.2 Energy spectrum . . . . .	3
1.2.3 Composition . . . . .	5
1.3 Sources and propagation . . . . .	6
1.3.1 Galactic cosmic rays . . . . .	9
1.3.2 Extragalactic cosmic rays . . . . .	10
1.4 High-energy cosmic rays . . . . .	11
1.5 Spectrum and composition models . . . . .	13
<b>2 Extensive air showers</b>	<b>17</b>
2.1 The atmosphere . . . . .	17
2.2 Air-shower development . . . . .	19
2.2.1 Heitler-Matthews model . . . . .	21
2.2.2 Particle distributions and air-shower observables . . . . .	25
2.3 Measurement techniques . . . . .	30
2.3.1 Particle-detector arrays . . . . .	30
2.3.2 Other detector types . . . . .	31
2.4 Air-shower simulations . . . . .	33
2.4.1 CORSIKA . . . . .	34
2.4.2 Hadronic interaction models . . . . .	35
2.4.3 Simulation uncertainties and the Muon Puzzle . . . . .	37

---

<b>3</b>	<b>The IceCube Neutrino Observatory</b>	<b>45</b>
3.1	Detection principle . . . . .	45
3.1.1	Cherenkov effect . . . . .	46
3.1.2	Energy loss in matter . . . . .	47
3.2	The detector . . . . .	49
3.2.1	Digital optical modules . . . . .	49
3.2.2	The in-ice array: IceCube . . . . .	52
3.2.3	The surface array: IceTop . . . . .	55
3.3	Data acquisition . . . . .	58
3.3.1	Local Coincidence . . . . .	59
3.3.2	Triggers . . . . .	60
3.3.3	Cosmic-ray filters . . . . .	61
3.4	Calibration . . . . .	62
3.4.1	VEMCal . . . . .	63
3.5	Detector simulation . . . . .	65
<b>4</b>	<b>Cosmic-ray physics with IceCube</b>	<b>69</b>
4.1	Coincident detection of air showers . . . . .	69
4.2	Cosmic-ray standard processing . . . . .	72
4.2.1	IceTop processing . . . . .	74
4.2.2	IceCube processing . . . . .	78
4.2.3	Quality cuts . . . . .	80
4.2.4	Performance . . . . .	82
4.3	Selected results from IceTop & IceCube . . . . .	84
4.3.1	Energy spectrum and composition . . . . .	86
4.3.2	Hadronic interaction models . . . . .	88
4.3.3	Density of GeV muons . . . . .	88
4.3.4	Atmospheric muons and seasonal variations . . . . .	90
4.4	Simulated datasets . . . . .	93
4.4.1	CORSIKA simulation specifics . . . . .	93
4.4.2	Detector simulation specifics . . . . .	96
<b>5</b>	<b>High-energy muon production profiles</b>	<b>99</b>
5.1	Introduction . . . . .	99
5.2	Parameterization of muon production profiles . . . . .	101
5.2.1	Simulations . . . . .	101
5.2.2	Formula . . . . .	102

---

5.2.3	Fit parameters . . . . .	107
5.2.4	Performance tests . . . . .	110
5.3	Example applications for IceCube . . . . .	119
5.3.1	Estimate of seasonal variations of TeV muon bundles . . . . .	119
5.3.2	Application to muon bundle energy loss . . . . .	123
5.4	Discussion . . . . .	127
<b>6</b>	<b>Testing hadronic interaction models</b>	<b>129</b>
6.1	Introduction . . . . .	129
6.2	Event selection . . . . .	132
6.3	Composition-sensitive observables . . . . .	133
6.3.1	Muon bundle energy loss . . . . .	134
6.3.2	IceTop LDF slope . . . . .	134
6.3.3	Density of GeV muons . . . . .	137
6.3.4	Systematic uncertainties . . . . .	139
6.4	Internal consistency of models . . . . .	140
6.5	Discussion . . . . .	142
<b>7</b>	<b>Multiplicity of TeV muons analysis</b>	<b>147</b>
7.1	Introduction . . . . .	148
7.2	Event selection . . . . .	149
7.3	MC expectations . . . . .	151
7.4	Neural-network reconstruction . . . . .	155
7.4.1	MC dataset . . . . .	155
7.4.2	Input . . . . .	155
7.4.3	Method . . . . .	158
7.4.4	Training . . . . .	161
7.4.5	Performance . . . . .	162
7.5	Calibration . . . . .	167
7.5.1	Correction factors . . . . .	171
7.5.2	Simple correction procedure . . . . .	173
7.5.3	Iterative correction procedure . . . . .	174
7.6	Simulation tests . . . . .	180
7.6.1	Analysis performance tests . . . . .	180
7.6.2	Data-MC comparison . . . . .	182
7.7	Results . . . . .	184
7.7.1	Application to data . . . . .	185

7.7.2	Systematic uncertainties . . . . .	186
7.7.3	Final results . . . . .	190
7.8	First look at seasonal variations . . . . .	195
7.9	Discussion & outlook . . . . .	200
<b>8</b>	<b>Summary and discussion</b>	<b>205</b>
	<b>Appendix A Low-energy muon profile parameterization</b>	<b>209</b>
	<b>Appendix B TeV muon multiplicity supplementary material</b>	<b>213</b>
B.1	Neural network architectures . . . . .	213
B.2	MC tests . . . . .	215
B.3	Data-MC comparison . . . . .	219
	<b>References</b>	<b>223</b>

# Nederlandstalige samenvatting

*“Die Ergebnisse der vorliegenden Beobachtungen scheinen am ehesten durch die Annahme erklärt werden zu können, daß eine Strahlung von sehr hoher Durchdringungskraft von oben her in unsere Atmosphäre eindringt (...)”*. Dit schreef Victor Hess in 1912 na een reeks experimenten, uitgevoerd tijdens ballonvaarten, die het begin van de astrodeeltjesfysica betekenden. De straling waarvan sprake wordt tegenwoordig “kosmische straling” of “cosmic rays” genoemd, en bestaat uit atoomkernen die zich aan zo goed als de lichtsnelheid voortbewegen, en ons vanuit elke richting uit de ruimte bereiken. Sinds de ontdekking van kosmische straling is de hoofdvraag in dit onderzoeksveld welke extreme processen in het universum de bronnen zijn van deze hoog-energetische deeltjes. Om deze vraag te beantwoorden kunnen bepaalde belangrijke eigenschappen van kosmische straling gemeten worden, zoals hun energie, massa, en bewegingsrichting. Zo’n metingen kunnen uitgevoerd worden op grote hoogte met ballonvaarten of satellietexperimenten. Bij de hoogste energieën is de flux van kosmische straling echter zo laag dat deze experimenten, met hun beperkte omvang, niet voldoende deeltjes meer opvangen voor gedetailleerde studies. Hiervoor bestaat weliswaar een oplossing, waarbij de atmosfeer van de Aarde als deel van de detector gebruikt wordt.

Wanneer een kosmisch stralingsdeeltje met hoge energie de atmosfeer bereikt, zal het botsen met een atmosferische atoomkern, zoals zuurstof of stikstof. In deze interactie worden allerlei nieuwe deeltjes geproduceerd, die op hun beurt opnieuw kunnen interageren. Op deze manier ontstaat er een cascade, waarin miljoenen deeltjes geproduceerd kunnen worden die zich kunnen verspreiden over een oppervlakte van meerdere vierkante kilometers. Deze deeltjeslawine wordt gewoonlijk een “air shower” genoemd. Air showers geven ons een manier om kosmische straling indirect te bestuderen, door op het aardoppervlak grote detectoren te bouwen die de deeltjes uit de air showers meten. Op basis van deze metingen kunnen bepaalde eigenschappen van het kosmisch stralingsdeeltje dat de atmosfeer binnendrong gereconstrueerd worden. Deze techniek laat dus toe om de eigenschappen van kosmische straling bij de hoogste energieën te bestuderen. De interacties in de atmosfeer kunnen verder ook gebruikt worden om deeltjesfysica te doen

bij energieën buiten het bereik van deeltjesversnellers gebouwd op Aarde, zoals de Large Hadron Collider in Zwitserland. Het experiment waarmee air showers gemeten worden in dit werk is het IceCube Neutrino Observatory.

Het IceCube Neutrino Observatory is een experiment dat zich bevindt op Antarctica, aan de geografische Zuidpool. Het experiment bestaat uit twee grote onderdelen: IceTop, een detector op het oppervlak voor kosmische stralingsfysica, en IceCube, een detector diep in het ijs met als voornaamste doel neutrinofysica. Beide werken op basis van het Cherenkov-effect, waarbij blauw licht uitgestraald wordt wanneer deeltjes met een snelheid door een medium bewegen die hoger is dan de snelheid van licht in dat medium. Het Cherenkov-licht kan opgevangen worden met optische sensoren om zo de passage van een deeltje te gaan waarnemen. IceTop bestaat uit cilindrische tanks verspreid over een oppervlakte van  $1 \text{ km}^2$  gevuld met ijs en twee zo'n optische sensoren, die het licht veroorzaakt door air-showerdeeltjes opvangen. De IceCube detector bestaat uit ongeveer 5000 optische sensoren verspreid over  $1 \text{ km}^3$  in het Antarctische ijs, op een diepte van 1.5 tot 2.5 km, en meet energetische deeltjes die zich door het ijs voortplanten.

IceTop en IceCube vormen samen een unieke kosmische-stralingsdetector. Met IceTop kan kosmische straling gemeten worden in het energiegebied van 1 PeV tot 1 EeV. De grote hoogte van de detector, ongeveer 2.8 km boven zeeniveau, zorgt dat air-showerdeeltjes gemeten worden voor er veel absorptie plaatsvindt in de atmosfeer. Het signaal in IceTop is gedomineerd door elektromagnetische deeltjes – elektronen, positronen, en fotonen – wat toelaat om de energie van de primaire kosmische straling nauwkeurig te bepalen. Van groot belang in air-showermetingen zijn ook muonen, die in mindere mate dan de elektromagnetische deeltjes voorkomen. Muonen bevatten veel informatie over de hadronische interacties in de atmosfeer, en het aantal muonen is afhankelijk van de massa van het primaire deeltje. IceTop meet ook muonen, typisch met energie rond 1 GeV. Muonen met hogere energieën, in de buurt van 1 TeV, kunnen diep in het ijs doordringen en, wanneer de geometrie van de air shower het toelaat, ook een signaal achterlaten in IceCube. Zo kunnen IceTop en IceCube verschillende componenten van de air shower meten om zowel informatie over de energie als de massa van kosmische straling te verkrijgen.

Een probleem eigen aan het bestuderen van kosmische straling via air showers is dat het een heel indirecte manier is om informatie te vergaren: er wordt een fractie van de miljoenen deeltjes gemeten die ontstaan in talloze interacties in de atmosfeer. Bij gevolg zijn gedetailleerde simulaties van de ontwikkeling van de air shower en van de detector nodig voor de interpretatie van de metingen. Simulaties van air showers beschrijven in detail hoe de verschillende deeltjes zich voortplanten in de atmosfeer en interageren.

De eigenschappen van deze interacties worden bepaald door modellen die hadronische interacties beschrijven, afgestemd op metingen die gedaan werden in experimenten met deeltjesversnellers. De interacties in air showers liggen echter niet altijd binnen het bereik van versnellerexperimenten, bijvoorbeeld wat betreft hun energie. Hierdoor moeten de hadronische interactiemodellen gaan extrapoleren, wat grote onzekerheden met zich meebrengt. Zo zijn er significante verschillen tussen air-showersimulaties gedaan op basis van verschillende modellen. Ook is het duidelijk vastgesteld dat in sommige gevallen de simulaties geen goede beschrijving geven van de observaties in experimentele data, wat het bepalen van belangrijke eigenschappen als de massacompositie, en bijgevolg het bepalen van de bronnen van kosmische straling, moeilijk maakt. Omdat deze verschillen vooral in de muon component van air showers voorkomen, wordt dit probleem “the Muon Puzzle” genoemd. Door metingen van muonen te doen met IceTop en IceCube proberen we bij te dragen aan het begrijpen en uiteindelijk oplossen van dit probleem.

Dit werd in dit werk in eerste instantie gedaan door observabelen die gevoelig zijn aan de massa van de kosmische straling in experimentele data te vergelijken met verwachtingen uit simulatie. Als de simulatie een goede beschrijving van de data geeft, dan zouden alle observabelen tot dezelfde conclusie moeten leiden over de massacompositie van kosmische straling. We observeren echter inconsistenties in simulaties gebaseerd op elk van de drie verschillende hadronische interactiemodellen gebruikt in de studie. Dit betekent dat analyses van de massacompositie sterk afhankelijk zijn van welke component van de air shower gemeten wordt, wat duidelijk een probleem vormt. In twee state-of-the-art modellen zijn er indicaties voor een inconsistentie tussen de GeV-muonen gemeten in IceTop en de TeV-muonen gemeten in IceCube. Dit vormde de motivatie om de TeV-muon component in meer detail te gaan bestuderen, complementair aan een reeds bestaande gedetailleerde meting van de GeV-muonen.

Dit werk bevat een aantal verschillende studies rond de hoog-energetische muon component van air showers. In de eerste plaats werd de muonproductie in de atmosfeer bestudeerd in simulaties, en werd een parameterizatie bekomen die gebruikt kan worden om op eenvoudige wijze berekeningen te maken over hoe de muonproductie afhangt van de atmosferische temperatuur. Zo kan de grootte van eventuele seizoenseffecten in muonmetingen in ondergrondse detectoren ingeschat worden.

Daarnaast werd een eerste analyse van data ontwikkeld die probeert het aantal muonen met een energie hoger dan 500 GeV in air showers te bepalen, als functie van de energie van de primaire kosmische straling. Hiervoor werden verticale air showers gebruikt die zowel in IceTop als IceCube gezien worden. Informatie uit het IceTop signaal en het signaal in IceCube, achtergelaten door de hoog-energetische muonen, werd gebruikt als

input voor neurale netwerken, die getraind werden om de energie van het kosmische deeltje te bepalen, alsook de multipliciteit van de hoog-energetische muonen in de air shower. De reconstructies die bekomen werden met de neurale netwerken werden daarna gecorrigeerd met correctiefactoren die bepaald werden uit simulatie. Op deze manier werden resultaten voor de muonmultipliciteit bekomen onder de veronderstelling van verschillende hadronische interactiemodellen. We vinden dat de resultaten overeenkomen met verwachtingen volgens recente modellen van de massacompositie. De inconsistentie met de metingen van GeV-muonen wordt echter bevestigd wanneer recente hadronische interactiemodellen gebruikt worden. Dit heeft belangrijke implicaties voor het begrijpen van de ontwikkeling van air showers in de atmosfeer en kan bijdragen tot een volgende generatie van hadronische interactiemodellen en verbeterde simulaties. Dit is cruciaal om tot een gedetailleerde beschrijving van kosmische straling te komen.

# Introduction

In the early 20th century, a series of measurements performed during balloon flights revealed that the Earth's atmosphere is constantly impinged by energetic radiation from outer space [1]. This radiation turned out to consist predominantly of positively charged particles with relativistic energies. More specifically, they are fully ionized atomic nuclei, and are referred to as cosmic rays. Soon, it was discovered that cosmic rays can induce large particle cascades through interactions in the atmosphere, called extensive air showers [2]. Measurements of air showers confirmed the existence of cosmic rays with extremely high energies, reaching, as we know today, even beyond 100 EeV [3]. The study of cosmic rays and the secondary radiation produced in the atmosphere has from the early days been of interest to both particle physics and astrophysics communities, and their discovery marked the beginning of the field called astroparticle physics.

Before human-made particle accelerators became the norm, the study of cosmic-ray interactions in the atmosphere was the main source of particle-physics discoveries. Examples are the discovery of anti-matter with the observation of the positron [4], the discovery of the muon [5], and the discovery of various mesons, such as the pion [6] and the kaon [7]. While particle accelerators are a powerful tool for precision measurements in particle physics, cosmic rays remain of interest, especially as the center-of-mass energy reached in the first interaction of ultra-high-energy cosmic rays can exceed that reached in the largest accelerator, the Large Hadron Collider [8], by two orders of magnitudes.

For astrophysics, cosmic rays are of interest as messengers from the extreme processes in the universe where they are produced. While many properties of cosmic rays have been measured in detail and various plausible source candidates are known from optical and other observations, a complete picture of sources and propagation of cosmic rays has not been established. This is especially true at the highest energies, where populations of Galactic cosmic-ray accelerators are expected to run out of power and to be overtaken by extragalactic sources.

One difficulty is that cosmic rays at the highest energies, beyond 100 TeV, can only be studied in sufficient numbers by detecting extensive air showers with large detectors at

the Earth's surface. To reconstruct from these indirect measurements properties such as the energy and mass of the primary cosmic ray, accurate calculations of the development of the shower are needed. This makes most measurements model dependent, and comes with large uncertainties, as models rely on extrapolations from accelerator measurements to describe the high-energy interactions in the shower. These uncertainties are currently the main factor limiting detailed measurements of cosmic rays at the highest energies, preventing the full potential for both particle physics and astrophysics to be fulfilled.

This work contributes to the understanding of the development of air showers through the measurement of muons with the IceTop and IceCube detectors at the South Pole. IceTop is a surface particle detector array, which measures the abundant electromagnetic air-shower particles accompanied by muons. The IceCube detector is a large multi-purpose detector embedded in the Antarctic ice sheet at a depth of 1.5 km, and can detect bundles of muons from air showers that are sufficiently energetic to propagate all the way through the ice. By measuring different components of the air shower, one can obtain insight into the shower development and perform tests of the performance of the hadronic interaction models used in simulations. The focus in this work is to do this using the high-energy muons seen in IceCube. The structure of this thesis, from background information to analysis results, is as follows.

Chapter 1 provides an introduction to cosmic rays and their experimental properties. After a brief discussion of their sources and how they propagate before reaching the Earth, various relevant experimental results for high-energy cosmic rays, at energies of 1 PeV and above, are discussed.

The indirect detection of cosmic rays through extensive air showers is the subject of Chapter 2. The development of air showers in the atmosphere is illustrated through a simple model and various simulation results. Afterwards, different ways of observing air showers are discussed. As air shower observations rely on detailed simulations for interpretation, the simulation software and different hadronic interaction models used are introduced afterwards. This is followed by a discussion on the so-called Muon Puzzle, which refers to various discrepancies that have been observed between data and simulation in the muon component of air showers specifically [9]. This is the context in which the work presented in this thesis is performed.

In Chapter 3, the IceCube Neutrino Observatory, including the IceTop and IceCube detectors are described. This includes a description of the detection principle, different detector components, and the data-acquisition system, as well as a brief discussion on how detector simulations are performed.

How the IceTop and IceCube detectors are used specifically for cosmic-ray physics is the subject of Chapter 4. A description of the type of event that is the focus of this work, namely near-vertical air showers triggering both IceTop and IceCube, is given, followed by different steps of cleaning and reconstruction performed on the data. A selection of results from IceTop and IceCube, relevant to this work, is discussed. The specifics of the different simulated datasets used in the analysis of experimental data are also described in this chapter.

A first study of high-energy muons is performed based on dedicated simulations in Chapter 5. Here, a parameterization of the longitudinal production profiles of high-energy muons is given. It includes the dependence of the production on the atmospheric density, and can therefore be used to estimate the seasonal variations expected in muon measurements for different detectors. Example applications relevant for IceCube are discussed.

In Chapter 6, an analysis testing the quality of the description of experimental data by simulations based on different hadronic interaction models is presented. This is done by comparing composition-sensitive observables based on different shower components, requiring that they lead to consistent conclusions about the underlying mass composition. As it will turn out, inconsistencies are observed, even for state-of-the-art models.

An analysis of the high-energy muon content in IceTop-IceCube coincident events is developed in Chapter 7. The goal is to obtain the average number of muons with an energy higher than 500 GeV as a function of the primary cosmic-ray energy. To this end, existing reconstructions of the air-shower in IceTop and the muon bundle energy loss in IceCube are combined with neural networks. The final result is obtained after the application of correction factors derived from simulation using different hadronic interaction models. The results are compared to expectations from composition models and a measurement of the density of GeV muons in IceTop. A first look at seasonal variations in the muon multiplicity and a comparison with the predictions from Chapter 5 are also presented.

In Chapter 8, we give a summary and final discussion of the results.



# Chapter 1

## Introduction to cosmic rays

Cosmic rays are fully ionized nuclei originating from the most violent processes in our Universe, spanning many orders of magnitude in energy. In this work, we will study properties of the large particle cascades cosmic rays induce in the Earth's atmosphere, introduced in Chapter 2. Measurements of such cascades are important to learn about the nature of cosmic rays, their sources, and how they propagate. Given that cosmic rays exist with energies far beyond the reach of particle accelerators on Earth, they are also of interest to study high-energy particle interactions.

In this chapter, we give an overview of concepts related to cosmic rays relevant to this work. After briefly introducing cosmic rays in Section 1.1, several basic experimental properties are discussed in Section 1.2. Subsequently, some considerations related to the possible sources and the propagation of cosmic rays are given in Section 1.3. In Section 1.4, more attention is given to the high-energy cosmic rays in the energy range of the IceTop detector (Chapter 3) and beyond, including experimental results and how they relate to possible source scenarios. Finally, several models describing the cosmic-ray flux and mass composition, relevant for calculations performed later in this work, are discussed in Section 1.5.

### 1.1 What are cosmic rays?

Early in the twentieth century, it was established that the Earth's atmosphere is constantly impinged by energetic radiation. Of crucial importance were the measurements of Victor Hess, who performed balloon flight experiments where he measured ionizing radiation in the atmosphere with electrometers. He found that the amount of radiation increased with altitude, indicating an extraterrestrial origin [1]. As the same intensity of radiation was observed at night as during the day, it was concluded that the Sun could not be

the source. The radiation was given the name *cosmic radiation*, or *cosmic rays*. For his discoveries, Hess was awarded the Nobel Prize in Physics in 1936.

We know today that cosmic rays are electrically charged particles moving through space at mostly relativistic velocities. While sometimes used to include different types of particles, e.g. electrons, the term cosmic rays will in this work always refer to ionized nuclei. Such nuclei enter the Earth's atmosphere at a rate of about 1000 per square meter per second. The majority of them, about 90%, are single protons. Second in abundance are helium nuclei, also called alpha particles. Heavier elements make up only a small fraction of all cosmic rays, although their relative contribution becomes more important at high energy, as will be discussed in Section 1.4.

Since the discovery, the central question of cosmic-ray physics has been what the origin of cosmic-ray nuclei is, and how they are accelerated to the extreme energies at which they can be observed. While large advancements have been made, the answer to these questions is not yet fully known. This is especially the case at the high-energy end of the spectrum, where cosmic rays are thought to be of extragalactic origin.

In addition to the astrophysical information they carry, cosmic rays are of interest for particle physics. In the 1930s and '40s, the positron [4], muon [5], pion [6], and kaon [7] were discovered by studying the secondary particles produced in interactions of cosmic rays in the atmosphere. While particle-accelerator experiments have been the main source of new and precise information about elementary particles in the last decades, cosmic-ray experiments still prove useful for particle physics, for example in studying interactions at energies beyond the reach of the accelerators. This will be illustrated in Chapter 2.

## 1.2 Cosmic-ray measurements

Cosmic-ray particles have several properties that contain valuable information about their origin and propagation. Of particular importance are the energy, charge, mass number, and arrival direction of the nucleus. Numerous experiments, varying strongly in nature, have provided us with data over the decades. In this section, we will briefly discuss the different classes of cosmic-ray experiments as well as the information they have collected on the energy spectrum and composition of cosmic rays. Measurements of the anisotropy in arrival direction, while important for uncovering the sources of cosmic rays, are beyond the scope of this work and will not be discussed further.

### 1.2.1 Type of experiment

Cosmic rays span a vast range of energies, from below a GeV to more than 100 EeV. As shown in Fig. 1.1, the flux drops steeply with increasing energy. As a result, different types of detectors are needed to cover different parts of the spectrum.

At energies below a few 100 TeV, the flux is sufficiently high to perform direct measurements of cosmic rays with satellite- or balloon-based experiments, which capture cosmic rays before they interact in the atmosphere. Detectors equipped with magnetic spectrometers, such as AMS02 [10], or calorimeters, such as CREAM [11], can provide measurements of the spectra of individual elements at high precision. The proton-only data shown in Fig. 1.1 was obtained from these two experiments.

At higher energies, the small flux of cosmic rays demands detectors with large areas exposed for long periods of time. This is achieved by building detectors at the surface of the Earth which measure the particles resulting from the cascades initiated by cosmic rays interacting in the atmosphere, so-called *extensive air showers*. Such experiments, called *air shower arrays*, can provide sufficient statistics to study cosmic rays up to the end of the energy spectrum because of their enormous dimensions and exposure times of several years or decades. These indirect measurements, however, come with the complication that they give only limited information about the cosmic-ray primaries, making it difficult to precisely determine their properties. Extensive air showers and the different techniques used to detect them are the subject of Chapter 2.

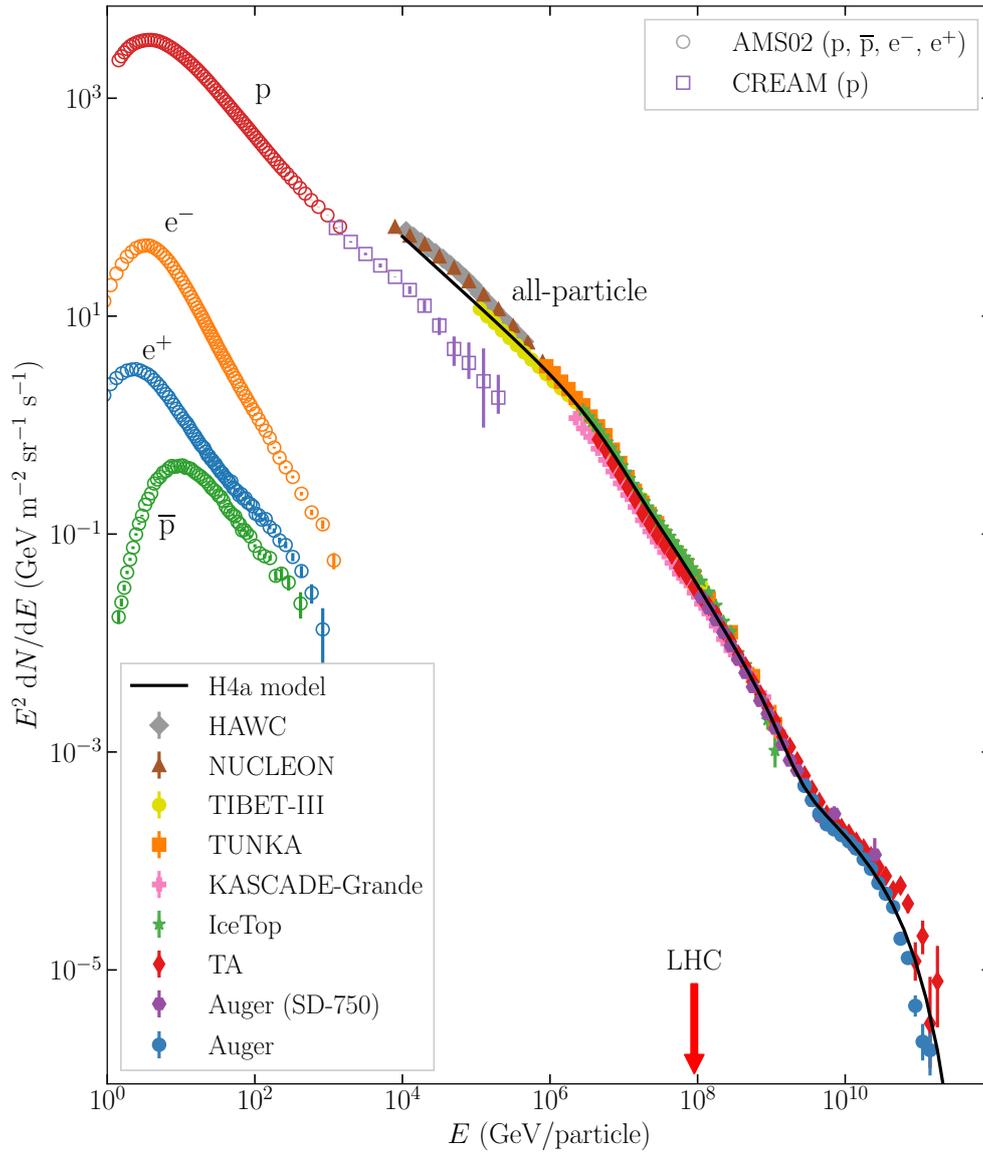
### 1.2.2 Energy spectrum

The cosmic-ray energy spectrum is expressed as the flux of nuclei of all types differential in energy, written as

$$J(E) \equiv \frac{dN}{dE dA d\Omega dt} \text{ (GeV}^{-1} \text{ m}^{-2} \text{ sr}^{-1} \text{ s}^{-1}) \text{ .} \quad (1.1)$$

Because of the enormous range of values the flux spans and how steeply it falls – over thirty orders of magnitude spanning about twelve orders of magnitude in energy – it is customary to multiply the flux with a power of the energy to better visualize its structures, as was done in Fig. 1.1. A notable characteristic of the spectrum is that it can be described in several large intervals of energy by inverse power laws of the form

$$J(E) \propto E^{-\alpha} \text{ ,} \quad (1.2)$$



**Fig. 1.1:** All-particle energy spectrum of cosmic rays at the Earth measured by a variety of experiments. At low energy, the proton-only spectrum measured by AMS [12] and CREAM [13] is shown, as well as the smaller fluxes of electrons, positrons and antiprotons measured by AMS. The red arrow indicates the cosmic-ray energy at which interaction with a stationary target has a center-of-mass energy of 13 TeV, the energy reached at the Large Hadron Collider (Run 2). Also shown is the total flux of the H4a model (see Section 1.5).

References for the all-particle measurements: HAWC [14], NUCLEON [15], TIBET-III [16], TUNKA-133 [17], KASCADE-Grande [18], IceTop [19], Telescope Array [20], Pierre Auger Observatory (SD-750) [21], Pierre Auger Observatory [22]; partially obtained from CRDB [23, 24].

where  $\alpha$  is the so-called *differential spectral index*. From 10 GeV to 1 PeV, it is  $\approx 2.7$ . At lower energies the spectrum is modified by solar modulation, the effect where the wind of ionized plasma emitted by the Sun prevents the lowest-energy Galactic cosmic rays from reaching the inner solar system. At energies beyond 1 PeV, a steepening of the spectrum can be observed, known as *the knee*, followed by a region of more complex structure, as discussed in Section 1.4. For energies between 100 PeV and 5 EeV,  $\alpha \approx 3.1$ . The transition to a harder spectrum<sup>1</sup> at higher energies, where  $\alpha$  is about 2.6, is referred to as *the ankle*. Above  $5 \times 10^{19}$  eV, data from the HiRes [25] experiment, Telescope Array [26], and the Pierre Auger observatory [27] suggest a suppression of the spectrum, sometimes called *the toe*.

### 1.2.3 Composition

The relative abundance of different nuclei in cosmic rays contains important information about their origin and propagation. In Fig. 1.2, we show the cosmic-ray composition measured on Earth, compared to the relative abundances of elements in the solar system. These show similar features, such as the odd-even effect, where the nuclei with an even number of protons  $Z$  are more stable and therefore more abundant. The abundance of elements heavier than iron is low; a result of the fact that the production of such elements requires an input of energy, so that they are not produced in stellar thermonuclear reactions. There are, however, also some important differences between the Solar-system and cosmic-ray abundances that are worth highlighting.

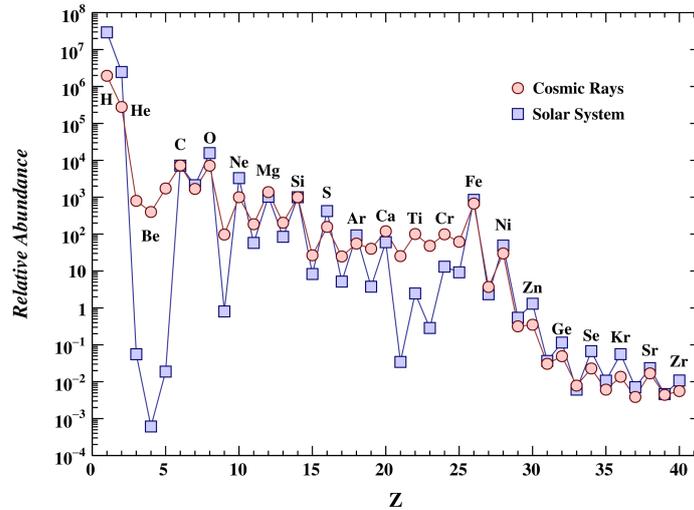
The groups of elements Li-B and Sc-Mn are rare in the solar system, as these are virtually absent as end-products of stellar nucleosynthesis. In cosmic radiation, however, they are present as the result of collisions of the cosmic-ray nuclei, mainly the abundant elements C, N, O, and Fe, with the interstellar medium. The relative abundance of different elements therefore contains important information about how the cosmic rays propagate between production and observation [29].

A second difference, which is less understood, is the lower relative abundance of protons in cosmic rays compared to hydrogen in the solar system. This could reflect a difference in composition at the sources of cosmic rays, or be related to the fact that hydrogen is relatively hard to ionize and therefore to accelerate.

Note that because of the steep spectrum, the measured composition is dominated by low-energy cosmic rays. The composition is, however, energy dependent. Measurements

---

<sup>1</sup>The spectrum of cosmic rays is said to be softer/harder if the flux decreases faster/slower with energy. A harder spectrum has more particles at higher energies than a softer spectrum, corresponding to a smaller spectral index.



**Fig. 1.2:** The elemental abundances of cosmic rays measured on Earth compared to the solar system abundances, relative to Silicon(=10<sup>3</sup>). Figure taken from Ref. [28].

of the composition of high-energy cosmic rays, and how they relate to their acceleration and propagation, are discussed in Section 1.4.

### 1.3 Sources and propagation

Before discussing the properties of high-energy cosmic rays in more detail in Section 1.4, we discuss in this section some general properties of the acceleration and propagation of cosmic rays at different energies. This is not meant to be an exhaustive description of this complex field, but rather an introduction to some concepts relevant to cosmic rays at energies around the knee and beyond. For a more complete discussion, we refer to Ref. [30].

The sources of cosmic rays are not yet fully understood. It is, however, clear that nearly all but the ones at very low energy ( $\lesssim 100$  MeV) come from outside the solar system [30, p. 16]. The majority of them are expected to come from local sources in the Galaxy, with a smaller contribution from extragalactic sources.

A possible mechanism for the acceleration of cosmic rays is Fermi acceleration in shock waves or magnetic clouds [31]. In this mechanism, kinetic energy of moving magnetized plasma gets transferred to individual charged particles. If a process is considered where a test particle increases its energy by an amount  $\Delta E = \xi E$  proportional to its original energy  $E$  through some proportionality constant  $\xi$ , and the particle has multiple encounters with a probability  $P_{\text{esc}}$  of escaping the acceleration region per encounter, it can be shown

that the number of particles  $N$  above a certain energy produced by the accelerator is given by

$$N(> E) \propto \frac{1}{P_{\text{esc}}} \left( \frac{E}{E_0} \right)^{-\gamma}, \quad (1.3)$$

with

$$\gamma \approx \frac{P_{\text{esc}}}{\xi}. \quad (1.4)$$

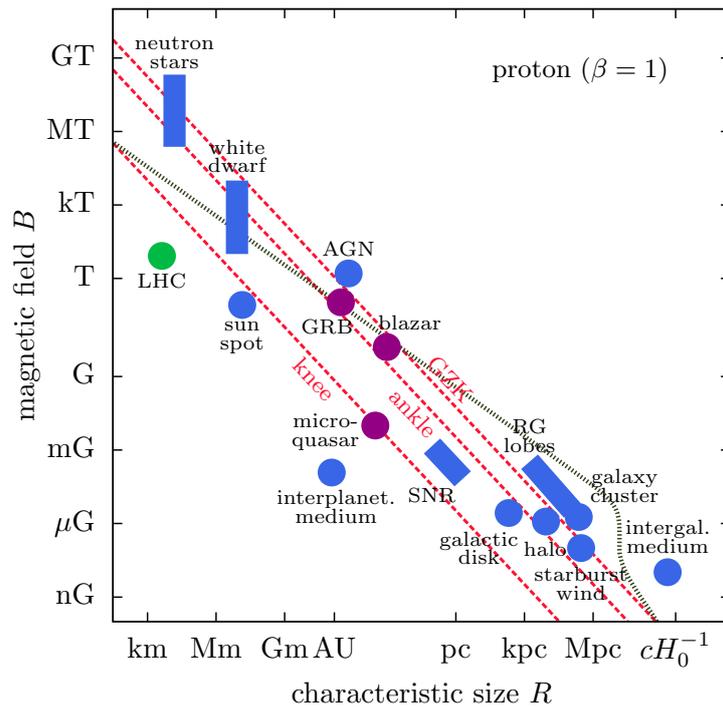
Such a mechanism therefore leads to the desired behavior of a power-law energy spectrum. A more detailed treatment for the case of strong shocks shows that the obtained spectral index has a value close to the one needed to describe the observed cosmic-ray spectrum. Note that a difference between the spectral index at Earth and at the source is expected due to propagation effects, such as an energy-dependent escape probability from the galaxy [30, p. 242].

The maximal energy that can be obtained is limited by the gyroradius of a particle inside the acceleration region. For a magnetic field  $B$  perpendicular to the direction of the particle, the gyroradius is given by  $r_L = R/B$ , where  $R = pc/Ze \approx E/Ze$  is the magnetic rigidity for a high-energy nucleus with charge  $Ze$ , momentum  $p$ , and energy  $E$ . The maximal energy at which a particle can be confined in the acceleration region is therefore limited by the dimensions of the accelerator to

$$E < ZeBR, \quad (1.5)$$

with  $R$  the radius of the acceleration region. The implications of this condition were pointed out by Hillas, who compared the size and magnetic field strength of potential sources of cosmic rays in a diagram known as *the Hillas plot* [32]. A more recent adaptation of this diagram is shown in Fig. 1.3. The diagonal lines show a lower limit for accelerators of protons; objects to the left of the diagonals cannot accelerate protons to such energies due to loss of confinement in the acceleration region. For example, typical supernova remnants (SNR) can be responsible for cosmic rays up to and inside the knee region, but are unlikely to be the main source of particles beyond the ankle.

The condition in Eq. (1.5) furthermore implies that nuclei with a higher charge can be accelerated to higher energies, as they are confined more easily in the acceleration region. If a dominant cosmic-ray source reaches its maximum energy for protons, this could cause an evolution of the composition with primary energy toward heavier nuclei, with successive cutoffs for helium, carbon, etc. This scenario is often referred to as a *Peters cycle* [33]. Note that also propagation effects, e.g. a rigidity-dependent leakage out of the galaxy, can cause a similar evolution in the composition.



**Fig. 1.3:** A recent version of the Hillas plot, showing the magnetic field strength and size of the acceleration region for different potential sources of cosmic rays. Red diagonal lines are drawn at the energy of the knee and ankle (Section 1.2.2) and the GZK cutoff (Section 1.3.2); objects to the left of them cannot accelerate protons up to this energy due to loss of confinement in the acceleration region. Also shown by the dotted gray line is an upper limit obtained for realistic sources including energy loss processes in the source (synchrotron losses and interaction with the cosmic photon background). Image taken from Ref. [34].

### 1.3.1 Galactic cosmic rays

When discussing potential Galactic sources of cosmic rays, it is useful to estimate the power required to maintain the Galactic density of cosmic rays in equilibrium. To quantify this equilibrium state, an understanding of the propagation of cosmic rays in the Galaxy is necessary.

Cosmic rays diffuse in the Galactic magnetic field of about  $3\ \mu\text{G}$  and may interact with the interstellar medium. Assuming that nuclei such as boron are mainly produced by spallation of primary nuclei and are themselves not produced in cosmic-ray sources, as discussed in Section 1.2.3, one can calculate from the relative abundance of such nuclei that cosmic rays in the GeV range traverse an effective grammage of on average  $5\text{--}10\ \text{g}/\text{cm}^2$  before observation, and that this grammage decreases with increasing energy [29]. The disk of the Milky Way, with an average density of 1 hydrogen atom per  $\text{cm}^3$  and a height of 200 pc, has a thickness of about  $1 \times 10^{-3}\ \text{g}/\text{cm}^2$ . This implies that cosmic rays travel much larger distances than this thickness, which suggests they diffuse in a volume including the Galactic disk.

The transport of cosmic rays in the Milky Way is governed by complex diffusion equations. Standard, simplified models exist, such as the *leaky box* model [30, p. 191]. In this model, the cosmic rays propagate freely in a containment volume, and have a small probability of escape each time they reach its boundary. This leads to an exponential distribution of path lengths with an energy dependent escape length. Together with the assumption that the local cosmic-ray energy spectrum is typical for the Milky Way, it can be calculated that the power required to maintain the Galactic cosmic rays in equilibrium is  $7 \times 10^{40}\ \text{erg}/\text{s}$  [30, p. 194].

This power requirement is suggestive of SNRs. A rough estimate learns that  $10\ M_{\odot}$  ejected with a velocity of  $\sim 5 \times 10^8\ \text{cm}/\text{s}$  from a type II supernova every 30 years, equates to a power of  $3 \times 10^{42}\ \text{erg}/\text{s}$ . Therefore, an efficiency of a few percent would suffice for supernova blast waves to energize all Galactic cosmic rays.

More direct evidence in favor of the acceleration of cosmic rays in SNRs can be obtained from gamma-ray observations. In contrast to charged particles which get deflected by magnetic fields, neutral particles like neutrinos and photons point back to their origin. Protons interacting with matter in sources or the interstellar medium can produce neutral pions which decay into two photons. In the rest frame of the  $\pi^0$ , the photons are produced back to back with energy  $E_{\gamma} = m_{\pi}/2$ , producing a characteristic peak in the gamma-ray spectrum. The Fermi-LAT collaboration reported evidence for the presence of this feature in the spectrum of two SNRs, which cannot be explained by gamma rays produced in leptonic processes [35].

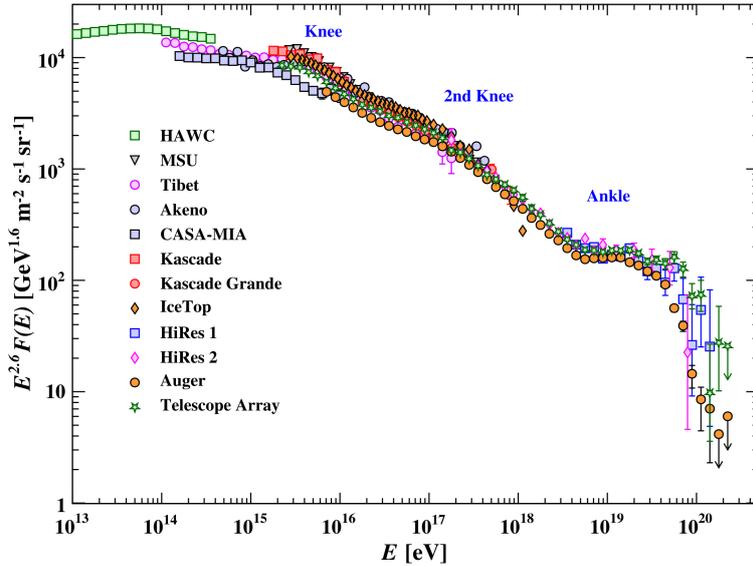
The maximal energy that can be reached in SNRs is uncertain, and it is possible that they are not the source of the highest-energy Galactic cosmic rays. Indications in measurements of the energy spectrum and composition of high-energy cosmic rays for a possible contribution from a different, higher energy, source are discussed in Sections 1.4 and 1.5. A promising candidate for the acceleration of PeV cosmic rays is, according to gamma-ray observations by H.E.S.S, the Galactic center [36]. Promising results have recently also been published by the LHAASO collaboration in Ref. [37], where they announced the discovery of 12 PeVatrons in our Galaxy.

### 1.3.2 Extragalactic cosmic rays

Particles with energies around the ankle and higher are referred to as ultra-high energy cosmic rays (UHECR). They are generally considered to originate from sources outside the Milky Way, as they show no sign of anisotropy related to the Galactic plane [38]. The situation is, however, less clear than for Galactic cosmic rays. From the Hillas plot, we see that typical SNRs are unlikely to accelerate particles up to the ankle. Potential sources that have a product of magnetic field strength  $B$  and size  $R$  sufficiently large to accelerate particles all the way up to  $10^{20}$  eV and are able to fulfill the power requirement to maintain the observed density of UHECR include galaxy clusters, active galactic nuclei (AGN), and gamma-ray bursts [30, p. 352].

AGNs are observed sources of high-energy gamma-rays. The observed spectrum can, however, also be explained through leptonic production. Along with gamma rays, cosmic-ray interactions will also produce neutrinos from the decay of charged pions and kaons. One method to discover extragalactic accelerators of cosmic-ray nuclei is therefore the simultaneous observation of neutrinos and gamma rays from the same source. In 2017, IceCube detected an astrophysical neutrino candidate from the direction of a blazar, a class of AGN, which was also in a gamma-ray flaring state [39]. The significance of the two being correlated was estimated to be  $3\sigma$ . This measurement has been an important milestone towards future multi-messenger detections of cosmic-ray sources, which together with future large-scale cosmic-ray experiments are expected to strongly constrain the potential source scenarios for UHECR.

For the propagation of UHECR over extragalactic distances, several energy loss processes have to be considered. The most important interaction target in this case is the cosmic microwave background (CMB), with a photon density of about  $410 \text{ cm}^{-3}$ . The energy threshold for photopion production  $\gamma p \rightarrow \pi^0 p$  with a CMB photon is  $E_{\text{th}} \sim 7 \times 10^{19}$  eV. Greisen, Zatsepin and Kuzmin realized that this effect makes the universe opaque to very high-energy cosmic-ray protons, leading to a suppression of the local flux



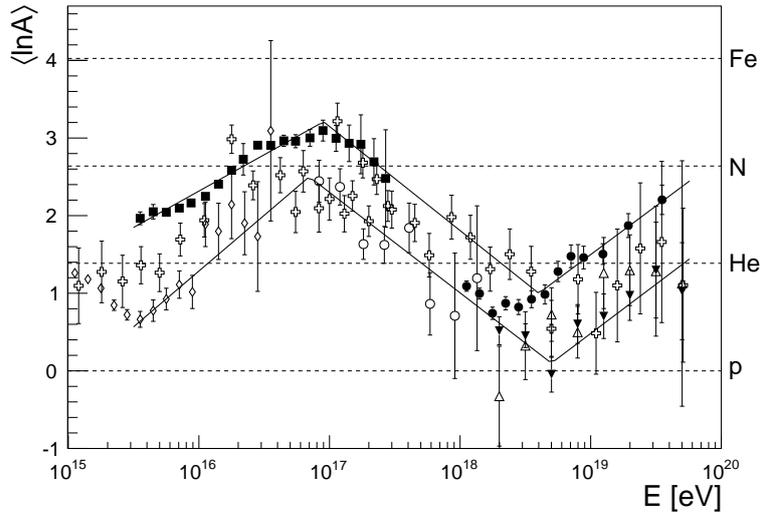
**Fig. 1.4:** The all-particle cosmic-ray energy spectrum as measured by a variety of air-shower experiments. Figure taken from Ref. [28].

above  $E_{\text{th}}$  known as the GZK cutoff [40, 41]. Other energy loss processes such as the electromagnetic production of  $e^+e^-$  pairs, and interaction with e.g. infrared background photons, have smaller contributions to the total proton energy loss. For heavier nuclei, the dominant interaction process at the highest energies is photodisintegration [42].

## 1.4 High-energy cosmic rays

As the main focus of this work will be indirect measurements of cosmic rays performed with the IceTop and IceCube detectors (Chapter 3), we give in this section some more details on the properties of cosmic rays with energies above 100 TeV.

Fig. 1.4 shows the all-particle energy spectrum of cosmic rays measured by a collection of air-shower experiments. It can be seen that the transition region where the spectrum steepens has a complex structure, rather than a single knee. In fact, there is a first knee-like structure at  $\sim 3$  PeV, followed by a hardening of the spectrum, after which the spectrum steepens again at the second knee around 100 PeV. The KASCADE and KASCADE-Grande experiments have shown that the first knee corresponds to a steepening of the spectra of light elements, while the heavy component of their data showed a knee-like structure at the energy of the second knee [43, 44]. This is supported by data collected by several experiments measuring the depth of shower maximum  $X_{\text{max}}$  (see Section 2.2), shown in Fig. 1.5. Because of the indirect measurement it is difficult



**Fig. 1.5:** Average logarithmic mass of cosmic rays as a function of energy derived from  $X_{\max}$  measurements with respect to simulations using the hadronic interaction model Sibyll 2.1 (see Section 2.4.2). Figure taken from Ref. [45].

to precisely determine the type of primary nucleus on an event-by-event basis, and the data is instead interpreted in terms of the mean logarithmic mass  $\langle \ln A \rangle$  of the cosmic rays. While there is a significant spread between the experiments, the results appear to show a monotonic increase in average mass from around the knee up to the second knee, after which it decreases again up to the ankle. Interpreting the knee as the point where a population of sources, e.g. SNRs, reach their maximal energy for protons, this is suggestive of a Peters cycle as described in Section 1.3.

As stated in Section 1.3.2, particles with energy above  $\sim 5$  EeV are generally assumed to be from extragalactic origin. The hardening of the ankle is therefore often interpreted as the transition from Galactic to extragalactic cosmic rays, although this is not the only possible explanation for this feature [30, p. 341]. Furthermore, data from KASCADE-Grande shows a hardening in the light component around 100 PeV, which may indicate that the transition to extragalactic cosmic rays already happens at this energy. Note that in the scenario where the knee region is the result of a Peters cycle of SNRs, a second higher-energy population of Galactic sources is needed to explain the presence of Galactic protons up to the energy where the extragalactic component becomes dominant.

The data in Fig. 1.4 shows a clear suppression of the flux above  $6 \times 10^{19}$  eV, which is consistent with the GZK cutoff, discussed in Section 1.3.2. It is, however, important to note that this is not the only possible explanation. For example, recent measurements by Auger show an increasingly heavy composition approaching the energy of the suppression,

raising the possibility that the suppression is a result of the maximum energy of the accelerators being reached [27].

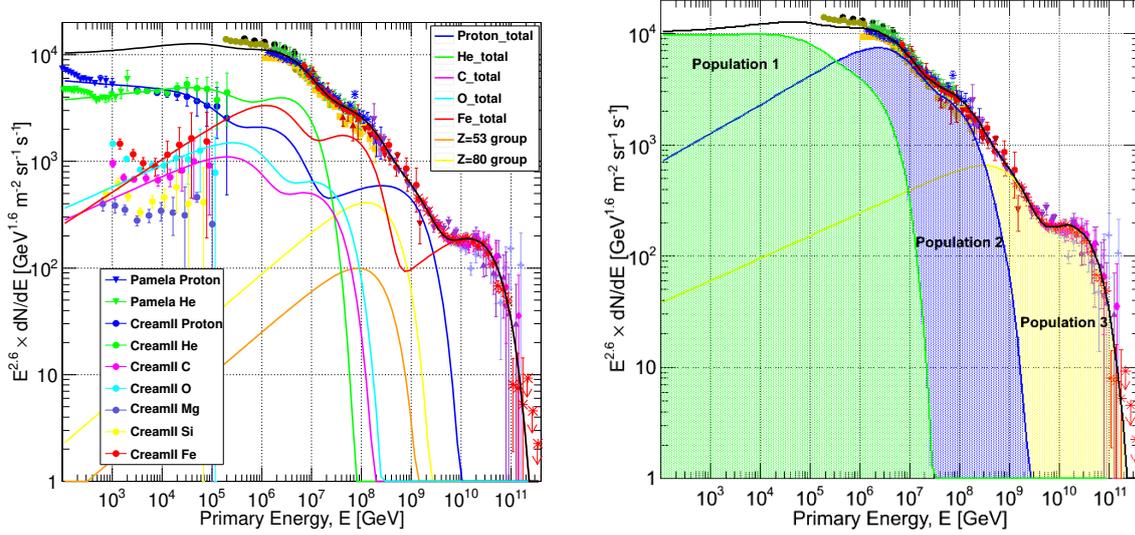
From the discussion above and in Section 1.3, we can conclude that there are still many open questions for high-energy Galactic and extragalactic cosmic rays, concerning both their sources and propagation. This is in part a consequence of the steeply falling flux and the difficulty of interpreting indirect measurements. Large uncertainties in the simulations of extensive air showers, necessary for the interpretation of the measurements, have prevented an unambiguous determination of the cosmic-ray composition, which is crucial for understanding where cosmic rays come from. It is among the goals of this thesis to study the effect of uncertainties in air-shower simulations on composition measurements.

## 1.5 Spectrum and composition models

Measurements of the shape of the all-particle energy spectrum and the nuclear composition of the primaries place constraints on the different contributing sources and their spectra. Models fitted to the observations may prove useful in understanding the underlying processes. Models of the flux of individual elements are also important for calculations of several derived quantities relevant for experiments, such as atmospheric lepton fluxes. In this section we introduce some recent models describing the energy spectrum and composition for high-energy cosmic rays.

As the determination of the mass of a cosmic-ray nucleus with indirect air-shower experiments is difficult, mainly because of large fluctuations from shower to shower, the composition is conventionally described in terms of four or five major groups of nuclei. Most commonly, these are p, He, CNO, Mg-Si and Fe.

To describe the measurements from the knee region to the end of the spectrum, Hillas proposed to assume three populations of cosmic rays [46]. The first population can be associated with acceleration by SNRs, with the knee being related to the cutoff of this population. The second population results from a hypothetical higher-energy source in our Galaxy. The third population is associated with particles of extragalactic origin. Note that in reality there could be many more components related to different classes of sources with different characteristics; the three population model is a minimal assumption describing the case where the transition between Galactic and extra-galactic cosmic rays occurs around the ankle. The models below assume that the features of the spectrum depend on magnetic rigidity, following Peters as discussed in Section 1.3. The sources are assumed to have an exponential cutoff at a characteristic rigidity  $R_c$ . The contribution



**Fig. 1.6:** Overview of the spectrum from below the knee to the ankle of the three-population GST model fit to experimental data. Left: Lines show contributions from individual groups of nuclei from all populations. Right: Shaded regions show total contributions from the three populations. Figures taken from Ref. [49].

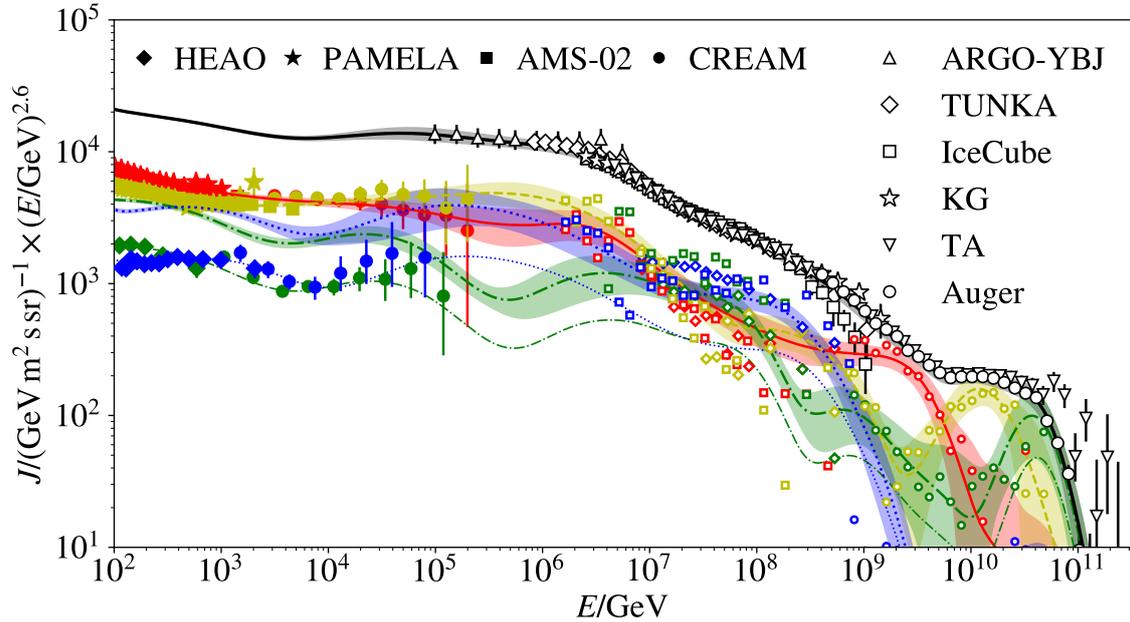
to the all-particle spectrum of mass group  $i$  can then be written as

$$\phi_i(E) = \frac{dN_i}{d \ln E} = \sum_{j=1}^3 a_{i,j} E^{-\gamma_{i,j}} \times \exp \left[ -\frac{E}{Z_i R_{c,j}} \right], \quad (1.6)$$

where the sum runs over the three populations and  $\gamma$  and  $a$  are the integral spectral index and a normalization constant respectively. An example of a three population, five mass group model is the H3a parameterization from Ref. [47]. A variant called the H4a parameterization assumes the extragalactic population to be protons only<sup>2</sup>.

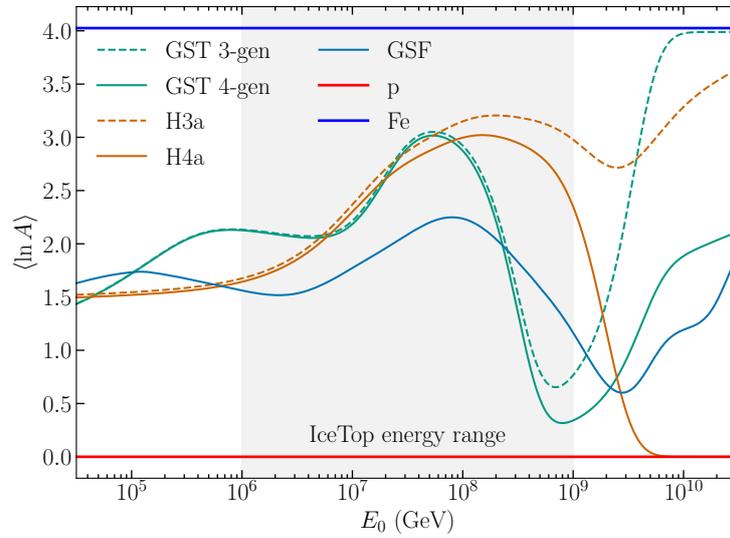
One difficulty in air shower measurements is the determination of the primary energy of the cosmic ray, for which experiments rely on air-shower simulations, which come with significant uncertainties. This leads to an offset in the energy spectra measured by different experiments. However, features like the knee which have a similar shape in the data of most experiments can be used to rescale the measurements and bring them in better agreement with each other. This was done in Ref. [49], after which a model similar as described above was fitted to the data. The resulting model, referred to as the GST model, is shown in Fig. 1.6.

<sup>2</sup>Note that a single-component proton-only extragalactic flux was found to be at strong variance with Auger data [48].



**Fig. 1.7:** The GSF model. Data points show energy-scale adjusted experimental data, with error bars representing combined statistical and systematic uncertainties. The lines show the all-particle flux (black solid line), the contribution from protons (red solid line), helium (yellow dashed line), the oxygen group (green dash-dotted line) and the iron group (blue dotted line). Bands around the model lines show a variation of one standard deviation. Figure taken from Ref. [50].

A different approach was recently used for the Global Spline Fit (GSF) model [50]. It does not try to explain the data in terms of theoretical models, but rather describes an average of the measurements of various experiments, taking into account both statistical and systematic uncertainties. As many experiments measure the cosmic-ray flux in four mass groups instead of five, this is also the approach taken by this model. The groups are p, He, O and Fe, and these have the attractive quality of spanning similar ranges in logarithmic mass  $\ln A$  (see Chapter 2). The differential flux of each mass group is parametrized with smooth spline curves. The energy scale is fixed by direct measurements and corrections for the energy-scale offsets between indirect experiments are derived as part of the global fit. In this way, the GSF model provides a world-average of the cosmic-ray spectrum and composition with a cross-calibrated common energy scale. It also provides a covariance matrix containing the experimental uncertainties, which can be propagated to derived quantities. The results are shown in Fig. 1.7.



**Fig. 1.8:** Mean logarithmic mass as a function of primary cosmic-ray energy for different composition models discussed in the text. The shaded region shows the energy region most relevant for IceTop/IceCube (see Chapters 3 and 4).

An overview of the average logarithmic mass given by the models discussed in this section is shown in Fig. 1.8.

# Chapter 2

## Extensive air showers

When a cosmic ray enters the Earth's atmosphere, it will eventually interact with an air nucleus. In this interaction, a variety of new particles is produced. Through subsequent interactions and decays, these particles will themselves produce more new particles, and so on. In this way, a large cascade of secondary particles forms, called an *extensive air shower* (EAS). Air showers provide a way to indirectly study cosmic rays at energies above 100 TeV, where direct measurements are not feasible. They also allow us to probe fundamental interactions at energies that are beyond the reach of modern particle accelerators.

In this chapter, we give an introduction to air-shower physics and related concepts that are relevant to this thesis. As air showers develop in the Earth's atmosphere, some properties of the atmosphere are introduced in Section 2.1. In Section 2.2, we discuss the development of different air-shower components and some general relations between air-shower observables and the primary cosmic ray. Various ways of detecting EAS are discussed in Section 2.3. Since EAS are very complex, their properties are best studied through detailed simulations. In Section 2.4, we describe modern air-shower simulations and discuss some of their known problems and uncertainties.

### 2.1 The atmosphere

The development of an EAS is determined by the propagation of particles in the atmosphere and their interactions with atmospheric nuclei. Of crucial importance is the density of the atmosphere, as it is a measure of the number of interaction targets that are present. Both the density at a certain altitude and the integrated density or the amount of traversed matter are relevant for understanding the development of the shower.

Atmospheric data, such as obtained from the AIRS instrument aboard the NASA Aqua satellite [51], usually provide the temperature of the atmosphere at different pressure levels. The atmosphere can be approximated to behave as an ideal gas, where the density  $\rho$  is related to the pressure  $P$  and temperature  $T$  as

$$\rho = \frac{MP}{RT}, \quad (2.1)$$

where  $R$  is the molar gas constant and  $M$  is the mean molar mass of air. For dry air which has a composition of 78% nitrogen, 21% oxygen and 1% argon,  $M = 0.02897 \text{ kg mol}^{-1}$ . This composition is stable up to heights of about 100 km [52].

The amount of matter traversed by a particle travelling vertically down the atmosphere is called the *vertical depth*  $X_v$ . At an altitude  $h$ , it is given by

$$X_v(h) = \int_h^\infty \rho(h') dh' = \frac{P(h)}{g}, \quad (2.2)$$

where  $g$  is the gravitational constant, and the density vanishes for large altitudes. The depth is usually expressed in units of  $\text{g/cm}^2$ . For a particle that enters the atmosphere at a zenith angle  $\theta$  relative to the vertical direction, more mass will be traversed and one can calculate the so-called *slant depth*  $X$  by integrating along the trajectory. In the approximation that the Earth is flat, which works well for  $\theta \leq 65^\circ$ , the relation between the vertical altitude  $h$  and the distance upwards along the trajectory  $l$  is simply  $l = h/\cos\theta$  and

$$X = \frac{X_v}{\cos\theta}. \quad (2.3)$$

Equivalently to Eq. (2.2), the density is  $\rho = -dX_v/dh$ . Combining this with Eq. (2.1), it is possible to write

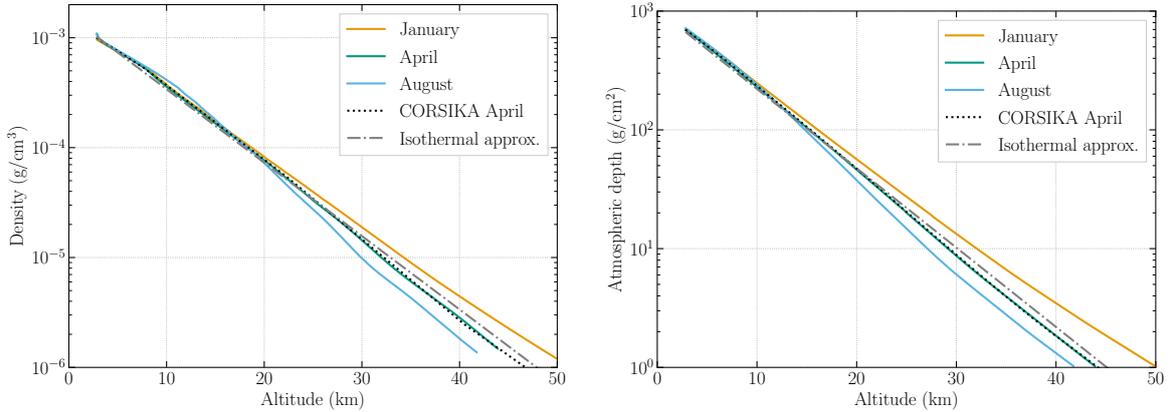
$$\frac{dX_v}{X_v} = -\frac{Mg}{RT} dh, \quad (2.4)$$

which leads to an exponential solution for an isothermal atmosphere

$$X_v = X_0 e^{-\frac{h}{h_0}}, \quad (2.5)$$

with a scale height  $h_0 = RT/Mg$ . For a typical lower stratospheric temperature of 220 K, this is  $\approx 6.5 \text{ km}$ . The vertical depth at sea level is  $X_0 \approx 1030 \text{ g/cm}^2$ .

The isothermal approximation is useful to get a sense of how the atmospheric depth relates to altitude. In reality, however, the atmosphere is not isothermal. Its temperature varies with altitude and has seasonal variations. Examples of the atmospheric density and depth as a function of altitude are shown in Fig. 2.1. In Chapter 5, we will discuss



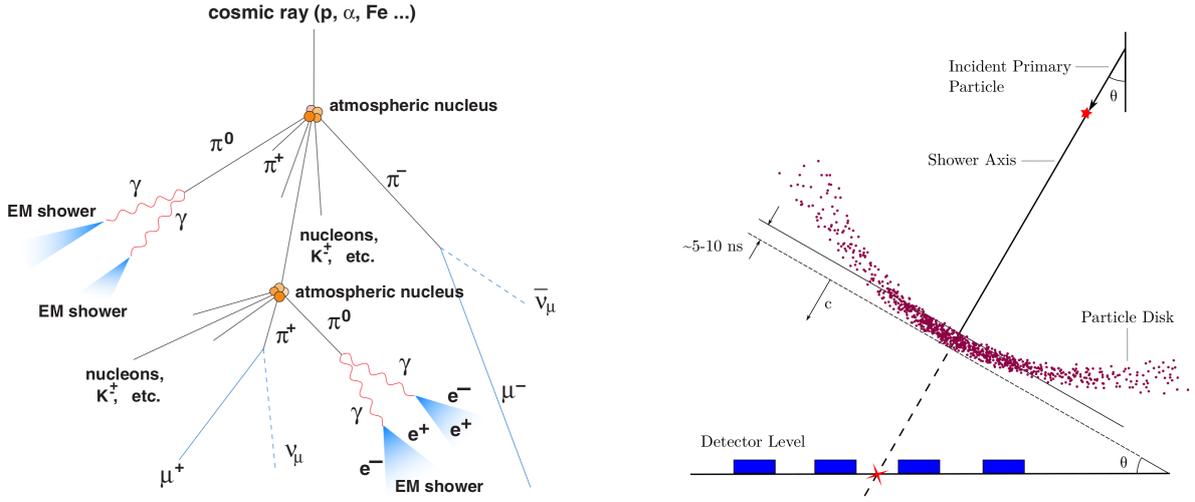
**Fig. 2.1:** Atmospheric density (left) and vertical depth (right) as a function of altitude. Shown are profiles based on AMRC balloon measurements [53] and AIRS satellite measurements [51] at the South Pole for three different months in 2013, kindly provided by Takao Kuwabara [54]. Also shown is the atmospheric model for April used in air shower simulations in this work, performed with CORSIKA (see Sections 2.4 and 4.4), and the isothermal approximation of Eq. (2.5).

how variations in the atmospheric temperature influence the production of high-energy muons in an air shower.

## 2.2 Air-shower development

Fig. 2.2 shows a schematic representation of an air shower. When a cosmic ray enters the Earth’s atmosphere and traverses more and more matter, it will eventually interact with an atmospheric nucleus. Subsequently, a cascade of particles will develop along the original trajectory of the primary cosmic ray, the so-called *shower axis*. The number of particles multiplies, reaches a maximum, and then attenuates, as the energy of the particles falls below the threshold for the production of new particles. Finally, the resulting particles arrive at the Earth’s surface as a thin disk of a few meters thick in the center, increasing up to a few hundred meters at large lateral distances.

An EAS induced by a cosmic ray has three main components: a hadronic, electromagnetic and muonic component. The hadronic component consists of a large variety of high-energy mesons and baryons, produced in every hadronic interaction including the first interaction of the cosmic ray in the atmosphere, and governs the development of the shower. The high-energy hadrons continually feed the electromagnetic component, primarily through the decay of neutral pions and eta particles into photons, which is their dominant decay channel. Each photon with sufficient energy will generate electro-



**Fig. 2.2:** Schematic representation of an extensive air shower. Left: Development of an EAS and its different components. Figure taken from Ref. [55]. Right: Propagation of the thin disk of EAS particles in the direction of the shower axis with zenith angle  $\theta$  with respect to the vertical direction. The shower core is marked by the red X. Figure taken from Ref. [56].

magnetic sub-showers, which develop through alternate positron-electron pair production and bremsstrahlung. High-energy hadrons may also interact with atmospheric nuclei and further contribute to the hadronic cascade. Lower-energy charged mesons, mainly pions and kaons, feed the muonic component by decaying into muons and neutrinos.

The longitudinal development of the number of different species of particles in a cascade can be described by a system of equations of the form

$$\frac{dN_i(E, X)}{dX} = -\frac{N_i(E, X)}{\lambda_i} - \frac{N_i(E, X)}{d_i} + \sum_j \int_E^\infty \frac{F_{ji}(E, E_j)}{E} \frac{N_j(E_j, X)}{\lambda_j} dE_j, \quad (2.6)$$

called *cascade equations* [30, p. 107]. Here,  $N_i(E, X)dE$  is the flux of particles of type  $i$  at slant depth  $X$  with energies in the interval  $E$  to  $E + dE$ . The probability that a particle of type  $i$  interacts when traversing an infinitesimal slant depth element  $dX$  is equal to  $dX/\lambda_i(E)$ , where  $\lambda_i$  is the interaction length in air. The probability for a particle to decay in  $dX$  is  $dX/d_i(E)$ , with  $d_i$  the decay length. Both the interaction and decay length are expressed in  $\text{g}/\text{cm}^2$ , as is  $X$ . The function  $F_{ji}(E, E_j)$  is the particle yield for a particle of energy  $E_j$  colliding with an air nucleus to produce a particle  $i$  with energy  $E < E_j$ , and follows from the inclusive cross section for the process. For

simplicity, a source term from particle decays and a term for energy loss due to ionization and stochastic losses have been omitted; for more details see Refs. [57] and [58]. The initial condition can either be chosen as the entire spectrum of cosmic-ray nucleons, as e.g. discussed in Section 1.5, to obtain inclusive lepton fluxes in the atmosphere, or as a single particle of a particular type and energy to calculate the longitudinal development of an air shower.

It is interesting to note the different behavior of the interaction and decay lengths with respect to the density of the atmosphere. The interaction length in  $\text{g}/\text{cm}^2$  is obtained by multiplying the corresponding length in units of distance  $l_i$  with the density. This results in

$$\lambda_i = l_i \rho = \frac{\rho}{n_A \sigma_i^{\text{air}}} = \frac{A m_p}{\sigma_i^{\text{air}}}, \quad (2.7)$$

where  $\rho$  is the density of the atmosphere at a certain altitude,  $n_A$  is the corresponding number density of nuclei with average mass number  $A$ , and  $\sigma_i^{\text{air}}$  is the cross section for interaction with air. Thus, the density dependence cancels out of the interaction length when expressed in  $\text{g}/\text{cm}^2$ . In the decay length on the other hand, the density dependence does not cancel. It is

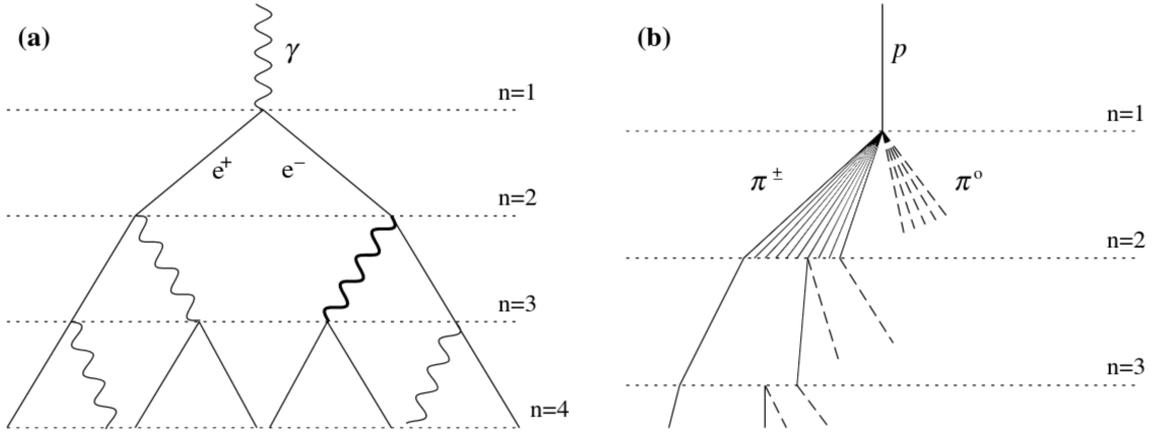
$$d_i = \rho \gamma c \tau_i, \quad (2.8)$$

with  $\gamma$  the Lorentz factor and  $\tau_i$  the lifetime of the particle. This is especially relevant when considering the production of muons, which depends on the probability of decay versus re-interaction of mesons, as discussed in Chapter 5.

In the remainder of this chapter, we will not study the properties of EAS by attempting to solve the cascade equations of Eq. (2.6). The principal way of examining the development of EAS is through detailed simulations, as described in Section 2.4. We will start by obtaining some insight in the development of an air shower through a simplified model in the following subsection. In the subsequent subsection, we discuss the air-shower properties in more detail through realistic distributions of particles of the different air-shower components.

### 2.2.1 Heitler-Matthews model

The basic physics behind several EAS observables can be understood through a simplified model of the development, which approximates the cascade by a simple deterministic branching model. We first introduce such a model for purely electromagnetic (EM) cascades, after which we discuss a generalization for cascades induced by cosmic rays.



**Fig. 2.3:** Heitler approximation of (a) an EM cascade and (b) a hadronic shower. In the EM cascade, at each interaction point a splitting happens due to bremsstrahlung or pair production. In the hadronic shower, at each interaction a number of pions is produced. Dashed lines indicate neutral pions, which quickly decay and yield EM sub-showers. Figure taken from Ref. [59].

### Electromagnetic cascades

An EM cascade can be initiated by a high-energy EM particle, either a photon, electron or positron<sup>1</sup>. Such a cascade can be approximated by a simple model introduced by Heitler [60]. As shown in Fig. 2.3 (left), the Heitler model assumes a particle to undergo a splitting after travelling a certain collision length  $\lambda_{\text{EM}}$ . For photons, the process is the production of an  $e^+e^-$  pair,  $\gamma \rightarrow e^+e^-$ . The other process is bremsstrahlung, where a charged particle scatters off a Coulomb field, producing a photon  $e^\pm \rightarrow e^\pm + \gamma$ . In each interaction, one new particle is produced, and the energy is divided equally over the two particles. After  $n$  collision lengths, a depth of  $X = n\lambda_{\text{EM}}$  has been traversed along the shower axis, and the number of particles is

$$N_{\text{EM}}(X) = 2^{X/\lambda_{\text{EM}}}. \quad (2.9)$$

The energy per particle at a certain depth is then

$$E(X) = E_0/N_{\text{EM}}(X). \quad (2.10)$$

The multiplication continues until the energies of the particles are too low for pair production or bremsstrahlung, and particles only lose energy through ionization or get

<sup>1</sup>In the context of air showers, we will often use the term electron to refer to both  $e^-$  and  $e^+$ . The same holds when talking about muons, which usually includes both muons and anti-muons.

absorbed. For EM cascades in air, this *critical energy* is  $E_c^{\text{EM}} \approx 87$  MeV. The maximal number of particles in the cascade is therefore

$$N_{\text{max}}^{\text{EM}} = N_{\text{EM}}(X_{\text{max}}^{\text{EM}}) = E_0/E_c^{\text{EM}}, \quad (2.11)$$

and  $X_{\text{max}}^{\text{EM}}$  is known as the *depth of shower maximum*. From Eqs. (2.9) and (2.11), this depth is

$$X_{\text{max}}^{\text{EM}} = \lambda_{\text{EM}} \frac{\ln(E_0/E_c)}{\ln 2}. \quad (2.12)$$

Therefore, this simple Heitler model for EM cascades predicts the following relation with primary energy for the number of particles at shower maximum and its depth:

$$N_{\text{max}}^{\text{EM}} \propto E_0 \quad \text{and} \quad X_{\text{max}}^{\text{EM}} \propto \ln(E_0). \quad (2.13)$$

### Hadronic showers

To gain insight into cosmic-ray induced EAS, the treatment of EM cascades above can be extended to hadronic cascades as presented by Matthews [59]. We first look at a cascade initiated by a single proton. Similar to before, hadrons are modeled in the Heitler-Matthews model as interacting after traveling a distance corresponding to their interaction length  $\lambda_{\text{had}}$ . Each interaction is assumed to produce  $n_{\text{tot}}$  pions, two-thirds of which are charged  $n_{\text{ch}}$  and one-third of which are neutral, with equal division of energy. The neutral pions decay immediately into two photons, which initiate EM cascades. Charged pions, on the other hand, will interact again with air nuclei if their energy is sufficiently high. When their interaction length becomes longer than their decay length, they decay into muons and neutrinos instead.

In each hadronic interaction, one-third of the energy is transferred to the EM shower component through neutral pions, so that after  $n$  interaction lengths, the energy in the hadronic and EM components is given by

$$E_{\text{had}} = \left(\frac{2}{3}\right)^n E_0 \quad \text{and} \quad E_{\text{EM}} = \left[1 - \left(\frac{2}{3}\right)^n\right] E_0. \quad (2.14)$$

For a typical  $n \approx 5$  [61], about 90% of the primary energy is eventually carried by EM particles. The shower maximum will therefore also be dominated by the EM sub-showers.

Considering only the EM sub-showers produced in the first hadronic interaction<sup>2</sup>, we find

$$X_{\max}^p \approx \lambda_{\text{had}} + X_{\max}^{\text{EM}}(E_0/(2n_{\text{tot}})) \approx \lambda_{\text{had}} + \lambda_{\text{EM}} \ln \left( \frac{E_0}{2n_{\text{tot}} E_c^{\text{EM}}} \right), \quad (2.15)$$

where the factor 1/2 accounts for the two photons in the decay of the neutral pions.

The number of muons in the shower follows from the number of charged hadrons, i.e.

$$N_{\mu}^p = n_{\text{ch}}^n = \left( \frac{2}{3} n_{\text{tot}} \right)^n. \quad (2.16)$$

The energy per pion is then equal to the critical energy  $E_c^{\pi}$  where decay and interaction have the same probability,

$$E_{\pi} = \frac{E_0}{n_{\text{tot}}^n} = E_c^{\pi}. \quad (2.17)$$

From these two equations we find that the number of muons depends on the primary energy  $E_0$  as

$$N_{\mu}^p = \left( \frac{E_0}{E_c^{\pi}} \right)^{\beta}, \quad \text{with } \beta = \frac{\ln n_{\text{ch}}}{\ln n_{\text{tot}}} = \frac{\ln \frac{2}{3} n_{\text{tot}}}{\ln n_{\text{tot}}}. \quad (2.18)$$

The number of EM particles at shower maximum can be estimated from Eq. (2.14) or from the share of the total energy that did not go into muons:

$$N_{e,\max}^p = \frac{1}{E_c^{\text{EM}}} \left[ E_0 - \left( \frac{E_0}{E_c^{\pi}} \right)^{\beta} \right] \approx \frac{E_0}{E_c^{\text{EM}}}. \quad (2.19)$$

## Nuclear primaries

The above derivation for hadronic cascades works for a single primary proton of energy  $E_0$ . Expectations for heavier nuclei can be obtained using the *superposition model*, which gives a simplified view of the interaction of a nucleus in the atmosphere. A primary nucleus with mass number  $A$  and total energy  $E_0$  is considered as  $A$  individual nucleons of energy  $E_0/A$  interacting independently. This approximation is justified because the binding energy of  $\sim 5$  MeV per nucleon is much smaller than the typical interaction energies [30, p. 320]. The resulting air shower is treated as the sum of  $A$  separate proton air showers all starting at the same point. For the shower maximum, this leads to the

---

<sup>2</sup>The shower maximum is dominated by EM sub-showers produced in the interaction with the largest inelasticity which is usually the first interaction.

predictions

$$N_{e,\max}^A(E_0) = A N_{e,\max}^p(E_0/AE_c^{\text{EM}}) \approx N_{e,\max}^p(E_0) \quad (2.20)$$

$$X_{\max}^A(E_0) = X_{\max}^p(E_0/A) \approx X_{\max}^p - \lambda_{\text{EM}} \ln A. \quad (2.21)$$

The number of EM particles in the Heitler-Matthews model is therefore expected to be independent of mass and only a function of the total energy. The depth of shower maximum is, on the other hand, a mass-dependent observable. The number of muons,

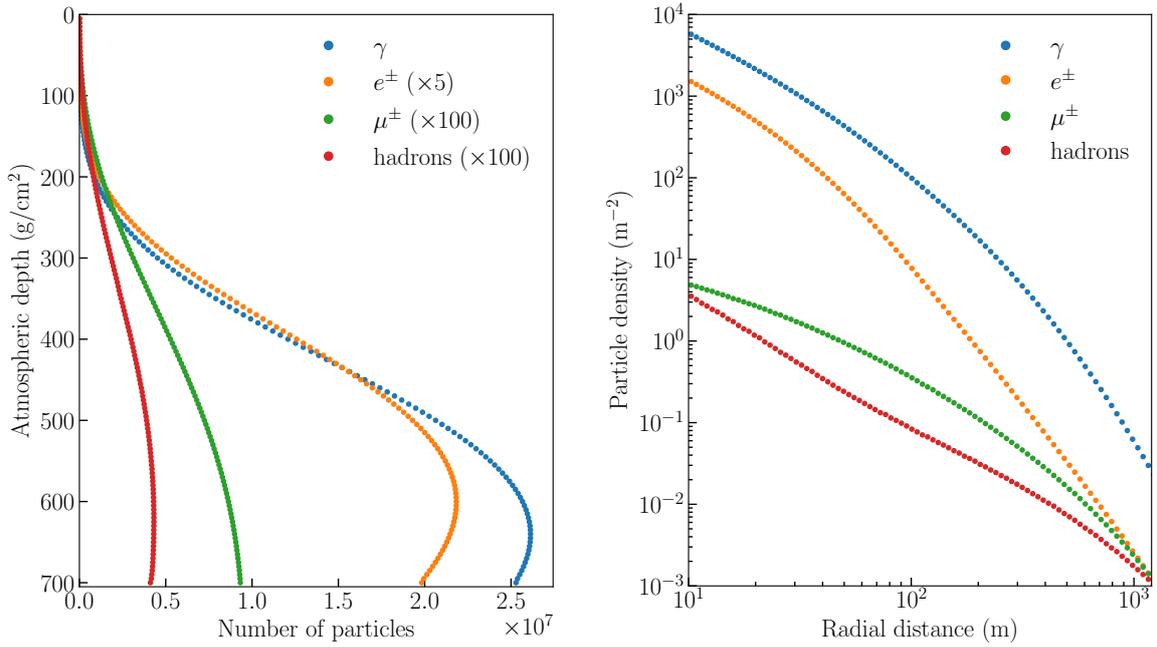
$$N_{\mu}^A(E_0) = A \left( \frac{E_0}{AE_c^{\pi}} \right)^{\beta} = A^{1-\beta} \left( \frac{E_0}{E_c^{\pi}} \right)^{\beta} = A^{1-\beta} N_{\mu}^p(E_0) \quad (2.22)$$

is also dependent on the primary mass, with a linear relation between  $\ln N_{\mu}$  and  $\ln A$ . Showers from heavier nuclei are thus expected to contain more muons, and reach their maximum higher in the atmosphere, compared to lighter nuclei with the same energy.

This model of course strongly simplifies the real development of an air shower. Quantities like the interaction length and the multiplicity of produced particles are to be seen as effective quantities in the model, while in reality they will depend on the energy of the interaction, among other things [62]. Other known phenomena that have been ignored can also have an important effect on the shower development, such as the fact that in an interaction of two hadrons a large fraction of the total energy is usually carried away by a single leading particle [63]. The impact of these properties is best evaluated through detailed simulations of air showers. Such simulations predict, for example, that  $\beta$  is in the range of 0.88–0.92, corresponding to effective multiplicities  $n_{\text{tot}}$  of 30–200 [64]. Although simple, the predictions of the Heitler-Matthews model regarding several air-shower observables and their dependence on primary energy and mass, have been shown to be in good accord with the qualitative behavior obtained from detailed simulations of air showers [45].

### 2.2.2 Particle distributions and air-shower observables

The Heitler-Matthews model of the previous section is useful to gain insight into the development of EAS, and the relations which may exist between the primary cosmic ray and different air-shower observables. A more realistic picture of EAS can be obtained through simulations of the full cascade. These provide us with detailed information on the distributions of different particles. Here, we introduce some of these basic features of air showers.



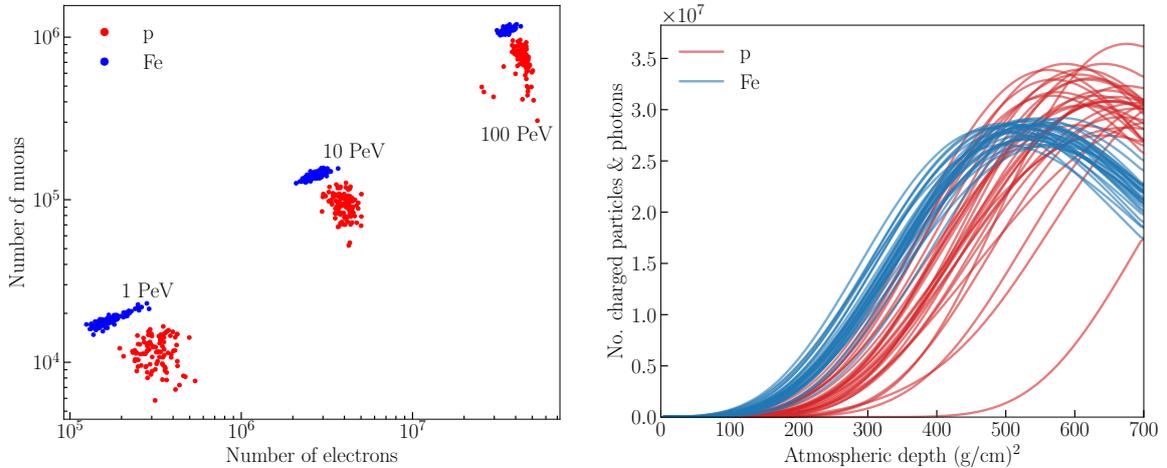
**Fig. 2.4:** Average longitudinal shower profile (left) and lateral particle distribution (right) for vertical proton-induced showers at 10 PeV. Showers were simulated with CORSIKA v73700 using Sibyll 2.1 and FLUKA as respectively the high- and low-energy hadronic interaction model (see Section 2.4). Conditions for IceTop have been used, i.e. a South Pole atmosphere (month April) and an observation level of 2835 m, corresponding to a depth of  $\approx 698$  g/cm<sup>2</sup>. The energy thresholds of the simulation were 0.05 GeV for muons and hadrons, 0.01 GeV for electrons and 0.002 GeV for photons.

The left panel of Fig. 2.4 shows the longitudinal development of EAS initiated by a 10 PeV proton. The shower is dominated by EM particles, which reach their maximum number at a particular depth, followed by a decline. This is in accordance with expectations, where most shower energy ends up in the EM component, is subdivided quickly until the critical energy is reached, and the remaining charged particles lose their energy through ionization. The number of muons, on the other hand, reaches a plateau without declining. This is because they rarely interact and lose energy a lot slower than electrons. On top, muon decay only becomes relevant for muons with very low energies. In addition to the EM and muonic component, a small number of hadrons will also reach the ground. We ignore neutrinos as they are relevant for neither development nor detection of an EAS.

The right panel of Fig. 2.4, shows the lateral distribution of the different components at observation level. The density of particles is highest around the shower axis, and drops quickly with radial distance. In hadronic interactions, secondary hadrons are produced with a typical average transverse momentum of  $\sim 350\text{--}400$  MeV, with small energy dependence [65, 66]. The EM sub-showers are produced by high-energy  $\pi_0$ , which are nearly parallel to the shower axis. Their lateral spread is dominated by multiple Coulomb scattering [67]. The majority of muons, on the other hand, are produced in the decay of low-energy pions, which have large angles relative to the shower axis. Their lateral distribution is therefore wider than that of EM particles.

The total number of electrons and muons at ground for EAS initiated by protons and iron nuclei is shown in the left panel of Fig. 2.5. These quantities are valuable to obtain information on both the primary energy and mass when observing air showers with particle-detector arrays (see next section). It can be seen that the abundance of both muons and electrons rises with primary energy, approximately as expected from the Heitler-Matthews model. The number of muons is significantly larger for iron showers, as expected from the superposition principle. The number of electrons in iron showers is lower, partially because more energy goes into the muonic component, and partially because they reach their maximum earlier in the atmosphere and thus a larger fraction of electrons will have been absorbed before reaching the ground. This is confirmed in the right panel of Fig. 2.5, which shows longitudinal profiles for a number of proton and iron showers.

Both panels furthermore illustrate that there are large shower-to-shower fluctuations in EAS observables. These fluctuations are largely a result of fluctuations in the first interaction of the cosmic ray [65, 68]. In simulations, large variations are observed in both the depth of the first interaction and the properties of the interaction itself, such as



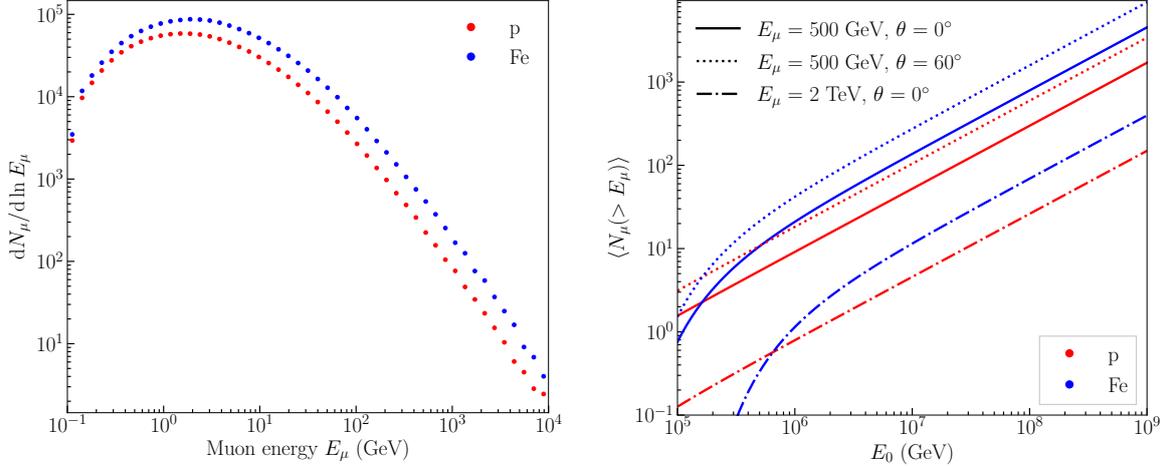
**Fig. 2.5:** Air-shower observables and their fluctuations in EAS simulated under the same conditions as Fig. 2.4. Left: Number of muons and electrons in 100 proton and iron showers at primary energies of 1 PeV, 10 PeV, and 100 PeV. Right: Evolution of the number of charged particles and photons throughout the atmosphere for 30 proton and iron showers at 10 PeV.

the types of particles produced, the multiplicities, the inelasticity, etc. Fluctuations in showers generated by heavy nuclei are smaller than those of light nuclei, a result of the shower developing approximately as a combination of sub-showers initiated by  $A$  different nucleons. Shower-to-shower fluctuations present a crucial difficulty in the determination of the properties of the primary cosmic ray, especially its mass, on an event-by-event basis.

### Muons in air showers

Of special importance for this thesis is the muonic component of EAS. Muons are mainly produced in the decay of low-energy charged pions and kaons, and the muon flux at ground level is dominated by muons with energies of several GeV and below. Mesons with high energies resulting from the first interactions in the atmosphere may also decay into muons, producing muons with energies above several hundred GeV, albeit with much lower probabilities. The average energy spectrum of muons in simulated 10 PeV proton and iron showers is shown in the left panel of Fig. 2.6. The number of muons indeed falls quickly with energy. At the lowest energies, the effect of muon decay becomes visible.

The majority of particles in an air shower gets stopped soon after reaching the ground, as they lose energy quickly propagating through dense media. The energy loss in matter for muons is, however, relatively low. As a result, high-energy muons can be observed



**Fig. 2.6:** Left: Average muon energy spectrum in 10 PeV proton and iron showers under the same conditions as Fig. 2.4. Right: Average number of high-energy muons expected as a function of primary energy, calculated with the Elbert formula (Eq. (2.23)) for different muon energy thresholds and zenith angles.

in deep underground detectors. As they result from the decay of high-energy mesons, they are highly collimated with respect to the shower axis and are commonly called a *muon bundle*. The minimum energy of muons that can reach a detector depends on the amount of matter that needs to be traversed (see Section 3.1.2). A formula based on an approximate solution of cascade equations, originally proposed by Elbert [69], gives the expected number of muons above a threshold energy in a shower initiated by a cosmic ray with energy  $E_0$ , mass number  $A$ , and zenith angle  $\theta$  as

$$\langle N_\mu(E_0, A, \theta, > E_\mu) \rangle \approx A \times \frac{0.0145 \text{ TeV}}{E_\mu \cos \theta} \left( \frac{E_0}{AE_\mu} \right)^{0.757} \left( 1 - \frac{AE_\mu}{E_0} \right)^{5.25}. \quad (2.23)$$

As one can see from the energy and mass dependence, the proportionality resembles that from the Heitler-Matthews model and the superposition approximation. The last factor of the equation adds a threshold factor for muon energies close to the energy per nucleon in the cosmic-ray nucleus. Some examples are shown in the right panel of Fig. 2.6. Note that the formula is approximate, and in reality the muon multiplicity depends, among other things, on the atmospheric conditions.

Minor contributions to the total muon content of an air shower also stem from electromagnetic muon pair production and from prompt muons. Prompt muons come from the decay of charm and bottom mesons, as well as unflavored vector mesons, which decay very rapidly. Because of this, they become dominant compared to the

conventional high-energy muon flux from the decays of light hadrons at muon energies around 1 PeV [70]. As their contribution is negligible for regular air-shower measurements, various generators used in simulation (Section 2.4) do not include charm production.

## 2.3 Measurement techniques

High-energy cosmic rays can be studied indirectly through measurements of the EAS they initiate in the atmosphere. A variety of detector types exists to achieve this. The common goal for most of them is to obtain information on the energy and mass of the primary cosmic ray, usually by measurements related to the observables introduced in the previous section: number of EM particles, number of muons, and depth of shower maximum. We introduce here the main classes of air-shower experiments, of which the particle-detector arrays are most relevant for this work.

### 2.3.1 Particle-detector arrays

A prevalent class of experiments are the surface arrays of particle detectors. The particle detectors are typically arranged in a regular pattern over a large surface area. Air showers are detected by looking for time coincidences of signals in neighboring detectors, as a result of the arrival of the disk of shower particles. The particle detectors are usually scintillation detectors or water Cherenkov tanks, which record signals produced by charged particles. The signals are therefore dominated by electrons, muons, and possibly photons, on the condition that the detectors are sufficiently thick so that the photons convert.

The arrival direction of a shower can be reconstructed from the arrival times of the particles in the detectors. The shower core, i.e. the point where the shower axis crosses the plane of detectors, is found by fitting the signal strength in the detector stations as a function of lateral distance with a *lateral distribution function* (LDF). The choice of LDF may be motivated by theoretical expectations for the distributions of different particle types, or its shape may be determined directly from data.

To determine the energy of the shower, typically, a parameter is defined based on the measured signal strengths in the detectors, for example the signal strength at a specific lateral distance from the shower axis. As the strength of the measured signals is related to the number of EM particles and muons in the shower, such a parameter will be related to the primary energy, as explained in the previous section.

Information on the mass of the primary with a particle detector array is usually obtained through observables related to the muon number. A measurement of the electron and muon numbers gives a handle on both the primary energy and mass. The signal in particle detectors is, however, often dominated by EM particles, and it is difficult to extract the muonic signal. A useful approach is to combine them with shielded particle detectors, which measure the muons after the EM part has been absorbed by some material covering the detector. Such an approach has been used in, for example, AGASA [71], CASA-MIA [72], EAS-TOP [73], GRAPES [74], KASCADE [75], KASCADE-Grande [76], GAMMA [77], Maked-ANI [78] and Yakutsk [79]. Further information about the mass composition based on signals in detector arrays can also be obtained from the steepness of the LDF, as done at Volcano Ranch [80]. Particles from showers which develop earlier in the atmosphere will be more spread out at the observation level. A larger number of muons compared to EM particles may also cause a flatter LDF. Therefore, showers from heavier primaries will have on average an LDF that is less steep than those initiated by lighter primaries. Also the rise time of the signal waveform recorded in the detectors has been used, e.g. at Haverah Park [81], as the time spread of the particles contains information on the production height.

Surface arrays are often used in a hybrid array. They can for example be used in coincidence with deep underground detectors, which can only be reached by high-energy muons in the showers. This provides a complementary way of deriving composition information. Examples are EAS-TOP with MACRO [82], the Baksan Underground Laboratory [83], SPASE-AMANDA [84, 85], and of course IceTop and IceCube (Chapters 3 and 4), the detectors used in this work. Tibet-AS $\gamma$  combines a scintillator array with emulsion chambers and so-called *burst detectors* [16]. Surface arrays have also been combined with fluorescence telescopes and Cherenkov detectors, as described below.

### 2.3.2 Other detector types

A variety of other techniques to study air showers exist that are not based on the direct detection of shower particles. We introduce them here briefly.

#### Fluorescence telescopes

Charged particles traversing the atmosphere can excite nitrogen molecules, which will in turn produce fluorescence light as they de-excite. For showers with energy  $\gtrsim 1 \times 10^{17}$  eV, fluorescence light can be used to directly measure the longitudinal development of air

showers, as the light yields are proportional to the energy deposited in the atmosphere [30, p. 334].

It is common to fit the measured longitudinal profiles with a function, such as the Gaisser-Hillas [86] function given by

$$N(X) = N_{\max} \left( \frac{X - X_0}{X_{\max} - X_0} \right)^{(X_{\max} - X_0)/\lambda} e^{-(X_{\max} - X)/\lambda}, \quad (2.24)$$

where  $N_{\max}$  and  $X_{\max}$  are the fitted shower size and slant depth at maximum and  $X_0$  and  $\lambda$  are treated as free shape parameters. A calorimetric measurement of the ionization energy deposited in the atmosphere is then given by the integral of the energy deposit profile, and is a good estimator of the primary cosmic-ray energy. In addition, the observation of the longitudinal profile provides a direct measurement of the shower maximum  $X_{\max}$ , which is sensitive to the mass composition.

Examples of air-shower experiments using the fluorescence method are the Fly's Eye detector [87] and its successor HiRes [88]. A drawback of the fluorescence technique is that it requires dark and clear nights, limiting their duty cycle to only 10–15%. Two experiments combining a surface array with fluorescence telescopes are the Pierre Auger Observatory [89] and the Telescope Array [90]. This provides the possibility of cross-calibrating the energy reconstruction of the two techniques, and for continued data-taking when the fluorescence measurements are not possible.

### Atmospheric Cherenkov light detectors

Charged particles moving through a medium at a velocity higher than the speed-of-light in that medium, determined by its refractive index  $n > 1$ , emit a large number of Cherenkov photons (see also Section 3.1). These photons can be used to detect EAS in a wide range of energies.

The emission of Cherenkov photons happens at an energy threshold and with an emission angle that depends on the refractive index of the medium, which is dependent on the density. Together with multiple Coulomb scattering of the photons, a characteristic lateral distribution of photons is observed at the ground. Simulations show that the density of photons at a distance around 120 m from the core and beyond is closely related to the energy of the shower. In addition, the slope of the LDF is found to depend on the depth of the shower maximum and thus on the cosmic-ray mass [30, p. 330].

Arrays of non-imaging Cherenkov detectors can be set up, similar to arrays of particle detectors, to sample the lateral distribution of Cherenkov light. This technique has been applied by e.g. AIROBICC [91], EAS-TOP [92], CASA-BLANCA [93], Tunka [94], and

Yakutsk [95]. Imaging atmospheric Cherenkov telescopes (IACT) can be used to detect showers with thresholds as low as 30 GeV, but their reach at high-energies is limited because of their relative small effective area. IACTs are mainly applied in the detection of showers induced by high-energy gamma rays. A disadvantage of the Cherenkov technique, similar as for fluorescence telescopes, is the requirement of clear, moonless nights.

### Radio detectors

Over the past two decades, important breakthroughs have been achieved toward the radio detection of EAS. Emission of radio signal from air showers has several contributions. The main contribution stems from the geomagnetic effect [96]. It is caused by the deflection of the abundant electrons and positrons in the shower disk by the geomagnetic field, producing a transverse current moving through the atmosphere. The current is time dependent because of the initial increase and later absorption of the particles, and through the influence of the changing density of the atmosphere on the charge separation. Another important contribution is the charge excess radiation, also called the *Askaryan effect* [97], which accounts for 10–15% of the signal. This radio emission is a result of the fact that there is a  $\sim 20\%$  excess of electrons with respect to positrons in an air shower, which is also time-dependent.

For wavelengths larger than the typical thickness of the shower disk, the emission is coherent and the expected electric field is proportional to the number of electrons  $N_e$  [30, p. 338]. The radiated power therefore scales quadratically with the number of electrons, and thus also with the primary energy. Various other observables, such as the shape of the lateral distribution of measured power and the slope of the observed frequency spectrum, depend on the depth of shower maximum, and are therefore sensitive to the mass of the primary. The feasibility of these techniques has been demonstrated by experiments such as Tunka-Rex [98], LOFAR [99], and AERA at the Pierre Auger Observatory [100]. In contrast to fluorescence and Cherenkov observations, the duty cycle for the observation of the EM shower component with radio observation is close to 100%. An enhancement of the IceTop array combining scintillators and radio antennas is planned to be deployed in the future [101, 102].

## 2.4 Air-shower simulations

Indirect studies of cosmic rays through the observation of EAS depend on a good understanding of the development of the cascade in the atmosphere. To relate air-shower observables to the properties of the primary cosmic ray, experiments have to rely on

simulations. Among all EAS measurements, the reconstruction of the cosmic-ray mass composition is the most difficult, as it relies on detailed comparisons of experimental observables, such as the number of muons, with air-shower simulations. These are, however, subject to uncertainties in the modelling of hadronic interactions at high-energies.

Below, we introduce the most commonly used tool for EAS simulation, CORSIKA. In Section 2.4.2, we introduce different hadronic interaction models that can be used to handle the interactions in the shower simulation. These models include large uncertainties, which lead to discrepancies between experimental data and simulation for air-shower measurements, discussed in Section 2.4.3.

### 2.4.1 CORSIKA

The CORSIKA program [103] is a detailed Monte Carlo (MC) program for the simulation of EAS that was originally developed for the KASCADE experiment. It is currently the most widely used simulation package for studying the evolution and properties of EAS in the atmosphere.

CORSIKA allows the simulation of EAS for shower energies up to several  $10^{20}$  eV, taking into account nuclei, hadrons, photons, electrons and muons. It tracks particles through the atmosphere, taking into account ionization energy loss and deflection by the Earth's magnetic field. For unstable particles, also the decays are simulated. Hadronic interactions of nuclei and hadrons with air nuclei are treated using various external hadronic interaction model codes, discussed in the next section. The simulation gives the type, energy, location, direction and arrival time of all secondary particles that reach a certain observation level. Various options can be selected, for example to study the longitudinal development of the shower, or to obtain more information about the production of muons. It is also possible to include the production of Cherenkov photons and neutrinos. Furthermore, the radio emission of air showers can be simulated using CoREAS [104].

By default, the Earth's atmosphere is treated as a flat disk, which is a good approximation for showers with zenith angles below  $\approx 65^\circ$ . To accurately simulate air showers at the location of a particular experiment, a realistic model of the density of the atmosphere is necessary. Standard CORSIKA atmospheres are built out of five layers in altitude. In the lowest four layers, the mass overburden or atmospheric depth has an exponential dependence on altitude

$$X_v(h) = a_i + b_i \exp^{-h/c_i}, \quad i = 1, \dots, 4. \quad (2.25)$$

In the fifth layer, it increases linearly with height

$$X_v(h) = a_5 - b_5 h / c_5. \quad (2.26)$$

The atmospheric model used in simulations for IceTop is discussed in Section 4.4.

As the number of particles to be tracked in an air shower rises quickly with the energy of the primary cosmic ray, so does the computing time of the simulation. To shorten the computing times for high-energy showers, the *thin sampling* option exists, based on an idea of Hillas for the MOCCA simulation program [105]. In the thinning mechanism, the energy of secondary particles resulting from an interaction is considered. When the sum of the energy of the secondary particles falls below a threshold energy, only one of them is followed, while the others are discarded. An appropriate weight is given to the surviving particle so that energy is conserved. When simulating a detector response, the simulated showers must be unthinned. As some amount of information is inevitably lost, both the thinning and unthinning must be done with caution.

A more detailed overview of the functionality and the underlying physics of CORSIKA is beyond the scope of this work and can be found in the CORSIKA documentation.

### 2.4.2 Hadronic interaction models

The development of an EAS is dominated by hadronic collisions with small momentum transfer. Because these interactions cannot be calculated with perturbative quantum chromodynamics, effective theories and phenomenology are used to predict their rates and the spectra of secondary particles produced in them. The software codes managing this are called *hadronic interaction models* or *generators*. Such generators are used by CORSIKA to simulate the hadronic interactions in a shower.

The treatment of hadronic interactions below and above a threshold energy, typically  $\sim 100$  GeV, is handled by separate models. Several commonly used models exist for both regimes. They differ in their descriptions of the interactions and as a result also in their predictions. This is especially the case for models describing high-energy hadronic interactions, which need to extrapolate to phase space that is not covered by measurements at accelerators such as the Large Hadron Collider (LHC). These high-energy models are therefore the main focus in this work. We introduce here briefly the most commonly used models. For a more exhaustive overview, see Ref. [9].

### Low-energy models

Low-energy hadronic interactions can be handled in CORSIKA by three different models. FLUKA [106, 107] is a package of Monte Carlo routines to propagate particles through matter. In combination with CORSIKA, only the part describing low-energy hadronic interactions is used. It is considered to be the state-of-the-art model, and is usually the preferred choice for detailed simulations. An alternative model which has a good description of the low-energy hadronic interactions is UrQMD [108, 109], designed specifically to treat low-energy hadron-nucleus interactions. The third option, GHEISHA [110], has been used to describe hadronic collisions up to 100 GeV in many experiments, but it is currently disfavored [9].

### High-energy models

Specialized generators exist to suit the needs for the simulation of EAS (and more). They need to be able to handle a variety of projectiles (protons, nuclei, charged pions and kaons) and targets (nitrogen, oxygen, argon). An important quality of the high-energy hadronic interaction models is their predictive power, as they need to extrapolate to energies that are outside the reach of human-made accelerators. As accelerator experiments usually measure in the mid-rapidity region<sup>3</sup>, EAS generators also need to extrapolate to cover more forward interaction regions, which are crucial for the shower development. The common models are all based on Gribov-Regge field theory, but have different ways of handling different aspects relevant to the EAS development, such as the treatment of diffraction or the description of nuclear collisions. The models are tuned to measurements from various accelerator experiments. The latest iteration of models has been tuned to data obtained by LHC-based experiments, and are sometimes referred to as *post-LHC models*.

The most commonly used post-LHC hadronic interaction models are Sibyll 2.3d [111], QGSJet-II.04 [112, 113], and EPOS-LHC [114]. QGSJet and Sibyll are focused on air-shower simulation; they only implement physics relevant for shower development and have a limited set of parameters. They focus on a small dataset for tuning, mainly from  $pp$  collisions at various energies and only a subset of final-state particles. The EPOS generator has a broader focus, including heavy-ion collisions. Correspondingly, it has

<sup>3</sup>The mid-rapidity region is often defined as  $|\eta| < 0.5$ , where  $\eta$  is the pseudorapidity. For a particle with momentum  $|p|$  with longitudinal component  $p_{\parallel}$ , i.e. the momentum along the beam axis, it is equal to

$$\eta = \frac{1}{2} \ln \frac{|p| + p_{\parallel}}{|p| - p_{\parallel}} = -\ln \tan \frac{\theta}{2},$$

where  $\theta$  is the polar angle of the particle with respect to the beam axis.

more parameters and tunes to a larger dataset, including data from  $e^+e^-$  collisions. We also mention the DPMJet-III model [115, 57], which has a more general high-energy physics scope, but can also be used in EAS simulation. Some older versions of models, such as Sibyll 2.1 [116], are still commonly used in analyses performed by experiments, because of the large computational cost of producing new air-shower datasets.

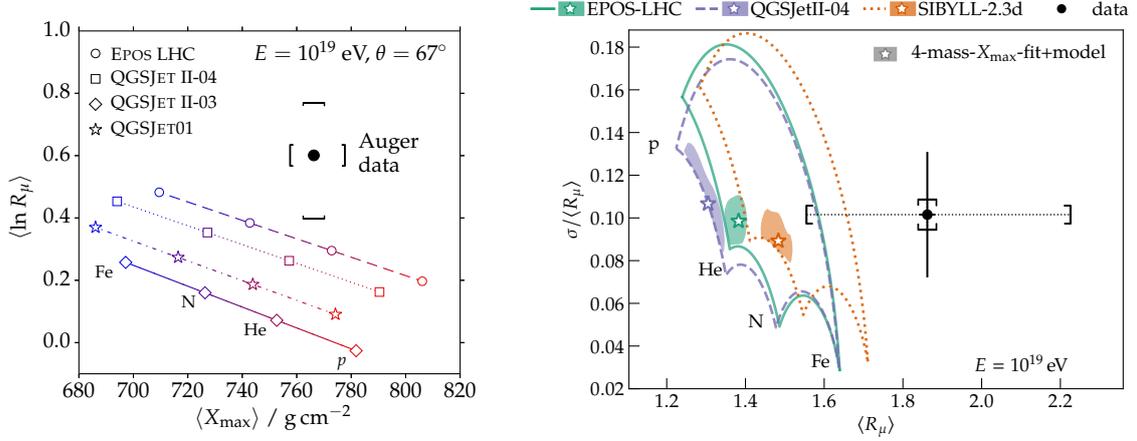
A detailed overview of the theory behind the models and how they differ is beyond the scope of this work. It is, however, important to mention that their predictions for several air-shower observables may disagree, as a result of uncertainties in the extrapolations. In Chapters 6 and 7, we discuss the predictions of observables relevant to IceTop and IceCube for the models Sibyll 2.1, QGSJet-II.04, and EPOS-LHC.

### 2.4.3 Simulation uncertainties and the Muon Puzzle

Cosmic-ray experiments rely on EAS simulations to derive the mass composition from experimental data. The two main air-shower features used for this are the depth of shower maximum  $X_{\max}$  and the number of muons  $N_\mu$  produced in the shower. The results for the mass composition derived from muon measurements show very large differences between different experiments [45]. Most of this uncertainty has been shown to be from theoretical nature, i.e. from the uncertainty in the description of the air-shower development [9]. While modern high-energy hadronic interaction models are tuned to LHC data, they still need to extrapolate towards higher center-of-mass energies, forward rapidities, and towards hadron-nuclear collision systems. This causes disagreement between the predictions of EAS observables using different models. Furthermore, significant discrepancies between the observed and predicted  $N_\mu$  have been observed, as will be discussed below. This is not the case for  $X_{\max}$ , where the measurements are bracketed by simulations of proton and iron showers over a wide range of shower energies.

#### The Muon Puzzle

Discrepancies between predictions and observations of the number of muons in EAS were first reported by the HiRes-MIA [117] and NEVOD-DECOR [118] experiments. A detailed comparison of measurements with predictions by the Pierre Auger Observatory established in 2015 in a nearly model-independent way that the post-LHC models consistently predict a lower production of muons with energies in the GeV range in air showers than is observed [119]. Their result is shown in Fig. 2.7, together with a recent update of the measurement including shower-to-shower fluctuations of  $N_\mu$  [120]. While there are also uncertainties in the predictions of  $X_{\max}$ , there is no such striking disagreement with the



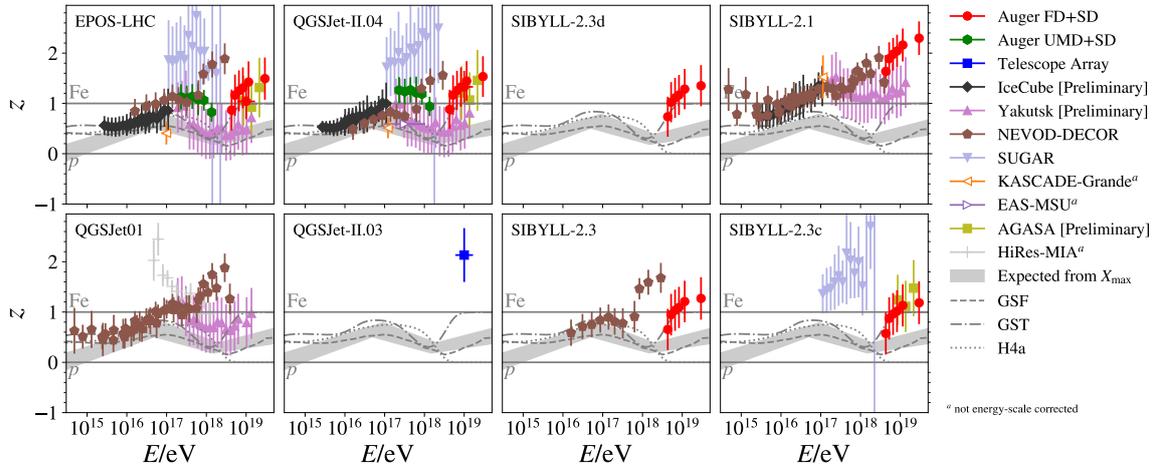
**Fig. 2.7:** Discrepancy in the normalized muon number  $R_\mu$  in inclined showers with shower energy of  $10^{19}$  eV as measured by the Pierre Auger Observatory. Left: Average logarithmic muon number as a function of the average shower depth in data compared to predictions of various models. Figure taken from Ref. [119]. Right: Fluctuations and average number of muons in data compared to post-LHC models. Contour lines contain all possible composition mixtures of the four components  $p$ , He, N, Fe. The shaded areas give the regions allowed by the Auger  $X_{\max}$  measurement [121]. Figure taken from Ref. [120].

observed values. This measurement was followed up by muon measurements from many leading air-shower experiments, under various different conditions. These results are, however, mixed, with some experiments also reporting evidence for a deficit of muons in simulations, while others did not seem to show a clear discrepancy.

The large amount of muon data led to a meta-analysis performed by the Working group of Hadronic Interactions and Shower Physics (WHISP) [122, 123], formed by members of eight leading air-shower collaborations. An abstract muon scale was defined which allowed to compare muon measurements of the different experiments, regardless of differences in zenith angle, muon energy threshold, etc. This was named the  $z$ -scale, defined as

$$z = \frac{\ln \langle N_\mu \rangle - \ln \langle N_\mu \rangle_{\text{p}}}{\ln \langle N_\mu \rangle_{\text{Fe}} - \ln \langle N_\mu \rangle_{\text{p}}}, \quad (2.27)$$

comparing the muon number  $\langle N_\mu \rangle$  or a quantity proportional to it in data, to the corresponding values  $\langle N_\mu \rangle_{\text{p}}$  and  $\langle N_\mu \rangle_{\text{Fe}}$  from proton and iron showers simulated using a certain hadronic interaction model. It should have a value  $0 < z < 1$ , as pure proton and iron are the expected extreme cases for the cosmic-ray mass composition. After cross-calibrating the energy scale of the different experiments, a fairly coherent picture in  $z$  is obtained, as shown in Fig. 2.8. Many measurements have  $z$ -values above 1, indicating



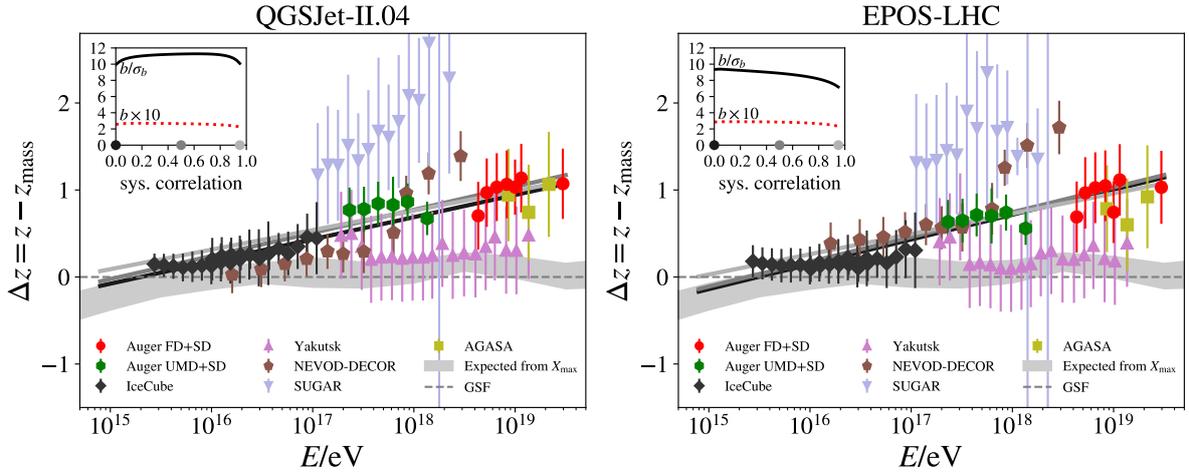
**Fig. 2.8:** Compilation of muon measurements by various experiments, expressed in the  $z$ -scale (Eq. (2.27)), after cross calibration of the energy scales. Also shown are the  $z$ -values based on various composition models introduced in Section 1.5, and those corresponding to  $X_{\max}$  measurements. Figure taken from Ref. [123].

a muon excess with respect to simulations. Also shown are the expected  $z$ -values obtained from composition measurements based on the shower maximum and different composition models, as described in Section 1.5. Subtracting the expected values corresponding to a model like GSF, which gives a world-average of composition data, from the values obtained from muons, should give a result  $\Delta z = z - z_{\max}$  consistent with zero. However, as can be seen in Fig. 2.9, this  $\Delta z$  shows a logarithmic increase with the shower energy  $E$  with a non-zero slope of  $8\sigma$  significance. The deviation starts around 40 PeV. This discrepancy is referred to as *the Muon Puzzle*, because authors of the hadronic interaction models have not been able to resolve it by tuning the model parameters.

The Muon Puzzle refers specifically to the number or density of GeV muons, produced in the late stages of a hadronic cascade. There are of course also other properties of muons in EAS that can be analyzed, and discrepancies have been reported in various other measurements. A summary is given in Ref. [9]. In Chapters 6 and 7, we will present measurements indicating an inconsistency in the description of GeV and TeV muons.

### Possible solutions for the Muon Puzzle

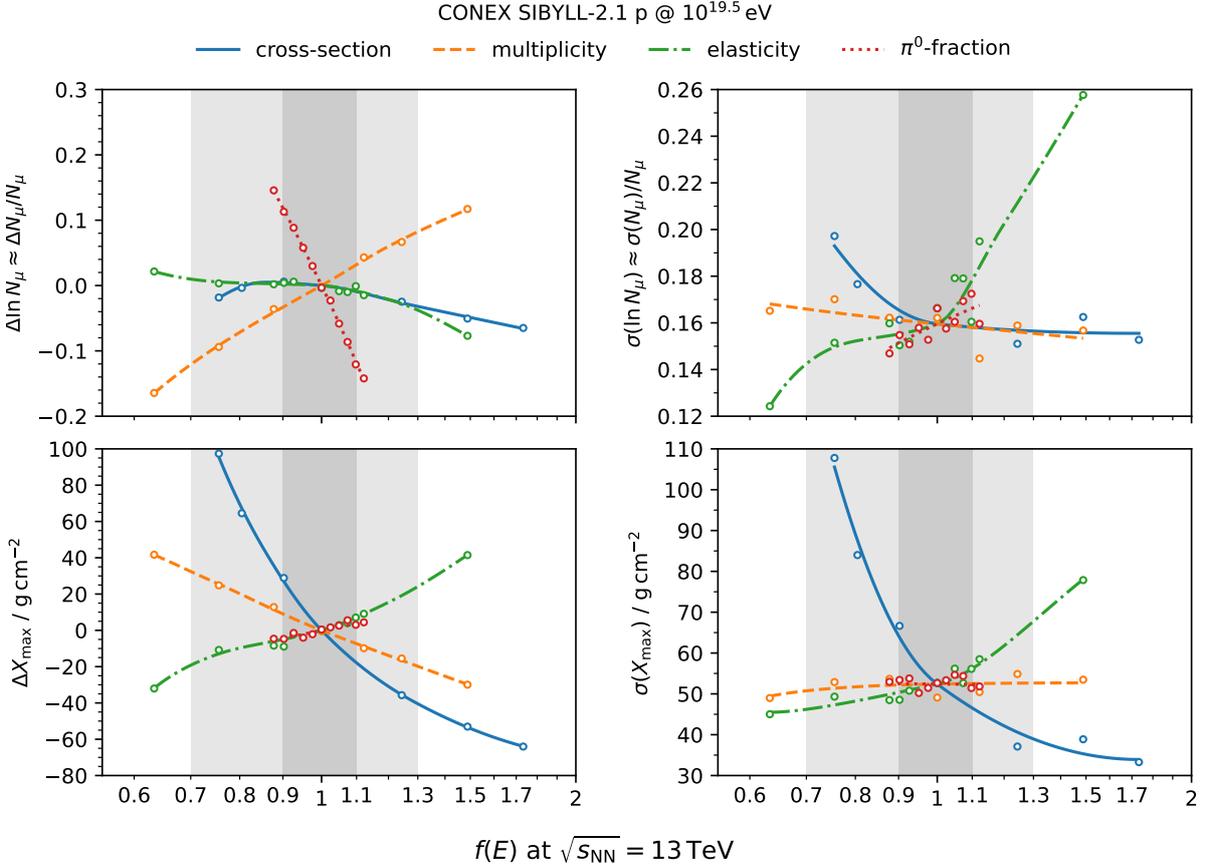
The origin of the Muon Puzzle is at this time not known. There are, however, various results indicating what a solution could look like [9]. The WHISP result discussed above showed that the deviation in the GeV muon number starts around 40 PeV, implying that the late shower stages where the hadrons have a lower energy are sufficiently well



**Fig. 2.9:** Muon data from Fig. 2.8 after subtracting  $z_{\text{mass}}$ , the expected variation in  $z$  due to changes in the primary cosmic-ray composition given by the GSF model. The lines represent fits with different assumptions for the correlation of the systematic errors. A positive slope is found with large significance. Figure taken from Ref. [123].

described by simulations. The origin is therefore likely to be found in the earlier stages of the shower. The relative deficit in muons increases approximately logarithmically with the shower energy, similarly to the number of stages in the hadronic cascade. This points towards a compounding effect, where a small change in the properties of the hadronic interactions results after several interactions in a larger effect in  $N_\mu$ .

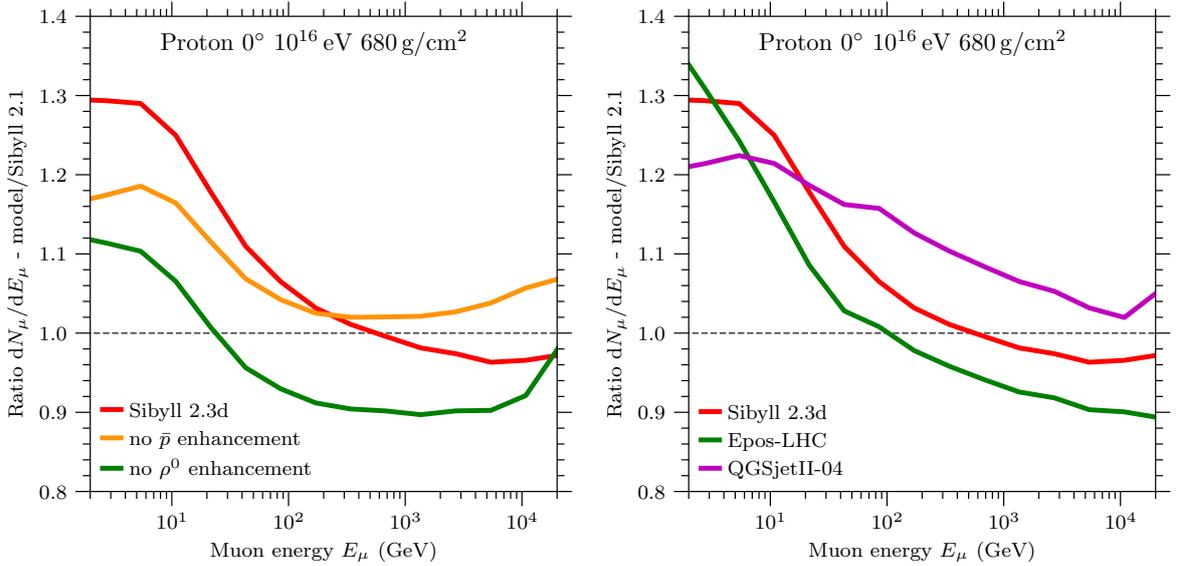
Four important parameters of hadronic interactions are the inelastic cross section, the multiplicity of secondary particles, the elasticity or energy fraction carried by the most energetic particle, and the energy ratio going into photons from decay of e.g.  $\pi^0$  and  $\eta$  versus into long-lived hadrons. In Ref. [124], these parameters were changed in an ad-hoc way during simulations to test the effect on various EAS observables. The parameters were modified with a factor  $f(E)$ , which grows logarithmically with the energy of the colliding hadron. Below 1 PeV,  $f$  was chosen to be 1, because at these energies the generators are fairly constrained by accelerator data. The results for the mean and standard deviation of the muon number  $N_\mu$  and depth of shower maximum  $X_{\text{max}}$  are shown in Fig. 2.10 for a  $10^{19.5}$  eV proton shower as a function of the size of the modification at a center-of-mass energy of 13 TeV, corresponding to the LHC. The easiest way to increase  $N_\mu$ , without having a large impact on the fluctuations on  $N_\mu$  or  $X_{\text{max}}$  (see Fig. 2.7), is by decreasing the  $\pi^0$  fraction, or equivalently, the amount of energy going into the EM cascade in each hadronic interaction. This corresponds to increasing the factor  $2/3$  in the Heitler-Matthews model. Changes to the cross section and the elasticity have a negligible effect. In Ref. [125], it was furthermore shown that



**Fig. 2.10:** Impact of changes to properties of hadronic interactions in Sibyll 2.1 on the mean and standard deviation of the logarithm of the muon number  $N_\mu$  and the depth of shower maximum  $X_{\text{max}}$  for a  $10^{19.5}$  eV proton shower. Changes are shown as a function of the modification factor at a center-of-mass energy of 13 TeV in the nucleon-nucleon system. Shaded regions indicate modifications of  $\pm 10\%$  and  $\pm 30\%$ . Figure taken from Ref. [9], where it was adapted from the original in Ref. [124]

the discrepancy in the left panel of Fig. 2.7 can be solved by modifying the ratio of energy going into neutral pions, but cannot be solved by changing the multiplicity. Note that various exotic explanations of the Muon Puzzle have been proposed as well, usually through an extreme change to the physics of the highly-energetic first interaction. Such models tend to reduce the fluctuations in  $N_\mu$  by a lot, in disagreement with the result from Fig. 2.7 (right).

Various ways in which the distribution of energy between the hadronic and EM shower components can be modified are already known. Ref. [111] shows for example how the muon yield in Sibyll 2.3d changes compared to Sibyll 2.1, in part by an enhancement of  $\rho^0$  and baryon production. The production of the neutral vector mesons  $\rho^0$  competes

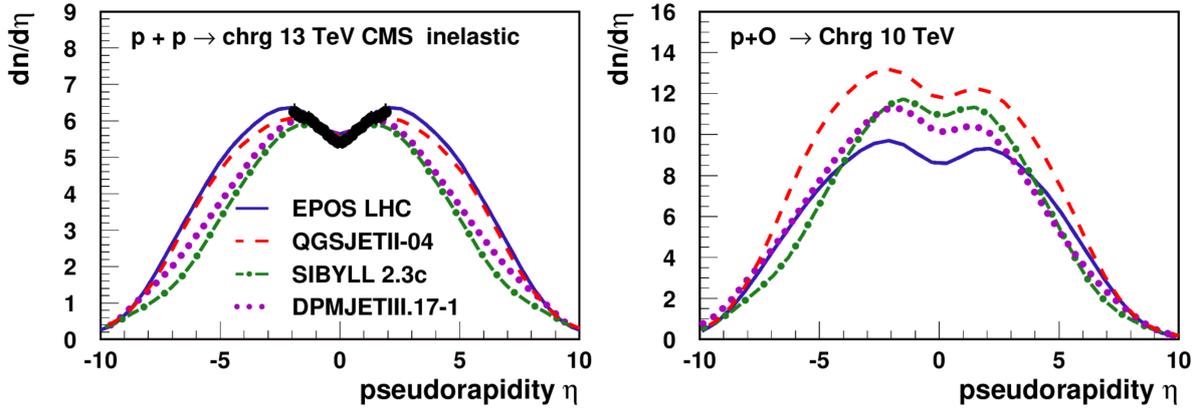


**Fig. 2.11:** Ratio of the muon energy spectrum between different interaction models and Sibyll 2.1. Left: Different versions of Sibyll 2.3d, where the extensions for enhanced  $\rho^0$  and baryon production have been switched off. Right: Comparison of different post-LHC interaction models. Figures kindly provided by the author of Ref. [111].

with the production of  $\pi^0$ . In contrast to  $\pi^0$ , it does not decay to photons, but to two charged pions  $\rho^0 \rightarrow \pi^+\pi^-$ . An increase in baryon-pair production will result in more re-interaction producing more pions and kaons and hence ultimately more muons. The influence of these two effects on the muon spectrum in Sibyll 2.3d is shown in Fig. 2.11.

As shown before in Fig. 2.8, the Muon Puzzle seems to be present in different Sibyll 2.3 versions and so additional effects are probably needed. The ALICE collaboration has recently discovered on top of a baryon enhancement also a universal strangeness enhancement in  $pp$ ,  $p$ -Pb, and Pb-Pb collisions [126], i.e. an increase in the yields of strange hadrons compared to pions as a function of the multiplicity of the interaction. This enhancement is currently regarded as the potential key to the Muon Puzzle. As the observation by ALICE was in the mid-rapidity region, its presence has to be demonstrated also in the forward region before the mechanism can be accepted as a solution. Such a measurement can be performed by LHCb [127].

There is a variety of other measurements that could be performed at the LHC and fixed-target experiments which could crucially improve the simulation of EAS. They are mainly related to average properties of light hadron production at low momentum transfer, such as the inelastic and diffractive cross sections, the charged particle spectra, the forward energy flow, etc. Fixed-target experiments such as NA61/SHINE [128] and



**Fig. 2.12:** Pseudorapidity distribution of charged hadrons predicted by different hadronic interaction models. Left: in  $pp$  at 13 TeV, together with data from CMS. Right: In  $p$ -O at 10 TeV. Figure taken from Ref. [9].

LHCb-SMOG [129] can furthermore improve our knowledge of hadronic interactions in later stages of an EAS, using a variety of target nuclei to study nuclear effects. Very promising is the prospect of a run with oxygen beams at the LHC to study  $p$ -O and O-O collisions, as there is currently considerable theoretical uncertainty in the extrapolation of  $pp$  to  $p$ -O, as shown in Fig. 2.12. A comprehensive overview of measurements at colliders and their connection to air-shower physics and the Muon Puzzle is given in Ref. [9].

The existence of the Muon Puzzle was established for muons with energies in the GeV range. Various explanations to increase the muon yield at this energy will also have an impact on the production of muons at higher energies. As discussed in the following chapters, the IceCube Neutrino Observatory has the unique capacity to measure muons in the GeV range at the surface with IceTop, and in the TeV range with IceCube. As explained in Ref. [111] and demonstrated in Fig. 2.11, the combination of measurements in both energy ranges has the potential to discriminate among interaction models and provide constraints to mechanisms of muon production in EAS. This motivates the work presented in Chapters 6 and 7.



# Chapter 3

## The IceCube Neutrino Observatory

The IceCube Neutrino Observatory is a hybrid particle detector located at the geographical South Pole, near the Amundsen–Scott South Pole Station. Its primary scientific goal is the observation of astrophysical neutrinos with a cubic kilometer photomultiplier array embedded in the glacial ice. Other science objectives of the in-ice detector include searches for exotic particles, studies of the oscillations of atmospheric neutrinos, and the detection of neutrino bursts from Galactic supernovae. At the surface it is supplemented by a square kilometer particle detector array called IceTop. The IceTop detector observes air showers from cosmic rays in the energy range between 100 TeV and 1 EeV, and can be used in coincidence with the in-ice detector to study high-energy muons in the shower. Construction and commissioning of IceTop and IceCube finished in 2011. In this chapter, we introduce the design of these detectors and how they are used for detecting various particles, with a focus on their application as air-shower detectors.

In Section 3.1, we introduce the mechanisms that allow the detection of charged particles propagating through matter. This detection is done using photomultiplier detectors in both the surface and in-ice array; these are introduced in Section 3.2. The data acquisition (DAQ) system is introduced in Section 3.3. Section 3.4 describes how the detectors are calibrated. Finally, a brief discussion on the simulation of the detector response is given in Section 3.5.

### 3.1 Detection principle

Charged particles travelling at relativistic velocities through a medium can, under several conditions, produce Cherenkov light. While propagating, they will also lose energy through processes like ionization and bremsstrahlung. Energetic particles resulting from such processes may themselves result in an important contribution of Cherenkov

light [130]. Both IceTop and IceCube utilize this light production to detect particles moving through ice.

### 3.1.1 Cherenkov effect

The Cherenkov effect refers to the emission of a bluish light by charged particles moving through a medium at velocities larger than the phase velocity of light in that medium. A full theoretical treatment and an overview of the experimental discovery and early applications are given in Ref. [131]. In this section, we give a brief introduction.

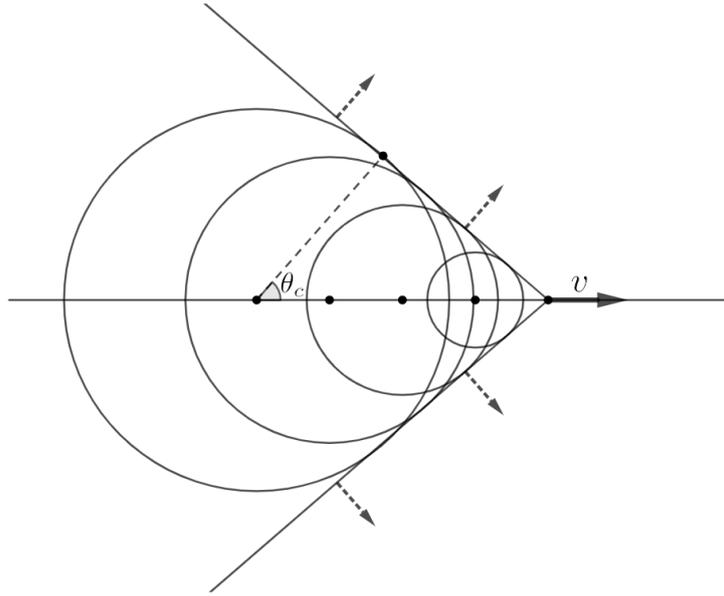
Consider a dielectric medium with refractive index  $n$ . The phase velocity of light in this medium is equal to  $c/n$ , where  $c$  is the speed of light in vacuum. When a charged particle moves through this medium, the charge of the particle distorts the atoms in the medium so that it becomes polarized. A dipole field temporarily appears along a segment of the track of the particle, leading to a brief electromagnetic pulse. In general, the resultant field intensity at a distant point will be zero. However, if the velocity of the particle is higher than the phase velocity of light in the medium, it will outrun any electromagnetic waves it emits. The waves from all portions of the track will then combine to form a wavefront of coherent radiation. Using the Huygens principle, one finds that this resulting wavefront propagates at a particular angle  $\theta_c$  with respect to the track of the particle. As demonstrated in Fig. 3.1, it is given by

$$\cos \theta_c = \frac{c/n}{v} = \frac{1}{\beta n}, \quad (3.1)$$

where  $v$  is the velocity of the particle and  $\beta = v/c$ . The Cherenkov radiation is thus emitted in a typical cone shape. Below the threshold velocity  $\beta = 1/n$ , no radiation takes place, and the possible wavelengths  $\lambda$  are restricted by  $n(\lambda) > 1/\beta$ . The spectrum of the produced photons is given by the Frank-Tamm formula [132, 133] and follows  $1/\lambda^2$  when ignoring the wavelength dependence of  $n$ . The radiation occurs mainly in the visible and near-visible regions of the spectrum.

For a wavelength of 400 nm, the refractive index of ice is 1.32 [134]. The threshold velocity is therefore  $0.76c$ . For a relativistic particle with a velocity close to  $c$ , the Cherenkov angle is given by  $\theta_c \approx 41^\circ$ .

The energy loss of a particle due to Cherenkov radiation is of the order of only several keV  $\text{cm}^{-1}$  [131], which is negligible compared to the energy loss processes introduced in the next section.



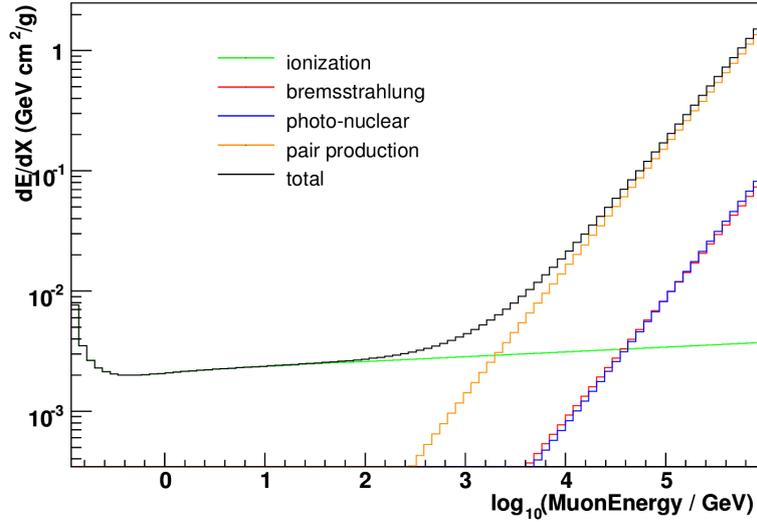
**Fig. 3.1:** Schematic view of the Cherenkov effect. The direction of the wavefront propagation is determined by the velocity of the particle and the phase velocity of light in the medium.

### 3.1.2 Energy loss in matter

Particles propagating through matter undergo various interactions, causing them to lose energy. Most important for IceCube is the case of charged leptons propagating through ice.

Charged particles moving through a medium interact with its electrons and as a result ionize the medium, losing energy. A charged particle can also undergo so-called *radiative losses*, i.e. bremsstrahlung, pair production, or photonuclear interaction. In bremsstrahlung, a charged particle loses energy in the field of a nucleus, thereby emitting a photon. Charged particles can also produce a pair of leptons in the field of a nucleus, most importantly  $e^+e^-$  pairs. The photonuclear interaction is the inelastic interaction of a lepton with a nucleus, leading to the breakup of the nucleus. The resulting particles may furthermore initiate a cascade in the medium. Secondary particles that are sufficiently energetic will produce Cherenkov light and contribute to the total light yield of a relativistic charged particle moving through a medium. While ionization energy loss is almost continuous, the radiative losses exhibit a stochastic behavior with large fluctuations [133].

Losses due to ionization vary only slowly with the energy of the particle. The average energy loss due to radiative processes, on the other hand, shows a clear increase with



**Fig. 3.2:** Muon energy loss  $dE/dX$  from ionization and radiative loss processes in ice as a function of muon energy. Figure taken from Ref. [136].

energy. The contributions to the average energy loss by the different processes for muons in ice is shown in Fig. 3.2. Ionization is the dominant contribution below  $\sim 1$  TeV. Muons with energies slightly below 1 GeV, around the broad minimum of the energy loss, are often called *minimum ionizing*. Above  $\sim 1$  TeV, radiative energy loss becomes dominant, with the largest contribution from pair production. For electrons, the situation is different. Above 100 MeV the energy loss is dominated by bremsstrahlung, and is considerably higher than for muons. This is in part due to an inverse mass-squared factor in the bremsstrahlung cross-section, suppressing bremsstrahlung for muons compared to electrons. For a detailed overview of the energy loss processes of charged leptons in matter and their cross-sections, see Ref. [135].

The average energy loss of muons is often parametrized as

$$\frac{dE}{dX} = -a - bE, \quad (3.2)$$

where  $E$  is the muon energy,  $b = b_{\text{br}} + b_{\text{pair}} + b_{\text{ph}}$  accounts for the three radiation processes, and  $a$  accounts for ionization losses [30, p. 165]. Values for muons in ice are given in Ref. [135] as  $a = 0.249$  GeV/mwe and  $b = 0.422 \times 10^{-3}$  /mwe<sup>1</sup>.

<sup>1</sup>1 mwe = 1 meter water equivalent. Expresses an amount of traversed matter in the equivalent amount of water. As the density of water is approximately 1 g/cm<sup>3</sup>, one meter of water corresponds to a depth of 100 g/cm<sup>2</sup>.

The solution for Eq. (3.2) is

$$\langle E(X) \rangle = (E_0 + \frac{a}{b})e^{-bX} - \frac{a}{b}, \quad (3.3)$$

and is to be interpreted as the mean energy of a beam of muons of initial energy  $E_0$  after penetrating a depth  $X$ . To estimate the minimum energy that a muon requires at the surface to reach an underground detector at a slant depth  $X$ , the left-hand side of Eq. (3.3) is set to zero, which gives

$$E_0^{\min} = \frac{a}{b} (e^{bX} - 1). \quad (3.4)$$

A quick estimate for IceCube, located under 1450 m of ice which has a density of about  $0.92 \text{ g cm}^{-3}$ , gives an energy threshold of 446 GeV. In reality, the range of muons of a certain energy will of course fluctuate due to fluctuations in the energy loss interactions. The best way to study muon propagation is therefore through Monte Carlo simulations.

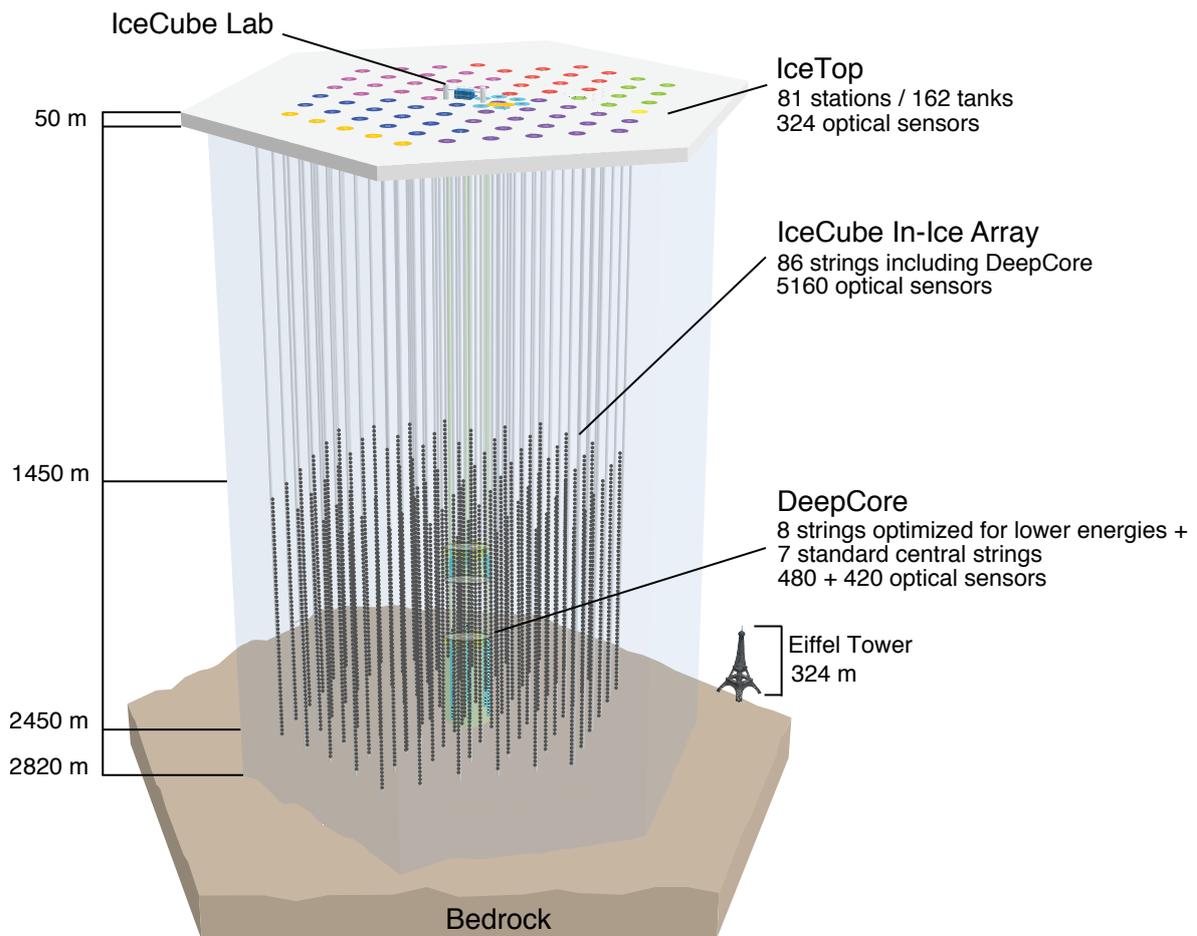
## 3.2 The detector

An overview of the IceCube Neutrino Observatory is shown in Fig. 3.3. It consists of the in-ice array IceCube and the surface array IceTop. At the center of the array on the surface sits the central operations building for the experiment, the IceCube Laboratory (ICL). This is where all cables arrive and where the computers for DAQ and online processing are located.

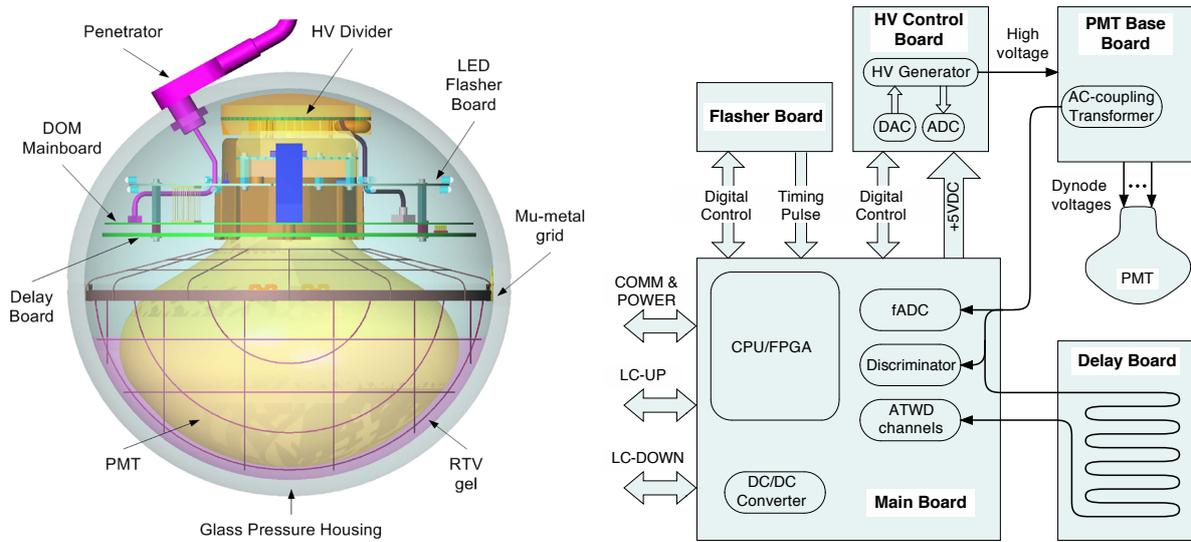
IceTop and IceCube both work by detecting the light produced by charged particles moving through ice. For IceCube, the medium is the deep glacial ice. For IceTop, it is the ice inside the IceTop tanks. In both cases, the light is detected using Digital Optical Modules (DOMs). Below, we introduce the design and functionality of the DOMs and how they are used in the two detectors.

### 3.2.1 Digital optical modules

The DOMs are the fundamental detector units allowing the detection of photons [137]. The first steps of data acquisition are also performed inside the DOM by the onboard electronics. Fig. 3.4 shows a schematic overview of a DOM and its main components. The insides of the DOM are housed in a glass sphere with a diameter of 13" and a thickness of 0.5", offering protection against the 250 bar pressure at the bottom of IceCube. The



**Fig. 3.3:** Overview of the IceCube Neutrino Observatory with the in-ice array, its sub-array DeepCore, and the surface air-shower array IceTop. Figure taken from Ref. [137].



**Fig. 3.4:** The IceCube Digital Optical Module (DOM). Left: Schematic view of the DOM components. Right: Functional connections of the electronic components. Figure taken from Ref. [138].

housing is made from borosilicate glass with good optical transmission and very low amounts of radioactive trace elements, which could cause noise.

Each DOM contains a downward-facing 10"-diameter Hamamatsu photomultiplier tube (PMT). The PMT is optimal for the detection of photons with wavelengths in the 300–650 nm range, with a peak quantum efficiency around 25% near 390 nm, which gives the average number of electrons ejected per photon hitting the photocathode. A full characterization of the PMT is given in Ref. [139].

Through a stage of dynodes, the initial photoelectron gets multiplied with a large amplification factor. The resulting signal waveforms can cover a large range of amplitudes, depending on how many photons reach the DOM simultaneously. The gain at which the DOMs operate is  $10^7$  for IceCube. For IceTop, DOMs run at two different gains of  $10^5$  and  $5 \times 10^6$ , as described in Section 3.2.3.

The PMT is secured in the DOM with a silicone gel surrounding the photocathode area, providing mechanical support to both the PMT and the assembly of circuit boards. The gel is optically clear and provides good coupling with the glass, with less than 0.1% reflection as light passes from the glass into the gel and then into the PMT. The PMT bulb is surrounded by a mu-metal grid to reduce effects of the geomagnetic field on the collection efficiency.

The PMT signals are compared to a discriminator threshold by the Mainboard electronics. When the discriminator is crossed, high-speed waveform capture and digitization

begins. The signals are sent to an Analog Transient Waveform Digitizer (ATWD) chip, sampling the input voltage at 300 Msps for a total of 128 bins, corresponding to a sampling time of about 427 ns. The signal is read from the Delay board, which delays the signal by 75 ns, so that also the part of the waveform before the discriminator crossing can be read out. The ATWD has three different channels which are used for data taking, with different gains to cover the large possible range of voltages without loss of information due to saturation. Every DOM is equipped with two ATWDs to reduce deadtime. In parallel, a fast-Analog-to-Digital Converter (fADC) samples continuously at a rate of 40 Msps and its data is recorded for 6.4  $\mu$ s after the discriminator crossing. This allows to capture photons that have been produced far away and which may arrive over a broad time interval due to scattering in the ice.

Every digitizer launch results in a *hit* record. The DOM stores hit information for about one second before sending the data up to the DAQ system in the ICL. The information which is sent for a hit depends on a coincidence condition with nearby DOMs, which is different for IceCube and IceTop, as described in Section 3.3. In case of an isolated signal (no coincidence), only a timestamp and a brief summary of the measured charge are included. In case the coincidence condition is met, the full waveform is compressed and included in the hit record.

A DOM also contains an LED flasher board, which is used to provide a light source for a variety of calibration purposes [140].

DOMs are in contact with the outside world with three wire pairs which enter through a penetrator at the top. One wire pair carries power and a digital communication stream to the ICL. The other wire pairs are connected with neighboring DOMs to check for coincident hits.

An important part of the design of the DOMs was their long-term reliability, as they are stuck in the ice with no way of fixing or replacing them. This led to a very low failure rate, currently estimated at only  $4.1 \pm 1.2$  out of more than 5000 DOMs per year.

For more details about the DOM and its role in the detector and DAQ system, see Ref. [137].

### 3.2.2 The in-ice array: IceCube

The ice cap at the South Pole is about three kilometers thick and provides a detector medium with excellent optical qualities. A combination of the long optical attenuation length of South Pole ice and large-area PMTs allow the instrumentation of a large volume of ice with a rather sparse spacing of detectors.

The IceCube in-ice array [137] consists of 5160 DOMs deployed between 1450 m and 2450 m below the surface on 86 vertical strings, with 60 DOMs per string. The primary array consists of 78 standard strings with a vertical separation of the DOMs of 17 m. The strings are positioned on a triangular grid with a horizontal spacing of about 125 m. All strings together form a hexagonal footprint, as shown in Fig. 3.3, instrumenting a volume of about one cubic kilometer of ice. The DOMs are attached to the strings via a harness of steel rope, with the main cable bending around the DOM, ensuring that the DOM axis is vertically aligned with the string, with the PMT facing downward. The design of IceCube was chosen in order to accomplish the primary scientific goal of detecting astrophysical neutrinos with TeV to PeV energies.

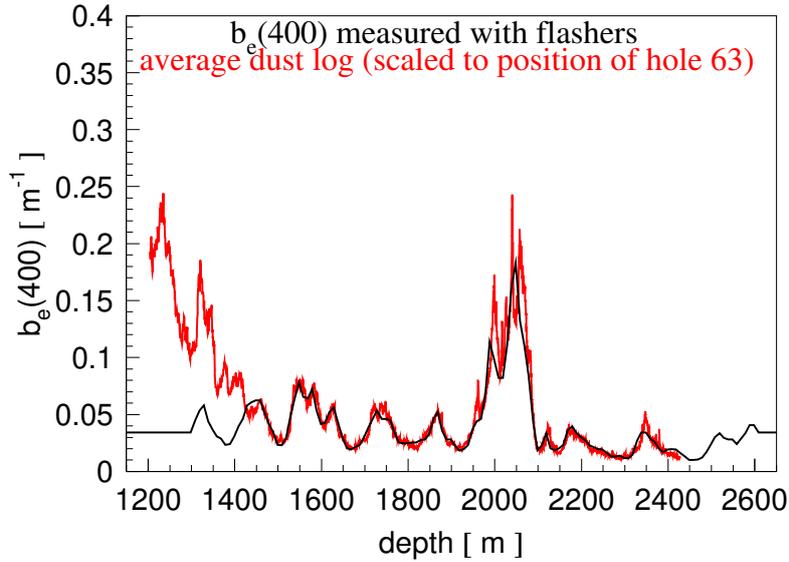
A subset of in-ice DOMs is deployed on eight specialized, closely-spaced strings in the center of the array. Together with the seven central standard strings, they make up the more densely instrumented DeepCore array. The inter-string spacing in DeepCore is on average 72 m. The specialized DeepCore strings have a vertical spacing of 7 m between DOMs for the bottom 50 DOMs, deployed at depths of 2100 m to 2450 m. The remaining 10 DOMs are deployed above 2000 m with a spacing of 10 m. The majority of DOMs on the DeepCore strings is instrumented with DOMs using PMTs with a 35% higher quantum efficiency than the standard IceCube DOMs. This design is optimized for the detection of neutrinos with energies from 10 to 100 GeV, relevant for e.g. the study of atmospheric neutrino oscillations and Galactic supernovae.

The total amount of light generated by a particle propagating inside the in-ice detector is a measure for its deposited energy. The timing and charge information of hits in individual DOMs can be exploited to reconstruct the direction of the particle. Reconstructions are however dependent on a good knowledge of the detector medium.

### Ice properties

The South Pole ice is the interaction medium for particles moving through the detector, and its optical properties are important for the propagation of Cherenkov photons. A good understanding of these properties, mainly scattering and absorption, is necessary for reliable reconstructions and realistic simulations. A measurement of the optical properties of the ice using the LED flashers of the DOMs is described in Ref. [140]. A recent overview of ice modelling is given in Ref. [141]. We highlight here some general characteristics.

The geometrical scattering coefficient  $b$  determines the average distance between successive scatters of a photon, given by  $1/b$ . The effective scattering coefficient is defined



**Fig. 3.5:** Effective scattering coefficient  $b_e$  for 400 nm photons as derived from flasher data. Overlaid is the variation in the average dust concentrations as function of depth. Figure taken from Ref. [140].

as

$$b_e = b(1 - \langle \cos \theta \rangle), \quad (3.5)$$

where  $\theta$  is the deflection angle after scattering. The average distance traveled by a photon before it is absorbed is similarly described by an absorption coefficient  $a$ . They are largely determined by the presence of impurities in the ice, especially dust and air bubbles. Below 1400 m, air bubbles have over long periods of time converted to non-scattering air hydrate crystals as a result of the large pressure [142], and the effective scattering and absorption lengths are on the order of 25 m and 100 m respectively for photons with a wavelength of 400 nm [140]. Fig. 3.5 shows the values of  $b_e$  as function of depth in the ice compared to measurements of the dust concentration, as determined with dust loggers used during the deployment of the string. It can be seen that around 2000–2100 m below the surface, large concentrations of dust are present. This causes a significant increase of the scattering and absorption in this so-called *dust layer*. This is the reason that the DeepCore strings are not instrumented at these depths.

IceCube was deployed using hot water drilling to melt a column of ice for the strings to be lowered in. After deployment, the water in the column froze again locking the DOMs in place. The resulting ice has a higher concentration of air bubbles than the surrounding ice, resulting in a decreased local scattering length of about 50 cm [140]. The ice in these columns of approximately 30 cm diameter is referred to as *hole ice*.

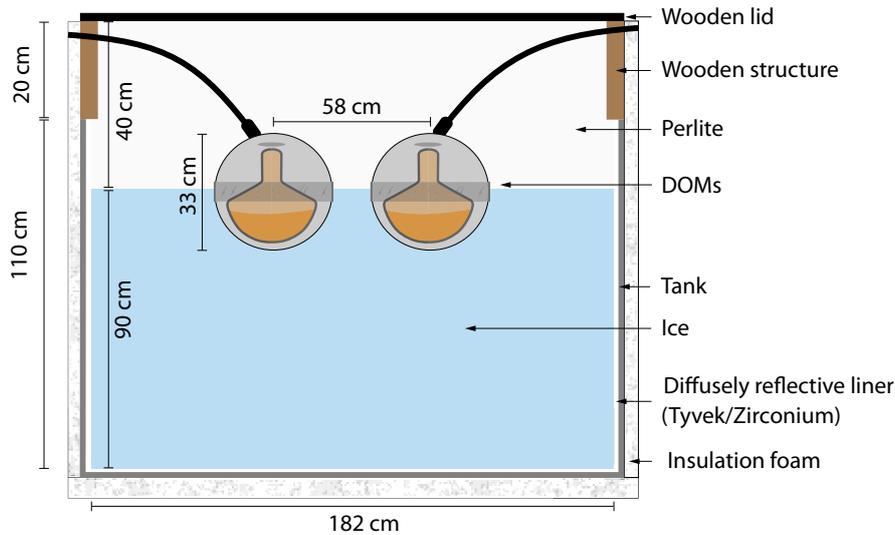
Other important effects have been discovered throughout the years. An anisotropy in the light propagation in the ice has been reported, aligned with the local ice flow direction [143]. Furthermore, layers of ice formed at the same time and with similar properties have found to be tilted, slowly varying in depth throughout the detector [140]. The inclusion of a birefringence effect has strongly improved recent models [141].

An accurate model describing the South Pole ice is important to understand the light yield in the detector. Various measurements of the ice have been performed over the years, and multiple models have been constructed including (a subset of) these effects in different ways. Although the knowledge has improved significantly, ice properties still result in an important systematic uncertainty in many analyses.

### 3.2.3 The surface array: IceTop

The air shower particle detector array IceTop [138] is located above the in-ice IceCube detector. It sits at an average elevation of 2835 m above sea level, corresponding to an atmospheric depth around  $680 \text{ g/cm}^2$ . It consists of 162 Cherenkov tanks filled with ice, placed at 81 stations over an area of  $1 \text{ km}^2$  on approximately the same grid as the in-ice array. Each station is made up of two tanks with a separation of 10 m. The altitude of the tanks varies by about 6 m throughout the array. The layout is ideal for the study of cosmic rays with primary energy in the PeV to EeV range. A denser infill array formed by eight stations in the center of IceTop, similar to DeepCore in the in-ice array, lowers the energy threshold to about 100 TeV.

The IceTop tanks are cylindrical with a height of 1.3 m and a diameter of 1.82 m, and are filled with ice to a height of 0.9 m, as shown in Fig. 3.6. In each tank, two DOMs are frozen into the ice with their bottom halves, PMTs facing down. In this way, they detect the light that shower particles produce when propagating through the tanks. IceTop tanks are not only sensitive to the number of charged particles. The tank ice is about two radiation lengths thick, so most photons in the shower front convert to  $e^+e^-$  pairs, and they provide the dominant fraction of the EM signal. The light generated by shower particles is a measure of the energy deposited in the tanks. The two DOMs operate at different PMT gains to allow better coverage of the huge range of possible energy depositions in the tank. The high-gain (HG) DOM operates at  $5 \times 10^6$ , the low-gain (LG) DOM at  $10^5$ . The tanks are lined on the inside with a reflective material to obtain optimal detection of photons produced in the tank. The space between the ice surface and the lid of the tank is filled with expanded perlite, providing thermal insulation and light protection.



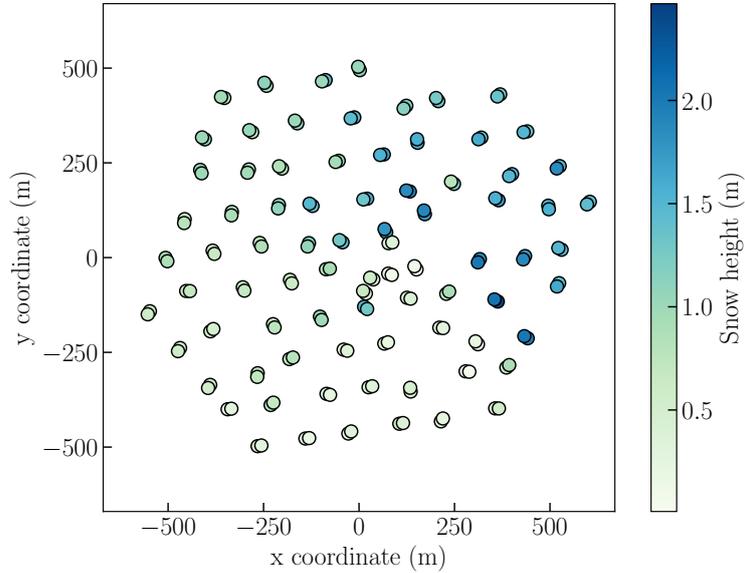
**Fig. 3.6:** Layout and dimensions of an IceTop tank in a cross-sectional view. Wall thicknesses are not to scale. Figure taken from Ref. [138].

For deployment, trenches were dug for the tanks so that their lids were at the same level as the surrounding snow. A special *Freeze Control Unit* was designed to allow the water in the tank to freeze in a controlled manner, ensuring clear and crack-free ice. After deployment, the trenches were filled in again, to minimize temperature variations in the tank. For more details about the tank design and the deployment of the IceTop array, we refer to Ref. [138].

### Environmental conditions

As the Earth's atmosphere can be considered part of the IceTop detector medium, atmospheric variations affect air-shower measurements [144]. The South Pole atmosphere undergoes an annual cycle with large variations. In the winter months (April to September) the sun is below the horizon, and the surface temperature ranges from  $-20^{\circ}\text{C}$  down to  $-70^{\circ}\text{C}$ . The atmosphere is then much denser than in summer. As described in Section 2.2, this influences the shower development. The changes in the density profile of the atmosphere cause a  $\pm 5\%$  variation in the IceTop DOM hit rates, and a  $\pm 8 - 9\%$  variation in the high-energy muon rate in IceCube.

The ground pressure shows no cyclic variation throughout the year but changes on much shorter time scales. As an increase in air pressure causes a stronger attenuation, there is a strong correlation between pressure and IceTop DOM rates. Furthermore, an increase in air pressure causes a decrease of the observed signal strength in IceTop for a given shower.



**Fig. 3.7:** Snow height measurements performed on October 21st, 2012. Each circle represents a tank, and its color shows the corresponding snow height. Positions are given in the detector coordinate system.

While in-ice DOMs function at a constant temperature, the temperature of the Mainboard of IceTop DOMs varies between  $-20^{\circ}\text{C}$  and  $-40^{\circ}\text{C}$  over the year. This has an impact on the shape of the waveforms that needs to be corrected during signal extraction [138].

Another important environmental effect on the data is snow accumulation. While the tanks were deployed with their lid level with the surface, they are now covered in snow. Precipitation only causes about 2 cm of snow per year, but there is a significant amount of snow drift due to wind, causing the snow height on top of tanks to increase by on average 20 cm per year. The accumulation is influenced by surrounding terrain and buildings, and it is therefore different for each tank. The snow height for a tank is measured two times per year. The measurements from October 2012 are shown in Fig. 3.7, and have an average value of  $\sim 1$  m. By March 2021, the average went up to 2.6 m.

We come back to the influence of several environmental effects on analyses in Chapter 4.

### 3.3 Data acquisition

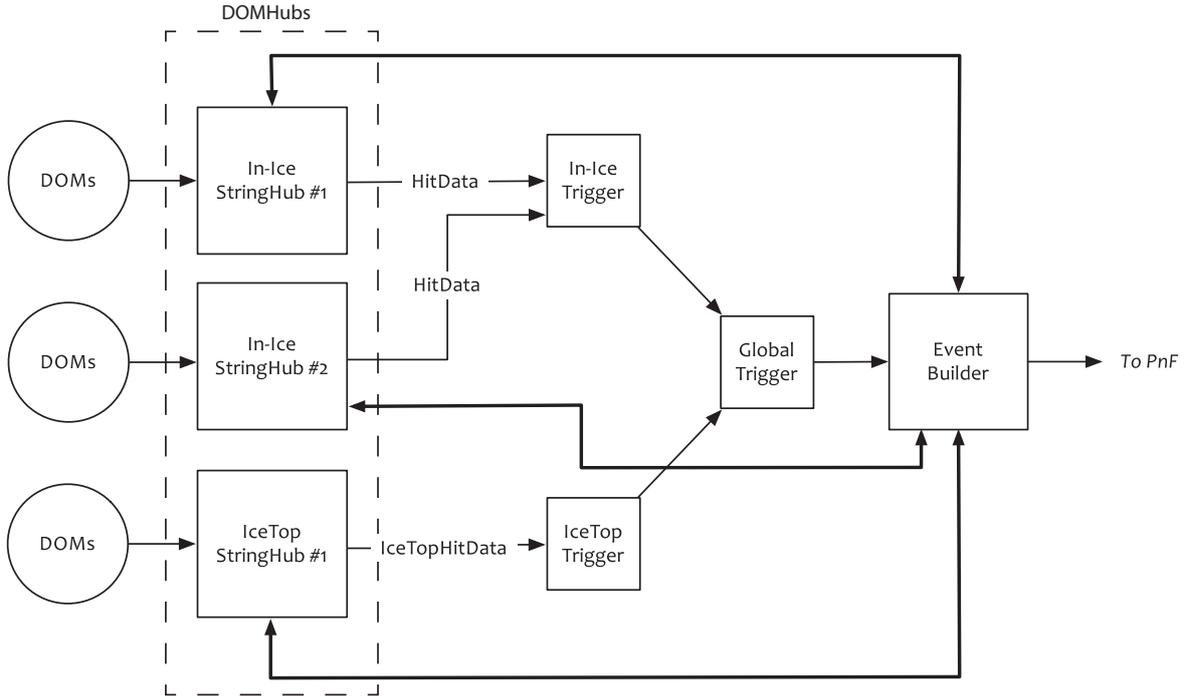
The selection of physically interesting events in the detector happens through a multi-level trigger system including both hardware and software decisions. A schematic overview is shown in Fig. 3.8.

As introduced in Section 3.2.1, the DAQ process starts in the DOMs, which act as autonomous data collectors by capturing data using their ATWDs and fADC when a pulse passes a discriminator threshold. For in-ice DOMs, the discriminator is typically set to a voltage threshold corresponding to 0.25 PE, the typical charge recorded for a single-photoelectron pulse (Section 3.4). For air-shower detection with IceTop, the HG DOM has a threshold corresponding to a signal charge of about 23 PE, and the LG DOM to 270 PE. The discriminator passing initiates the process of capture and digitization of waveforms, resulting in a *hit*.

All DOMs are connected to computers called *DOMHubs* in the ICL, which regularly read out the DOM hits stored in the internal memory of the DOMs. The amount of information sent to the DOMHubs depends on if a Local Coincidence condition is fulfilled or not, as described in Section 3.3.1. The next level of selection happens by the DAQ system, which runs software trigger algorithms on the hits, looking for temporal and optionally spacial patterns that suggest a causal relationship. A trigger defines a time window where all hits are read out. An air shower or a particle moving through the detector will often satisfy more than one trigger condition, with overlapping readout windows. To avoid overlapping events, the triggers and their readout windows are merged into a *Global Trigger*. The *Event Builder* subsequently reads out all hits within the resulting window, and bundles them into an event: the fundamental output of the DAQ.

All events subsequently go through the online Processing and Filtering (PnF) system. The treatment includes application of various calibrations and rudimentary reconstructions. For in-ice DOMs, the waveforms are deconvolved using the known DOM response to a single photoelectron (SPE) to extract the photon arrival times and amplitude information [145]. A subset of events (about 15%), selected by filter algorithms, are transferred to the Northern Hemisphere over satellite. All processed events are archived locally on tapes, which are shipped to the north during the Antarctic summer. More information on the different trigger and filter algorithms is given in Sections 3.3.2 and 3.3.3.

Data-taking runs are divided into 8-hour periods and assigned a unique run number. Detector configuration parameters that may affect physics analysis, such as parameters for certain triggers or filters, are changed at most once per year. This happens typically in May, and indicates the start of a new *physics run*, also often called *season*.



**Fig. 3.8:** Overview of the data flow in the IceCube DAQ system, as described in the text. Figure taken from Ref. [138].

The operational performance is important, especially to remain sensitive to rare astrophysical transient events. Due to the implementation of many redundancies and fail-safes in all detector systems, an average detector uptime over 99% has been achieved. A complete overview of the instrumentation and online systems can again be found in Ref. [137].

More high-level selection and processing of events that are likely induced by cosmic-ray air showers is discussed in Chapter 4.

### 3.3.1 Local Coincidence

DOM hit rates are dominated by various sources of dark noise for in-ice DOMs and by single muons from small showers passing a tank in IceTop. To reduce the influence of these hits, which are most likely not related to a physics event of interest, a first step of data reduction takes place inside the DOM using the Local Coincidence (LC) condition. To achieve this, every DOM has two wires that can connect to other DOMs. For in-ice DOMs, these are the DOM below and above on the same string. For IceTop, the high-gain DOM of each tank is connected to the high- and low-gain DOMs of the other tank in the station. When a DOM launch happens, a signal is sent through this

dedicated wiring to look for coincident hits in a certain time window. This window is set to  $\pm 1 \mu\text{s}$  for both in-ice and IceTop DOMs. When a hit is found in this time window, the LC condition is fulfilled. The hits are then saved in full detail, including the full waveform. If the LC condition is not fulfilled, only abbreviated information is saved, namely a timestamp and a chargestamp. For in-ice DOMs, the chargestamp consists of three samples of the fADC waveform centered around the peak value. For IceTop DOMs, it is the sum of all bins of the ATWD waveform, after pedestal subtraction. When a HG DOM in IceTop launches (which usually happens before the LG DOM passes the discriminator threshold) and the LC condition is fulfilled, each DOM in the other tank of the station which had a discriminator trigger within the time window is read out. Station data therefore always consists of at least signals from both HG DOMs, and optionally additional signals from LG DOMs.

Hits which pass the LC condition are often called *hard local coincidence* (HLC) hits, while isolated hits where no LC is found are called *soft local coincidence* (SLC) hits. This is because two operation modes were foreseen for IceCube: an HLC mode, where only those DOMs which have an LC transfer data, and an SLC mode, where all DOMs with hits transfer data. In practice however, IceCube is always running in SLC mode, so all hits are considered. The trigger algorithms however only use HLC hits. When the Event Builder combines hits into an event, both HLC and SLC hits are included.

### 3.3.2 Triggers

The DAQ system is a set of software components running on the DOMHubs and dedicated servers in the ICL. Its goal is detecting patterns of hits that are likely caused by physically interesting events, such as air showers hitting IceTop or muons propagating through IceCube, and storing the corresponding hits as events. The trigger algorithms work by looking for clusters of HLC hits in a certain trigger time window, sometimes with an additional geometric requirement. When a trigger is formed, longer readout windows are appended before and after the trigger window, to include possible early and late hits belonging to the event. Triggers are generally restricted to a subset of DOMs, e.g. IceTop DOMs, DeepCore DOMs, or all in-ice DOMs.

The most basic trigger for all subsets is the Simple Multiplicity Trigger (SMT). It requires  $N$  or more HLC hits within a time window depending on the subset, without locality conditions. When the condition is met, it is extended until there is a time period of equal length of the initial trigger without any HLC hits. The multiplicities and time windows for the different SMTs are given in Table 3.1, as well as an approximate trigger rate typical for 2013 (the IceTop trigger rate decreases with time because of snow

**Table 3.1:** Simple Multiplicity Trigger (SMT) parameters together with typical trigger rates. The trigger rate varies seasonally with the atmospheric muon flux. The rate for IceTop decreases over time because of snow accumulation. Given values are determined for October 2012.

DOM set	$N$ HLC hits	Window ( $\mu\text{s}$ )	Rate (Hz)
IceTop	6	5	32
in-ice	8	5	2220
DeepCore	3	2.5	270

accumulation on the tanks, see Section 3.2.3). The readout window for the IceTop SMT trigger starts  $10\ \mu\text{s}$  before the trigger window and ends  $10\ \mu\text{s}$  after the last of the HLC hits causing the trigger. For the in-ice SMT trigger, it starts  $4\ \mu\text{s}$  before the trigger window and ends  $6\ \mu\text{s}$  after the trigger window.

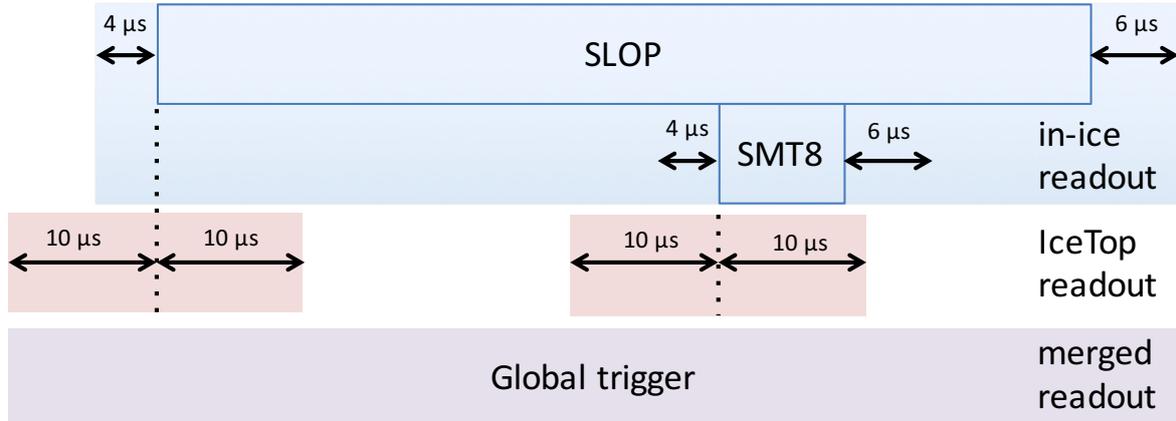
For IceCube, triggers like the Volume and String triggers require a certain number of HLC hits in a cylindrical volume or along the same string. A Slow Particle (SLOP) trigger has also been developed to search for track-like signatures of particles that travel at velocities  $\ll c$ , mainly the hypothetical magnetic monopoles [146]. Other special purpose triggers related to calibration exist for both IceTop and IceCube.

As many events will satisfy multiple trigger conditions, overlapping triggers and their associated readout are merged to a Global Trigger, while retaining information about the separate triggers. All hits from the entire detector, both HLC and SLC from all DOM subsets, are then requested and built into events. An example is shown in Fig. 3.9. As long global events may contain several causally independent physics events, they are typically re-split before reconstruction and analysis.

### 3.3.3 Cosmic-ray filters

Among other functions, the online PnF system filters triggered events as quickly as possible after collection by the DAQ system. The filters tag events, based on simple conditions or basic reconstructions, as possibly interesting for specific analyses and for transmission via satellite. We discuss here the ones of interest for cosmic-ray analyses with IceTop.

The rudimentary IceTop filter is called *IceTopSTA3* and requires that at least three stations have an HLC hit. The cosmic-ray energy threshold for this filter is around 300 TeV. The events passing are further broken down by the *IceTopSTA5* and *IceTop\_InFill\_STA3* filters. The former requires at least five stations to have HLC hits. The latter requires



**Fig. 3.9:** Example of how a Global trigger and merged readout window is formed based on the IceTop and in-ice readout for a long event satisfying the SLOP and in-ice SMT triggers. Figure taken from Ref. [137].

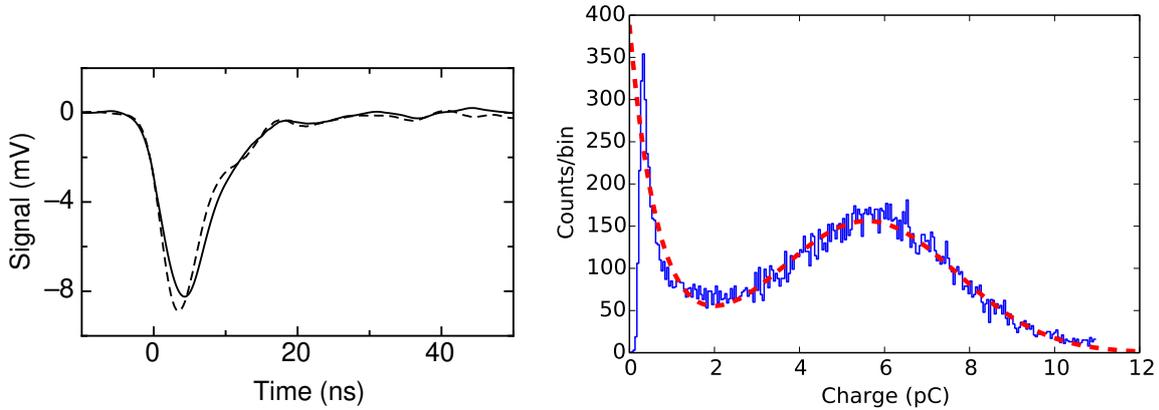
at least three stations from the infill tanks to have HLC hits, with the aim of selecting events with energies down to 100 TeV.

An additional IceTop filter, *InIceSMT\_IceTopCoin*, requires the in-ice SMT to be fulfilled together with the presence of at least one HLC hit in IceTop.

### 3.4 Calibration

To translate the recorded signals into a measurement of Cherenkov light incident at a particular time, a collection of calibration procedures is in use. Various properties such as the optical efficiency of the DOMs, have been measured in the lab [139]. Other properties need to be calibrated on a regular basis in the detector itself. For the in-ice detector, this means regularly calibrating each individual DOM. For IceTop, also a calibration of signals on tank-level is performed, besides the calibration of the DOMs.

Some procedures do not allow simultaneous data taking, as they include for example the flashing of LEDs in the DOMs. They are therefore performed only at fixed times to limit detector downtime. An example is the DOMCal routine, which provides calibration constants used to convert DOM waveforms into physical units. Other calibration procedures are continuously ongoing. An example is the determination of the baseline of the values of the ATWDs, i.e. the average value of the digitizers when the input voltage is zero, which needs to be subtracted from captured waveforms. This is determined during data-taking by launching DOMs at a rate of 1 Hz to capture “empty events”. Together with corrections for other effects, such as temperature dependent changes in the shape of



**Fig. 3.10:** Left: Average SPE waveform for a PMT at gain  $10^7$ , measured with two different transformers. Individual waveforms may differ in amplitude due to fluctuations in the amplification process in the PMT dynodes, but have shapes similar to within a few percent. Right: SPE charge spectrum at a gain of  $10^{7.5}$  as recorded by DOMCal, showing a clear Gaussian SPE peak. Figures taken from Refs. [139] and [137].

the waveform, this results in an absolute calibration of charges in units of PE, which represent the charge measured for a single-photoelectron pulse. A typical SPE waveform and SPE charge spectrum are shown in Fig. 3.10. This is crucial for the unfolding of the in-ice waveform in terms of SPE pulses [145].

Another continuously-running routine is the RAPCal timing calibration. It is used to translate hit timestamps based on the individual DOM clocks to the synchronized surface clocks in the ICL. This establishes a common time base for the array to  $\mathcal{O}(\text{ns})$  accuracy, which can be converted to UTC. It is based on the transmission of a pulse from the ICL to a DOM and back, allowing the time from the two clocks to be matched independent of prior knowledge of cable length.

An extensive overview of all DOM calibration procedures is beyond the scope of this work; more details can be found in Ref. [137]. In the following section, we discuss briefly the calibration of charges measured in the IceTop tanks, which are expressed in units of Vertical Equivalent Muons (VEM).

### 3.4.1 VEMCal

In contrast to in-ice waveforms which are unfolded in terms of single photon pulses, only the timing and total calibrated charge of IceTop waveforms are relevant for typical cosmic-ray analyses. As different tanks may have different optical properties and DOM efficiencies, the DOM response to the same energy deposition in a tank varies from tank

to tank. Therefore, the tank signals are calibrated by expressing them in terms of a common reference: the charge of a vertical muon passing through a tank, defining the VEM unit [138].

This calibration is performed using the natural flux of atmospheric muons. The average muon energy at the detector level is  $\mathcal{O}(\text{GeV})$ , where muons are minimum ionizing (Section 3.1.2). Therefore, they always lose approximately the same amount of energy for a given track length through the tank. The total number of Cherenkov photons emitted scales with the energy deposit and thus with the track length of the muon. The distribution of track lengths peaks at 90 cm, mainly due to muons passing vertically through the top and bottom of a tank. As a result, the charge spectrum recorded in a tank exhibits a corresponding peak, which can be used to define the VEM unit.

Calibration data is taken with special calibration DOM launches, called VEMCal launches, which require that the LC condition is not fulfilled. This enhances contributions of single muons. They are only enabled for HG DOMs, as LG DOMs do not allow triggering on single muons.

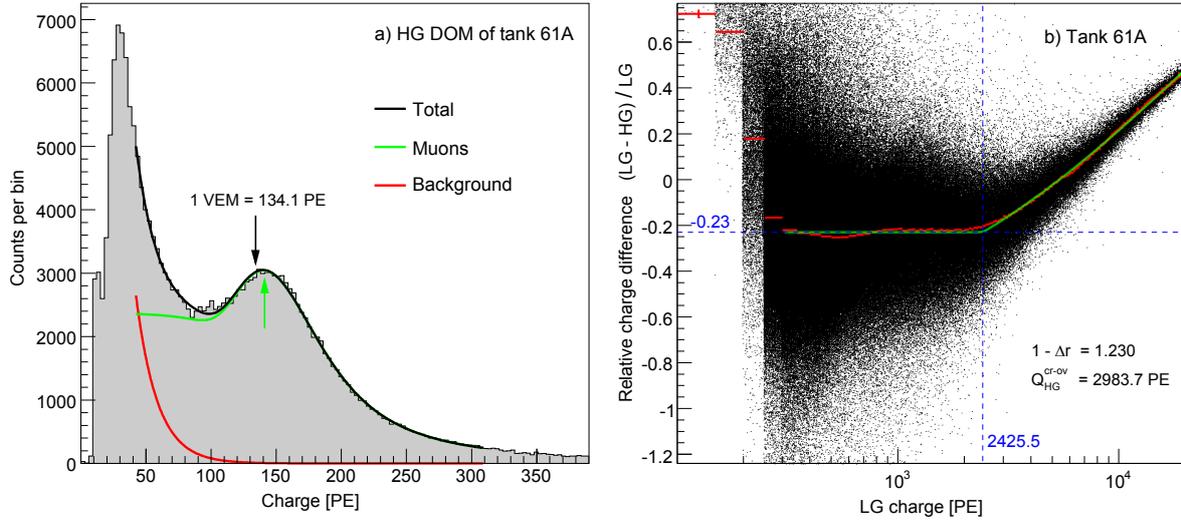
The raw ATWD waveform is calibrated using the charge calibration procedure described earlier, resulting in the charge in PE units. VEMCal hit data are processed offline in the North on a biweekly basis. The charge spectrum in PE of the HG DOMs is fit to extract the muon peak, as demonstrated in Fig. 3.11 (left). The unit of VEM is defined as the charge value at 95% of the muon peak<sup>2</sup>. It corresponds for most HG DOMs to about 125 PE.

As this procedure can only be performed for HG DOMs, the LG DOMs are cross-calibrated assuming that the two DOMs in a tank record on average proportional light yields per particle. The ratio of the signals of the LG and HG DOM should then be constant. This is the case in a large region of charge as shown in the right panel of Fig. 3.11; the region where saturation of the HG DOM does not play a role. The relation between LG and HG charges is fit, and the result can be used to calibrate the VEM unit for LG DOMs. It is found to be systematically lower, with typical charges of 105 PE. This difference is not fully understood, but by the nature of the VEM calibration it is accounted for in data-taking and simulation.

The merit of the VEM calibration lies in the fact that the details of a tank, which could lead to different charges in PE for identical energy deposits, become unimportant when talking about charges in units of VEM. As the shower energy can only be determined by

---

<sup>2</sup>The shift to 95% of the peak value was chosen as a result of studies performed with a muon tagger, used to select nearly vertical muons, placed on top of a tank. The measurement of these tagged vertical muons resulted in a 5% lower muon peak [138]. The choice of 5% does not affect physics results, as it also appears in simulations.



**Fig. 3.11:** Calibration monitoring plots of the DOMs in a tank of station 61. Left: Charge spectrum of the HG DOM, fit with a function that is a combination of a muon and EM contribution. The black arrow indicates the definition of 1 VEM. Right: Relative charge difference of the HG and LG DOMs in the tank. The fit (green) to the average behavior (red) is used for the calibration of the LG DOM. At high charges, saturation of the HG DOM is visible. Figures taken from Ref. [138]

comparison with simulations, it is important that the simulation uses the same definition of VEM to quantify the tank response. As a result of the VEM calibration, the details of each tank do not have to be simulated, but one typical tank behavior can approximately be applied to all tanks in simulation.

The VEM calibration is rather stable. Variations of the number of PE per VEM between two successive calibration do not change more than  $\pm 10\%$  with an RMS spread of 3%.

More information about tank calibration can be found in Ref. [138]. While originally the VEMCal procedure happened soon after events were transferred to the north, it was found that this led to unreliable results. The correct calibration now happens in the high-level cosmic-ray standard processing, described in Section 4.2.

### 3.5 Detector simulation

As discussed in the previous chapters, accurate simulations of cosmic-ray induced events are necessary to infer properties of the primary particle from actual measurements. We present a brief overview of what the simulation chain looks like for the simulation of EAS which have the possibility to trigger IceTop and may, depending on the geometry of

the event, have high-energy muons trigger IceCube. The first steps describing how the particles generate Cherenkov photons, and subsequently photoelectrons in the DOMs, differs for IceTop and in-ice. The simulation of the DOM response and the triggers is common.

The IceTop detector simulation is described in detail in Ref. [138]. It starts with the particle output of a CORSIKA EAS simulation (Section 2.4.1) at a certain observation level, a little above the snow on the tanks. Particles with a trajectory that falls within a 30 cm radius of a tank's volume are considered for further simulation. Propagation through the full structure of a tank and its surrounding snow, including the individual snow height on top<sup>3</sup>, as well as the air above the snow, is realistically simulated using Geant4 [147, 148]. To save on computation time, the Cherenkov photons emitted in the tanks are not tracked. The number of photons in the wavelength interval 300–650 nm is counted<sup>4</sup>, and their propagation is modeled based on detailed simulation studies [149]. Their arrival times at the DOMs are modeled as being distributed according to an exponential distribution, tuned to match the waveform decay times in experimental data.

For in-ice simulation, only muons with an energy above 273 GeV are considered, based on energy loss arguments<sup>5</sup>. The propagation of muons through the ice is performed with PROPOSAL [135], which simulates in detail all the interactions with the medium which cause energy loss. The energy losses and the corresponding secondary particles are stored. The next step in the simulation is the generation and propagation of Cherenkov photons. This is done in a hybrid way. Cherenkov photons produced by the muons and their low-energy secondaries are directly propagated, tracking them through the ice [152, 153]. The Cherenkov photons from cascades resulting from large radiative losses are described using interpolations of photon look-up tables [154]. Such tables contain information about the number and arrival times of photoelectrons expected in a DOM. They are pre-produced from simulations, and are used to save time in the simulation of large cascades. For the photon propagation step, an accurate ice model is important, as described in Section 3.2.2. This results in a list of Cherenkov photons arriving at certain DOMs at a certain time, creating photoelectrons, as in the IceTop simulation. Before the PMT response is simulated, noise hits related to e.g. radioactive decays are added for the in-ice DOMs [155].

---

<sup>3</sup>As snow heights change throughout the years, simulation sets are made for specific periods of time. See also Section 4.4.

<sup>4</sup>Outside this interval, the acceptance is nearly zero, as shown in Ref. [140].

<sup>5</sup>The value of 273 GeV was derived as the energy for which 0.01% of muons reach the top of the AMANDA-II detector, or the depth of 1498.5 m [150]. AMANDA-II was a predecessor of IceCube [151]. For IceCube, this depth is actually  $\sim 50$  m below the top of the detector.

The next steps happen for both the IceTop and IceCube part of the simulated event. The generated photoelectrons (and possible noise hits) at the DOM are injected into a simulation of the PMT. This is done by superimposing Gaussian shaped waveforms for single photoelectrons, with a charge randomly chosen from the measured single photoelectron spectrum (e.g. Fig. 3.10). After this, the behavior of the DOM Mainboard is simulated, including the discriminator and launching. Thereafter, the LC conditions are evaluated. Finally, the digitization of the waveforms by the ATWDs and fADCs are simulated. After this, the different trigger algorithms are applied to the simulated hits.

The result of the entire process is simulated data in the same format as the experimental data. Starting from this, further processing and reconstruction is applied to both depending on the physics that one wants to study. The standard processing for cosmic-ray events is described in the next chapter.



# Chapter 4

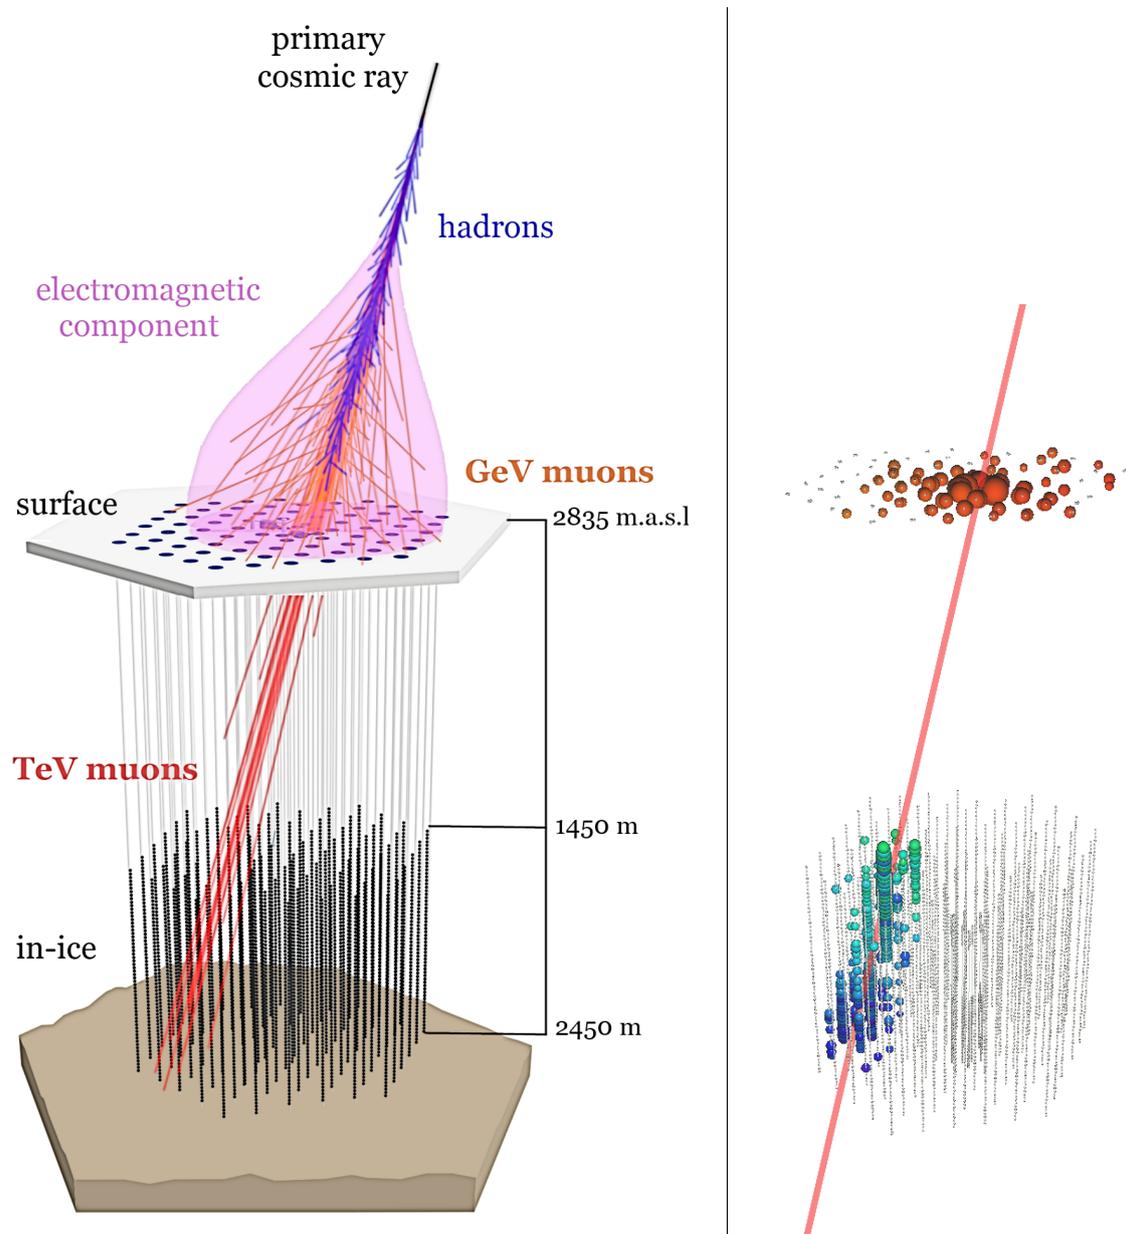
## Cosmic-ray physics with IceCube

With its combination of a surface air-shower array and a large-volume underground detector, the IceCube Neutrino Observatory is an excellent experiment for the indirect detection of cosmic rays. Chapter 1 of this work presented the topic of high-energy cosmic rays, and how studying their properties can provide valuable information about their origin and propagation. In Chapter 2, we introduced how cosmic-ray induced EAS develop in the atmosphere, and how they provide a way to study both cosmic rays and particle physics at very high energies. It was noted that the hadronic physics governing the shower development is not fully understood, motivating a detailed examination of muons in air showers. In Chapter 3, the technical design and functionality of the IceTop and IceCube detectors have been presented. In this chapter, we bring several of these aspects together and discuss how cosmic-ray and air-shower physics are studied with data obtained at the IceCube Neutrino Observatory.

We begin by giving an overview of the detection of EAS with both IceTop and IceCube in Section 4.1. Subsequently, the standard processing of events induced by EAS, including various reconstructions of important air-shower quantities, is discussed in Section 4.2. A selection of cosmic-ray results obtained by the IceCube collaboration relevant to this work is presented in Section 4.3. Finally, the properties of the simulation datasets used in the analyses of Chapters 6 and 7 are discussed in Section 4.4.

### 4.1 Coincident detection of air showers

The IceTop array is well-suited for the detection of EAS from cosmic rays with primary energy in the PeV to EeV range. The lower limit is determined by the distance between the stations, while the upper limit is a result of the total surface area the array covers.



**Fig. 4.1:** Air shower detected in coincidence by IceTop and IceCube. Left: Sketch of the shower development and the different shower components measured at the surface and in the ice. Right: Event display of a simulated 10 PeV proton-induced shower. The red line indicates the shower axis. Hit DOMs are represented by colored spheres. The color scale shows the time evolution, with red happening first and blue last. The size of the spheres is related to the charge measured in the DOMs. Pulses are selected based on the algorithms described in this chapter.

The infill array in the center, with a reduced spacing, can be triggered by showers as low as 100 TeV.

The elevation of the detector and corresponding depth close to  $700 \text{ g/cm}^2$  puts it in the proximity of the shower maximum, as could already be seen in Fig. 2.5. At 1 PeV,  $X_{\text{max}}$  is around  $550 \text{ g/cm}^2$  for proton and  $X_{\text{max}}$   $450 \text{ g/cm}^2$  for iron, while at 100 PeV the maximum for proton showers happens around the surface level. As a result, the shower still has an abundant EM component, as was shown in Fig. 2.4. Given the larger energy loss of electrons compared to muons (Section 3.1.2) and the fact that most photons in the shower convert in the tanks (Section 3.2.3), the total charge measured in the IceTop array in an event is also strongly dominated by the EM component of an EAS, with smaller contributions from hadrons and muons. As will be shown in the following sections, the measured signals in IceTop allow a precise reconstruction of the shower direction and the core position. The abundant EM signal also allows a fairly accurate and model-independent estimate of the primary energy. As shown before in Fig. 2.6, the muonic component is dominated by muons with  $\mathcal{O}(\text{GeV})$  energy, which we refer to as *GeV muons*. While their signal contribution is usually outweighed by the EM signal, their flatter LDF and typical signal of 1 VEM allows them to be observed far from the shower axis.

The detection of an air shower with IceTop can be combined with a detection of high-energy muons, when the geometry of the event permits it. As estimated in Section 3.1.2, muons need several hundred GeV at the surface to cross the 1450 m of ice before reaching IceCube. These muons, which we collectively refer to as *TeV muons*, arrive in a collimated bundle along the shower axis, as they are produced in the first interactions in the atmosphere, when the transverse momentum of the parent hadrons is still small. A sketch of an EAS and its different components measured in coincidence between IceTop and IceCube is shown in Fig. 4.1, together with a visualization of a simulated event. As shown in Fig. 2.6, the number of TeV muons goes from  $\mathcal{O}(10)$  at 1 PeV to  $\mathcal{O}(1000)$  at 1 EeV and is dependent on the mass of the primary. The small radius of the bundle, combined with the sparseness of the detector and the scattering of Cherenkov photons, does not allow the identification of individual muons. However, the light emission of the bundle allows to estimate the total energy loss, which is related to the number of muons. In this way, the TeV muons provide an excellent way of studying the mass composition and hadronic physics, while an independent estimate of the primary energy can be obtained from IceTop.

In the next section, we will describe different reconstructions that are used to quantify the event properties. Afterwards, different examples of analyses that have been performed

based on measurements of different combinations of shower components are discussed. The focus for past measurements has been mostly on *contained* events, where the shower core lies within the boundaries of the IceTop array. This results in the most accurate shower reconstructions, and will therefore also be the focus in this work. Several quality cuts will be introduced to obtain a set of contained events with clear muon bundle signature in IceCube. Because of the geometry of the detectors, the zenith angle range of such events is limited to  $\lesssim 35^\circ$ . Reconstruction methods for more inclined events where a bundle goes through IceCube and which trigger IceTop, but where the shower core lies outside IceTop, are currently being investigated [156].

## 4.2 Cosmic-ray standard processing

Events usually undergo an extensive processing chain before they are used in analysis. Many low-level processing steps are common to all events. After this standard processing, analyzers add various other steps specifically designed for the type of analysis they want to perform. A standard high-level processing to obtain analysis-quality events specifically for EAS events recorded with IceTop, and possibly in coincidence with IceCube, has been devised, called *cosmic-ray Level3 (L3) processing*. In this section, we give an overview of the important steps in data processing, with a focus on reconstructions relevant for the remainder of this work.

The first processing for events is performed by the PnF system at the South Pole, as already described in Chapter 3. All data that passes the filtering and arrives in the data warehouse via satellite transfer, undergoes further processing there. After this point, data is said to be at *Level2*. We highlight the most important steps up to this point in the processing below.

- Hits in DOMs that were known to behave unreliable during the run are cleaned out, and the DOMs are excluded from further reconstructions.
- Hits are calibrated by the procedures that were briefly discussed in Section 3.4, followed by the extraction of *pulses*.
  - For IceTop, a pulse is described by the total charge of the recorded waveform in PE, later calibrated to VEM<sup>1</sup>. For HLC hits, the pulse time is determined by fitting the leading edge of the waveform between 10% and 90% of its

---

<sup>1</sup>An initial VEM calibration happens early in the data processing chain. Some undesired behavior was, however, observed, which is why it is corrected in the L3 processing after manual verification of the VEMCal values.

maximum charge, and extrapolating it down to the baseline. The trailing edge is defined as the time when the waveform drops below 10% again after the peak. This is illustrated in the left panel of Fig. 4.2. For SLCs, there is no waveform available and charge and time were determined by the DOM electronics. To make up for differences with the more precise HLC processing, further SLC calibration is performed later in the L3 processing.

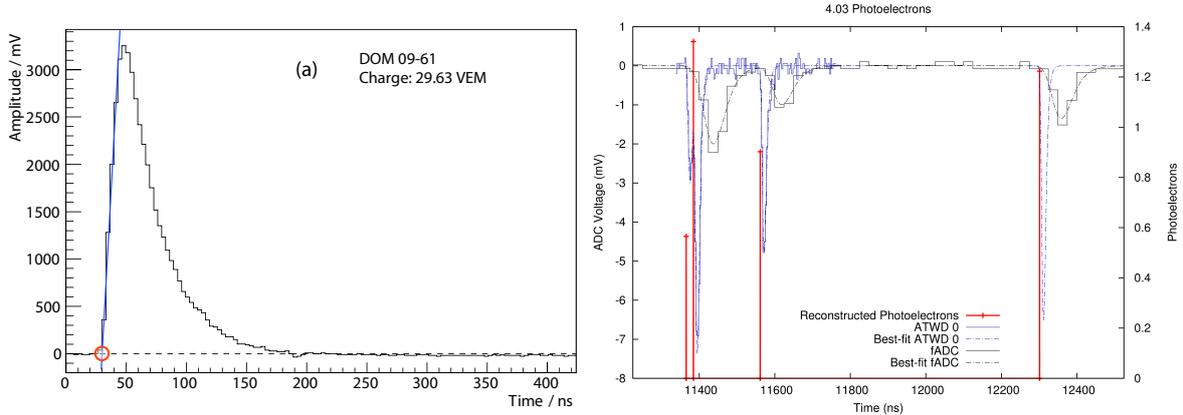
- For IceCube, waveforms are deconvolved in terms of SPE pulses [145]. This leads to a series of pulses with a certain charge (in PE) and arrival time, corresponding approximately to the arrival times of the recorded photo-electrons. An example is shown in the right panel of Fig. 4.2.
- In case both the low-gain and high-gain DOM in a tank were hit, their information is combined into a global tank pulse. By default, the high-gain DOM is used to represent the tank, as it provides the most detailed waveform. If the high-gain DOM was saturated, the low-gain pulse is used if the time difference does not exceed 40 ns. If there is no such pulse, the high-gain pulse is marked as saturated, and this is taken into account when performing later reconstructions.
- As hits in a global readout window (Section 3.3.2) may originate from different physics events with no causal relation, they are split based on different algorithms. For IceTop, a pulse cleaning and clustering algorithm is used based on HLC tank pulses to look for spatial and temporal correlations between pulses. This is done by ordering the pulses in time, and requiring that the time difference  $\Delta t$  between subsequent hits fulfills

$$\Delta t < \Delta x/c + 200 \text{ ns}, \quad (4.1)$$

where  $\Delta x$  is the distance between the two tanks in which the hits were recorded. This results in hits grouped in clusters that could belong to the same shower, and removes a large fraction of noise. In-ice pulses that could be related to the event seen by IceTop are later selected in the L3 processing.

- The cosmic-ray filters, listed in Section 3.3.3, are run on the events resulting from the splitting procedure described above.

The L3 processing improves on the standard processing in several ways. For example, snow measurements are performed only twice per year. These values are afterwards interpolated on a run-by-run basis and inserted in the data to be taken into account during the shower reconstruction algorithm. It also takes care of accurate VEM calibration



**Fig. 4.2:** Examples of IceTop and IceCube waveforms. Left: Typical IceTop waveform with a steep rise, single peak, and an exponential tail. The pulse time is defined by extrapolating the slope of the leading edge to the baseline. Figure taken from Ref. [138]. Right: Unfolding of an IceCube waveform in terms of photoelectrons. The vertical lines represent the times and charges of the unfolded pulses relative to the right-hand axis. As a result of PMT amplification variance, the SPE pulse charges fluctuate. The dashed lines show the corresponding predictions for the various digitizer read-outs (ATWD 0 in blue, fADC in black). Figure adapted from Ref. [145].

(Section 3.4.1), after verification of the values, and performs a basic calibration for the timestamp and chargestamp of SLC pulses which have been determined by the onboard DOM software. Furthermore, it provides cleaning of noise pulses with various algorithms, as well as detailed reconstructions of different EAS observables. Some of these reconstructions are extensively used in the following chapters. Therefore, we discuss the most relevant steps in the processing in the following subsections. A detailed overview of the data processing chain for cosmic-ray events is given in Ref. [157]. Its development was strongly inspired by the analysis presented in Ref. [158].

### 4.2.1 IceTop processing

All events that fail to pass an IceTop filter are omitted. If they pass the three station filter but not the five station filter, it is checked if the hit stations are located next to each other. If not, the event is also removed.

Before the air-shower reconstruction procedure, a pulse cleaning algorithm is used which is stricter than the simple splitting algorithm described above. It is only run on HLC pulses, as the shower reconstruction does not use SLC pulses. The algorithm first looks for a cluster of 3 HLC hits within a certain radius and time interval. It then selects HLC hits in a radius of 180 m within a time of 450 ns around these hits. Three iterations

of this selection are performed. This process keeps the hits related to the real EAS event and removes most of the noise hits, e.g. from coincident low-energy showers which only hit one station.

The resulting set of HLC pulses is used as input for the air-shower reconstruction procedure, called *Laputop*, which aims to obtain the core position, shower direction, and two parameters describing the lateral distribution of measured charges: the shower size, i.e. the signal strength at a certain lateral distance, and the slope. This is done by fitting the measured charges with a lateral distribution function (LDF) and the signal times with a function describing the geometric shape of the shower front.

An LDF specific for IceTop was derived from simulations [159, 160]. The expected charge  $S$  at a distance  $R$  from the shower axis is described by

$$S(R) = S_{\text{ref}} \cdot \left( \frac{R}{R_{\text{ref}}} \right)^{-\beta - \kappa \log_{10} \left( \frac{R}{R_{\text{ref}}} \right)}, \quad (4.2)$$

with free parameters  $S_{\text{ref}}$ ,  $\beta$ , and  $\kappa$ . It is referred to as a double logarithmic parabola (DLP), as the logarithm of  $S$  is a second order polynomial in the logarithm of  $R$ . In that representation,  $\beta$  and  $\kappa$  are respectively the slope and the curvature of the parabola at  $R = R_{\text{ref}}$ . As suggested by simulation studies, the curvature is kept at a constant value of  $\kappa = 0.303$  [160]. As will be discussed in more detail in Chapter 6,  $\beta$  is composition dependent, which is to be expected from the discussion in Section 2.3.1. The signal  $S_{\text{ref}}$  at the reference distance  $R_{\text{ref}}$  is known as the shower size. As will be discussed later, it is strongly correlated to the primary energy. The current standard fit uses  $R_{\text{ref}} = 125$  m, which was found to minimize the composition dependence of the energy estimator.

The effect of snow accumulation on the tanks, discussed in Section 3.2.3, is approximately corrected for during the reconstruction. This is done by modifying the expected charge from Eq. (4.2) assuming an exponential absorption model:

$$S_i^{\text{att}} = S_i \exp \left( \frac{d_{\text{snow}}}{\lambda_{\text{eff}} \cos \theta} \right). \quad (4.3)$$

The expected attenuated signal  $S_i^{\text{att}}$  depends on the snow height  $d_{\text{snow}}^i$  on tank  $i$  and the zenith angle  $\theta$ . The attenuation length  $\lambda_{\text{eff}}$  determines the magnitude of the effect, and is an *effective* parameter because EM particles and muons are affected differently by the snow. Its ideal value was obtained as 2.1 m from studies on experimental data in 2010 [161], but changes throughout the years as the muon signal tends to become overcorrected. For the 2012/2013 data used in this work, a value of 2.25 m is used [19]. Ideally, the snow correction would depend on the distance to the core and on the primary

mass, as these both change the ratio of muonic to EM signal, but this is difficult to achieve in practice. As a result, a fairly large systematic uncertainty of  $\pm 0.2$  m is assigned to  $\lambda_{\text{eff}}$ . Studies for improved treatment of the snow absorption are currently being performed [162].

The expected signal time  $t$  at position  $\mathbf{x}$ , determined by the shape of the shower front, can be described as

$$t(\mathbf{x}) = t_0 + \frac{1}{c}(\mathbf{x} - \mathbf{x}_c) \cdot \mathbf{n} + \Delta t(R), \quad (4.4)$$

where  $t_0$  is the time the shower core reaches the ground at position  $\mathbf{x}_c$  and  $\mathbf{n}$  is the unit vector along the direction of the shower<sup>2</sup>. The term  $\Delta t$  describes how the curved shower front deviates from a plane perpendicular to the shower axis. From experimental data, it was found that it can be described by the sum of a parabola and a Gaussian function symmetric around the shower axis:

$$\Delta t(R) = aR^2 + b \left( 1 - \exp\left(-\frac{R^2}{2\sigma^2}\right) \right), \quad (4.5)$$

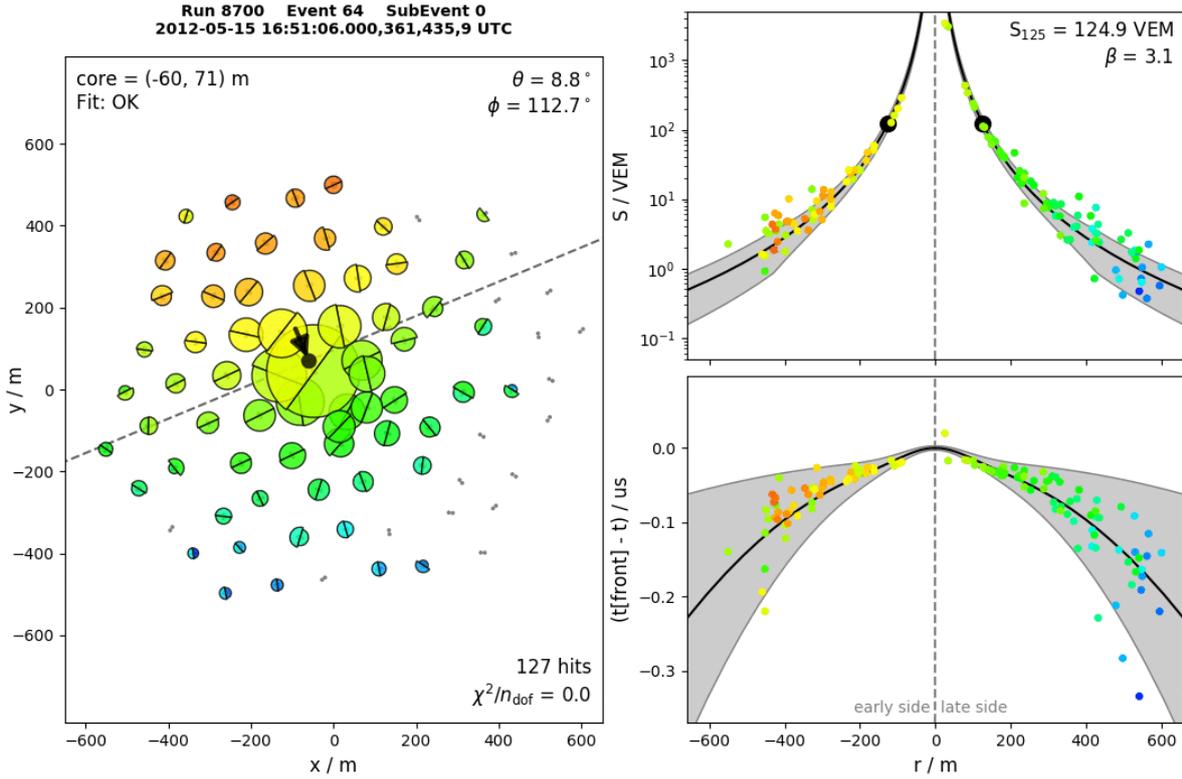
with constants  $a = 4.823 \times 10^{-4} \text{ ns m}^{-2}$ ,  $b = 19.41 \text{ ns}$ , and  $\sigma = 8.35 \text{ m}$  [160].

Eqs. 4.2 to 4.5, describing the expectations for the charge and time of air shower signals in IceTop, are fitted to measured data with a maximum likelihood method. It is therefore necessary to also know the magnitude of the fluctuations of signal charge and time for each tank. These have been studied in experimental data by looking at local fluctuations between the two tanks of a station, as well as in simulation. The likelihood also includes a term taking into account stations that did not trigger, and a term that treats saturated tanks close to the shower core. A detailed description of the full likelihood is given in Ref. [158].

The minimization procedure happens in a number of steps. It starts with calculating a first-guess for the shower core position and  $S_{125}$  based on the measured tank charges, for the direction based on describing the arrival times with a plane shower front, and assumes  $\beta = 2.6$ . Three iterations of the minimization are performed, varying the boundaries of the different free parameters. The procedure improves the stability of the fit and was optimized for showers with their core inside or close to the IceTop array and zenith angles up to  $40^\circ$ . An example is shown in Fig. 4.3. A number of quality cuts to obtain an event sample with accurate reconstructions are discussed in Section 4.2.3.

---

<sup>2</sup>The altitude of ground level is defined as the average of the altitudes of the participating tanks, weighted with  $\sqrt{S}$ , where  $S$  is the measured charge.



**Fig. 4.3:** Event view of a simulated proton shower at 97 PeV with a true zenith of  $9.0^\circ$ , azimuth of  $113.4^\circ$ , and core position  $(-60 \text{ m}, 68 \text{ m})$ . Shown are the HLC pulses with charge corrected for snow attenuation after the cleaning algorithm that happens before the shower reconstruction, as described in the text. The timing is represented by the color, with red happening first and blue last. Left: Visualization of the detected pulses, where each tank is represented by a half circle, oriented along the line connecting the two tanks in a station. Larger circles indicate a larger detected signal. Top right: Reconstructed LDF as described by Eq. (4.2) compared to the detected charges. Bottom right: Deviation of the pulse times from a plane shower front compared to Eq. (4.5).

Based on the air-shower reconstruction obtained from the above procedure, a third pulse cleaning algorithm is implemented. It starts from the uncleaned tank pulses and selects them based on their arrival time and location being in accordance with expectations from the track reconstruction, using a plane shower front approximation. Hits are accepted within a window of  $-200$  ns to  $+800$  ns around this plane. This is a further improvement from the previous algorithms, as now also isolated pulses further away from the shower core can be selected. This is especially important for the inclusion of SLC hits, that may be caused by muons at large lateral distances belonging to the shower.

### 4.2.2 IceCube processing

The L3 processing of in-ice data has the goal of extracting the (HLC and SLC) hits resulting from the high-energy muon bundle corresponding to the air shower that triggered IceTop, and to reconstruct its energy loss.

As a first step, an in-ice SMT trigger which could be causally connected to an IceTop event is sought after. As the muons travel at  $\sim c$ , the muon bundle should arrive approximately  $5 \mu\text{s}$  after the IceTop signals for a vertical shower. To allow for more inclined showers and showers that have their core outside IceTop, in-ice SMT triggers in a window ranging from 2 to  $7 \mu\text{s}$  after the IceTop trigger are considered. When a corresponding trigger is found, the pulses are selected from  $0.3 \mu\text{s}$  before until  $0.4 \mu\text{s}$  after the trigger, or until the maximum length of  $6.5 \mu\text{s}$  is reached<sup>3</sup>.

Next, the shower axis reconstructed with IceTop is used to clean the pulses, as it describes how the muon bundle propagates. Hits that happened more than  $400$  ns before the muon bundle is expected to pass through the array are removed. Furthermore, a cylinder with a radius of  $500$  m around the track is defined outside of which all hits are removed, as they are most probably noise hits. This also effectively removes random coincidences, where e.g. an unrelated atmospheric muon passes through the detector at the same time as the muon bundle from the air shower.

The resulting pulses are used to perform a reconstruction of the energy loss of the muon bundle along the track with the so-called *Millipede* algorithm. As described in Section 3.1.2, the energy loss of a muon is a combination of quasi-continuous ionization losses and large radiative losses. Due to the small radius of the muon bundle compared to

<sup>3</sup>The length of  $6.5 \mu\text{s}$  is chosen to exclude pulses which are likely a result of PMT afterpulses, and which typically have a delay around  $6 \mu\text{s}$  [137]. Afterpulses are a common feature of PMTs. They follow true signal pulses after a short delay period and can be caused by various mechanisms, such as the ionization of residual gases in the PMT by electrons accelerated between the dynodes [163].

the spacing of the strings, only the combined light yield of the muons in the bundle can be seen. The Cherenkov photons undergo scattering in the ice before being detected by a DOM. As a result, the DOM waveform is a combination of light emitted along different parts of the track. The energy loss procedure applied to the in-ice data unfolds the contributions from different segments of 20 m along the track with a maximum likelihood algorithm [145]. For each segment, a template of a cascade is fitted. The expected light yield from a segment where a certain amount of energy is lost, is obtained from tabulated MC simulation of light propagation in the ice [164]. This results in a series of fitted energy losses along the track, between slant depths of 1300 m to 2600 m<sup>4</sup>.

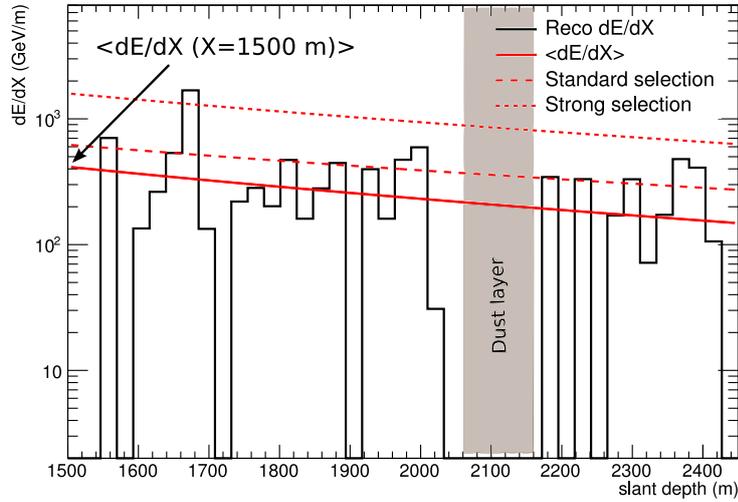
As a next step, the standard processing extracts various features relevant for studying the high-energy muon content from the reconstructed energy deposits. First, segments which lay outside the detector or inside the dust layer are removed, as the reconstruction is inaccurate there. Next, the energy loss per distance interval is calculated by dividing the reconstructed energy deposit of a cascade in a segment by the length of the segment (20 m). A physically motivated function [158] is then fitted to the energy loss values, as shown in Fig. 4.4. From this, the energy loss at a slant depth of 1500 m, referred to as  $dE/dX_{1500}$ , is extracted, which is expected to be related to the number of muons in the bundle. The fit is also used to define two different thresholds, which are used to count large stochastic losses, as these contain information about the energy spectrum of the muons and therefore of the mass composition<sup>5</sup>. A complete overview of this muon bundle processing and how the resulting variables are used for a composition analysis is given in Refs. [158] and [19].

For the procedure described above to accurately capture the true energy loss, a good track reconstruction is necessary. The quality cuts used in this work, presented in Section 4.2.3, ensure showers contained in IceTop with a good reconstruction. For uncontained showers, alternative track reconstructions are provided in the L3 processing based on only in-ice information. The accuracy is, however, significantly worse. As this case is not relevant for this work, we do not discuss it further. Studies for improved reconstructions of uncontained showers combining IceTop and IceCube information are ongoing [156].

---

<sup>4</sup>This includes possible energy losses above and below the detector for vertical showers. For easy comparison, the slant depth range was kept the same for more inclined showers that remain after the standard selections of Section 4.2.3 (only up to  $\sim 30^\circ$ ), as explained in Ref. [158]. As will be discussed below, segments outside the detector will be removed for further analysis, as the reconstruction is less accurate there.

<sup>5</sup>Due to the superposition principle (Chapter 2), muons in a proton-induced shower have on average higher energies than muons in an iron-induced shower of the same primary energy. As seen in Section 3.1.2, muons with higher energy are more likely to undergo radiative losses.



**Fig. 4.4:** Example of the energy loss reconstruction performed on the signal deposited by a muon bundle in IceCube. The black line shows the reconstructed energy loss in segments along the track. The red line is a fit to the energy loss, from which the average energy loss at 1500 m is obtained. The dashed lines define thresholds for counting large stochastic losses. Segments in the dust layer were removed. Figure taken from Ref. [19].

### 4.2.3 Quality cuts

The L3 processing provides a number of event-quality conditions which are checked for every event and which can be used as quality cuts to obtain an event sample which ensures accurate reconstructions for high-level analyses. We describe them below and use them for the work presented in Chapters 6 and 7. They are summarized in Tables 4.1 and 4.2.

#### IceTop cuts

The goal of the standard IceTop cuts is to obtain showers above 1 PeV that have their core inside the array and have zenith angles typically below  $40^\circ$ , which can be reconstructed with good accuracy with the procedure of Section 4.2.1.

As a first basic requirement, events are required to pass the IceTopSTA5 filter, which requires at least five stations to have HLC hits (Section 3.3.3), after the standard hit cleaning.

To ensure that the shower core lies inside the array (and is not misreconstructed as being inside), it is required that the tank with the highest measured charge is not on the edge of the array. This takes care of showers with cores close to the array, where the EM component is still important. If the core is further away, muon signals may result

**Table 4.1:** Summary of the standard IceTop quality cuts.

Quality Cut	Requirement
IceTop STA5 filter	Passed
Station density parameter (text)	$> 0.2$
Maximum signal in event (after snow correction)	$\geq 6$ VEM
Signal in tank next to maximum signal	$\geq 4$ VEM
Station with maximum signal	Not on edge
Air-shower reconstruction (Laputop)	Succeeded
LDF slope at 125 m	$1.4 < \beta < 9.5$
IceTop containment fraction (text)	$< 0.96$

in the largest recorded signal to lie inside the array, even though the core is outside. To remove such events, the highest signal in a tank is required to be at least 6 VEM, and the other tank in the same station is required to have at least 4 VEM. For these cuts, the effect of snow on the signal is taken into account using Eq. (4.3). These cuts furthermore effectively remove events with low energy and/or high inclination. Finally, a cut is applied directly based on the reconstructed shower core. For this, a polygon is defined by connecting the outer tanks of the array, and a line is drawn between the center of the array and the reconstructed core. The ratio of the distance of the center to the core over the total length of the line toward the polygon, sometimes referred to as the IceTop size or IceTop containment fraction, is required to be smaller than 0.96.

A “station density” condition is also included to filter events which have their hits sparsely distributed over the array, leading to unreliable reconstructions. For this, the average location of the hit tanks weighted with the signal is used as the center of a circle which extends to the furthest hit tank. When the ratio of the number of hit tanks to the total number of tanks inside the circle is less than 0.2 for an event, it is removed.

Lastly, the slope of the LDF  $\beta$  is required to be larger than 1.4, as lower values tend to indicate lower quality reconstructions. An upper bound of 9.5 is also implemented, although values above 6 are rarely observed.

### IceCube cuts

Several cuts are also applied on the in-ice pulses and reconstructions, to ensure a good muon bundle energy loss reconstruction along the track reconstructed with IceTop.

First, more than 8 HLC hits are required to remain after the cleaning described in Section 4.2.2. This is basically a repetition of the SMT8 trigger requirement (Section 3.3.2) after the pulse cleaning.

**Table 4.2:** Summary of the standard in-ice quality cuts.

Quality Cut	Requirement
# HLC hits	$\geq 8$
Millipede reduced log-likelihood	$< 10^2$
Millipede forward folded charge / measured charge	$> 10^{-0.03} (\approx 0.93)$
# track segments with non-zero energy loss	$> 3$
Fit to energy loss	Succeeded

Two cuts are applied to output parameters of the energy loss reconstruction algorithm to ensure its quality. The  $\log_{10}$  of the reduced log-likelihood is required to be less than 2. A forward folded estimate of the charge corresponding to the reconstructed energy loss profile is ensured to be of similar magnitude as the actual measured charge by requiring that their ratio is greater than 0.93 [158].

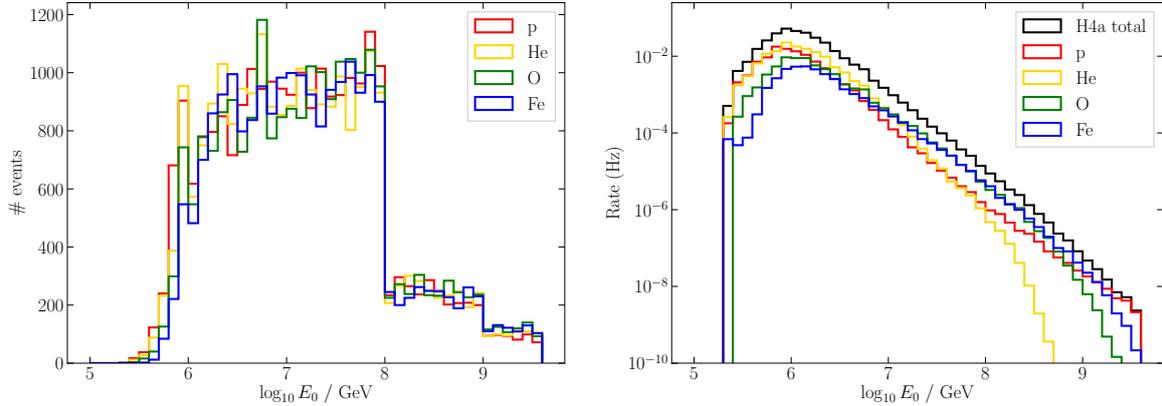
The reconstruction is furthermore required to have at least three cascades with a non-zero energy deposit after removal of the track segments inside the dust layer and outside the detector. Also the fit of a smooth function to the resulting energy loss profile mentioned in Section 4.2.2 can be required to succeed.

The combination of these cuts results in a sample of events where the track of the muon bundle intersects significantly with the IceCube detector, removing corner-clipping trajectories or muon bundles that skim the edge of the array.

#### 4.2.4 Performance

After running the L3 processing and applying the standard quality cuts as described above, we are left with a sample of well-reconstructed showers coincident between IceTop and IceCube. This is the event sample we use in nearly identical form in Chapters 6 and 7. One more cut that is applied in both chapters is  $\cos \theta > 0.95$  or  $\theta \lesssim 18^\circ$  on the reconstructed zenith angle  $\theta$  to obtain near-vertical showers, while the zenith distribution after the L3 cuts extends to  $\cos \theta \approx 0.8$ . We will also limit the primary energy range, as will be discussed later.

Below, we show some plots summarizing some relevant properties of the remaining event sample after the extra zenith cut. These plots are based on the simulated datasets described in Section 4.4, and are weighted according to the H4a flux model, as described in Section 1.5. The energy distribution of simulated events after all cuts is shown for the four simulated primary masses combining low- and high-energy sets based on Sibyll 2.1 in the left panel of Fig. 4.5, while the spectrum after weighting is shown in the right

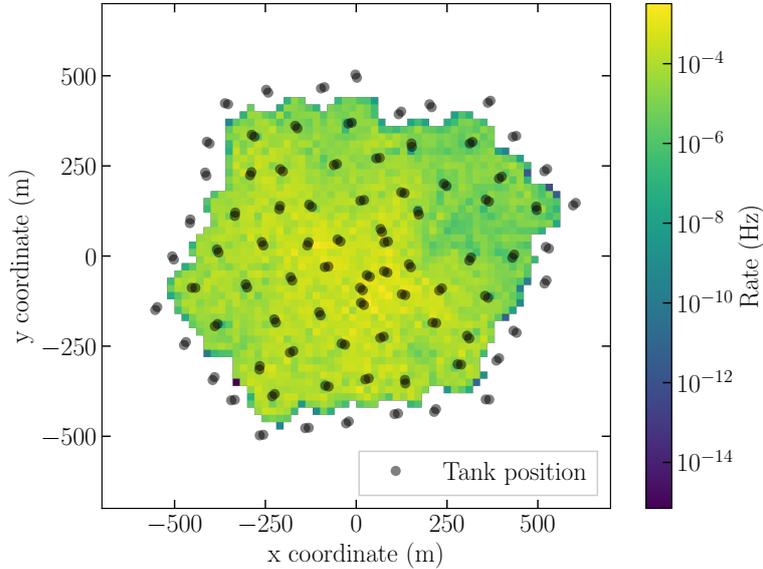


**Fig. 4.5:** Left: Energy distribution of simulated events of the low- and high-energy Sibyll 2.1 sets given in Table 4.3. Right: Simulated events weighted according to the H4a flux model.

panel<sup>6</sup>. At the low-energy end of the plot, a threshold behavior can clearly be observed. This is because these low-energy showers are too small to pass the IceTop triggers or further processing. Because of the increasing snow coverage on top of IceTop, and the increased absorption of mainly the EM signal, the energy threshold where full detection efficiency is reached increases over time. The effect of snow on the threshold was studied in Ref. [165]. For the snow coverage of the period of data used (called IC86.2012, see also Fig. 3.7 and Section 4.4), the threshold for showers of any primary type with a zenith angle  $\cos \theta \geq 0.8$  is found to be about  $\log_{10} E_0 / \text{GeV} \approx 6.4$ . This is also the threshold which will be used for the analyses presented in Chapters 6 and 7. In Section 7.2, we explicitly confirm that this is the correct threshold. It is also demonstrated there that the selection does not lead to any mass dependent bias.

Fig. 4.6 shows the true core positions of showers that pass all quality cuts. It can be seen that the goal of contained showers is obtained. The resolution of the core position reconstruction, defined by the 68% percentile of distances between the true and reconstructed core, is shown for proton and iron simulations in the left panel of Fig. 4.7. Around 1 PeV energies, it is around the order of 12 m for proton showers, and slightly worse for heavier primaries. Above  $\sim 10$  PeV, the resolution has improved to around 5 m with no obvious composition-dependence. The resolution of the directional reconstruction is shown in the right panel of Fig. 4.7. It is defined as the 68% percentile of the space angle between the true and reconstructed direction. The resolution is below

<sup>6</sup>As the H4a model uses five primary mass groups (p, He, CNO, MgSi, Fe), and IceTop simulations are only produced for four primaries (p, He, O, Fe), the Oxygen fraction is modeled to describe the combination of the CNO and MgSi groups.



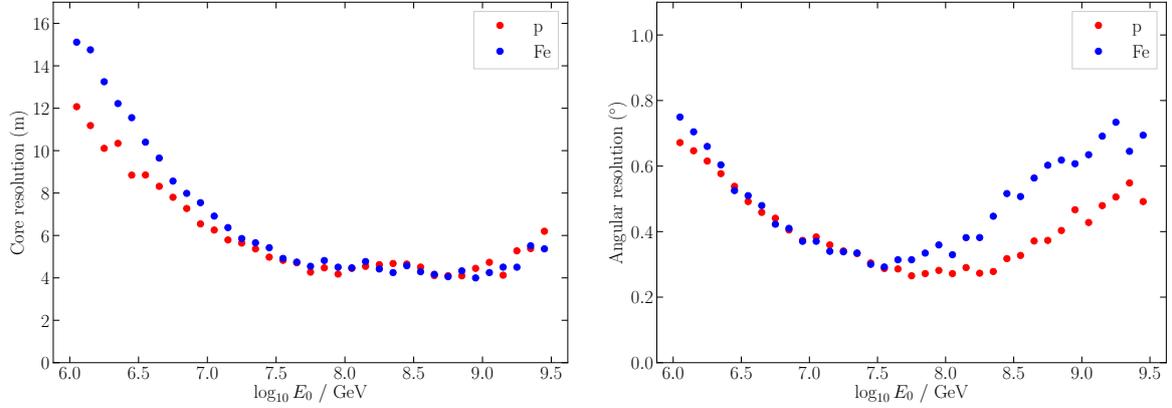
**Fig. 4.6:** Distribution of true core positions in near-vertical, H4a-weighted simulation after L3 processing and the application of the quality cuts as described in the text.

$1^\circ$  at all relevant primary energies for both proton and iron. The optimal resolution is obtained around 30 PeV, and worsens with decreasing and increasing energy. Towards higher energies, a slight composition dependence is visible, with iron-induced showers being reconstructed worse than proton-induced showers.

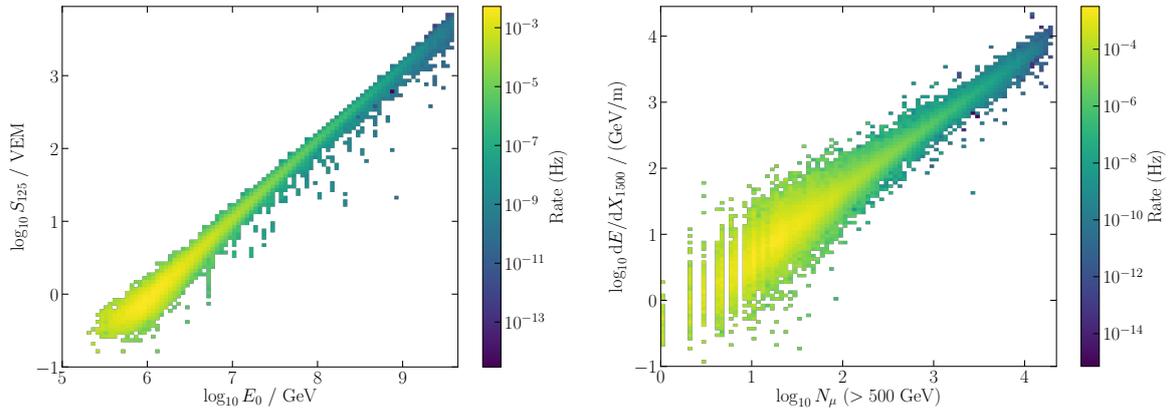
Two variables that have been important in previous analyses with IceTop and IceCube and that are important to this work are  $S_{125}$ , the reconstructed shower size, and  $dE/dX_{1500}$ , the reconstructed muon-bundle energy loss. As explained in Section 2.2, the fact that  $S_{125}$  is mainly determined by the abundant EM shower particles makes it a good estimator for the primary energy.  $dE/dX_{1500}$  on the other hand is determined only by the TeV muon component of the shower and is strongly correlated to the multiplicity of these muons. As a result, it is sensitive to the mass of the primary cosmic ray. These properties are illustrated in Fig. 4.8.

### 4.3 Selected results from IceTop & IceCube

Over the years, the IceCube collaboration has published a variety of results on cosmic-ray and air-shower physics using the surface detector, the in-ice detector, or a combination of both. In this section, we highlight some results which are specifically of interest for this work.



**Fig. 4.7:** Performance of the air-shower reconstruction performed using IceTop signals as derived from simulation. The definition of the resolution is given in the text. Left: Resolution of the reconstruction of the position of the shower core. Right: Resolution of the reconstruction of the direction of the shower.



**Fig. 4.8:** Important observables resulting from the L3 processing shown for 4-component simulation weighted according to H4a. Left: Correlation of  $S_{125}$  with the primary cosmic-ray energy. Right: Correlation of  $dE/dX_{1500}$  with the number of muons with an energy above 500 GeV in the shower.

### 4.3.1 Energy spectrum and composition

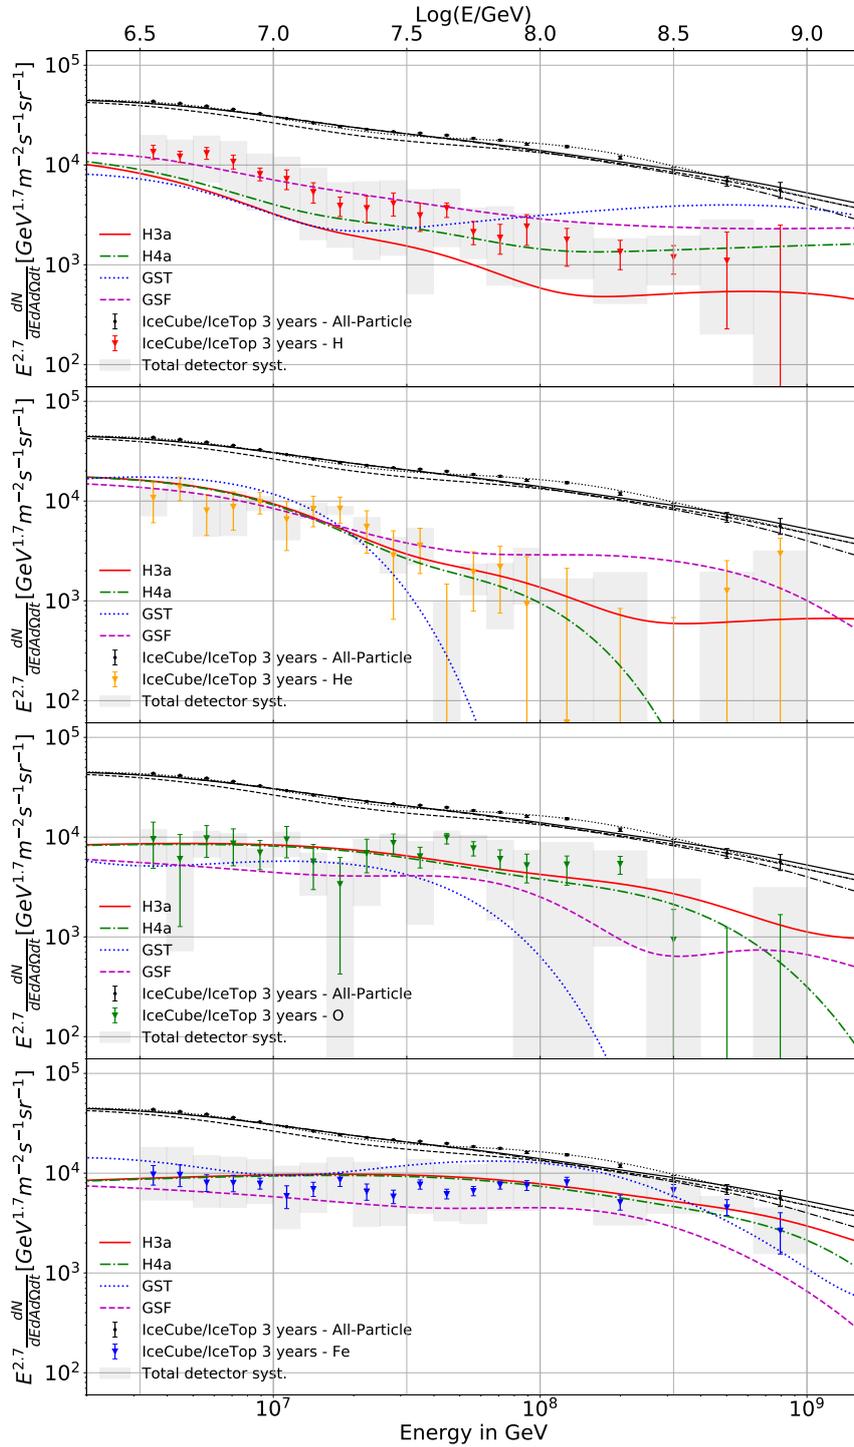
The all-particle cosmic-ray energy spectrum has been measured between 2.5 PeV and 1.3 EeV with IceTop alone. The analysis used contained events and a zenith angle range of  $\cos \theta > 0.8$  or  $\theta \lesssim 37^\circ$ . It is based on a conversion function from  $S_{125}$  to  $E_0$  of the form

$$\log_{10} E_0/\text{GeV} = p_1 \log_{10} S_{125}/\text{VEM} + p_0 \quad (4.6)$$

as can be fit to Fig. 4.8, for different zenith angle bins of size 0.05 in  $\cos \theta$ . The analysis is explained in detail in Ref. [161] using the at that time incomplete IceTop detector with 73 stations, and was repeated using three years of data in Ref. [19]. The results, which clearly show a double knee structure, were included in Fig. 1.4. The energy spectrum measurement has been extended below the knee down to 250 TeV in a dedicated analysis using the IceTop Infill array in Ref. [166].

Next to the all-particle measurement, Ref. [19] also presents a composition analysis which obtains energy spectra for four different mass groups: p, He, O, and Fe. Events are used which have undergone the processing and satisfy the quality cuts for coincident events described in Section 4.2. In this analysis,  $S_{125}$  from the IceTop reconstruction was used as main energy-dependent observable, while the muon bundle energy loss observable  $dE/dX_{1500}$  was used for its mass dependence. These variables were combined in a neural network with the reconstructed zenith angle from IceTop and two variables related to stochastic energy losses as described in Section 4.2.2. The 5-input feedforward neural network was trained to estimate the primary energy and mass. In bins of primary energy, the mass-output distribution of the neural network for experimental data was fit using templates obtained for simulation of the different mass groups. In this way, the fractional contribution of each mass group was determined. The resulting spectra, which were already included without uncertainties in Fig. 1.7, are now shown in Fig. 4.9 including statistical and systematic uncertainties related to the detector. The detector systematics are related to the snow accumulation on IceTop, the VEM calibration, and the ice model properties, which has the largest contribution to the total uncertainty. The results are consistent with an increase in the average mass between the first and second knee, as discussed in Section 1.4. Some features are, however, different from predictions of the models discussed in Section 1.5. The IceCube result also shows disagreements with some other experiments. This is, however, not surprising, recalling the discussion of uncertainties in air-shower observables in Chapters 1 and 2.

It is important to note that the spectrum and composition analyses were based on simulations with Sibyll 2.1, respectively for the determination of the  $S_{125}$  conversion func-



**Fig. 4.9:** Mass composition in the PeV–EeV range measured with IceTop and IceCube using hadronic interaction model Sibyll 2.1. The different panels show the individual spectra for the four mass groups including detector systematic uncertainties, compared to various models discussed in Section 1.5.

tion and the training and of the neural network and subsequent template fit. The results are therefore model dependent, as is often the case in indirect cosmic-ray measurements. As we will discuss in Section 4.4, smaller simulation sets are usually produced using other hadronic interaction models for comparison. Using a different model has a stronger impact on the composition analysis than the all-particle spectrum measurement, due to the differences in the muonic component. As the relative difference in the number of muons between models varies only slowly with energy (see Chapter 7), this mainly causes a global shift in the average composition, rather than a change in the evolution of the composition with primary energy [19].

### 4.3.2 Hadronic interaction models

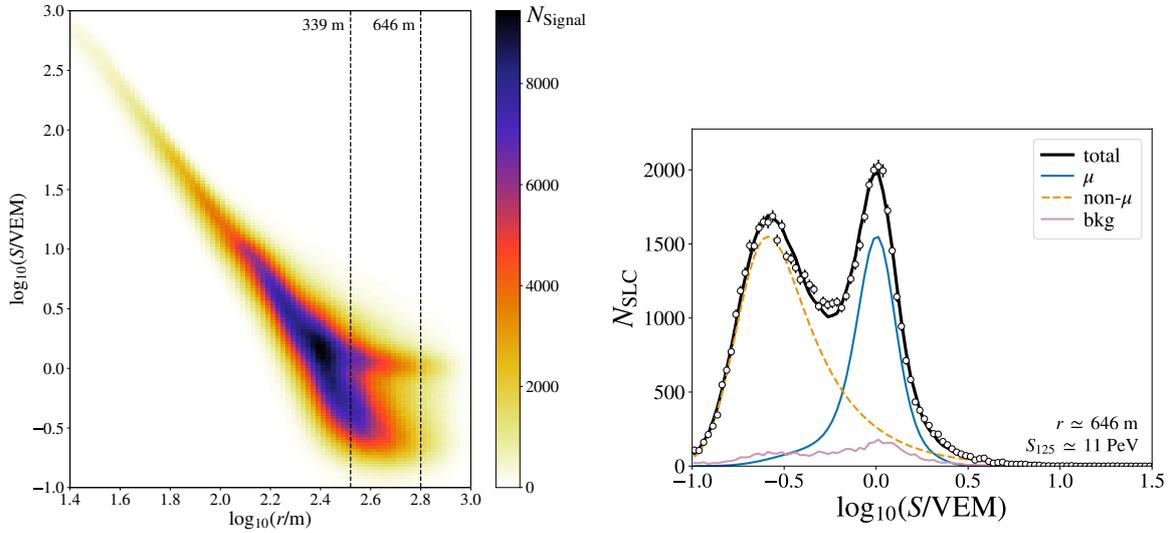
Uncertainties in the determination of the mass composition result not only from differences between hadronic interaction models, but also from internal inconsistencies in the predictions of different air-shower observables in individual models. Observations based on different mass-sensitive EAS observables should lead to consistent conclusions about the properties of the primary cosmic rays, while inconsistencies indicate deficiencies of the interaction model [45]. As discussed earlier in this chapter, IceTop and IceCube provide the possibility to measure observables based on different air-shower components and perform unique tests of hadronic interactions.

This was explored in Refs. [157] and [167] using the composition-dependent slope of the IceTop LDF  $\beta$  together with  $dE/dX_{1500}$ . The latter obtains its composition sensitivity from the number of high-energy muons, while for  $\beta$  it is the result of a more complex interplay of the number of low-energy muons and the depth of shower maximum, as discussed in Section 2.3.1 and further explored in Chapter 6. The average values obtained for  $\beta$  and  $dE/dX_{1500}$  in data were compared to the values in proton and iron simulation to obtain their average composition interpretation, as a function of the energy-proxy  $S_{125}$ . This was performed with a similar sample of coincident events as described above.

For the simulations using QGSJet-II.04, a consistent interpretation was found. Inconsistencies were found for EPOS-LHC and Sibyll 2.1. For the latter,  $\beta$  furthermore showed unphysical behavior in simulation, as it led to a composition interpretation of data heavier than iron. This work is extended in Chapter 6.

### 4.3.3 Density of GeV muons

The density of muons far from the shower core has been measured with IceTop in near-vertical showers with  $\cos\theta > 0.95$  [168]. The dominant contribution comes from



**Fig. 4.10:** Distributions of HLC and SLC signals in IceTop, showing the typical muon signal around 1 VEM that becomes visible at large distance. Left: Distribution of signals as a function of lateral distance and charge in near-vertical showers for an  $S_{125}$  bin corresponding approximately to energies between 10 PeV and 12.5 PeV. At large distances, the muon thumb is visible. Right: Signal distribution at lateral distances around 646 m, corresponding to a vertical slice in the left-hand figure. The lines show the fit muon signal model, the distribution of signals with no muons, and the distribution of accidental signals. Figures taken from Ref. [168].

muons with  $\mathcal{O}(\text{GeV})$  energy. As the technique used in this analysis will also be used in Chapter 6, we summarize here the important steps and results.

While close to the core, the signal is dominated by EM particles, at further distances the contribution of muons with their typical 1 VEM signal deposition increases. This is exploited by making charge versus distance histograms of large numbers of events in different  $S_{125}$  bins, as shown in Fig. 4.10. Here, a typical structure caused by the muonic signal component, called *the muon thumb*, becomes apparent.

This analysis is different from most other IceTop analyses, as it includes not only HLC but also SLC hits, as they are important for the detection of single muons at the outer edge of the shower. These hits are, however, more likely to result from accidental coincidences of low-energy showers. As no background is included in the simulation of IceTop events, the simulation chain is enhanced with an ad-hoc addition of background hits with a rate of 1470 Hz, as obtained from data. Furthermore, a possible smearing of the charge-distance histograms as a result of the 3% fluctuations in the VEM definition

obtained in bi-weekly intervals (Section 3.4.1) is included by a Gaussian smearing of the simulated charges.

The hits used in the analysis are the HLCs and SLCs that have passed the cleaning based on the IceTop shower reconstruction described in Section 4.2.1. To limit background further, pulses are required to arrive within  $-4\ \mu\text{s}$  to  $+0.5\ \mu\text{s}$  around the reconstructed shower front. The resulting charge distributions per distance and  $S_{125}$  bin show two clear peaks at large distance, which can be fit with individual signal contributions from the EM and muon component, together with a small background component. The energy distribution of the EM component is modeled as a power law. The muon signal distribution is described by a physically motivated function which depends on the average number of muons observed per tank, as given in Ref. [168]. From these fits, of which an example is shown in Fig. 4.10, the average density of muons in a sample of showers can be estimated.

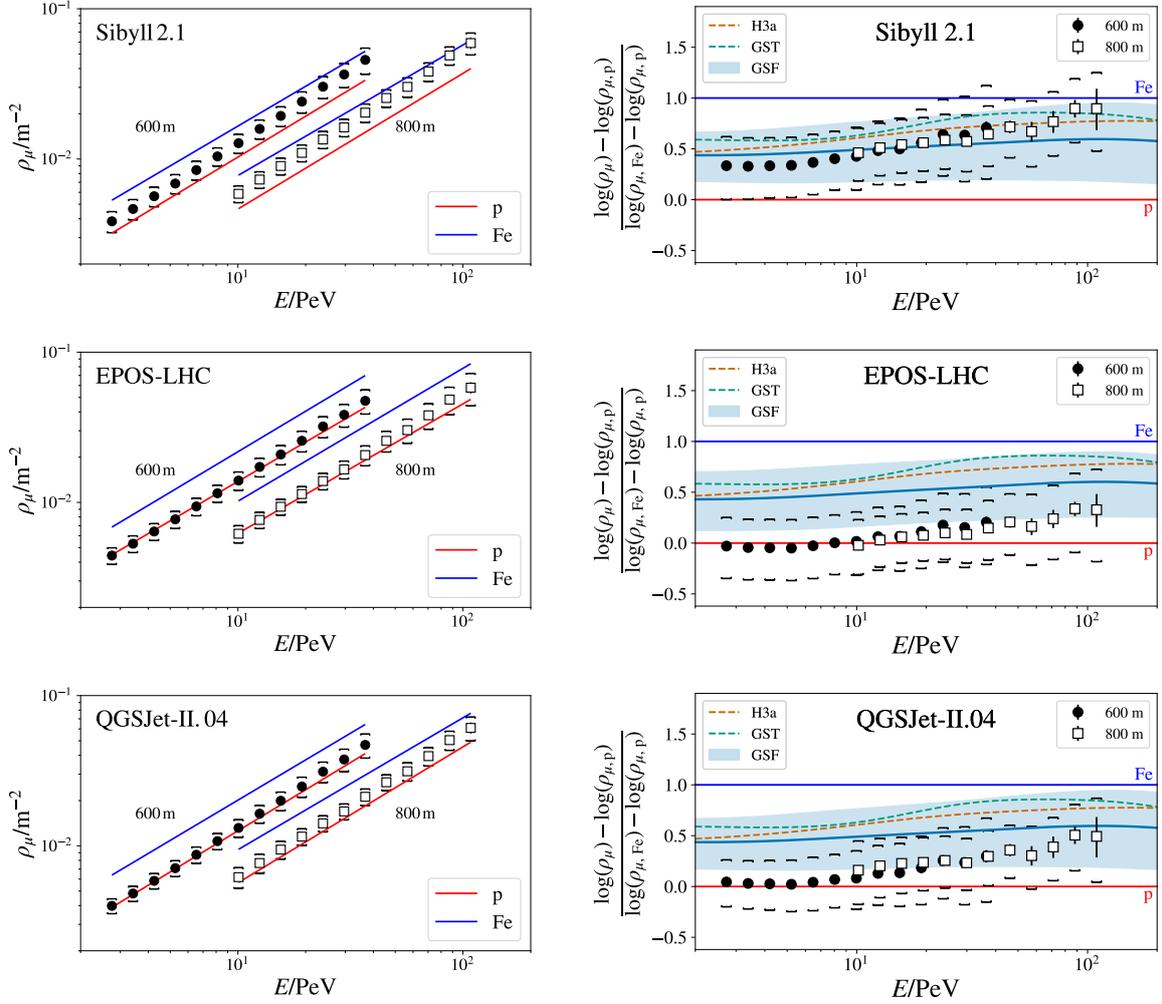
The resulting values of average muon density as a function of distance are then fit to interpolate to the muon densities  $\rho_\mu$  at two reference distances of 600 and 800 m. The central values of the  $S_{125}$  bins are converted to primary energy with a simple conversion function (Eq. (4.6)). This results in reconstructed values of  $\rho_\mu$  as a function of energy.

Different effects, such as imperfect signal models and misreconstructions in shower direction, cause the results to be systematically biased. This effect is quantified by comparing reconstructed  $\rho_\mu$  values in simulation to the true values expected from CORSIKA simulation (without detector simulation). The ratio of the two is taken and fitted with a linear function to obtain a correction factor. As the correction factor differs slightly for proton and iron simulation, the average of the two is taken, and the difference is included as an uncertainty. The resulting correction factor can then be applied to the reconstructions from experimental data to obtain a measurement of the muon density. The correction factor is derived from simulation and therefore adds a model dependence to the measurement.

The results for three hadronic interaction models are shown in Fig. 4.11. The result using Sibyll 2.1 is close to the expectations of the models discussed in Section 1.5, while those using the post-LHC models QGSJet-II.04 and EPOS-LHC indicate a lighter composition. The systematic uncertainties are, however, large. These results were also included in Fig. 2.8, related to the discussion of the Muon Puzzle.

#### 4.3.4 Atmospheric muons and seasonal variations

Muons produced in air showers account for the dominant part of the event rate in large underground detectors. In IceCube, the in-ice SMT trigger rate is dominated by



**Fig. 4.11:** Measured muon densities at 600 m and 800 m lateral distance after applying model-dependent correction factors. Left: Muon densities in data obtained assuming a certain hadronic interaction model compared to the predicted values for a pure proton and iron composition obtained from simulations with the same model. Error bars indicate the statistical uncertainty, brackets the systematic uncertainty. Right: Muon densities from the left-hand plots expressed in the z-scale (Eq. (2.27)). Shown for comparison are the expectations from the cosmic-ray flux models discussed in Section 1.5.

single muons with an energy around the threshold necessary to reach the detector, i.e. several hundred GeV (see Section 3.1.2). These muons are produced predominantly by showers with primary energies between 10 and 100 times this energy [30, p. 136]. As explained in Section 2.2, the competition between decay and re-interaction of parent mesons will result in a seasonal variation in the muon yield, as a result of the variation in the density/temperature of the atmosphere. This causes a seasonal variation in the SMT trigger rate of  $\pm 8\%$ . These variations are strongly correlated with the temperature of the stratosphere [144].

There is a long history of measurements of underground muon rates and their seasonal variations. This temperature effect was first discussed in detail in 1952 in Ref. [169], based on observations in a detector in a deep cavity near Cornell University. It is common to describe the atmospheric temperature by means of a single number, called the *effective temperature*. It is the average of the atmospheric temperature  $T(X)$  along a path weighted by the inclusive production spectrum of atmospheric muons, e.g.

$$T_{\text{eff}}(\theta) = \frac{\int dE_\mu \int dX P_\mu(E_\mu, \theta, X) A_{\text{eff}}(E_\mu, \theta) T(X)}{\int dE_\mu \int dX P_\mu(E_\mu, \theta, X) A_{\text{eff}}(E_\mu, \theta)} \quad (4.7)$$

where  $A_{\text{eff}}$  is the effective area of the detector and the muon production spectrum  $P_\mu$  is the number of muons produced per energy and slant depth along a trajectory at zenith angle  $\theta$  [170]. The exact definition of effective temperature adopted for analyses varies, however, between works [171]. The relation between the relative variation in the measured muon rate  $R$  and the atmospheric temperature is commonly quantified by a correlation coefficient  $a_T$ , defined as

$$\frac{\delta R}{R^{\text{av}}} = \alpha_T \frac{\delta T_{\text{eff}}}{T_{\text{eff}}^{\text{av}}}, \quad (4.8)$$

where  $\delta$  gives the deviation from the average value indicated by the superscript. As the magnitude of the variations relate to the relative contributions of pions and kaons, this can be used to constrain the atmospheric  $K/\pi$  ratio [172].

Seasonal variations of atmospheric neutrinos are also analyzed in IceCube, using a sample of upward-going neutrinos from the Southern sky with zenith angles of  $90^\circ$  to  $120^\circ$  [173].

A seasonal variation has furthermore been observed in the value of the muon bundle energy loss  $dE/dX_{1500}$  [174], related to variations in the multiplicity of high-energy muons in individual showers. A correction for this effect was devised and applied to the

$dE/dX_{1500}$  values in the composition analysis of Section 4.3.1. We will come back to this effect in more detail in Chapters 5 and 7.

A variety of other measurements related to atmospheric muons has been performed with IceCube. A broad characterization of the atmospheric muon flux measured in IceCube is given in Ref. [70].

One analysis that has been performed is a measurement of the flux of high-energy muons up to  $\sim$  PeV energy, based on large stochastic energy losses in the ice [70, 175]. These measurements indicate the presence of a prompt component (Chapter 2), albeit with limited significance. Improved analysis methods and larger statistics are expected to allow to strongly constrain the prompt muon flux [176].

Another analysis targets a rare class of air shower events, where two parallel tracks originating from the same shower are observed in IceCube [177, 178]. Single muons with a large lateral separation from the muon bundle in the shower originate predominantly from the first interaction of the cosmic ray in the atmosphere. The lateral separation is a direct probe of the transverse momentum of the parent hadron, and the measured distribution of lateral separations can therefore be used to constrain hadronic interaction models.

## 4.4 Simulated datasets

As discussed in Chapter 2, the interpretation of EAS measurements is dependent on simulations. The analyses of experimental data obtained with IceTop and IceCube presented in Chapters 6 and 7 will therefore also make use of various simulated cosmic-ray datasets, which are available as resources of the Cosmic Ray Working Group of the IceCube collaboration. We list here the properties of these simulations, which we will refer to as IceTop simulation. The simulations happen in two main steps: the simulation of the air shower development in the atmosphere with CORSIKA (Section 2.4.1), and the simulation of the detector response (Section 3.5). A summary of important properties of the different datasets is given in Table 4.3.

### 4.4.1 CORSIKA simulation specifics

EAS simulations for IceTop analysis used in this work were produced with CORSIKA version 7.3700. For hadronic interactions below 80 GeV, FLUKA2011.2b was used. The high-energy hadronic model used for most simulation is Sibyll 2.1. More limited sets are also available for EPOS-LHC and QGSJet-II.04.

**Table 4.3:** Overview of the cosmic-ray simulation datasets used in this work. All included datasets are simulated according to the IC86.2012 season. The first column gives the internal IceCube dataset identifier.

ID	HE model	Primary	$\log_{10} E_0/\text{GeV}$	Zenith	# Showers	Thinned
12360	Sibyll 2.1	$p$	5.0–8.0	$0^\circ\text{--}65^\circ$	20000	No
12630	Sibyll 2.1	He	5.0–8.0	$0^\circ\text{--}65^\circ$	20000	No
12631	Sibyll 2.1	O	5.0–8.0	$0^\circ\text{--}65^\circ$	20000	No
12362	Sibyll 2.1	Fe	5.0–8.0	$0^\circ\text{--}65^\circ$	20000	No
20143	Sibyll 2.1	$p$	7.0–9.6	$0^\circ\text{--}65^\circ$	24000	Yes
20145	Sibyll 2.1	He	7.0–9.6	$0^\circ\text{--}65^\circ$	24000	Yes
20146	Sibyll 2.1	O	7.0–9.6	$0^\circ\text{--}65^\circ$	24000	Yes
20144	Sibyll 2.1	Fe	7.0–9.6	$0^\circ\text{--}65^\circ$	24000	Yes
12634	EPOS-LHC	$p$	5.0–8.0	$0^\circ\text{--}65^\circ$	6000	No
12635	EPOS-LHC	Fe	5.0–8.0	$0^\circ\text{--}65^\circ$	6000	No
20196	EPOS-LHC	$p$	5.0–8.0	$0^\circ\text{--}40^\circ$	6000	No
20197	EPOS-LHC	Fe	5.0–8.0	$0^\circ\text{--}40^\circ$	6000	No
12636	QGSJet-II.04	$p$	5.0–8.0	$0^\circ\text{--}65^\circ$	6000	No
12637	QGSJet-II.04	Fe	5.0–8.0	$0^\circ\text{--}65^\circ$	6000	No
20194	QGSJet-II.04	$p$	5.0–8.0	$0^\circ\text{--}40^\circ$	6000	No
20195	QGSJet-II.04	Fe	5.0–8.0	$0^\circ\text{--}40^\circ$	6000	No

Simulations are produced using four primary types, as it is too computationally intensive to produce all possible primaries. IceTop simulation uses proton, helium, oxygen, and iron. These elements are roughly equidistant in  $\ln A$  and should be thought of as mass groups, including their neighboring, less abundant, elements with similar shower properties.

The particles in the shower are simulated down to a certain cutoff energy, below which they are discarded. Cutoff values are chosen sufficiently low so that particles at those energies cannot make a significant signal contribution. For muons and hadrons, it is 0.05 GeV. For electrons and photons, they are 0.01 GeV and 0.002 GeV respectively.

Datasets are produced with two different energy ranges. One goes from 0.1 to 100 PeV, or 5 to 8 in  $\log_{10}(E_0 / \text{GeV})$ , with full simulation of all shower particles. Higher-energy datasets go from 10 PeV to 4 EeV, or 7 to 9.6 in  $\log_{10}(E_0 / \text{GeV})$ , and use thinning to keep the runtime in check. For each dataset, a fixed number of showers is produced per bin of width 0.1 in  $\log_{10} E_0$ , distributed according to  $E^{-1}$ . This is reweighted to a realistic flux model during analysis. The overlap between the energy regions allows to verify the thinning and unthinning procedures.

As discussed in Section 2.4.1, the shower thinning is done by keeping only one particle out of all secondary particles originating from a particle interaction with energies below a threshold energy  $E_{\text{thin}} = \epsilon \times E_0$ , where  $\epsilon$  is the thinning level and  $E_0$  the primary cosmic-ray energy. The simulations use  $\epsilon = 10^{-6}$  for primary energies up to  $E_0 = 10^{8.4}$  GeV. Above this, the thinning level is set to  $\epsilon = 273 \text{ GeV}/E_0$ , so that muons which can propagate toward IceCube are never assigned thinning weights. To reduce the fluctuations induced by thinning, no more thinning is applied to daughters of particles if the weight exceeds  $w_{\text{max}} = \epsilon \times E_0$ .

The zenith and azimuth angles are generated corresponding to an isotropic distribution of primary cosmic rays. For the azimuth, this corresponds to sampling from a uniform distribution between 0 and  $2\pi$ . The distribution of the zenith angle  $\theta$  of events is proportional to  $\sin \theta \cos \theta$ , as a result of the  $\sin \theta$  factor in the solid angle element in spherical coordinates, and the projection onto a flat detector.

The atmospheric model used in the simulations is a custom model describing a realistic South Pole atmosphere for the month April. This was chosen as it is close to the average of the South Pole atmosphere over the entire year. The exact model was derived based on atmospheric temperature measurements between 2007 and 2011 in Ref. [157]. CORSIKA has included it as a standard atmosphere since version 7.7400. The relation between depth and altitude for the model was included in Fig. 2.1. The observation level is set at 2837 m above sea level, which is at least 1 m above the maximum snow level on the

**Table 4.4:** Resampling radius used in the simulation as a function of the primary cosmic-ray energy.

$\log_{10} E/\text{GeV}$	Radius (m)
5.0–6.0	800
6.0–7.0	1100
7.0–8.0	1700
8.0–9.0	2600
9.0–9.6	3800

tanks<sup>7</sup>. From the observation level, the particles are further simulated through the snow and tanks with Geant4. The observation level corresponds to an atmospheric depth of 698 g/cm<sup>2</sup>.

#### 4.4.2 Detector simulation specifics

A large fraction of the simulated CORSIKA showers, when placed at a certain point on the detector, would not cause a trigger nor make it to the final analysis level. To save on computational resources, each shower is *resampled* 100 times. This is done by randomly placing the shower core inside a circle around the IceTop array. The radius is chosen so that it is at least the distance up to which a shower of a certain energy could still trigger the detector. The radius thus depends on the shower energy, which determines the typical size of the footprint of the shower. The values are given in Table 4.4. The number of resamples is chosen so that every shower remains on average only once or a handful of times in the final sample after typical quality cuts. This limits possible artificial fluctuations due to the resampling procedure, while still making effective use of the generated showers.

For shower simulations that used thinning, the particles need to be "dethinned" in some way before the detector response is simulated. This is done similarly to the dethinning algorithm described in Ref. [179]. In essence, it works by taking a weighted particle which hits a certain sampling area around a tank and generating a number of clones distributed uniformly over the sampling area, each with the same energy as the parent particle. More information on the calculation of the size of the sampling area and how it was optimized can be found in Refs. [180] and [181].

To account for the effect of snow accumulation on the IceTop tanks, the detector is simulated in different epochs. These are chosen to correspond to different data-taking

<sup>7</sup>For dataset 12360, the observation level was set at 2834 m, which is below the snow level for some tanks. This is remedied by shifting the particles up artificially by 3 m before the detector simulation.

seasons. The simulation used in this work corresponds to the data-taking period referred to as IC86.2012, which ran from May 15th, 2012 to May 2nd, 2013<sup>8</sup>. The snow heights used in the simulations are the ones measured in-situ in October 2012, as shown in Fig. 3.7. They roughly correspond to the average of the season.

The ice model used in the simulations used in this work is the SPICEMIE model [140].

---

<sup>8</sup>At the time of writing, this is the only season for which simulation sets are available using different hadronic interaction models.



# Chapter 5

## High-energy muon production profiles

*This chapter is dedicated to the memory of Thomas K. Gaisser.*

In Chapters 3 and 4, the fact that IceCube can detect high-energy muons deep in the Antarctic ice was introduced. Before performing an analysis on actual IceCube data, as in Chapters 6 and 7, we discuss here in a more general way the production of such muons in the atmosphere. We present a parameterization based on air-shower simulations that describes longitudinal profiles of muon production above a certain threshold energy, given a primary cosmic ray and an atmospheric temperature profile. This parameterization can be used to make estimates of the number of muons measured in underground detectors, and how this depends on the atmospheric conditions.

The work presented in this chapter was developed in close collaboration with Tom Gaisser and was published in Ref. [182]<sup>1</sup>. In this chapter, we will focus on the part of the work most relevant to IceCube.

A brief introduction on the topic of underground measurements of high-energy muons and how the idea for this work emerged is given in Section 5.1. Following that, the derivation of the muon production profile parameterization is given Section 5.2. In Section 5.3 some example calculations relevant to IceCube are discussed, followed by some concluding thoughts in Section 5.4.

### 5.1 Introduction

The yield of high-energy muons in air showers induced by cosmic rays interacting near the top of the atmosphere is relevant for understanding event rates and properties of muon

---

<sup>1</sup>Presentations related to this work were shown at the ICRC2021 conference; see Refs. [183] and [184].

bundles in underground detectors. The so-called Elbert formula, given in its original form in Eq. (2.23), has been used to estimate the average multiplicity of muons  $\langle N_\mu \rangle$  above a certain energy  $E_\mu$  in an air shower initiated by a cosmic ray with energy  $E_0$ , mass number  $A$ , and zenith angle  $\theta$ . We repeat it here in its more general form,

$$\langle N_\mu(E_0, A, \theta, > E_\mu) \rangle \approx A \times \frac{K}{E_\mu \cos \theta} \left( \frac{E_0}{A E_\mu} \right)^{\alpha_1} \left( 1 - \frac{A E_\mu}{E_0} \right)^{\alpha_2}, \quad (5.1)$$

where the normalization constant  $K$  and exponents  $\alpha_1$  and  $\alpha_2$  can be derived from simulations. As described in Section 3.1.2, the threshold energy for a certain detector can be derived from the energy loss of muons in matter and the mass overburden of the detector.

While the Elbert formula is useful for quick estimates, it does not include any information about the production height of the muons or the dependence on atmospheric conditions. The muon production depends on the density of the atmosphere through the competition between decay and re-interaction of the parent mesons, as can be understood from the decay and interaction lengths introduced in Eqs. (2.7) and (2.8). We therefore introduce in this chapter a generalization of the Elbert formula that describes the production of high-energy muons as a function of depth along the trajectory of the primary cosmic ray, including an explicit dependence on the atmospheric temperature. The integral of this profile is the mean number of muons per shower, similar to the Elbert formula. This topic was a shared interest between Tom Gaisser and me, which we approached from two different but complementary analysis goals.

My personal interest for studying the longitudinal muon production profiles lies in the application to coincident measurements between an underground and a surface detector, for example IceTop and IceCube. This was especially inspired by the study of seasonal variations of the muon bundle energy loss in IceCube in Ref. [174], in the context of the composition analysis described in Section 4.3.1. We will return to this analysis and perform related calculations as example applications in Section 5.3.

Tom's interest mainly came from understanding the seasonal variation of the IceCube SMT8 trigger rate, which is dominated by (single) atmospheric muons, as introduced in Section 4.3.4. The standard approach to seasonal variations in muon rates is an analytic approach, using formulas for the inclusive production spectrum of muons integrated over the primary spectrum [169]. Using the parameterization of muon production profiles and integrating over the primary spectrum offers an alternative to the analytic approach. It also allows one to consider the dependence on muon multiplicity and production altitude. For examples of such applications, we refer to Refs. [182] and [184].

## 5.2 Parameterization of muon production profiles

This section describes in detail how muon production profiles were derived from air-shower simulations and how they are fitted with a physically motivated formula. The behavior of the fit parameters is then parametrized in terms of the primary cosmic-ray energy, mass, and the muon energy threshold.

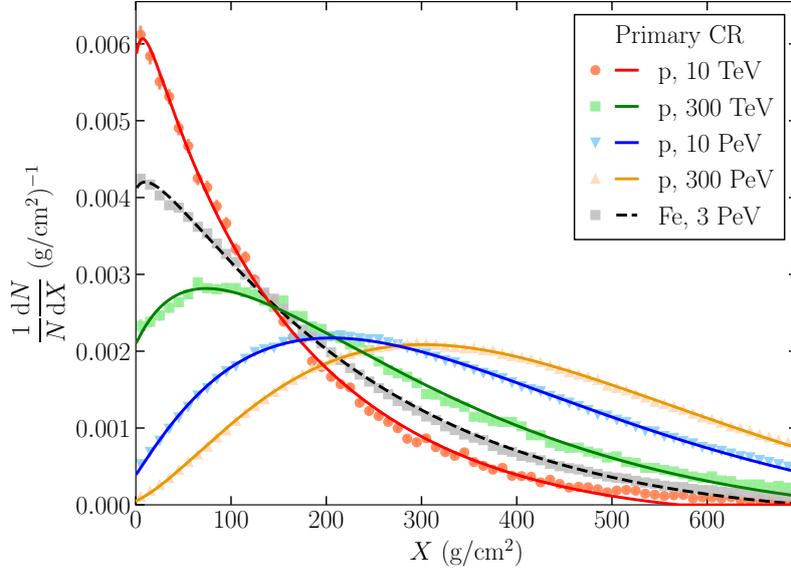
### 5.2.1 Simulations

The production of muons above a certain energy threshold, differential in slant depth throughout the atmosphere along the shower axis, is referred to as the longitudinal production profile. The idea is to perform a large number of air-shower simulations to obtain the average muon production profile for primary cosmic rays with energy  $E_0$ , mass number  $A$ , zenith angle  $\theta$ , and for muons with energy above  $E_\mu$ .

We have performed dedicated simulations with CORSIKA v7.7100 using Sibyll 2.3c [111] as the high-energy interaction model. The observation level and atmosphere used correspond to these of IceTop, i.e. 2835 m and an atmospheric density profile describing the average South Pole atmosphere in April between 2007 and 2011, as described in Section 4.4.1, although any other atmosphere could have been used.

Simulations were performed for muon threshold energies  $E_\mu$  between 300 GeV and 1 TeV over a large range of primary energies. As we are only interested in the number of muons above this energy, particles below  $E_\mu$  were not tracked further in the simulation, significantly speeding up the process. From the longitudinal CORSIKA output, which gives the number of particles of different types at regular intervals of slant depth, we derived the muon production profile as the increase in number of muons per step of slant depth. Because the muon energies are high, muon decay can be ignored. For the derivation of the parameterization, the simulations were limited to vertical proton showers. The performance for heavier primaries and inclined showers is discussed in Section 5.2.4.

The number of showers simulated was chosen sufficiently high to obtain an average with small statistical errors, ranging between  $10^3$  and  $10^6$ . This depends on the ratio of the primary energy to the muon threshold energy, as lower energy showers have a lower muon content and need more statistics to average out fluctuations. Examples of the obtained muon production profiles for  $E_\mu > 500$  GeV are given in Fig. 5.1, together with the fits described below.



**Fig. 5.1:** Normalized muon production profiles for vertical showers with  $E_\mu > 500$  GeV. The markers show the average profiles obtained from CORSIKA simulations. The lines are the result of fitting Eq. (5.2) to the simulated profiles.

An alternative approach that was not pursued here would be to estimate the muon production profiles from solutions of the cascade equations Eq. (2.6), using for example the MCEq toolkit [185].

### 5.2.2 Formula

The muon production profiles obtained from simulation, which give the distribution of slant depths over which muons above a certain energy threshold are produced in an air shower, are fitted with a physically motivated formula. The basic idea is to interpret the derivative of the Gaisser-Hillas function Eq. (2.24) as the rate of production of charged mesons per  $dX(\text{g}/\text{cm}^2)$  along the trajectory of a primary cosmic ray and then multiply by appropriate factors describing the decay of the mesons to muons, including a dependence on the atmospheric temperature  $T$  at the slant depth  $X$  (equivalent to the density). The

formula for the production of muons  $dN$  per slant depth interval  $dX$  is given by

$$\left\langle \frac{dN}{dX}(X, T, E_0, A, \theta, > E_\mu) \right\rangle = \quad (5.2)$$

$$N_{\max} \times \exp((X_{\max} - X)/\lambda) \times \left( \frac{X_0 - X}{X_0 - X_{\max}} \right)^{(X_{\max} - X_0)/\lambda} \times \frac{X_{\max} - X}{\lambda(X - X_0)} \quad (5.2a)$$

$$\times \left[ f_\pi \times \frac{r_\pi \lambda_\pi \epsilon_\pi}{1 + \frac{r_\pi \lambda_\pi \epsilon_\pi}{f E_\mu \cos(\theta) X}} + f_K \times \frac{r_K \lambda_K \epsilon_K}{1 + \frac{r_K \lambda_K \epsilon_K}{f E_\mu \cos(\theta) X}} \right] \times \frac{1}{f E_\mu \cos(\theta) X} \quad (5.2b)$$

$$\times \left( 1 - \frac{A E_\mu}{E_0} \right)^{\alpha_2}. \quad (5.2c)$$

Below, we will explain in detail the different factors that were combined in the formula.

### Eq. (5.2a): Production of charged mesons

The first line of Eq. (5.2) is the derivative of the Gaisser-Hillas function. This function gives a good phenomenological description of longitudinal shower profiles. It was already introduced in Section 2.3.2 in the context of fluorescence telescopes, as it is frequently used to fit the longitudinal shower profiles they observe. In this context,  $N_{\max}$  is the fitted shower size at the slant depth  $X_{\max}$  where the shower reaches its maximum, while  $\lambda$  and  $X_0$  are considered free shape parameters.

In the formula here, we use the derivative of this function with respect to  $X$ , to describe the production of charged mesons which could contribute to the production of muons through decay. The free parameters in the fit are again  $N_{\max}$ ,  $X_{\max}$ ,  $\lambda$ , and  $X_0$ . As we will see later on, the values of the fitted parameters will numerically be quite different from those of the original Gaisser-Hillas formula for air showers dominated by the EM cascade, and should not be interpreted as more than shape parameters.

### Eq. (5.2b): Decay to muons

In the second term, we multiply by the relative probability of decay of a meson to a muon relative to the total rate of decay and re-interaction. We consider two channels for muon production, namely decay of charged pions and kaons  $\pi^\pm/K^\pm \rightarrow \mu + \nu_\mu$ , with branching ratios of 100% and 63.5% respectively. For pions with decay length  $d_\pi$  and interaction length  $\lambda_\pi$ , the decay fraction is given by

$$\frac{1/d_\pi}{1/d_\pi + 1/\lambda_\pi}. \quad (5.3)$$

From the discussion on the atmosphere in Section 2.1, we find that the atmospheric density, assuming zenith angles below  $65^\circ$  and a simple exponential atmosphere, can be written as

$$\rho = \frac{-dX_v}{dh} = \frac{X_v}{h_0} = \frac{X \cos \theta}{h_0}, \quad (5.4)$$

with the scale height  $h_0 = RT/Mg$ . Combining this with the definition of the decay length, Eq. (2.8), results in

$$\frac{1}{d_\pi} = \frac{m_\pi c^2}{E_\pi X \cos \theta} \frac{1}{c\tau_\pi} \frac{RT}{Mg} \equiv \frac{\epsilon_\pi}{E_\pi X \cos \theta}, \quad (5.5)$$

where  $m_\pi$  and  $\tau_\pi$  are the mass and lifetime of the pion. The pion critical energy  $\epsilon_\pi$  is given by

$$\epsilon_\pi = \frac{m_\pi c^2}{c\tau_\pi} \frac{RT}{Mg} \approx 0.524 \text{ GeV} \times \frac{T}{\text{K}}, \quad (5.6)$$

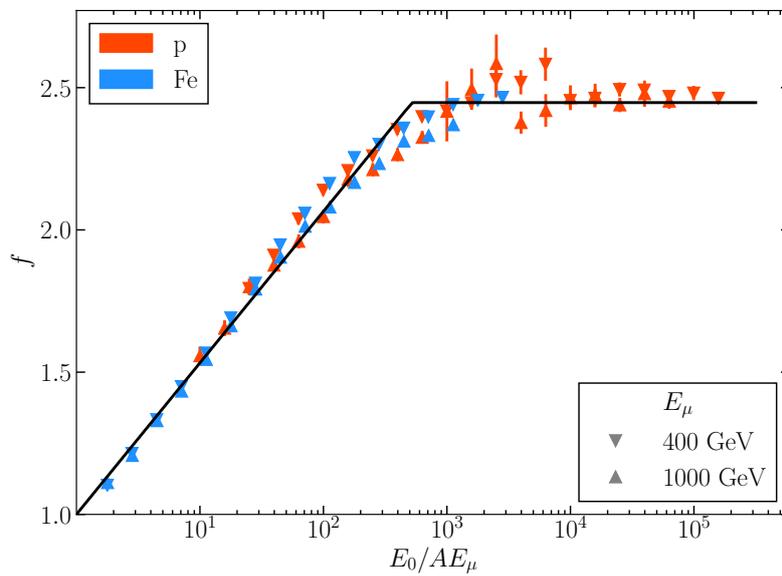
with  $T$  the atmospheric temperature in Kelvin [30, p. 122]. In pion decay, the muon carries an average energy of  $r_\pi \times E_\pi$ , with  $r_\pi \approx 0.79$  [30, p. 129]. The average energy of muons above the energy threshold  $E_\mu$  in an air shower can be written as  $f \times E_\mu$  and has been derived from simulations. The behavior of  $f$  as a function of the ratio of the energy per nucleon  $E_0/A$  and the muon energy threshold  $E_\mu$  is shown in Fig. 5.2. It increases from a value of 1 at the production threshold to an approximately constant value  $f \approx 2.45$  at higher energies, and we approximate it with a simple piecewise linear function (see Table 5.1). With this, we can rewrite Eq. (5.3) as a function of the muon threshold energy  $E_\mu$  as

$$\frac{1/d_\pi}{1/d_\pi + 1/\lambda_\pi} = \frac{r_\pi \lambda_\pi \epsilon_\pi}{f E_\mu \cos(\theta) X} \times \frac{1}{1 + \frac{r_\pi \lambda_\pi \epsilon_\pi}{f E_\mu \cos(\theta) X}}. \quad (5.7)$$

As the interaction length varies only slowly with energy, driven by the inelastic cross section, we use a constant value of  $111 \text{ g/cm}^2$  for  $\lambda_\pi$  [30, p. 123].

An analogous term to Eq. (5.7) can be obtained for decay of parent kaons to muons with some substitutions. The average energy of muons in kaon decay is  $r_K \approx 0.52$ . The critical energy is larger than the pion critical energy by a factor of 7.45 because of the larger mass and shorter decay length of kaons. For the interaction length, a value of  $122 \text{ g/cm}^2$  is used.

The pion and kaon terms are combined weighted by prefactors  $f_\pi$  and  $f_K$ . They are derived from the average momentum fraction carried by charged pions and kaons in proton-air interactions. Their numerical values are based on Fig. 5.2 of Ref. [30],



**Fig. 5.2:** Ratio  $f$  of the mean energy of muons above the threshold in an air shower over the muon energy threshold  $E_\mu$  itself. The markers show the values obtained from simulations of vertical proton and iron showers. The black line represents the approximation of  $f$  used in Eq. (5.2), and the parameters describing it are included in Table 5.1.

where the so-called *spectrum-weighted moments* for primary protons interacting with air at 1 TeV are plotted. These spectrum-weighted moments are quantities that are introduced when analytically solving the cascade equations of Eq. (2.6) for a primary power law spectrum, and they are a way to characterize the production of a certain type of particles in a cascade with a single number. While their exact definition is outside the scope of this work, it is important to highlight that evaluating the spectrum weighted moments  $Z_{ji}$  at a spectral index  $\gamma = 1$ , simply gives the average fraction of the interaction energy that goes into particles of type  $i$  in interactions of particles of type  $j$ . From the aforementioned Figure, these fractions for charged pions and charged kaons in p-air interactions are  $Z_{N,\pi} = 0.29$  and  $Z_{N,K} = 0.040$  respectively. Combining these with the branching fractions for decay to muons, we obtain  $f_\pi = 0.29/(0.29 + 0.635 \times 0.04)$  and  $f_K = 1 - f_\pi$ , i.e.

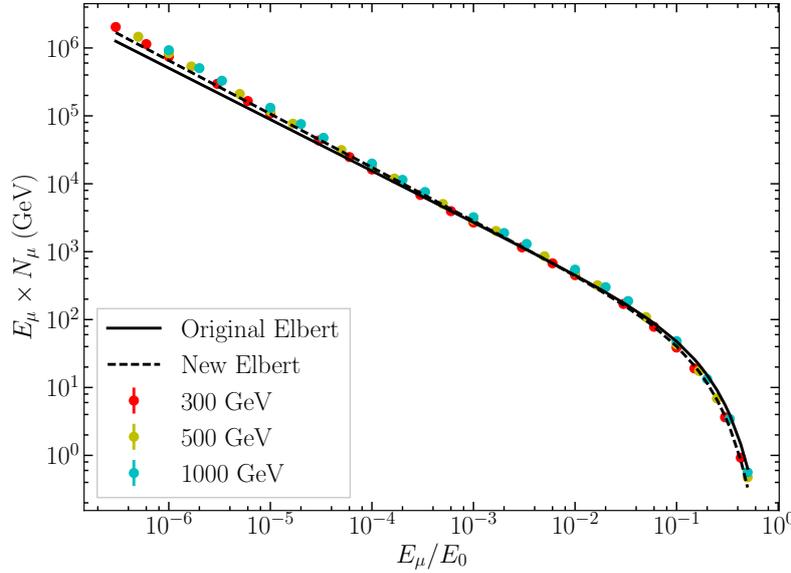
$$f_\pi = 0.92, \quad f_K = 0.08. \quad (5.8)$$

The total decay factor combining terms for charged pions and kaons is shown in Eq. (5.2b). This factor gives the formula its dependence on the atmospheric temperature through the critical energies  $\epsilon_\pi$  and  $\epsilon_K$  and is thus the feature that allows us to estimate seasonal variations in the muon production. As  $f_\pi$  and  $f_K$  are calculated based on a fixed interaction energy, it may be interesting to explore other ways of combining the pion and kaon decay factors, to check if they offer any improvement in the description of the temperature dependence in the final parameterization. However, as we will show below, the current approach already describes the relative variations in muon production in reasonable agreement with simulations.

### Eq. (5.2c): Threshold factor

The last factor of the formula can be recognized as the threshold factor of the Elbert formula, Eq. (5.1). It describes the threshold behavior in the muon production in the regime where the muon energy threshold is close to the energy per nucleon of the primary nucleus. In this regime, the typical power-law behavior between the muon multiplicity and the energy per nucleon does not hold.

The normalization constant  $K$  and exponents  $\alpha_1$  and  $\alpha_2$  of the Elbert formula depend on the hadronic physics in the shower. It is therefore customary to derive the exponents from simulation. Here, we fitted the Elbert formula to the total number of muons obtained in the simulations performed with Sibyll 2.3c, using different muon energies and a large range of primary energies. This is illustrated in Fig. 5.3, where we show the



**Fig. 5.3:** Comparison of the original Elbert formula Eq. (5.1) and an updated version as described in the text. The latter was fitted to muon multiplicity values obtained from vertical proton simulations at different values of  $E_\mu$ , as indicated by the markers.

simulations compared to both the original Elbert formula, Eq. (2.23), and our newly derived one. Values of  $K = 12.4$ ,  $\alpha_1 = 0.787$ , and  $\alpha_2 = 5.99$  are found in the fit.

The threshold factor of the newly derived Elbert formula, with exponent  $\alpha_2 = 5.99$ , is included in our formula Eq. (5.2) as is; the value of the exponent is fixed as part of the formula and is not a free parameter in the fits to the muon production profiles. Including this threshold factor will ensure the correct behavior of the total multiplicity predicted by the parameterization in the low-energy regime.

### 5.2.3 Fit parameters

The formula described in the previous section, Eq. (5.2), is used to fit the muon production profiles obtained from simulations. Some example fits are shown in Fig. 5.1. For these fits, the temperature in the decay factors was kept constant at a value of 217 K, a good approximation of the average South Pole atmosphere. While a temperature profile given by the simulated atmosphere may slightly improve the fits, this will have a minor impact on the final result.

A fit leads to values of four free parameters which, given a temperature profile or a constant temperature value, fully describe the shape of the muon profile:  $N_{\max}$ ,  $X_{\max}$ ,  $\lambda$ , and  $X_0$ . Repeating the fit procedure for a large number of muon and primary cosmic-ray

energies reveals that the optimal values of the parameters depend in leading order on  $E_0/AE_\mu$ . This is shown in Fig. 5.4 for simulations of vertical proton showers with  $E_\mu > 300, 400, 500, 700, 1000$  GeV. The behavior of these parameters is fitted with the following functions,

$$N_{\max} = c_i \times A \times \left( \frac{E_0}{AE_\mu} \right)^{p_i} \quad (5.9)$$

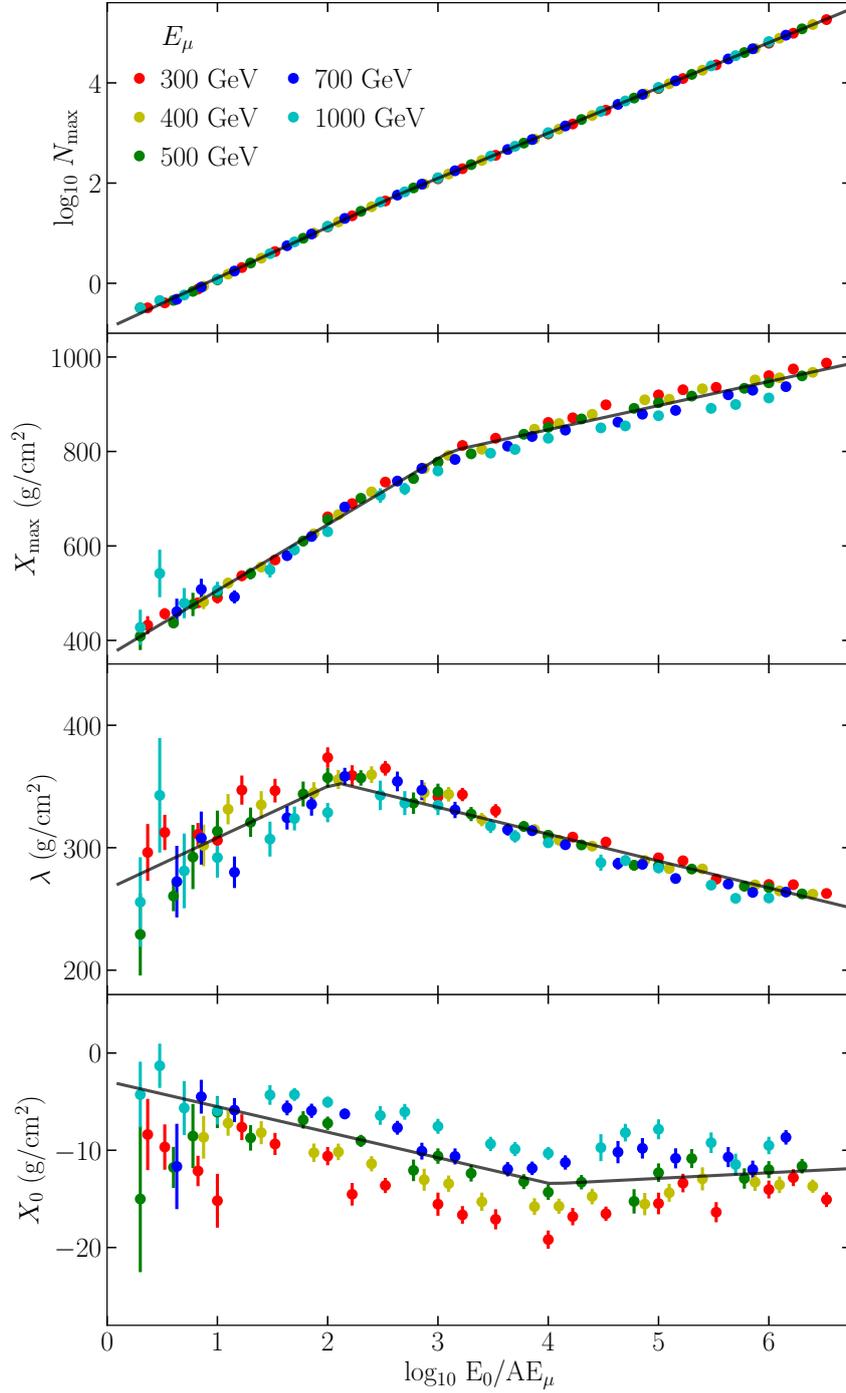
$$X_{\max}, \lambda, \text{ and } X_0 = a_i + b_i \times \log_{10} \left( \frac{E_0}{AE_\mu} \right),$$

where  $c_i, p_i, a_i,$  and  $b_i$  are defined for each parameter separately and have two regimes with a break at  $E_0/AE_\mu = 10^q$  with parameters  $i = 1$  below the break and  $i = 2$  above. The values obtained for these fit parameters are given in Table 5.1 and the fits are shown in Fig. 5.4. The table also includes the behavior of the  $f$ -factor.

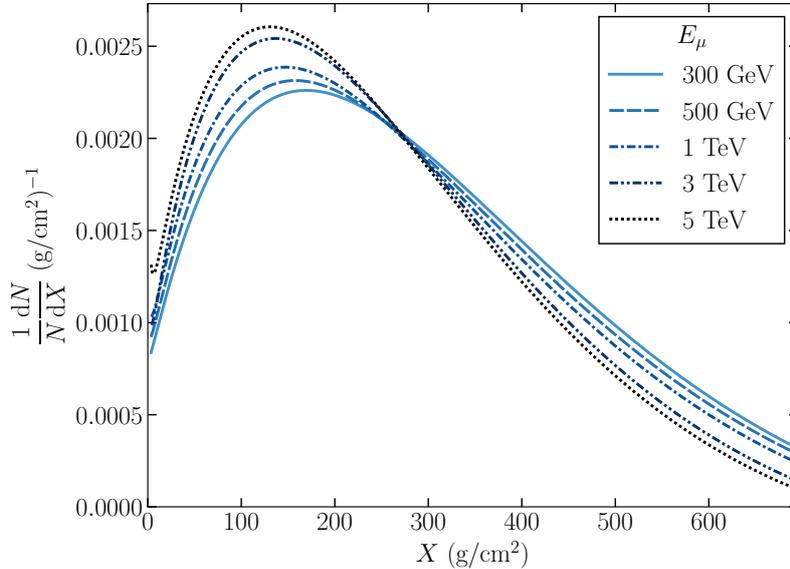
As noted in Ref. [186], the parameter  $X_0$  is often negative in fits and is to be considered simply as a shape parameter rather than the starting point of the cascade. As can be seen from Eq. (5.2a), depths  $X > X_{\max}$  give negative values. There is no further significant muon production at this point, so the muon profile is set to zero. It is also interesting to note that if we did not include a threshold factor in the formula, as given by Eq. (5.2c), the behavior of  $\log_{10} N_{\max}$  would deviate from a straight line at low energies and would show a suppression similar to the shape of the threshold factor.

**Table 5.1:** Parameter values for Eq. (5.9) for  $300 \text{ GeV} \lesssim E_\mu \lesssim 1 \text{ TeV}$ .

	$i$	$c_i$	$p_i$	$q$
$N_{\max}$	1	0.124	1.012	2.677
	2	0.244	0.902	
	$i$	$a_i$ (g/cm <sup>2</sup> )	$b_i$ (g/cm <sup>2</sup> )	$q$
$X_{\max}$	1	366.2	139.5	3.117
	2	642.2	51.0	
$\lambda$	1	266.0	42.1	2.074
	2	398.8	-21.9	
$X_0$	1	-2.9	-2.6	4.025
	2	-15.8	0.6	
$f$	1	1	0.53	2.72
	2	2.45	-	



**Fig. 5.4:** Optimal values of the four parameters in fits of the muon production spectrum Eq. (5.2) to simulations of vertical proton showers with different threshold muon energies  $E_\mu$  over a large range of primary energies. The black line shows the fits to these values given by Eq. (5.9) and Table 5.1.



**Fig. 5.5:** Normalized muon production profiles fitted to vertical proton simulations with primary energies chosen so that  $E_0/E_\mu = 10^4$  for different muon threshold energies  $E_\mu$ , showing that the shape of the profiles has a remaining dependence on  $E_\mu$ .

The scaling of the parameters with  $E_0/AE_\mu$  only holds approximately. This can already be observed from Fig. 5.4, where a remaining dependence on  $E_\mu$  is visible. This is further demonstrated in Fig. 5.5, where profiles fitted to proton simulations are shown with a constant  $E_0/E_\mu$  for different  $E_\mu$ . A difference in the shape of the profiles can be observed, with the peak of the production happening higher up in the atmospheres for higher  $E_\mu$ . As a result, the parameterization can only be used reliably in a limited range of  $E_\mu$ . As discussed in the next section, the parameterization presented in Table 5.1 works best around  $E_\mu \approx 500$  GeV, but worsens somewhat at lower and higher energies. For this reason, it is best to optimize the parameterization using simulations for a specific energy range for the application at hand. In Ref. [182], for example, we present calculations for the NOvA Near Detector [187] with parameters derived for a threshold energy of 50 GeV. The corresponding parameters are included in Appendix A. The parameterization with both sets of parameters has been made available online as a simple Python script<sup>2</sup>.

#### 5.2.4 Performance tests

The parameterization derived above is meant to be a useful tool for estimating the number of muons produced in a shower given a certain atmospheric temperature profile. As a

<sup>2</sup><https://github.com/verpoest/muon-profile-parameterization>

result of some approximations in the derivation, it is not expected to be a precision tool. To see what one can realistically expect from it, it is useful to check the performance of the parameterization compared to actual MC simulations. This may also provide ideas for how to improve this kind of parameterization in the future.

As described above, the parameterization is derived based on simulated vertical proton showers in a South Pole April atmosphere, over a large range of primary energies and for muon threshold energies  $E_\mu$  between 300 GeV and 1 TeV. From the fits in Fig. 5.4, we can expect it to perform best for  $E_\mu \approx 500$  GeV. In Fig. 5.6, a comparison is shown between muon production profiles derived from simulation, and the profiles derived from the parameterization with a constant temperature of 217 K as assumed during the profile fits. In general, there is a pretty good agreement between the two. For most primary energies, features like the normalization and the position of the peak are reproduced well. The worst performance in this case is seen for the case where  $E_0 = 500$  TeV. Likely, this is a result of the fact that this is in the energy range where the behavior of some parameters shows a break, which could be described better with a smooth function rather than the piecewise linear approximation used here.

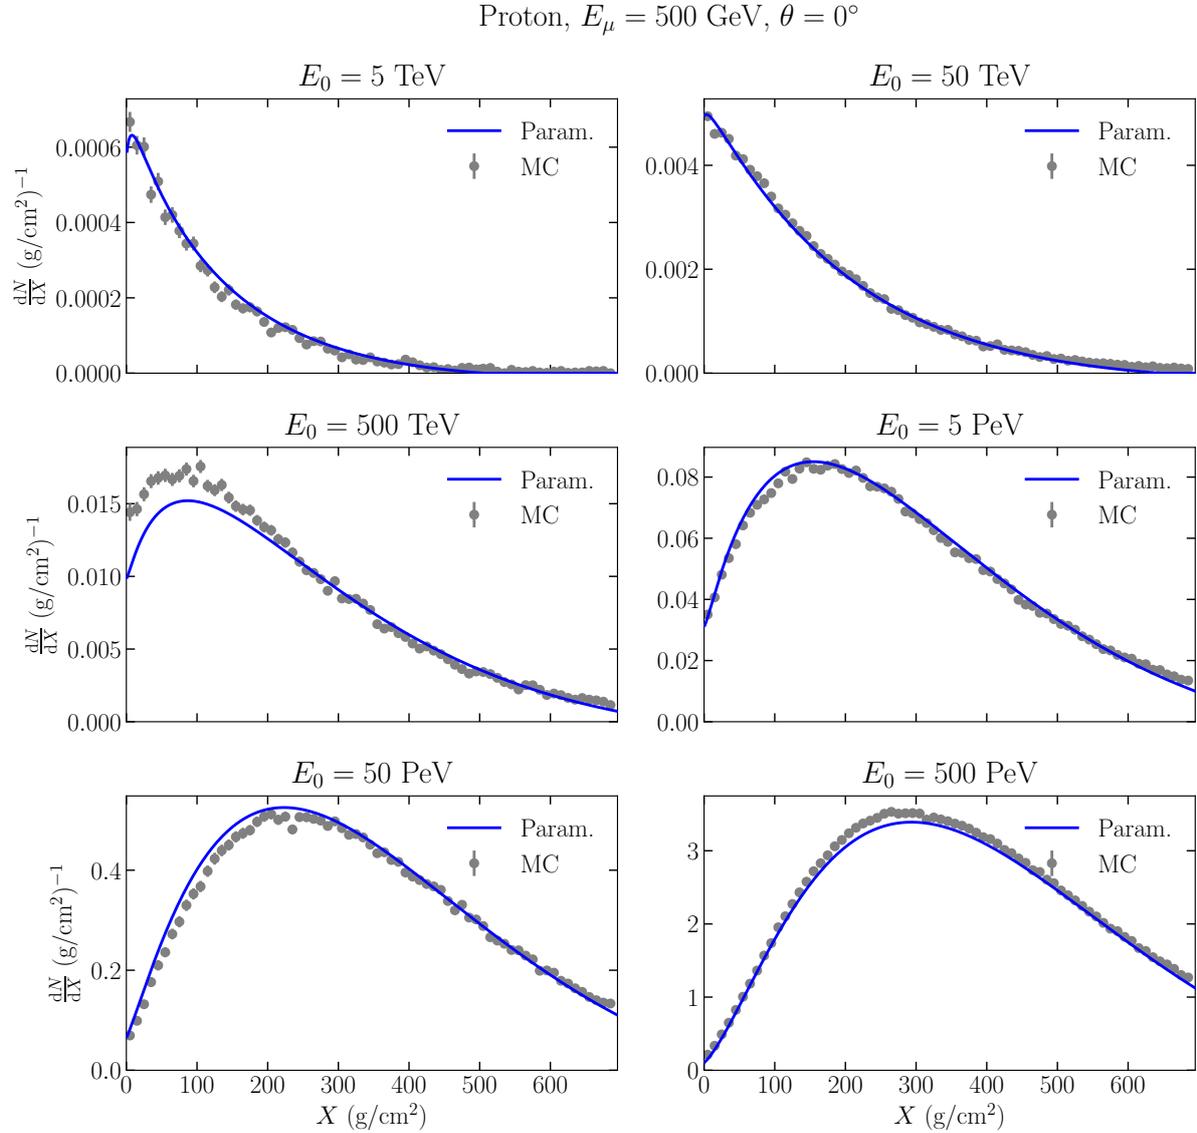
Below, we show a number of comparison plots considering different  $E_\mu$ , primaries, zenith angles, and atmospheric profiles, and discuss some takeaways from these plots.

### Dependence on $E_\mu$

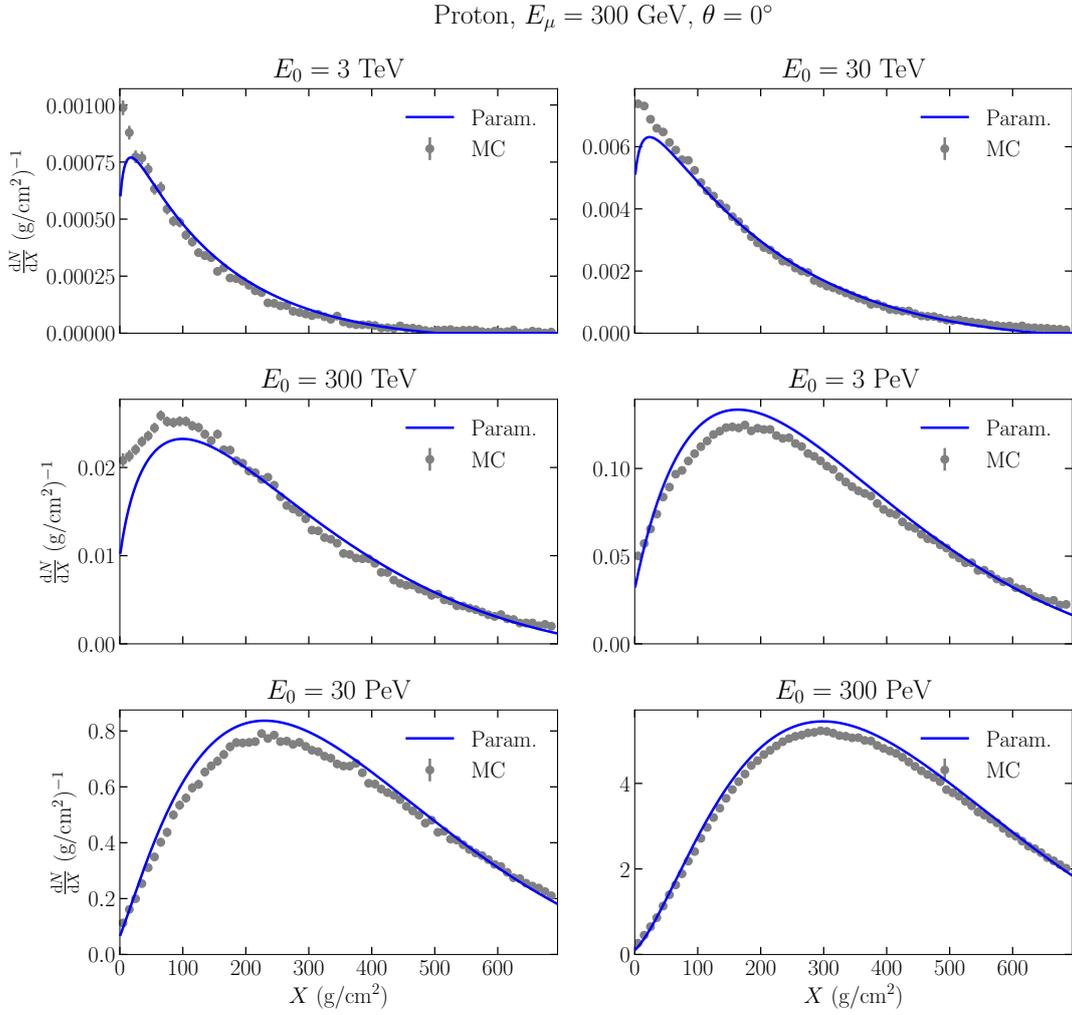
From the imperfect scaling of the fit parameters with  $E_0/AE_\mu$ , it is expected that the agreement between the parameterization and the simulated profiles becomes worse when moving away from the middle of the  $E_\mu$  range used in the fits. In Figs. 5.7 and 5.8, comparisons are shown for  $E_\mu = 300$  GeV and 1 TeV, analogously to Fig. 5.6. While the shape of the profile is in general still described reasonably well, one can see more systematic deviations from the true profiles. This affirms our previous statement that for calculations for muon energies in a different order of magnitude, it is recommended to re-optimize the parameters.

### Dependence on $A$

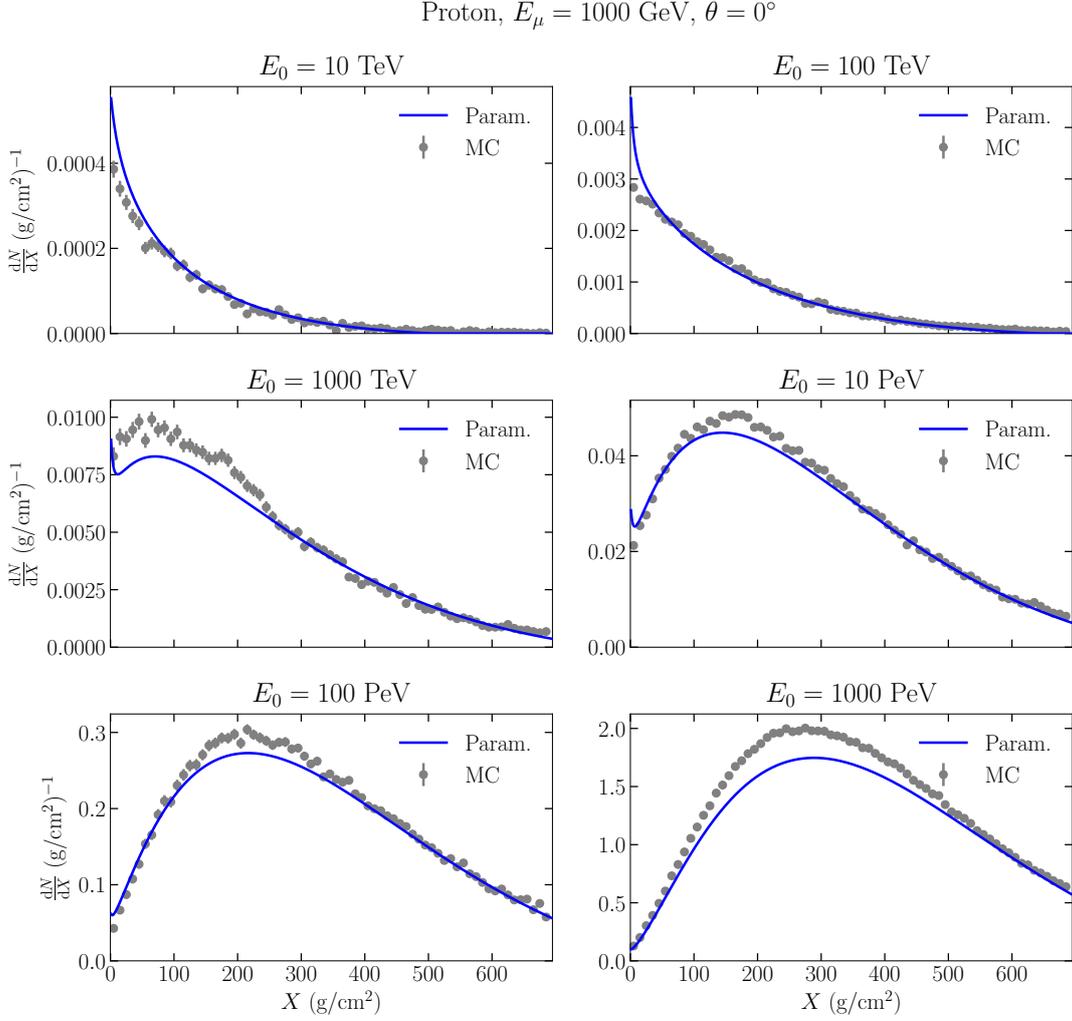
For reasons of simplicity, the parameterization derived in this chapter was fitted only to proton showers, assuming that it would also work for showers of heavier nuclei based on the superposition principle (see Chapter 2). In Fig. 5.9, we show a comparison of profiles derived from MC with iron primaries compared to the expectation of the parameterization. A good agreement, comparable to the one for proton primaries, is



**Fig. 5.6:** Muon production profiles obtained from simulations of vertical proton showers with  $E_\mu > 500 \text{ GeV}$  at various primary energies, compared to the profiles given by the parameterization of Eqs. 5.2, 5.9 and Table 5.1, assuming an isothermal atmosphere with  $T = 217 \text{ K}$ .



**Fig. 5.7:** Comparison of muon production profiles obtained from simulations and from the parameterization of Table 5.1 at different primary energies  $E_0$ . Considered are vertical proton showers with a muon threshold of 300 GeV. The simulated atmosphere is the April South Pole atmosphere. The parameterization is evaluated with a constant temperature of 217 K.



**Fig. 5.8:** Comparison of muon production profiles obtained from simulations and from the parameterization of Table 5.1 at different primary energies  $E_0$ . Considered are vertical proton showers with a muon threshold of 1 TeV. The simulated atmosphere is the April South Pole atmosphere. The parameterization is evaluated with a constant temperature of 217 K.

observed. We can therefore confidently use the parameterization for primary nuclei with different masses.

### Dependence on $\theta$

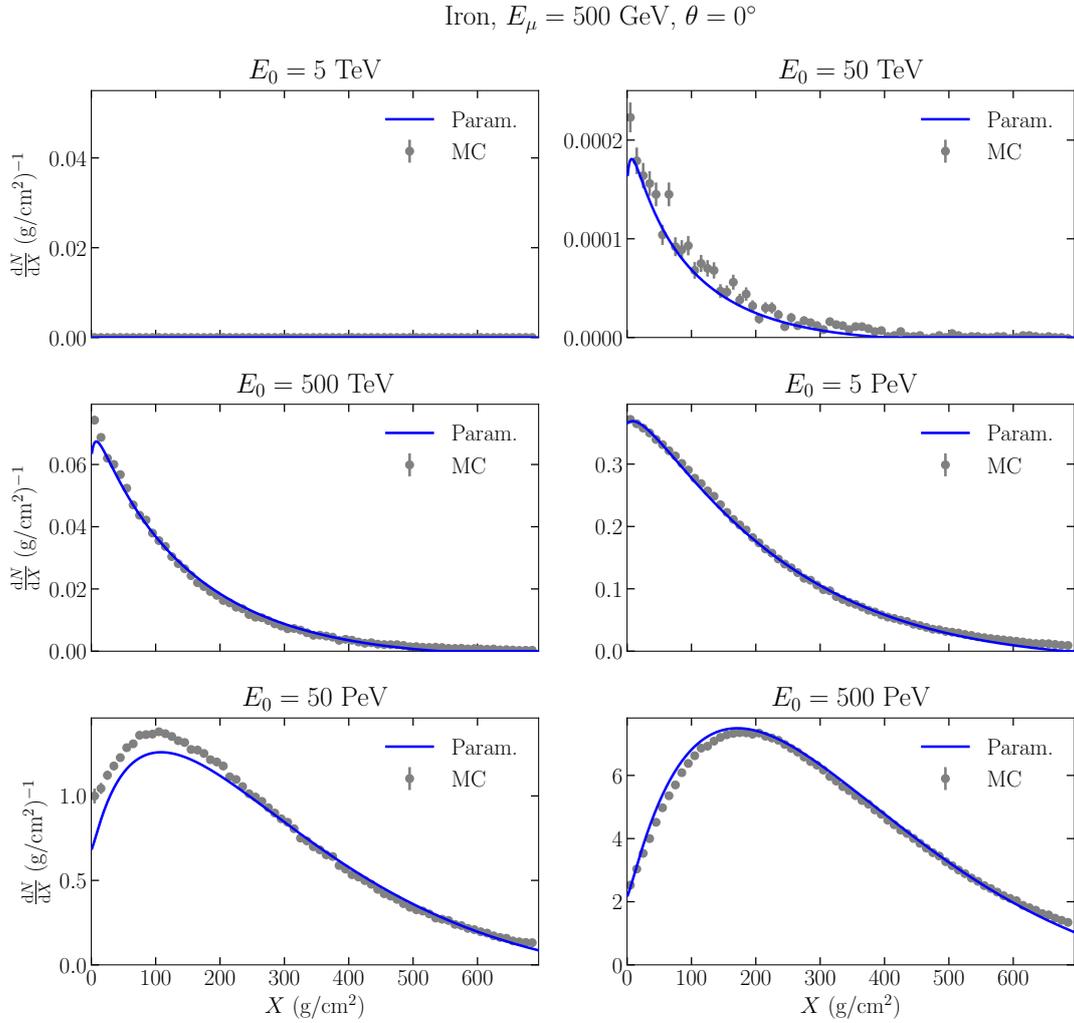
The parameterization was derived only from vertical showers. The production profiles are expected to look different for inclined showers, given the explicit  $\cos\theta$  terms in the formula of Eq. (5.2). In Figs. 5.10 and 5.11, results are shown for showers with zenith angles of  $30^\circ$  and  $60^\circ$  respectively. Good agreement is still observed for  $30^\circ$  degrees. For  $60^\circ$  however, the shape given by the parameterization deviates more from the truth. This is especially the case at large slant depth, where the production is underestimated by the parameterization. This is likely a result of the fact that the simulated profiles used in the fits do not go beyond  $700\text{ g/cm}^2$ , which could be remedied by simply lowering the observation level in the CORSIKA simulations.

### Dependence on atmospheric temperature

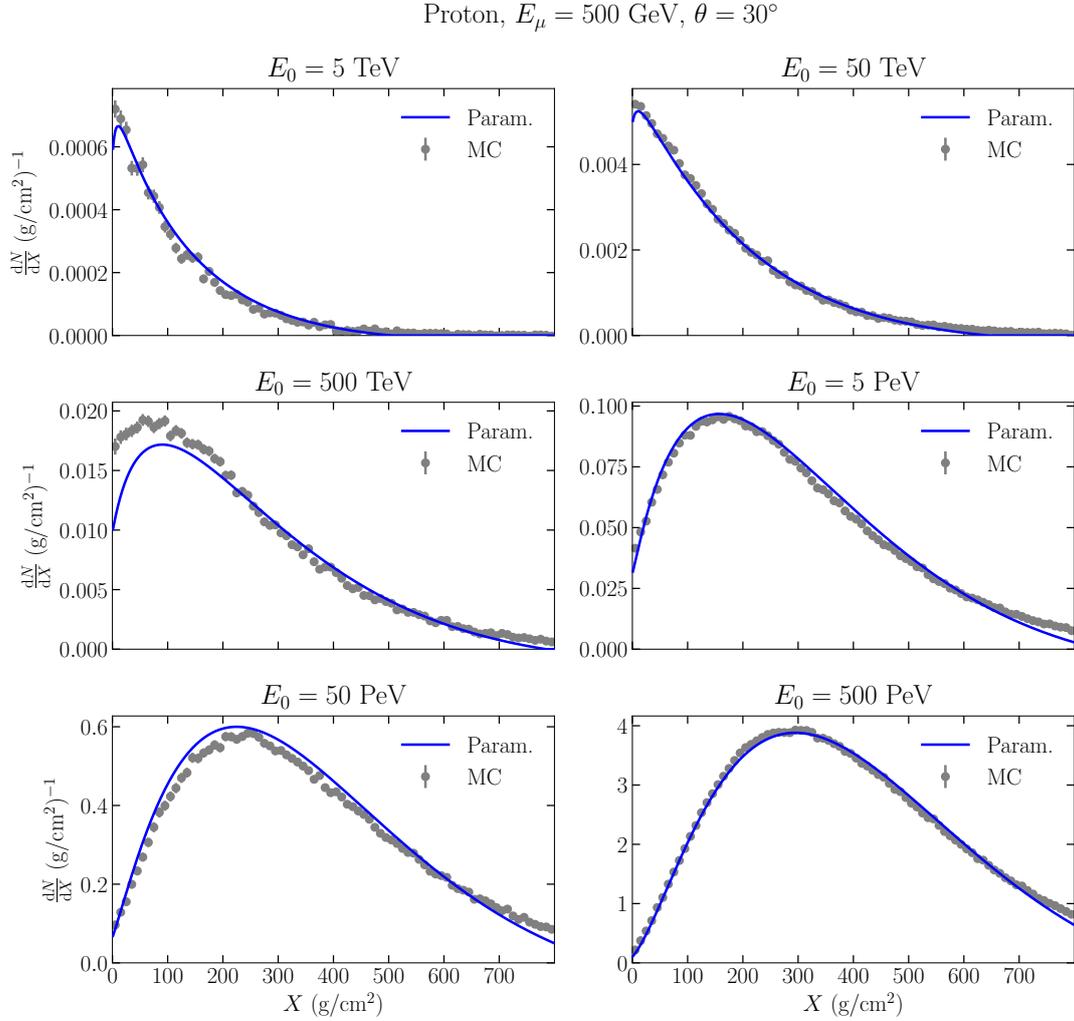
One of the important reasons for devising this parameterization, and an important difference to the Elbert formula, is the possibility of easily using it with different atmospheric profiles, allowing to estimate seasonal variations in the muon production. As the parameterization was fit to simulations using only one specific atmosphere, it needs to be tested if it accurately reproduces profiles for other atmospheres as well.

To test this, CORSIKA simulations with two different atmospheres were produced: an average January and August atmosphere, both for the South Pole, as given in Ref. [157]. The CORSIKA atmospheres are described by a five-layer model of the vertical atmospheric depth  $X_v$  as a function of altitude, as discussed in Section 2.4.1. To obtain an equivalent temperature profile that can be plugged into Eq. (5.2), we used the ideal gas law and the relation between depth and density described in Section 2.1.

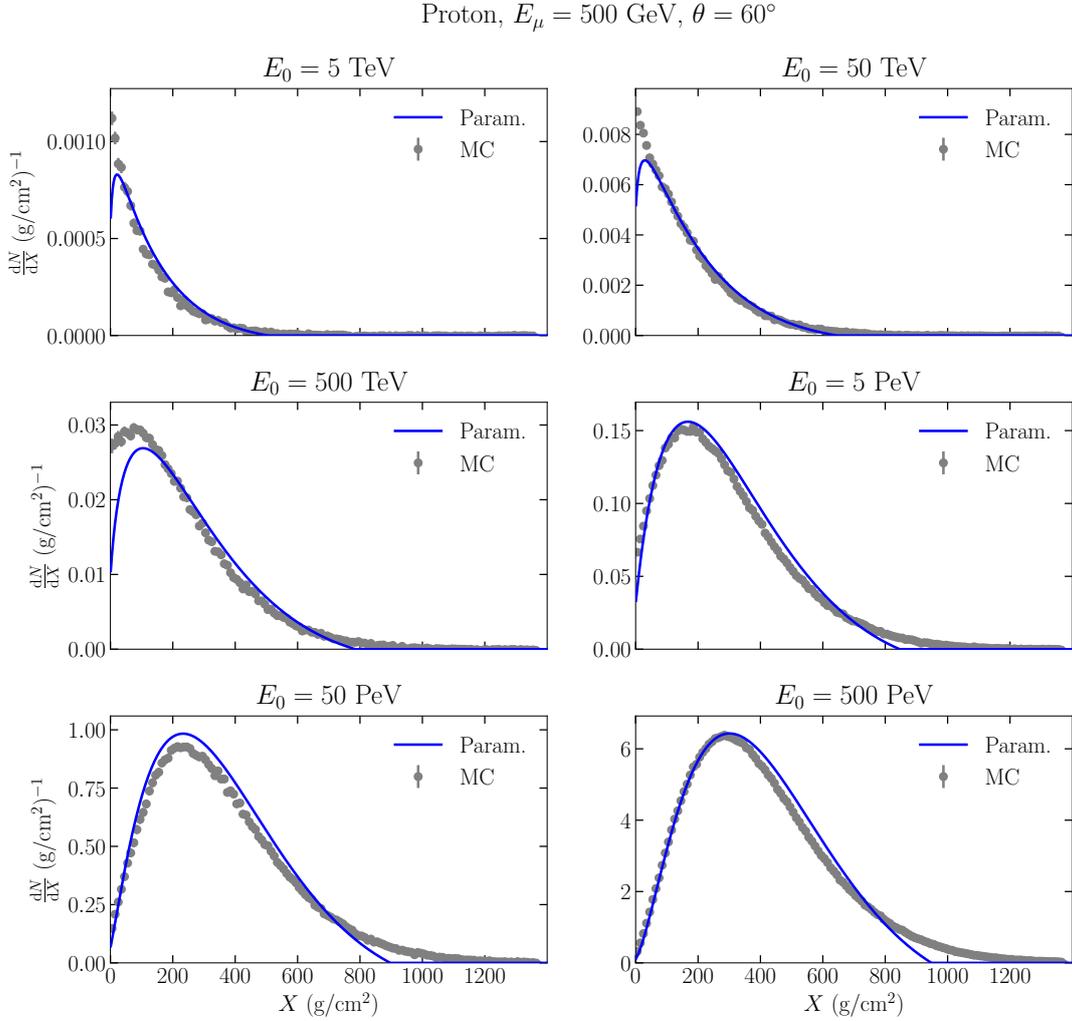
A comparison between the simulated muon production profiles and the profiles obtained from the parameterization is shown in Fig. 5.12, together with the total number of muons above  $E_\mu$  at the surface. First, there is a clear dependence on the atmosphere in the simulated profiles: in January, when the South Pole atmosphere is warmer, more muons are produced as a result of the larger decay probability, and vice versa for August. It can be seen that the profiles given by the parameterization show similar trends to the simulation. While the overlap with the simulated profiles is not equally good for all primary energies, the magnitude of the seasonal effect is described rather well. The total number of muons, obtained from the simulation or from integrating the profiles from the parameterization, agrees to within 10%. The relative variation between the months



**Fig. 5.9:** Comparison of muon production profiles obtained from simulations and from the parameterization of Table 5.1 at different primary energies  $E_0$ . Considered are vertical iron showers with a muon threshold of 500 GeV. The simulated atmosphere is the April South Pole atmosphere. The parameterization is evaluated with a constant temperature of 217 K.



**Fig. 5.10:** Comparison of muon production profiles obtained from simulations and from the parameterization of Table 5.1 at different primary energies  $E_0$ . Considered are proton showers with a zenith angle of  $30^\circ$  and a muon threshold of 500 GeV. The simulated atmosphere is the April South Pole atmosphere. The parameterization is evaluated with a constant temperature of 217 K.



**Fig. 5.11:** Comparison of muon production profiles obtained from simulations and from the parameterization of Table 5.1 at different primary energies  $E_0$ . Considered are proton showers with a zenith angle of  $60^\circ$  and a muon threshold of 500 GeV. The simulated atmosphere is the April South Pole atmosphere. The parameterization is evaluated with a constant temperature of 217 K.

is similar to the simulated ones, typically between 4-8%. Similar results were obtained from tests performed with the cascade-equation code `MCEq`.

Note that the profiles for April in Fig. 5.12 look slightly different compared to the ones shown in Fig. 5.6. This is because the former are calculated with an actual temperature profile corresponding to the simulated April atmosphere, rather than the constant value of 217 K. The slightly worse agreement in Fig. 5.12 between calculated and true indicates that taking the changing temperature profile into account during the fits from which the parameterization is derived may improve its performance.

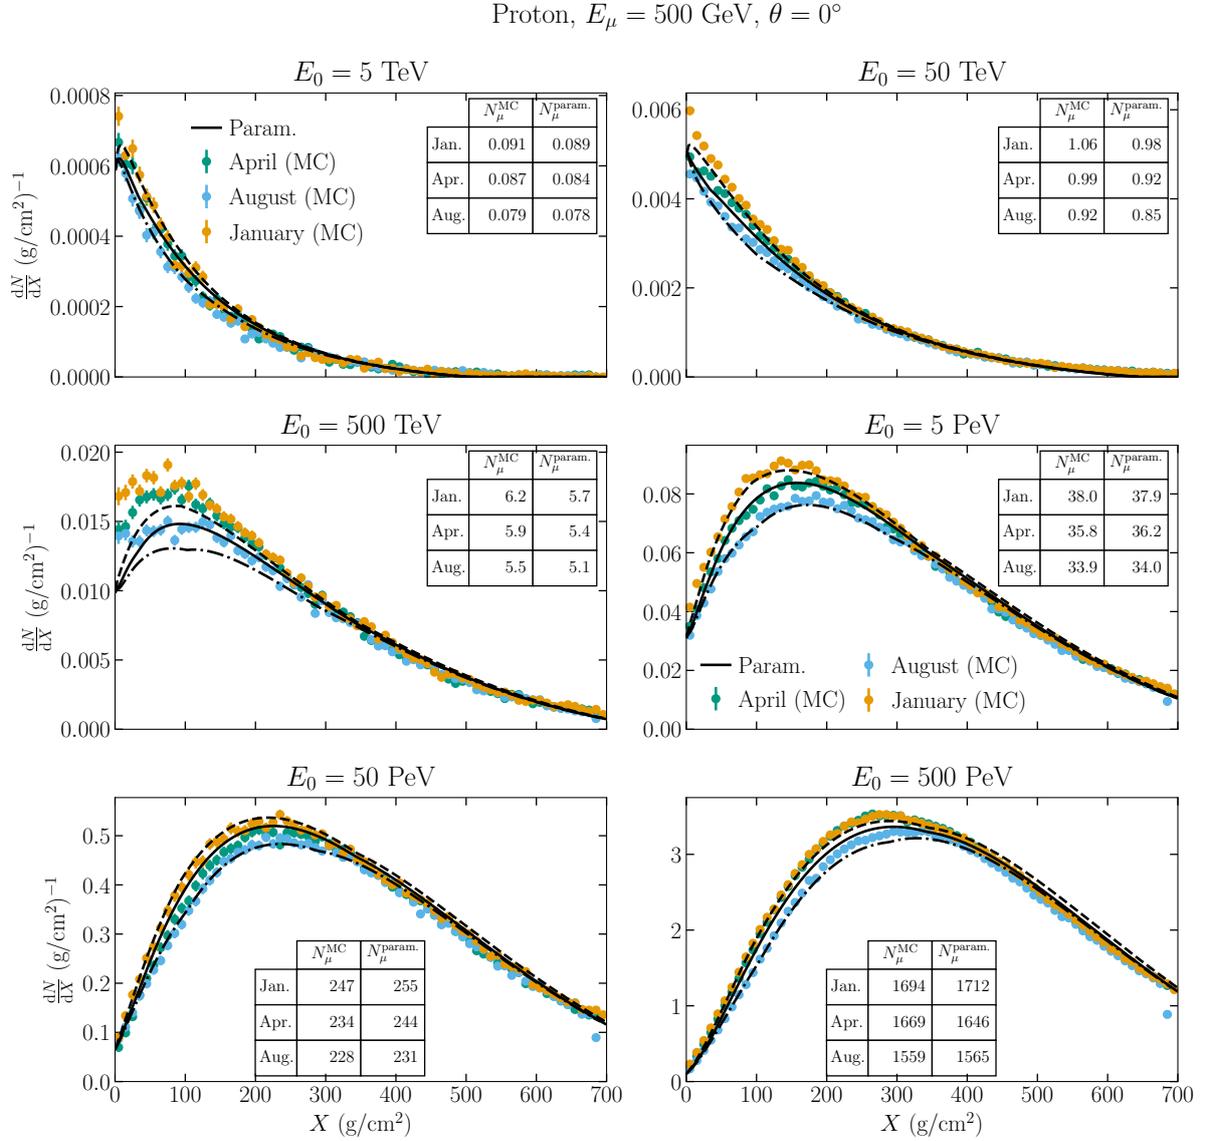
## 5.3 Example applications for IceCube

The parameterization presented in the previous sections can be used for making estimates of observables depending on the longitudinal production profiles of high-energy muons, typically related to muon bundles or inclusive muon rates in underground detectors. In experiments performing coincident measurements of air showers with a surface array and an underground muon detector, the primary cosmic-ray energy can approximately be fixed by the surface array. This allows to study the properties of muon bundles as a function of energy, and how they depend on the atmospheric conditions. Below, two simple examples of applications related to the detection of air showers with IceCube are given.

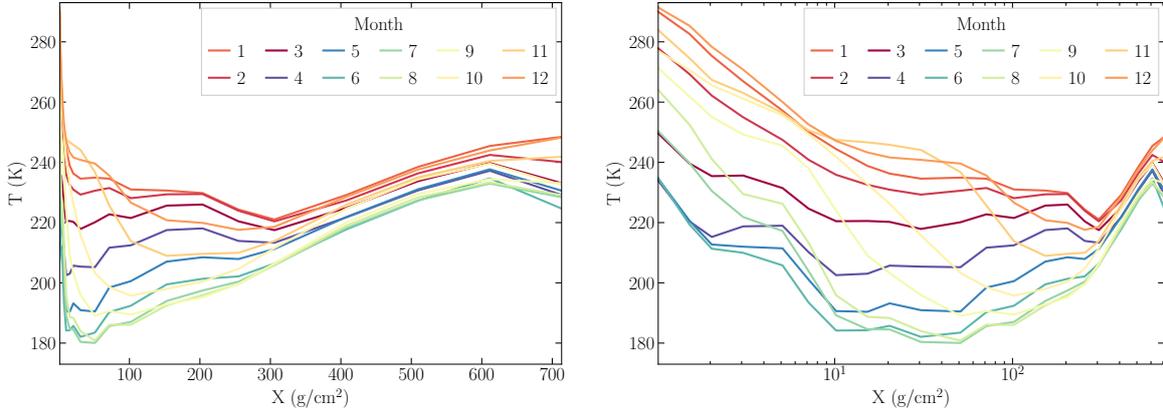
### 5.3.1 Estimate of seasonal variations of TeV muon bundles

A simple calculation of the seasonal variation of the number of high-energy muons in a shower can be performed by obtaining the production profiles for different atmospheric temperature profiles and integrating them. We show this using parameters relevant for IceTop and IceCube, specifically vertical showers, a primary energy of 10 PeV, a muon threshold of 500 GeV, and an atmospheric depth of about 700 g/cm<sup>2</sup>. We calculate the number of muons produced in the shower, as they would arrive at the surface, but do not take into account effects of propagation through ice to the detector.

The calculations are performed using atmospheric data for the South Pole obtained from the AIRS satellite, which provides the temperature at atmospheric pressure levels unevenly spaced between 1 hPa and 700 hPa. The pressure levels are converted to a grid of vertical depth and linearly interpolated to a regular grid. Monthly averages of such measurements are shown for the year 2017 in Fig. 5.13 as an example. The muon profiles obtained from the parameterization and daily temperature profiles are integrated



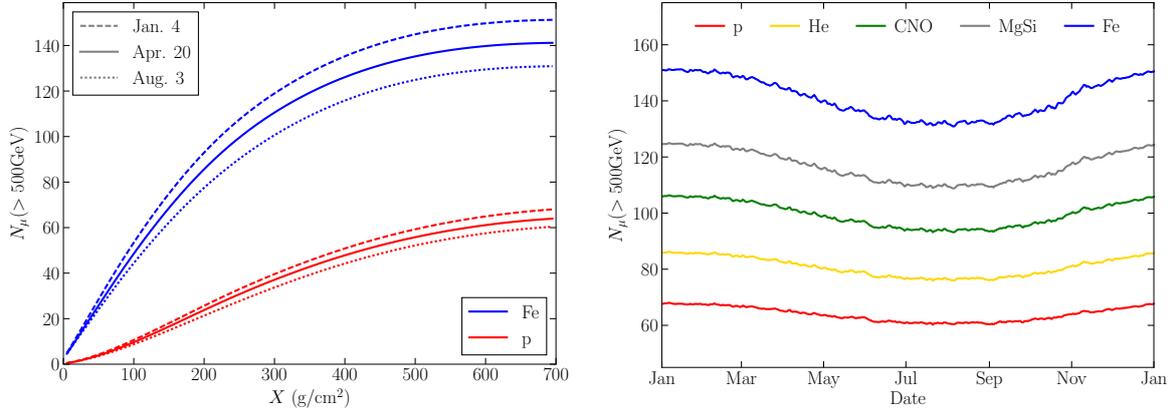
**Fig. 5.12:** Comparison between muon production profiles obtained from simulation and from the parameterization, for different atmospheric profiles. The simulations were produced with a depth profile corresponding to the South Pole atmosphere in January, April, and August. For the parameterization, these profiles were converted to a temperature profile assuming the ideal gas law. The tables show the total number of muons expected at the ground for the different cases.



**Fig. 5.13:** Atmospheric temperature at the geographical South Pole measured by the AIRS satellite for 2017, averaged by month, as a function of the vertical depth in linear (left) and logarithmic (right) scale.

to obtain the muon multiplicity. Fig. 5.14 shows the expected multiplicity throughout the year 2017 for different primary masses, as well as the integral profiles for three days corresponding roughly to the days with minimal, average, and maximal multiplicity. The multiplicity is maximal in austral summer, when temperatures are highest. The relative variations in multiplicity are around 6% and 7% for proton and iron respectively, and they slowly decrease with primary energy.

As the parameterization describes the muon production as a function of slant depth, it is possible to extract information about the altitude at which muons are produced. Ideally, the atmospheric data includes the altitude corresponding to a pressure or depth level at which the temperature is measured. If this is not the case, it can also be calculated assuming the ideal gas law. As an estimate, one can invert the relation of Eq. (2.5) to obtain the altitude as a function of depth, but seeing that this formula was derived for an isothermal atmosphere, it is preferred to numerically integrate Eq. (2.4) to obtain the altitude corresponding to a certain depth. We follow here this approach, and define  $h = 0$  as the surface level of the South Pole at an elevation of 2835 m with an atmospheric depth  $X_0 \approx 700 \text{ g/cm}^2$ . The left panel of Fig. 5.15 shows the differential muon production as a function of altitude for vertical proton and iron showers at the same average and extremal days as before. Muons are on average produced higher in the atmosphere for the iron showers. For a given primary mass, one observes that muon production happens at a higher altitude in the summer compared to colder periods, a result of the thermal expansion of the atmosphere.



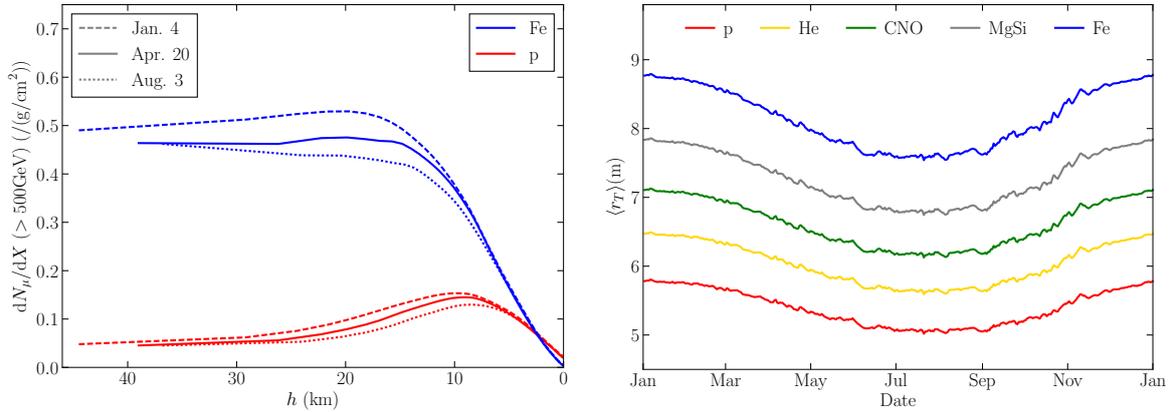
**Fig. 5.14:** Left: Integral production profiles for vertical 10 PeV proton and iron showers, calculated using temperature profiles obtained at the South Pole on three different days in 2017 where the expected muon multiplicity is approximately the yearly minimum, average, and maximum. Right: Variation of the expected muon multiplicity above threshold throughout the year for vertical showers at 10 PeV for five primary masses.

From this, the transverse distance of the muons from the shower axis can be estimated. A muon with energy  $E_\mu$  produced at an altitude  $h$  with a transverse momentum  $p_T$  will have a transverse distance from the shower axis given by

$$r_T = \frac{p_T}{E_\mu} \times \frac{h}{\cos \theta}, \quad (5.10)$$

where  $\theta$  is the zenith angle of the primary. We perform a simple estimate of the expected muon bundle size, assuming a mean value of transverse momentum for the muons of  $\langle p_T \rangle \approx 350$  MeV [188]. The expected muon bundle size  $\langle r_T \rangle$  is estimated by calculating the transverse distance for a muon with a transverse momentum  $\langle p_T \rangle$  from Eq. (5.10), and taking the weighted average by multiplying with the production profile and integrating over depth. The result of this calculation is shown in the right panel of Fig. 5.15. The bundle size is maximal in the summer as a result of the higher production altitude. The relative variations in this case are of the order 7-8% for proton and iron respectively. One could also obtain an estimate of the size based on the average muon energy in the shower, rather than the minimum energy, by replacing  $E_\mu$  by  $fE_\mu$ , where  $f$  is the ratio defined in the previous section.

The calculations in this section present only approximations of the muon bundle properties when it reaches the surface and the shower development stops. For a full calculation of the properties of the muon bundle as it reaches the underground detector, the propagation through matter needs to be taken into account. Important effects are



**Fig. 5.15:** Left: Differential muon production as a function of altitude for vertical 10 PeV proton showers for three different days at the South Pole. The altitude is measured relative to the surface above IceCube, and is calculated as described in the text. Right: Seasonal variation of the estimated transverse size of the muon bundle as a result of the changing production altitude for vertical showers with  $E_0 = 10$  PeV and five mass groups.

multiple Coulomb scattering in the overburden [189], and charge separation of muons caused by deflection in the geomagnetic field [190], which would both increase the transverse distance of the muons. A detailed treatment of these effects is beyond the scope of this work.

A full simulation in Ref. [191] obtained comparable conclusions regarding the average size and multiplicity of high-energy muon bundles in underground detectors.

### 5.3.2 Application to muon bundle energy loss

The second example discussed is directly related to the composition measurement performed with IceTop and IceCube as discussed in Section 4.3.1. We will therefore use some jargon introduced in Chapter 4, and use a small amount of IceCube data for illustrative purposes.

As introduced in Section 4.3.4, the main mass-dependent observable in the composition analysis  $dE/dX_{1500}$ , related to the energy loss of the muon bundle in the detector, shows a clear seasonal variation. As the energy loss is related to the number of muons in the bundle, this effect is not surprising, considering the discussion above. In Ref. [174], the magnitude of the seasonal variation was found to be 10-15% of the difference between  $dE/dX_{1500}$  in proton and iron simulation. For an analysis which assigns a mass estimate on an event-by-event basis using a neural network that was trained on a single atmosphere, this is an important effect to take into account, and an event-by-event correction for the

variation was devised. We illustrate here briefly how such a correction could be performed in the future using the muon production profiles obtained from the parameterization introduced in this chapter.

The seasonal variation can be described in terms of variations in the atmospheric temperature by defining an effective temperature that reduces the atmospheric temperature profile to a single number. The atmospheric temperature can be defined by weighting the temperature profile  $T(X)$  by the production profile of muons above a threshold  $E_\mu$  as

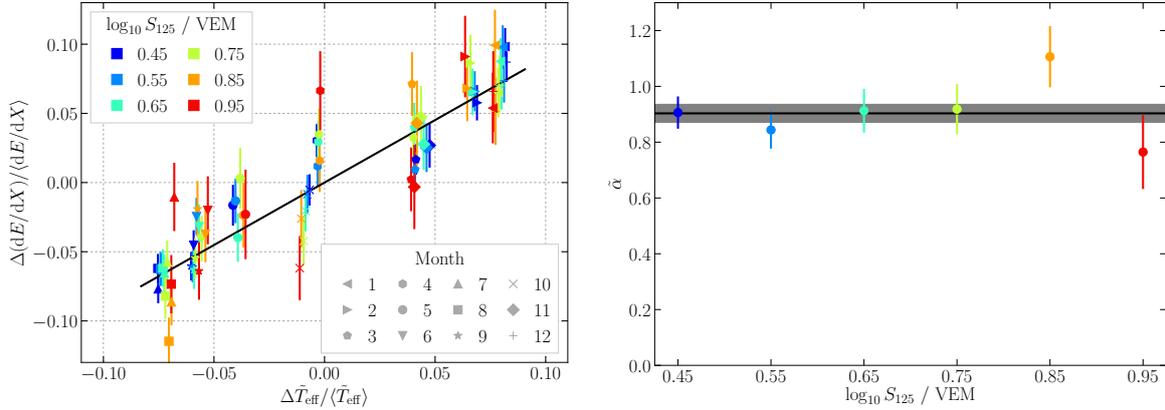
$$\tilde{T}_{\text{eff}}(E_0, A) = \frac{\int \frac{dN}{dX}(E_0, A, > E_\mu) T(X) dX}{\int \frac{dN}{dX}(E_0, A, > E_\mu) dX}, \quad (5.11)$$

where  $E_0$  and  $A$  are the primary cosmic-ray energy and mass. Note that this is different from the effective temperature Eq. (4.7) that is used when describing the seasonal variation of event rates; Eq. (4.7) considers the inclusive muon production in the atmosphere (i.e. integrated over the primary cosmic-ray spectrum), while in Eq. (5.11) we consider the muon production in single showers. The seasonal variation of  $dE/dX_{1500}$  observed in data can be described by a correlation coefficient  $\tilde{\alpha}$ , which captures the relation between relative variations in  $dE/dX_{1500}$  and  $\tilde{T}_{\text{eff}}$

$$\frac{\Delta dE/dX_{1500}}{\langle dE/dX_{1500} \rangle} = \tilde{\alpha} \frac{\Delta \tilde{T}_{\text{eff}}}{\langle \tilde{T}_{\text{eff}} \rangle}, \quad (5.12)$$

which may be dependent on the primary energy. In what follows, we bin the data by month and study the relative variations compared to the yearly average.

To obtain  $\tilde{\alpha}$  from data, we use 10% of IC86.2012 data with the quality cuts described in Section 4.2.3 applied, binned per month and in the primary energy proxy  $S_{125}$ . Atmospheric temperature profiles are obtained from balloon measurements from AMRC [53] and measurements from the AIRS satellite [51], merged and kindly provided by Takao Kuwabara [54], and are averaged by month. For each event, we use the  $dE/dX_{1500}$  value and calculate a  $\tilde{T}_{\text{eff}}$  value. For this, a muon production profile needs to be calculated from the parameterization. We use a threshold value of  $E_\mu > 500$  GeV, and estimate the primary energy using a simple  $S_{125}$ -to- $E_0$  conversion function. To deal with the mass dependence, an average based on the H4a composition model (Section 1.5) is used. With this, the average  $dE/dX_{1500}$  and  $\tilde{T}_{\text{eff}}$  can be calculated per month- $S_{125}$  bin. The relative variations for each month compared to the yearly average are shown in the left panel of Fig. 5.16. The correlation  $\tilde{\alpha}$  is derived by fitting a line through the origin to the points per  $S_{125}$  bin. The result is shown in the right panel of the figure. In the energy



**Fig. 5.16:** Left: Relative variation in the muon bundle energy loss variable  $dE/dX_{1500}$  used in the IceCube composition analysis, compared to variations in the effective temperature Eq. (5.11), binned by month and primary energy proxy  $S_{125}$ . The black line corresponds to the average correlation coefficient shown in the right panel. Right: Correlation coefficient  $\tilde{\alpha}$  derived from straight line fits to the variations plotted in the left panel. The black band shows the average value and its statistical uncertainty.

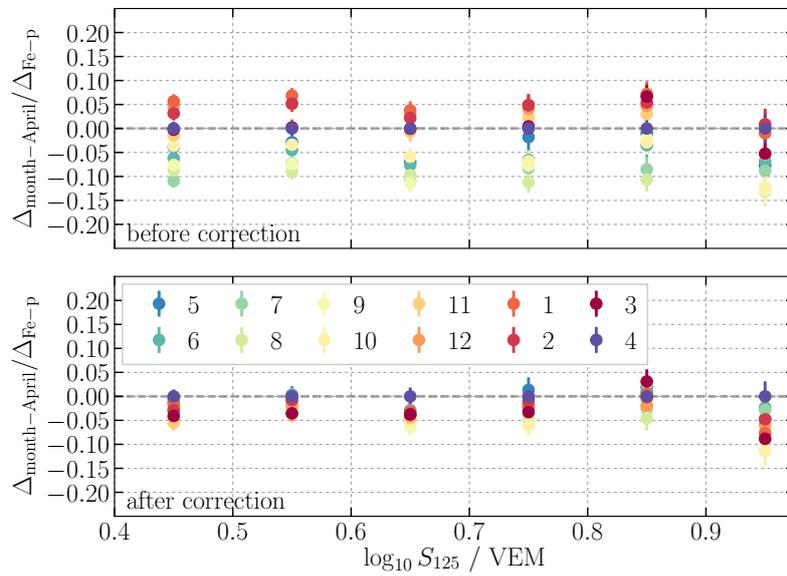
range accessible here with the limited statistics used, there is no evidence for an energy dependence of the correlation, so that a constant value of  $\tilde{\alpha} \approx 0.9$  is obtained.

The definition of  $\tilde{T}_{\text{eff}}$  and the value of  $\tilde{\alpha}$  can in turn be used to correct the  $dE/dX_{1500}$  values on an event-by-event basis to some reference effective temperature. The current IceTop simulations use an atmospheric profile that describes the average South Pole atmosphere in April. Events are therefore corrected to the reference value  $\tilde{T}_{\text{eff}}^{\text{April}}(E_0, A)$ , again by estimating the energy from  $S_{125}$  and averaging according to the H4a model composition, as follows:

$$\Delta\tilde{T}_{\text{eff}} = \frac{\tilde{T}_{\text{eff}}^{\text{month}} - \tilde{T}_{\text{eff}}^{\text{April}}}{\tilde{T}_{\text{eff}}^{\text{April}}}, \quad dE/dX_{1500}^{\text{corr}} = \frac{dE/dX_{1500}}{1 + \tilde{\alpha}\Delta\tilde{T}_{\text{eff}}}. \quad (5.13)$$

Fig. 5.17 demonstrates how this method effectively reduces the spread between the monthly averaged  $dE/dX_{1500}$  values from  $\sim 15\%$  of the proton-iron difference to about 5%. Any remaining spread or constant offset from the chosen reference can be included as a systematic uncertainty in further analysis, or could likely be improved upon by further optimizing the simple calculations presented here.

We will return to seasonal variations in the multiplicity of muon bundles as detected by IceCube in Chapter 7.



**Fig. 5.17:** Difference in the average  $dE/dX_{1500}$  between a certain month and the reference month April, relative to the difference between  $dE/dX_{1500}$  in proton and iron simulations based on the month April. Different months are indicated by the colors as shown in the legend. The bottom panel shows the remaining variations after correcting the  $dE/dX_{1500}$  values to the reference based on Eq. (5.13).

## 5.4 Discussion

Motivated by the understanding of observables accessible with underground detectors, this chapter investigated the longitudinal production profiles of high-energy muons in cosmic-ray air showers based on simulations. A physically motivated formula was presented that can be used to accurately fit such profiles, regardless of primary energy  $E_0$ , mass number  $A$ , zenith angle  $\theta$ , and muon threshold energy  $E_\mu$ . The results of fits to numerous simulated profiles were parametrized in terms of these variables. The result is a formula and a set of parameters, in usage comparable to the Elbert formula, but with the clear enhancement that it describes the entire muon production profile rather than just the total number of muons produced. Furthermore, it has an explicit dependence on the atmospheric temperature profile, a result from the competition between interaction and decay of parent mesons.

Example applications of the parameterization to the calculation of muon bundle properties were presented, clearly showing how it can be used to estimate the seasonal variations of these properties, which may be an important effect to take into account for various analyses.

Although the parameterization was derived from a set of simulations with some specific conditions, it was shown that it can be used in broader way, although with a somewhat decreased accuracy. In this respect it is important to note that the parameterization in its current form is not a precision tool. Rather, it is an instrument that can be used to estimate with a reasonable accuracy the properties of certain muon observables relevant for different experiments. It is a simple and especially very efficient way of obtaining insight into basic physical processes compared to a full simulation of the problem, which would be necessary for more detailed analyses.

Various suggestions to increase the precision and scope of the parameterization were mentioned throughout the chapter. Re-optimizing the parameters for an experiment with different characteristics than presented is one of them. Taking the simulated temperature profile into account during the profile fits is another. It is likely also a good idea to simulate showers down to sea level or at larger zenith angles, so that slant depths larger than the vertical depth of the observation level of a certain detector would be described well. Furthermore, a better description of the remaining dependence of the variables on  $E_\mu$ , after describing them in terms of  $E_0/AE_\mu$ , would be a significant step forward. One could also think about a different way of combining decay factors corresponding to different muon production channels. Lastly, a good estimate of the uncertainty on the obtained profiles would be very useful for more quantitative analyses.



# Chapter 6

## Testing hadronic interaction models

Indirect measurements of cosmic rays through the air showers they produce in the atmosphere require accurate simulations to be interpreted in terms of properties of the primary nucleus. As discussed in Chapter 2, uncertainties in state-of-the-art hadronic interaction models used in such simulations still hinder the accurate prediction of various air-shower observables. This is especially true for the muonic component of air showers.

As explained in Chapter 4, the combination of the surface array IceTop and the deep in-ice detector IceCube forms a cosmic-ray detector with the unique capability of measuring a shower's electromagnetic component together with muons in two different energy regimes. By reconstructing different composition-dependent observables, one can provide strong tests of the self-consistency of hadronic interaction models and the related uncertainty in cosmic-ray composition measurements. In this chapter, a basic analysis of this kind is presented<sup>1</sup>.

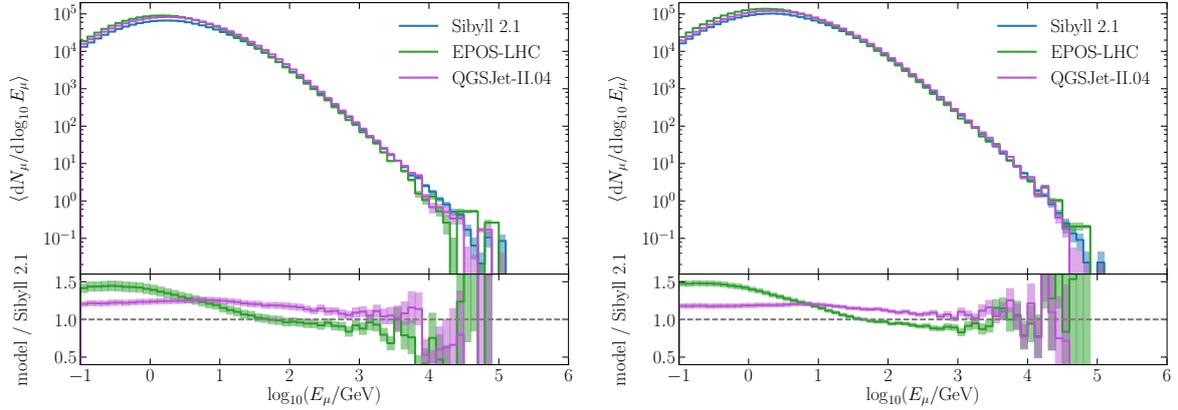
In Section 6.1, we introduce the motivation and goal of the analysis, as well as the approach. Details about the dataset used are given in Section 6.2. The different composition-dependent observables used are introduced in Section 6.3. In Section 6.4, the composition interpretation of these different observables in data is obtained and compared. The implications of the results are finally discussed in Section 6.5.

### 6.1 Introduction

As discussed in Chapters 1 and 2, cosmic-ray measurements above 100 TeV are performed indirectly by observing extensive air showers. Measurements of properties of the primary cosmic rays, such as their energy spectrum and mass composition, are dependent on

---

<sup>1</sup>The analysis described in this chapter was presented at the ICRC2021 conference. For a more concise summary of the work, see the corresponding proceeding, Ref. [192].

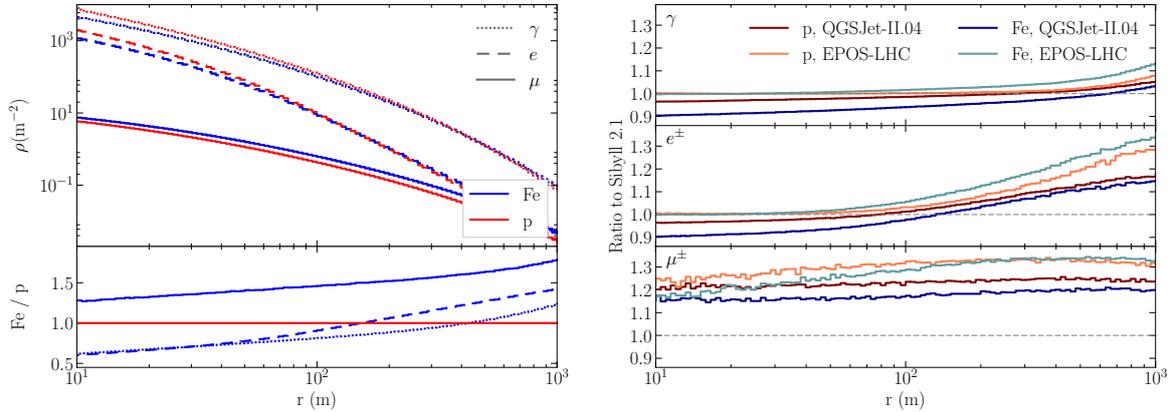


**Fig. 6.1:** Average muon energy spectrum in simulated vertical proton (left) and iron (right) showers with primary energy around 10 PeV for the different hadronic interaction models considered in this work.

air-showers simulations, which utilize effective models to describe hadronic interactions at high energies. The uncertainties in these models and in the resulting air-shower simulations make it difficult to unambiguously interpret the data, leading to conflicting results in analyses that use different models. Furthermore, the existence of discrepancies between data and simulation has been demonstrated in measurements related to the number of (GeV) muons in air showers, implying deficiencies in the modeling of hadronic interactions (Section 2.4.3).

The IceCube Neutrino Observatory provides a unique opportunity to study hadronic interactions and provide input to this problem by measuring multiple components of air showers at once: the EM component, the GeV muon component, and the TeV muon component. As summarized in Section 4.3.2, this was first used to study the consistency of hadronic interaction models by combining different composition-dependent observables and comparing their values in data to predictions in simulations in Ref. [157]. If the simulations describe data accurately, the composition interpretation obtained from the different observables should be consistent with each other. The observables used were the slope of the IceTop LDF  $\beta$ , and the in-ice muon bundle energy loss parameter  $dE/dX_{1500}$ , which were introduced already in Chapter 4. Here, we expand on this work by adding a measurement of the density of GeV muons at large lateral distance in air showers, following the method of the analysis described in Section 4.3.3. The different composition-sensitive observables used are discussed in more detail in Section 6.3.

Before further discussing the analysis, it is interesting to highlight some differences in air-shower properties between simulations performed with different high-energy hadronic



**Fig. 6.2:** Left: Average lateral distribution of photons, electrons, and muons in vertical proton (red) and iron (blue) showers around 10 PeV obtained from simulations using Sibyll 2.1. Right: Ratio of the different lateral distributions in QGSJet-II.04 and EPOS-LHC compared to Sibyll 2.1.

interaction models, relevant for measurements with IceTop and IceCube. These are obtained from the available IceTop simulations, described in Section 4.4. Fig. 6.1 shows the average muon energy spectrum in air showers for the models under consideration: Sibyll 2.1, QGSJet-II.04, and EPOS-LHC<sup>2</sup>. The post-LHC models predict a larger number of GeV muons compared to Sibyll 2.1. For muons in the TeV energy range, simulations with QGSJet-II.04 contain a larger number than Sibyll 2.1, while EPOS-LHC has less. In Fig. 6.2, the lateral distributions of photons, electrons, and muons is shown for Sibyll 2.1, as well as how they change for the post-LHC models. The increase in the number of muons in the post-LHC models is clear over the entire distance up to 1 km away from the core. The plots furthermore imply that the LDFs of the EM component are less steep in these models compared to Sibyll 2.1.

In what follows, we will discuss experimental quantities related to these shower properties. The impact of the changes between the models on the composition interpretation of a measurement will be studied, as well as the agreement between observables within individual models. This will reflect uncertainties in indirect cosmic-ray composition measurements and may reveal where the model predictions deviate from experimental data.

<sup>2</sup>The post-LHC model Sibyll 2.3d is not included, as there were no IceTop simulations using this model available at the time of writing.

## 6.2 Event selection

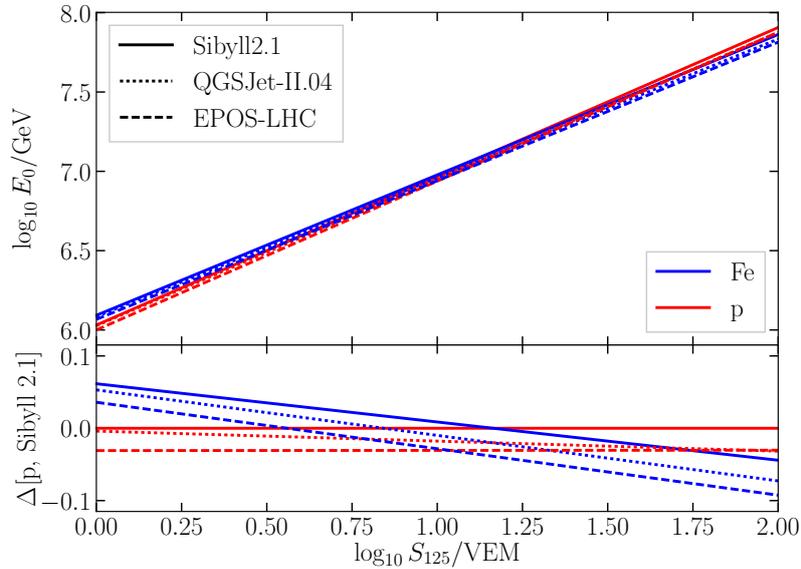
This analysis starts from the cosmic-ray Level3 processing that was described in Section 4.2. As the goal is to compare observables related to both low- and high-energy muons, we specifically want coincident events that trigger IceTop and have a muon bundle going through IceCube. The quality cuts described in Section 4.2.3 are applied, ensuring that events are contained in IceTop, are well reconstructed, and have a successful muon bundle energy loss reconstruction in IceCube.

Compared to the previous analysis (Section 4.3.2), a more strict zenith cut is applied, requiring  $\cos\theta > 0.95$ . This is so that the muon density reconstruction technique (Section 4.3.3), which has not been studied well outside this range, can be applied.

The data used will be limited to the IC86.2012 data-taking season, as this is the season for which simulation is available with different hadronic interaction models. The simulated datasets used are those listed in Table 4.3. As snow accumulation may cause disagreement between simulation and data of different years, no other data is included. The results presented in Section 6.4 are obtained using 10% of the experimental data collected during IC86.2012. The reason for this limited amount is that the analysis in its current form is still considered somewhat preliminary, for reasons discussed in Section 6.5. The 10% of data consists out of all runs ending in zero. Therefore, it is spread out approximately uniformly throughout the year. As discussed in Section 4.4, the yearly average of the atmosphere is close to a typical atmosphere of the month April, so that the data can be directly compared to the simulation, which uses an April density profile.

Average values of different composition-sensitive observables will be compared in bins of the energy proxy  $S_{125}$ , the reconstructed shower size in IceTop. The correlation between  $S_{125}$  and the primary cosmic-ray energy, illustrated already in Fig. 4.8, can be used to define a simple conversion function as in Eq. (4.6). Examples are shown in Fig. 6.3 for different hadronic interaction models and for proton and iron primaries. This can be used to estimate the primary energy of an event with a certain measured  $S_{125}$ . The high elevation of IceTop, putting it close to  $X_{\max}$  where the shower is strongly dominated by the EM component (Chapter 2), result in only a small model and mass dependence of the conversion functions.

As the analysis involves a direct comparison of data to simulation, the range of  $S_{125}$  values we can use is limited by the energy range in which simulations exist. The datasets for QGSJet-II.04 and EPOS-LHC are limited to 100 PeV, while for Sibyll 2.1, simulation is available up to approximately 4 EeV. To avoid composition-dependent threshold effects, the lower end of the energy range is placed above the threshold for full efficiency of IceTop. As discussed in Section 4.2.4, this is about  $\log_{10} E_0/\text{GeV} = 6.4$  for the season

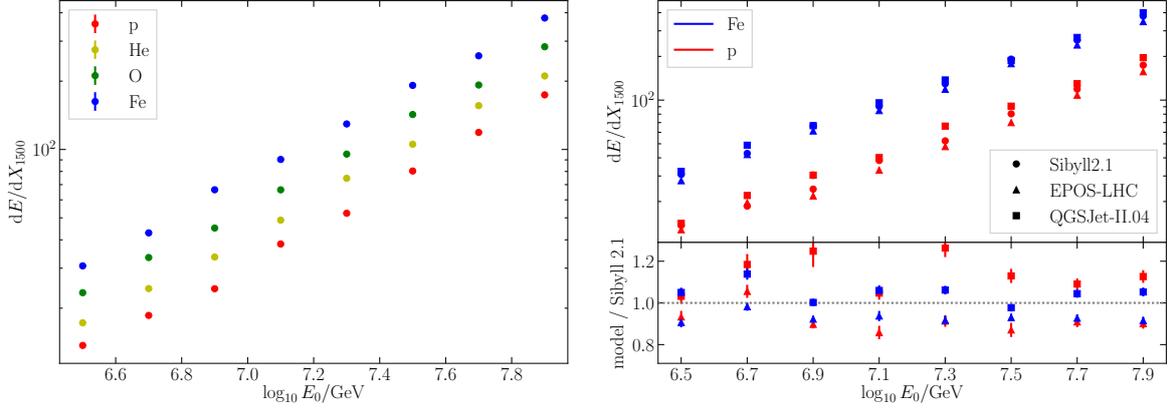


**Fig. 6.3:** Simple shower size  $S_{125}$  to primary energy  $E_0$  conversion function of the form of Eq. (4.6) for IC86.2012. Functions are obtained by fitting the average  $\log_{10} E_0$  in bins of  $\log_{10} S_{125}$  in simulated proton and iron showers with  $\cos \theta > 0.95$ . The bottom plot shows the difference between the lines in the top plot and the result obtained for proton and Sibyll 2.1.

considered here, corresponding to approximately  $\log_{10} S_{125}/\text{VEM} = 0.4$ . As not to be influenced by an incomplete description of bin migration effects at the energy where the simulated dataset stops, the  $\log_{10} S_{125}/\text{VEM}$  range is limited to 3.0 for Sibyll 2.1 and 2.0 for QGSJet-II.04 and EPOS-LHC. These correspond to a  $\log_{10} E_0/\text{GeV}$  of approximately 8.8 and 7.9 respectively.

### 6.3 Composition-sensitive observables

The combination of IceTop and IceCube can be used to reconstruct several composition-sensitive observables, based on different air-shower components. Below, three observables are introduced, namely the muon bundle energy loss parameter  $dE/dX_{1500}$ , the IceTop LDF slope  $\beta$ , and the muon density at large lateral distance  $\rho_\mu$ . We describe how they are obtained from the measurements and how they behave in simulations for different primaries and hadronic interaction models.



**Fig. 6.4:** Left: Average  $dE/dX_{1500}$  in bins of true primary energy obtained from simulations (Sibyll 2.1) of different primary nuclei. Right: Comparison of the average  $dE/dX_{1500}$  values obtained from simulations with different hadronic interaction models.

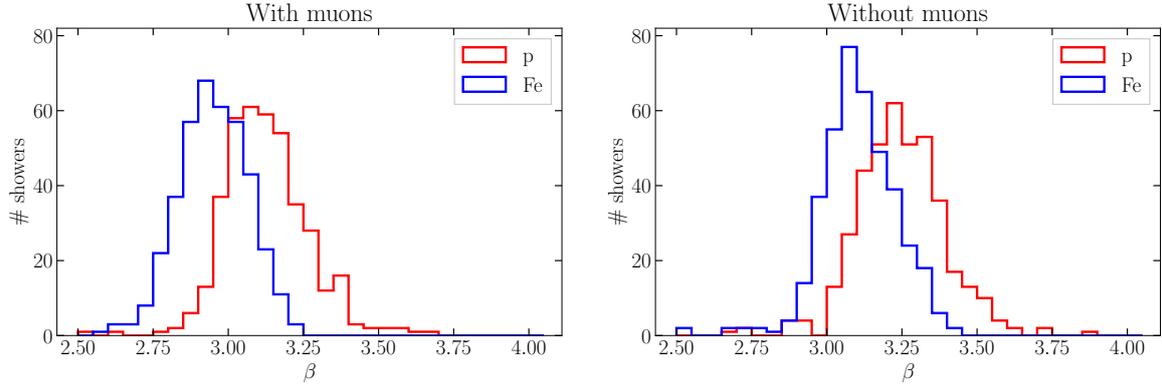
### 6.3.1 Muon bundle energy loss

We include a single in-ice observable in the analysis: the muon bundle energy loss parameter  $dE/dX_{1500}$ , introduced in Section 4.2.2. It is obtained by fitting the segmented energy loss reconstruction performed on the muon bundle signal in IceCube along the reconstructed shower axis. The energy deposit of a muon bundle can be expected to be approximately proportional to the number of muons in the bundle. In Fig. 4.8, it was shown that there is indeed such a correlation between  $dE/dX_{1500}$  and the number of muons above several hundred GeV in the air shower.

As the number of muons in a shower depends on the primary mass (Chapter 2), also  $dE/dX_{1500}$  will show such a dependence. This is illustrated with simulations of different primary nuclei in the left panel of Fig. 6.4, showing that heavier primaries have larger energy loss values. The right panel shows the changes in the average  $dE/dX_{1500}$  value for different hadronic interaction models. Simulations predict higher values when using QGSJet-II.04 compared to Sibyll 2.1 and lower values when using EPOS-LHC, qualitatively consistent with expectations from the muon spectra of Fig. 6.1.

### 6.3.2 IceTop LDF slope

One of the IceTop parameters used in this study is  $\beta$ , introduced in Section 4.2.1. It is the slope of the lateral distribution function that is fitted to the measured IceTop charges during the standard air-shower reconstruction, Eq. (4.2). Together with  $S_{125}$ , it fully defines this function.

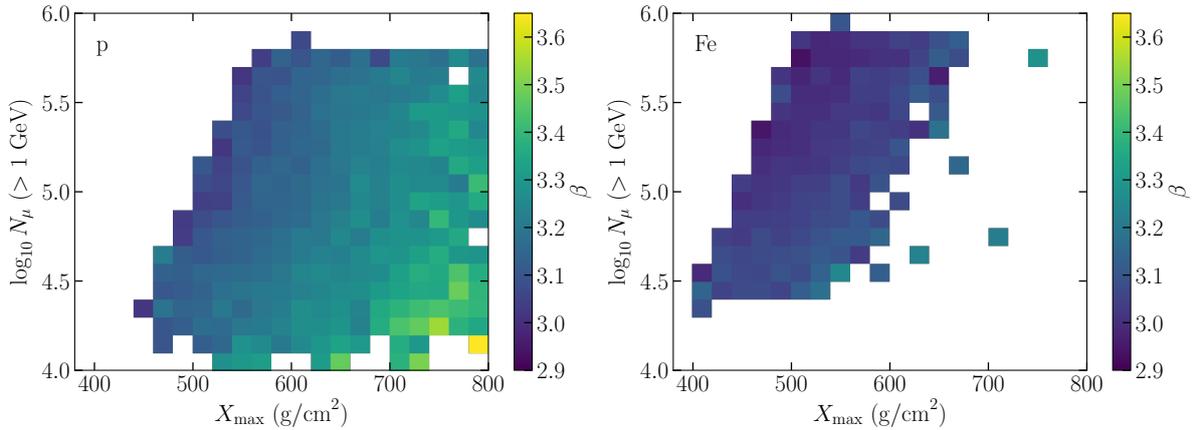


**Fig. 6.5:** Distribution of  $\beta$  values obtained from 400 simulated vertical proton and iron showers at 10 PeV with their core in the center of the IceTop array. Left: A regular detector simulation has been performed. Right: Muons have been ignored in the detector simulation.

The lateral distribution of charges can be expected to be composition dependent for multiple reasons. The dependence is not as clear as for e.g.  $dE/dX_{1500}$ , which can clearly only be influenced by high-energy muons. The observed charges in IceTop are dominated by EM particles close to the shower axis, but have a muonic contribution that increases with lateral distance. From simulations, it becomes clear that  $\beta$  is not a pure estimator of the low-energy muons, as it was sometimes interpreted in the past, but is also influenced significantly by changes in the EM shower component.

To illustrate this, dedicated simulations were produced of vertical showers at a fixed energy of 10 PeV, landing in the center of the IceTop array. The simulated showers were once processed with a normal IceTop detector simulation, and once with a detector simulation where all muons were ignored. The resulting distributions of  $\beta$ -values are shown in Fig. 6.5. It can be seen that when removing the muons,  $\beta$  shifts toward higher values, corresponding to a steeper LDF. However, even without muons, there is a clear difference in the distributions for proton and iron showers, which must be a result of differences in the EM component. We furthermore show the relation between  $\beta$ , the depth of shower maximum  $X_{\max}$ , and the number of muons in the shower  $N_{\mu}$ , based on standard IceTop simulation in Fig. 6.6. The plots indicate that at fixed  $X_{\max}$ ,  $\beta$  decreases with increasing  $N_{\mu}$ , while at fixed  $N_{\mu}$ ,  $\beta$  increases with increasing  $X_{\max}$ .

The composition dependence of  $\beta$ , demonstrated in the left panel of Fig. 6.7, can then be understood as follows. First, showers from heavier primaries develop earlier in the atmosphere so that at the observation level the particles have spread out more than for lighter primaries. Secondly, showers from heavier primaries contain more muons,



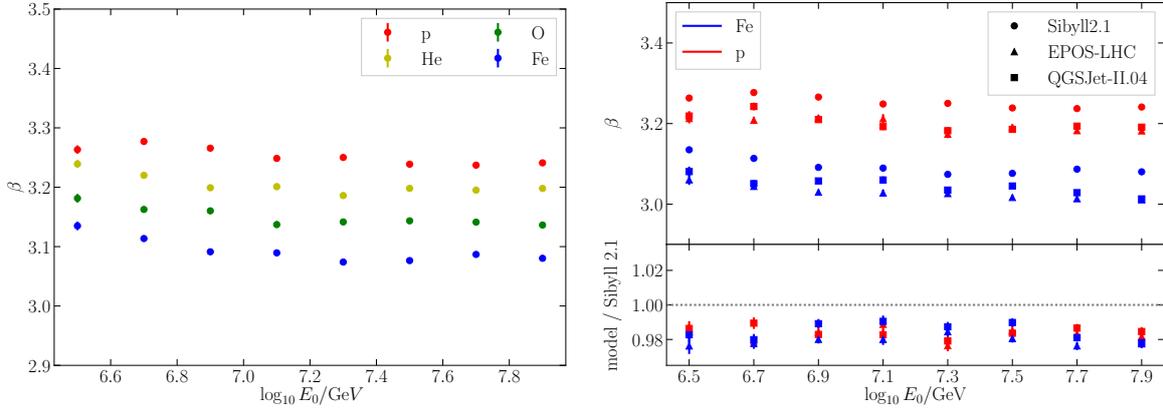
**Fig. 6.6:** 2D profile plots showing the average  $\beta$  value as a function of the depth of shower maximum  $X_{\max}$  and the number of muons above 1 GeV in the shower. Distributions were obtained from standard IceTop simulation with primary energy in the range  $6.4 \leq \log_{10} E_0/\text{GeV} \leq 8.0$  and zenith angles  $\cos \theta > 0.95$  for proton (left) and iron (right).

which have an LDF that is less steep than the LDF of EM particles. Both effects cause the lateral distribution of charges detected in IceTop to become less steep for heavier nuclei, corresponding to the smaller values of  $\beta$  observed in simulation. The right panel of Fig. 6.7 shows how the average value of  $\beta$  differs between hadronic models. The value is lower for both post-LHC models compared to Sibyll 2.1, consistent with the increase in muons and flattening of the EM LDFs shown in Fig. 6.2.

### The “ $\beta$ problem”

In past analyses,  $\beta$  was not used beyond energies of 100 PeV. This is because it was found that in the high-energy simulated datasets,  $\beta$  showed an unphysical behavior, abruptly increasing in value with increasing primary energy [158, 157]. This is illustrated in the left panel of Fig. 6.8. Given the energy at which the effect starts and the range over which it increases, together with the fact that the behavior is not seen in the low-energy simulation sets, this behavior was thought to be related to a problem in the thinning or dethinning of the shower [181].

Initiating the updated analysis presented here, it was observed that this problem was no longer present in the high-energy datasets of the newly produced simulations of the season IC86.2012. While the exact reason for the bug in the older simulation is not known, the smooth behavior with energy and the agreement between the low-energy and high-energy (thinned) datasets in the new simulation is reassuring. Note that the increase in average  $\beta$  as function of  $S_{125}$  that can be seen at the upper energy end of



**Fig. 6.7:** Left: Average  $\beta$  in bins of true primary energy obtained from simulations of different primary nuclei. Right: Comparison of the average  $\beta$  values obtained from simulations with different hadronic interaction models.

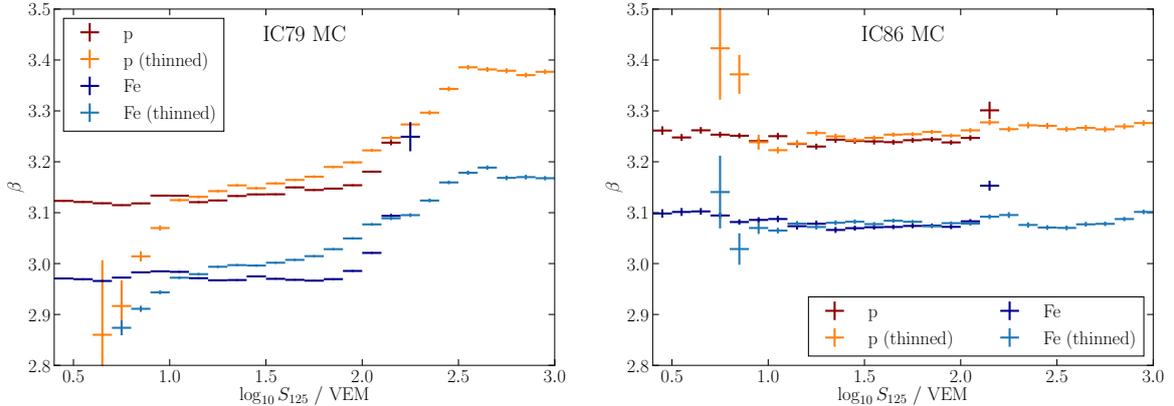
the datasets without thinning is a result of the dataset ending at a fixed energy, and motivates limiting the  $S_{125}$  range to values where events with energies both lower and higher than the average expected in the bin can contribute. Combining the low- and high-energy datasets gets rid of the related features seen in the right panel of Fig. 6.8 around 0.8 and 2.2 in  $\log_{10} S_{125}$ .

### 6.3.3 Density of GeV muons

As a second IceTop observable, the density of GeV muons  $\rho_\mu$  at large lateral distance is used. The density of muons in IceTop, dominated by muons in the GeV energy range, is reconstructed using the method of the analysis discussed in Section 4.3.3<sup>3</sup>. It relies on the typical VEM signal produced by muons passing vertically through a tank, which can be used to separate the muon signal from hits of other charged particles far from the shower axis. This is a statistical analysis, in which the average muon density in air showers is derived over the entire data sample, in contrast to the previously discussed observables which are reconstructed on an event-by-event basis. The reference distances at which the muon density is derived are 600 m and 800 m.

The muon densities  $\rho_\mu^{600}$  and  $\rho_\mu^{800}$  are, as expected, dependent on the primary mass, with showers from heavier nuclei possessing more muons. The relation between the true muon density and the primary energy is shown for different primary nuclei and

<sup>3</sup>The energy conversion function used in the original analysis was updated to one derived from the IC86.2012 simulations used here.

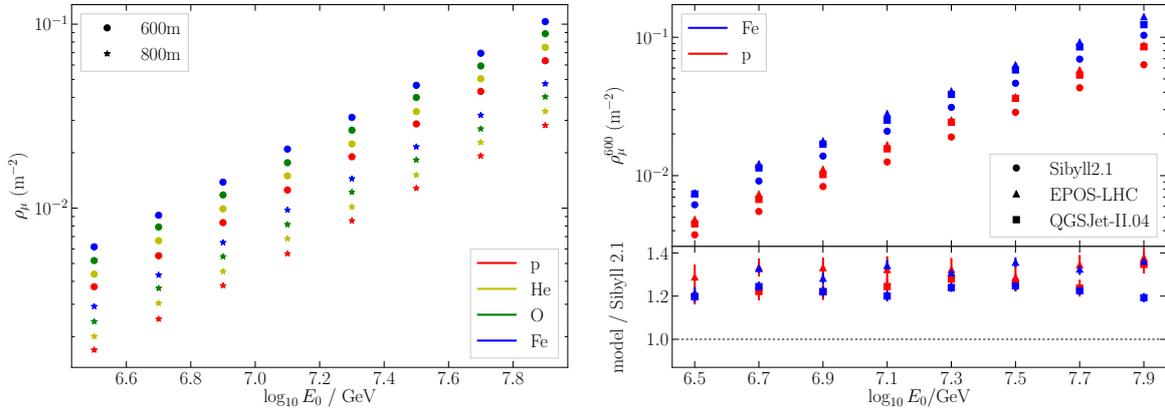


**Fig. 6.8:** Distribution of the average  $\beta$  as function of  $S_{125}$  in low- and high-energy IceTop datasets. The left panel, made with older IC79 simulation, shows the so-called “ $\beta$  problem” in the thinned datasets. In the newer simulation of Table 4.3, shown in the right panel, this issue is not present anymore.

different hadronic interaction models in Fig. 6.9. As could already be seen in Fig. 6.2, the post-LHC models contain more muons than Sibyll 2.1.

Note that this observable is conceptually different to  $dE/dX_{1500}$  and  $\beta$ . The density of muons is a property of an air shower, it exists without any notion of a detector, and can be defined directly from CORSIKA simulation. The other observables, on the other hand, are detector observables. They are determined by properties of IceTop and IceCube, and can only be derived from a full detector simulation. Therefore, the expectations shown in Fig. 6.9 are the true muon densities obtained straight from CORSIKA, rather than those that would be obtained from applying the reconstruction method to our detector-simulated signals. We will later on also compare the muon densities obtained from data to the true expectations from CORSIKA, rather than to the reconstructions obtained from proton and iron simulation. This can be done as a result of the correction factor that gets applied to the muon density reconstructions, which takes into account the difference between reconstructed and true muon density, as explained in Section 4.3.3<sup>4</sup>.

<sup>4</sup>The application of the correction factor in this study is somewhat different to the one in the original analysis. There, the  $S_{125}$  values at which muon densities are derived get converted to primary energy values, so that a direct comparison to CORSIKA simulation is possible to obtain a correction factor. In the analysis presented here, the observables are kept in bins of  $S_{125}$  and not primary energy. A correction factor is still determined by converting the  $S_{125}$  values to the corresponding energy. It gets, however, applied to the muon densities as a function of  $S_{125}$  rather than expressing the final result as function of primary energy. The uncertainty related to converting between  $S_{125}$  and primary energy is taken into account as a systematic uncertainty.



**Fig. 6.9:** Left: Average muon density  $\rho_\mu$  at 600 m and 800 m in bins of true primary energy obtained from simulations of different primary nuclei. Right: Comparison of the average  $\rho_\mu$  values at 600 m obtained from simulations with different hadronic interaction models.

### 6.3.4 Systematic uncertainties

In the next section, we will compare the composition interpretation obtained from experimental data using the different observables to each other. To be able to make a meaningful comparison, it is important to know what the systematic uncertainties on the measurements are. Given that this study includes only observables that were already used in past analyses, the systematic uncertainties used in those cases can simply be reused here.

For the determination of the average  $dE/dX_{1500}$  as function of  $S_{125}$ , the relevant systematic uncertainties are those used in the composition analysis described in Section 4.3.1. The main uncertainty comes from the ice model, as introduced in Section 3.2.2. Uncertainties on scattering, absorption, hole ice, and DOM efficiency are combined in quadrature to obtain a total light yield uncertainty of  $+9.6\%/-12.5\%$  [193]. The systematic uncertainty on the average  $dE/dX_{1500}$  values is obtained by rerunning the energy loss reconstruction, modifying the parameter that describes the light-collection efficiency in the algorithm. There are two uncertainties related to IceTop which can influence the  $dE/dX_{1500}$  result, as they may change the  $S_{125}$  value, or the track along which the energy loss reconstruction is performed. First, there is the uncertainty of 0.2 m on the snow attenuation length  $\lambda_{\text{eff}}$  from Eq. (4.3) [161]. The effect on the result is evaluated by rerunning the reconstructions with modified  $\lambda_{\text{eff}}$  values. Secondly, there is a 3% uncertainty on the VEM calibration [194]. This was taken into account by shifting

measured charges up and down by 3%. All shifts in the average  $dE/dX_{1500}$  value obtained from data will be added in quadrature to obtain the total systematic uncertainty<sup>5</sup>.

For  $\beta$ , we consider the same uncertainties on  $\lambda_{\text{eff}}$  and on the VEM definition. Similar to before, the systematic uncertainty in  $\beta$  is determined from rerunning the processing on data and adding the resulting shifts in quadrature. For more details on how the individual uncertainties affect  $\beta$  and  $dE/dX_{1500}$ , see Ref. [157].

The systematic uncertainties on the muon densities  $\rho_{\mu}^{600}$  and  $\rho_{\mu}^{800}$  are somewhat more complex. First, there is a systematic uncertainty associated with the function used to model signals without muons in the signal fits used to derive the muonic signal. There is also an uncertainty in the energy estimation related to several effects described in Ref. [161]. Furthermore, uncertainties in the derivation of the correction factor are included as systematic uncertainties, namely its statistical uncertainty and an uncertainty related to the primary composition. For more details, we refer to Ref. [195].

## 6.4 Internal consistency of models

Measurements of different air-shower observables sensitive to the primary mass should lead to agreeing conclusions about the composition of cosmic rays. Inconsistencies would be a sign of deficiencies of the hadronic interaction model used in the simulation. We test the internal consistency of air-shower simulations based on Sibyll 2.1, QGSJet-II.04, and EPOS-LHC by comparing the mass-sensitive observables described in the previous section between data and simulation.

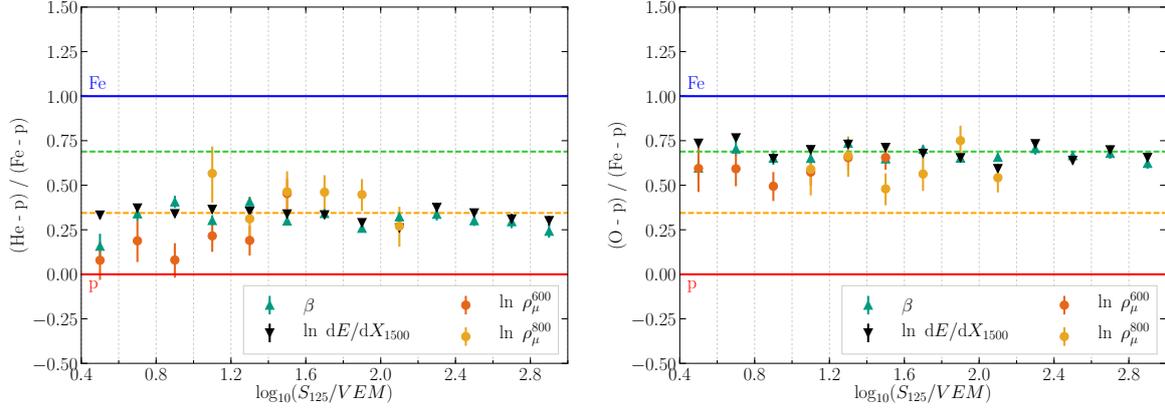
The observables are studied as a function of  $\log_{10} S_{125}/\text{VEM}$ , in bins of width 0.2, with a lower limit of 0.4 and an upper limit of 3.0 for Sibyll 2.1 and 2.0 for QGSJet-II.04 and EPOS-LHC, as described in Section 6.2. To compare the different observables, they are plotted in a “proton-iron space”, by calculating in each bin<sup>6</sup>

$$\frac{x_{\text{data}} - x_{\text{p}}}{x_{\text{Fe}} - x_{\text{p}}} \quad (6.1)$$

for each observable  $x \in \{\beta, \ln(dE/dX_{1500}), \ln(\rho_{\mu}^{600}), \ln(\rho_{\mu}^{800})\}$ . The terms  $x_{\text{p}}$  and  $x_{\text{Fe}}$  are derived from proton and iron simulation based on a specific hadronic interaction model. Eq. (6.1) is constructed so that if the values in data agree with a pure proton

<sup>5</sup>The same systematic uncertainties will be used for the analysis presented in Chapter 7, where we will discuss in more detail the impact of the different uncertainties on a measurement of the multiplicity of high-energy muons in air showers.

<sup>6</sup>Note that this is similar to the  $z$ -scale introduced in Eq. (2.27), which has been defined as a way to compare muon measurements from different experiments.



**Fig. 6.10:** Comparison of helium (left) and oxygen (right) simulation to proton and iron simulation using Sibyll 2.1. The calculated  $z$ -values are given for the different observables as function of  $S_{125}$ . The dashed horizontal lines are drawn at  $\ln 4 / \ln 56$  and  $\ln 16 / \ln 56$ .

composition, it gives a value of 0, while if it agrees with iron, it is 1, with a mixed composition somewhere in between. For the comparison of the different observables, it is important that they depend on the primary mass  $A$  in the same way. The reason to use these specific observables, and to take the logarithm for some of them, is that for variables  $x$  that are linear in  $\ln A$ , Eq. (6.1) reduces to

$$\frac{\ln A_{\text{data}}}{\ln 56}. \quad (6.2)$$

From the Heitler-Matthews model discussed in Section 2.2, this relation is expected for the logarithm of the number of muons in a shower. By evaluating Eq. (6.1) for helium and oxygen simulation instead of experimental data, we can see that this relation also holds in good approximation for the observables under consideration here. This is illustrated in Fig. 6.10.

Note that, as discussed in the previous section, the proton and iron values for  $\beta$  and  $\ln(dE/dX_{1500})$  are the average values in the IceTop simulation, while for  $\ln(\rho_\mu)$ , they are the true muon density values obtained from CORSIKA simulation. For the event-by-event observables  $\beta$  and  $\ln(dE/dX_{1500})$  the mean value in each bin is calculated, while the muon densities  $\rho_\mu$  are determined over the entire sample<sup>7</sup>.

<sup>7</sup>The muon density method finds the average in each bin, i.e.  $\langle \rho_\mu \rangle$ , so the values plotted are  $\ln \langle \rho_\mu \rangle$ . For  $dE/dX_{1500}$ , both  $\ln \langle dE/dX_{1500} \rangle$  and  $\langle \ln dE/dX_{1500} \rangle$  were studied, and it was found that the agreement with expectations for He and O simulations is marginally better for the latter. We therefore chose to include the observable in this way, although the difference in the results is small.

The analysis is performed with 10% of the data of the IC86.2012 season<sup>8,9</sup>. The results are shown in Fig. 6.11. The systematic uncertainties are calculated as discussed in the previous section; they are largest for the muon density results. Although the dataset used is limited, the results are clearly dominated by systematic uncertainties up to the energies accessible with the datasets simulated without thinning.

It is found that for all models the curves of the observables increase with  $S_{125}$ , consistent with a composition that becomes heavier with primary energy. Barring this general trend, there are however clear differences between the observables. For Sibyll 2.1,  $\beta$  indicates an average composition that is much heavier than for the low- and high-energy muon measurements. Its values even go beyond the predictions for iron simulation, which is physically implausible. The muon observables, on the other hand, overlap with each other for Sibyll 2.1, giving a consistent composition interpretation. For the post-LHC models, we see the biggest differences compared to Sibyll 2.1 in  $\beta$  and  $\rho_\mu$ , with a clear shift towards a lighter composition. In the case of QGSJet-II.04, there is little overlap between all observables, which cover a large fraction of the area between proton and iron. For EPOS-LHC, the  $dE/dX_{1500}$ -curve moves slightly upwards, leading to a better agreement with  $\beta$ , but a larger discrepancy with the low-energy muons.

Note that the results for the muon density are qualitatively consistent with the full analysis introduced in Section 4.3.3. The results for  $dE/dX_{1500}$  are also in agreement with those of the prior analysis introduced in Section 4.3.2. For  $\beta$ , however, a small shift toward iron-like values is seen in this analysis. This is likely a result of the fact that the previous analysis only used April data, which deviates further from the simulated snow heights compared to a yearly average. This effect should be studied in more detail when  $\beta$  is included in future analyses.

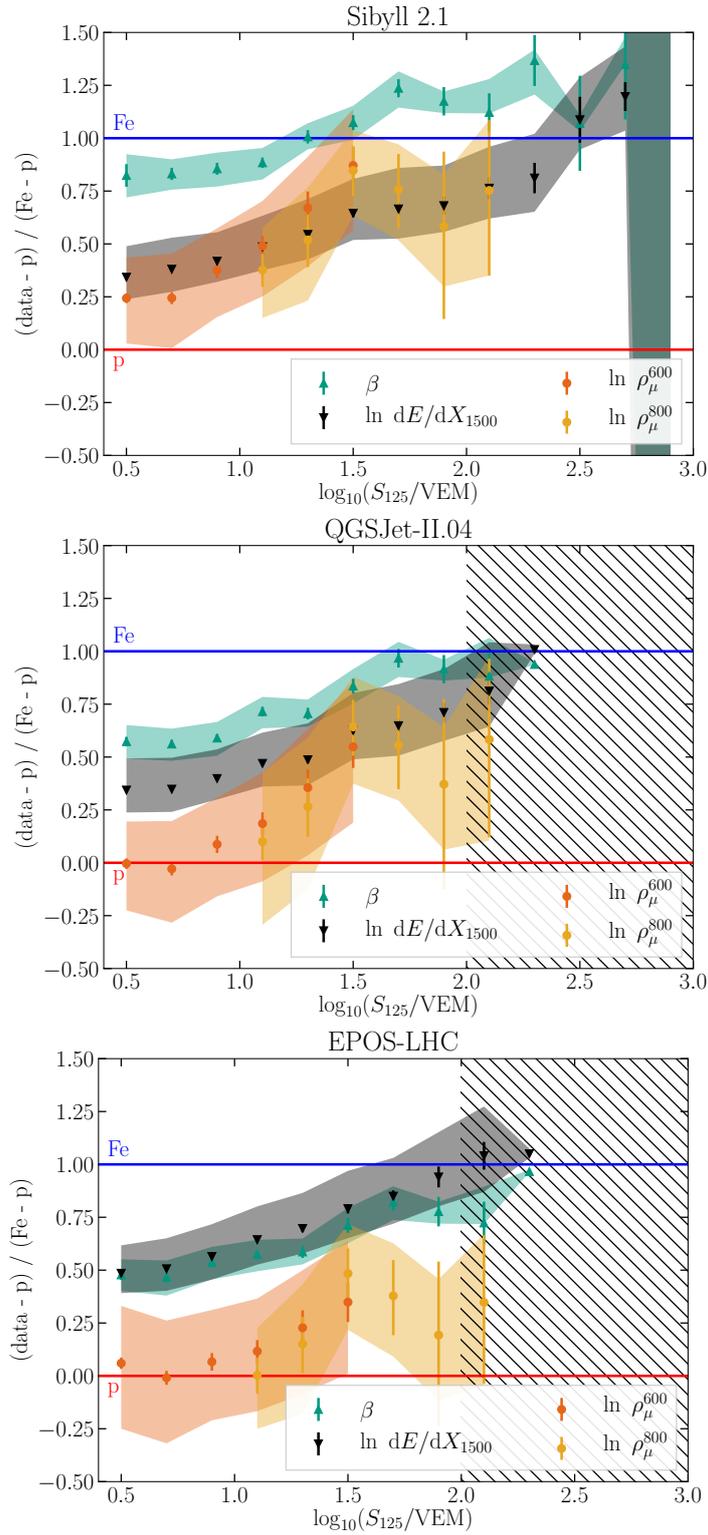
## 6.5 Discussion

In this chapter, the capability of IceTop and IceCube to measure different components of air showers at the same time has been applied to perform a test of hadronic interaction models. This was done by comparing different mass-sensitive observables in experimental

---

<sup>8</sup>The reason for the limited dataset is that this analysis in its current form is considered preliminary, due to the usage of detector-level variables, something we try to improve on with a dedicated high-energy muon analysis in Chapter 7.

<sup>9</sup>The 10% data consists out of all runs ending in zero. As a result, it is spread rather uniformly over the year. As discussed in Chapter 4, the atmosphere averaged over a year resembles a typical atmosphere of the month April rather well, so that the data can be directly compared to the simulation, which uses an April density profile. Similarly, the simulated snow coverage corresponds to measurements taken in October, roughly in the middle of the season.



**Fig. 6.11:** Different composition-sensitive observables as a function of the primary energy estimator  $S_{125}$  in the proton-iron space as represented by Eq. (6.1). The error bars give the statistical uncertainty, while the bands give the systematic uncertainty. Results are limited to  $\log_{10} S_{125}/\text{VEM} = 2$  due to a limited availability of simulations.

data to expectations from simulations of proton and iron showers. Given that all observables are obtained from the same experimental data, they should lead to consistent conclusions about the underlying mass composition.

From the result plots of Fig. 6.11, several conclusions about the current state of hadronic interaction models and indirect cosmic-ray composition measurements can be drawn. Comparing the same observable between different hadronic models shows clear shifts in the composition interpretation. As a result, a composition measurement based on one of these observables will show significant uncertainties related to the model used for the analysis. But the uncertainties go even further: for any of the models considered, discrepancies are observed in the composition interpretation between different observables. This implies, for example, that performing an IceTop-only composition analysis based on low-energy muons, one may obtain results that are in disagreement with the composition analysis using  $dE/dX_{1500}$  as described in Section 4.3.1. It also implies that one should be very careful when combining observables in a machine learning algorithm trained on air-shower simulations, as the distributions of the observables in simulations are different from those in data, or when comparing composition measurements of different experiments. The disagreements between observables indicate issues in the description of the hadronic interactions that govern the shower development, even when using state-of-the-art models.

While the study presented in this chapter already highlights clear problems in the considered models, it can still be considered somewhat preliminary. The detector observables  $\beta$  and  $dE/dX_{1500}$ , for example, are not easily interpreted in terms of a single property of the air shower, in contrast to the muon density  $\rho_\mu$ . This may limit the usefulness of this study as an input for future iterations of hadronic interaction models. For example, it is not clear if the discrepancies between  $\beta$  and  $\rho_\mu$  are a result of a wrong description of the lateral distribution of muons, an issue in the description of the EM shower component, or something else. The analysis could therefore be improved in several ways.

One possible improvement is to enhance the determination of high-energy muon content of air showers compared to  $dE/dX_{1500}$ . While  $dE/dX_{1500}$  is clearly related to the number of high-energy muons in an air shower, as demonstrated in Fig. 4.8, the resolution is not great, which makes it hard to interpret precisely. This motivates the analysis presented in the next chapter, where we will try to perform a precise analysis to obtain the number of muons above a certain energy threshold in an air shower, as a high-energy equivalent to the GeV muon density analysis discussed in Section 4.3.3. Another improvement could come from more clearly interpretable LDF-based observables

rather than  $\beta$ . For example, studying the possibility of performing a double-LDF fit to describe the EM and muon component separately would be useful. Other improvements could come from a reduction in systematic uncertainties. An improved reconstruction of GeV muons, which has the largest systematic uncertainty, could solidify the indications of discrepancies with other observables. For  $dE/dX_{1500}$ , improved systematic uncertainties should be in reach with the more modern ice models that have been developed compared to the one used in our current simulations.

One aspect that makes a relatively simple analysis of the kind presented here appealing is that it can easily be updated in the future. For example, if new mass-sensitive observables are introduced, they can be added to learn more about their behavior compared to previously used observables, as well as to further scrutinize the hadronic interaction models. In this regard, the planned upgrade of the surface array with scintillators and radio antennas [101, 102], as well as the proposed air-Cherenkov array IceACT [196], will bring promising opportunities. Of particular interest would be a measurement of  $X_{\max}$  accessible with the radio antennas. In addition to mass-sensitive observables, the new detectors could bring a larger accessible energy range and a better handle on uncertainties such as the snow accumulation on IceTop.



# Chapter 7

## Multiplicity of TeV muons analysis

Muons in air showers are tracers of the hadronic cascade and contain information about the primary cosmic-ray mass. Measurements from several observatories indicate that the production of these muons is not well understood, motivating further detailed measurements. With IceTop and IceCube, the unique possibility presents itself to measure muons in very different energy regimes in the shower, therefore probing the muon energy spectrum. Preliminary studies comparing several muon-related observables, described in Chapter 6, indicate possible discrepancies between low- and high-energy muons. Motivated by this, we present in this chapter a first analysis obtaining the average multiplicity of TeV muons in air showers as a function of primary energy using IceTop and IceCube coincident events, complementary to the existing IceTop measurement of the GeV muon density.

In Section 7.1, the idea for the analysis is introduced with some relevant background information. The event selection for the analysis is described in Section 7.2. In Section 7.3, we obtain predictions for the TeV muon multiplicity from simulations, to be used in various calculations. Next, different neural-network reconstructions for the primary energy and muon multiplicity are described in detail in Section 7.4. A method for correction for biases in the end result, caused by imperfections in the reconstructions, is described in Section 7.5, after which various tests of the method are given in Section 7.6. In Section 7.7, the method is applied to experimental data, which, after a discussion of the systematic uncertainties, leads to the final result plots of this work. These results are followed by a first look at seasonal variations in the high-energy muon multiplicity in Section 7.8. Finally, a discussion of the results and an outlook toward the future is given in Section 7.9.

## 7.1 Introduction

In Chapter 2, the relevance of muons in air-shower measurements was introduced. It was also shown that simulations based on different hadronic interaction models predict significant differences in the muon content of air showers, and that they do not necessarily agree with what is observed in experimental data. While most air-shower experiments, including IceTop, dominantly see muons with energies in the GeV range, the large ice sheet above IceCube specifically selects muons with energies above several hundred GeV. As was illustrated in Fig. 2.11, measurements of muons in these two energy ranges can provide important information toward a resolution of the problems with existing hadronic interaction models, as they provide unique tests of muon production models.

Several steps toward such tests have already been taken. A measurement of the density of GeV muons observed in IceTop was presented in Ref. [195] and summarized in Section 4.3.3. In Chapter 6, the same procedure was applied to a set of IceTop-IceCube coincident events as a first step in comparing measurements of the GeV muons with the TeV muons in IceCube. A comparison indicated inconsistencies between the composition interpretation obtained from the GeV muons and from the energy loss of the high-energy muon bundle for several models. This energy loss is, however, an indirect probe of the actual high-energy muon content, and is not necessarily easy to interpret<sup>1</sup>. Therefore, we develop in this chapter a first analysis aiming to obtain the average multiplicity of high-energy muons in air showers, as a complement to the IceTop GeV muon density analysis.

The idea for the analysis is to use IceTop-IceCube coincident air-shower events, and to obtain for each event an estimate of the primary cosmic-ray energy and the multiplicity of muons above a certain threshold in the shower. This estimate will be obtained using neural networks using as input information from IceTop for the energy and the muon bundle energy-loss profile in IceCube for the muon multiplicity. Using this, the average muon multiplicity as a function of energy will be determined, after which correction factors obtained from simulation will be applied when necessary. For simplicity, we start from existing tools to obtain quality coincident events, and restrict ourselves to near-vertical showers. The target we reconstruct is the number of muons in the shower

---

<sup>1</sup>For muons of a fixed energy, the average energy loss of a bundle can be expected to be proportional to the number of muons. However, as demonstrated in Fig. 3.2, higher energy muons have a higher average energy loss. Therefore, a degeneracy exists in the muon bundle energy loss between the number of muons in the bundle and their energies.

above an energy threshold chosen at 500 GeV, determined at the surface. At this energy, nearly all muons will reach the in-ice detector<sup>2</sup>.

In Eq. (2.27), a common way of representing muon data by comparing them to expectations from simulations using proton and iron primaries, was introduced. As it will frequently appear in the following sections, it is repeated here:

$$z = \frac{\ln\langle N_\mu \rangle - \ln\langle N_\mu \rangle_{\text{p}}}{\ln\langle N_\mu \rangle_{\text{Fe}} - \ln\langle N_\mu \rangle_{\text{p}}}, \quad (7.1)$$

where  $\langle N_\mu \rangle$  is the result obtained from data, and the subscripted values are the expectations from proton and iron simulation. Note that alternatively, one could define  $z$  replacing  $\ln\langle N_\mu \rangle$  with  $\langle \ln N_\mu \rangle$ . Differences between the definitions and related uncertainties are discussed in Refs. [197] and [198]. To be consistent with the IceTop GeV muon analysis and the meta-analysis of different muon measurements discussed in Section 2.4.3, we adopt the  $\ln\langle N_\mu \rangle$  definition, as given above. Therefore, the aim of the analysis is to obtain  $\langle N_\mu \rangle$  as a function of primary energy, where  $N_\mu$  refers to the number of muons above 500 GeV in the shower.

## 7.2 Event selection

The analysis uses events that are coincident between IceTop and IceCube. To obtain events that can be reconstructed well, various existing tools that have been described in Chapter 4 are used. We start from the standard Cosmic Ray Level 3 Processing, and apply the quality cuts described in Section 4.2.3, which have been developed in the past for the IceTop-IceCube coincident composition analysis of Section 4.3.1.

The IceTop quality cuts ensure that the shower core is located within the bounds of the IceTop array and leads to a set of events that have a good reconstruction of both shower core position and direction, and have a shower size parameter  $S_{125}$  that is well correlated with the primary energy.

The IceCube quality cuts select events with a clear muon bundle signature to which a successful energy loss reconstruction has been performed, an example of which was shown in Fig. 4.4. This energy loss reconstruction will be crucial to estimate the muon multiplicity.

A further cut is applied to the reconstructed zenith angle to obtain near-vertical showers with  $\cos\theta > 0.95$  or  $\theta \lesssim 18^\circ$ . Restricting the zenith range in this way significantly

---

<sup>2</sup>We refer to the muons 500 GeV collectively as the TeV muon component. From Fig. 5.2, it can be seen that the average energy of muons above this threshold is indeed of the order of  $\sim$ TeV.

simplifies the problem for a first measurement of this type. This is because the number of high-energy muons in an air shower is expected to increase approximately with  $1/\cos\theta$ , as can be seen in Eq. (2.23), while at the same time, the energy lost by muons before reaching the detector will increase with  $\cos\theta$ . With the restricted zenith range, any possible complications caused from these effects should be unimportant, allowing us to focus on establishing a detailed first measurement of muons in a small phase space, which can be expanded in future iterations of the analysis.

Distributions of several variables in simulation and the performance of different reconstructions after the cuts described above were already shown in Section 4.2.4.

One extra cut is applied to the events based on the reconstructed energy-loss profile in the detector, which is transformed into a fixed-length vector to be fed into a neural network, as will be described in detail in Section 7.4.2. This vector will be required to have at least three non-zero entries after the removal of segments outside the detector.

The analysis will also be limited to simulation and data from the IC86.2012 season, as this is the season for which most IceTop CORSIKA simulation is available, as shown in Table 4.3. As the neural network is trained on simulation with a simulated snow height of this season specifically, it may become unreliable when applied to experimental data for years when the snow height is significantly different. The distribution of simulated events remaining after the selection, on which the neural networks will be trained, will be shown in Section 7.4.1. Distributions in several variables obtained from data after the selection will be included in Section 7.6.2.

The lower limit of the energy range considered is defined by the threshold where all primary types reach maximal efficiency for triggering IceTop. As discussed in Chapter 4, this is about  $10^{6.4}$  GeV for vertical showers and the IC86.2012 snow coverage. The upper limit of the analysis will in practice be determined by running out of statistics due to the steep primary spectrum, but should in principle be able to extend up to  $\sim$ EeV energies. To confirm these statements, we calculate the so-called effective area  $A_{\text{eff}}$  of the detector at analysis level, defined as

$$A_{\text{eff}}(E) = A_{\text{sampling}}(E) \frac{\cos\theta_{\min} + \cos\theta_{\max}}{2} \epsilon(E). \quad (7.2)$$

Here,  $A_{\text{sampling}}$  is the area over which showers are generated, as discussed in Section 4.4, and the zenith range is in this case given by  $\cos\theta_{\min} = 1$  and  $\cos\theta_{\max} = 0.95$ . The efficiency  $\epsilon$  is defined as

$$\epsilon(E) = \frac{N_{\text{eff}}(E)}{N_{\text{gen}}(E)}, \quad (7.3)$$

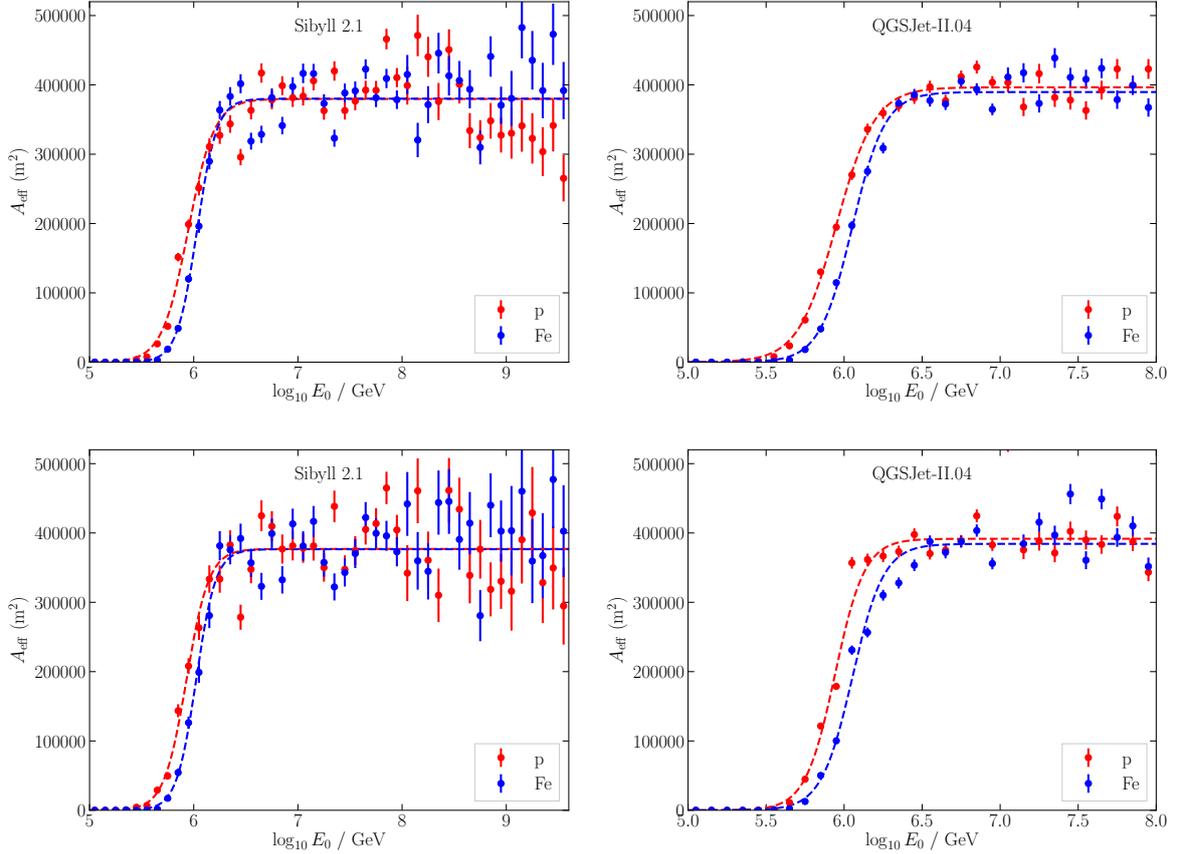
where  $N_{\text{gen}}$  is the total number of showers generated in an energy bin in the considered zenith range, and  $N_{\text{eff}}$  is the total number of events after the event selection with a reconstructed energy and zenith in the considered energy and zenith bins. Results obtained for the effective area from proton and iron simulation are shown in Fig. 7.1. The error bars come from assuming Poissonian errors on  $N_{\text{eff}}$ . We include plots based on the high-statistics QGSJet-II.04 datasets for a detailed examination of the threshold region, and plots based on Sibyll 2.1 to examine a larger energy range (see Table 4.3). In the top plots, we calculate  $N_{\text{eff}}$  based on the true energy of the events, and in the bottom plots, we use the energy reconstruction of the NENN method, which will be discussed below in Section 7.4.3. A clear threshold behavior can be seen: showers with energies below  $10^{5.6}$  GeV rarely make it through the event selection, however, with increasing energy, the effective area increases and becomes approximately flat. The behavior is fitted with a sigmoid function,

$$A_{\text{eff}} = \frac{p_0}{1 + \exp(-p_1 \times \log_{10}(E_0 / \text{GeV}) + p_2)}, \quad (7.4)$$

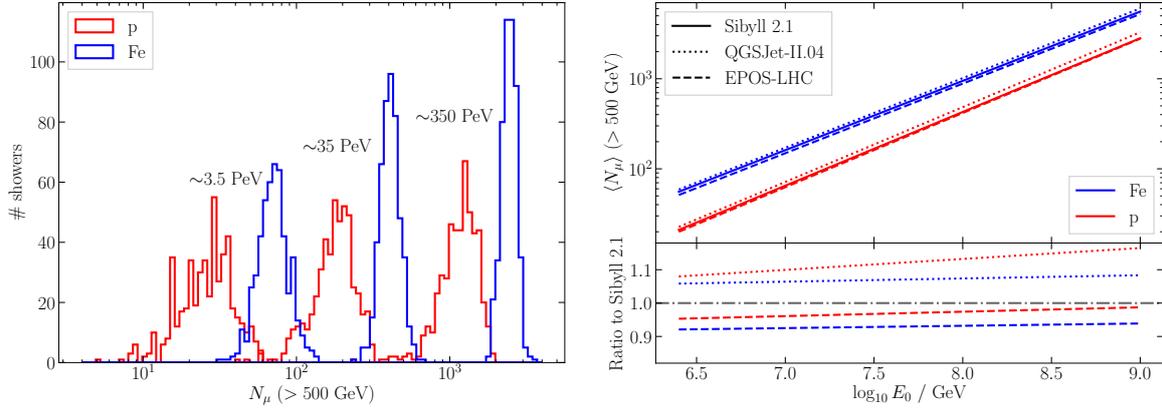
where  $(p_0, p_1, p_2)$  are free parameters. In all cases, the fitted curves reach full efficiency, which we define as 97% of the maximum value  $p_0$ , at an energy of  $10^{6.4}$  GeV or before. After this point, the effective area has a value that is about  $390 \times 10^3 \text{ m}^2$ , with a slight variation depending on the hadronic model and energy that is used. What is important, however, is that the effective areas obtained for proton and iron are in agreement after the threshold, so that no composition-dependent bias is present in the event selection. This is the case here, confirming that  $10^{6.4}$  GeV is a valid threshold for the analysis. The effective area remains approximately constant with energy, although a possible energy dependent effect can be seen in proton beyond  $10^{8.5}$  GeV. This should be studied in more detail when expanding the analysis to these energies. However, this is no issue in the results presented in this chapter, which are limited either by statistics in data or available simulation, as will be discussed later.

### 7.3 MC expectations

It is useful to obtain expectations from simulation for the quantity we try to measure in this analysis, namely the number of muons above 500 GeV in air showers. For this, dedicated simulations were produced using CORSIKA for proton and iron primaries at fixed energies of  $\log_{10} E_0 / \text{GeV} = 6.45, 6.55, \dots, 8.65$ . For simplicity, the showers were simulated at a fixed zenith angle representative for the dataset; not the average zenith



**Fig. 7.1:** Effective area as a function of primary energy for proton and iron calculated as in Eq. (7.4). Results are shown for simulated datasets of Sibyll 2.1 (largest energy range) and QGSJet-II.04 (improved statistics), listed in Table 4.3. In the top plots, the effective area is calculated using the true energy for the events remaining after the event selection, in the bottom plots the reconstructed energy from the NENN neural network is used (Section 7.4.3). The behavior is fitted with a sigmoid function. The statistical uncertainty is larger in the bottom left plot because when performing the calculation using reconstructed energy, only the test set for Sibyll 2.1 was used, which is 40% of the full Sibyll 2.1 set included in the top left plot (see again Section 7.4.3).



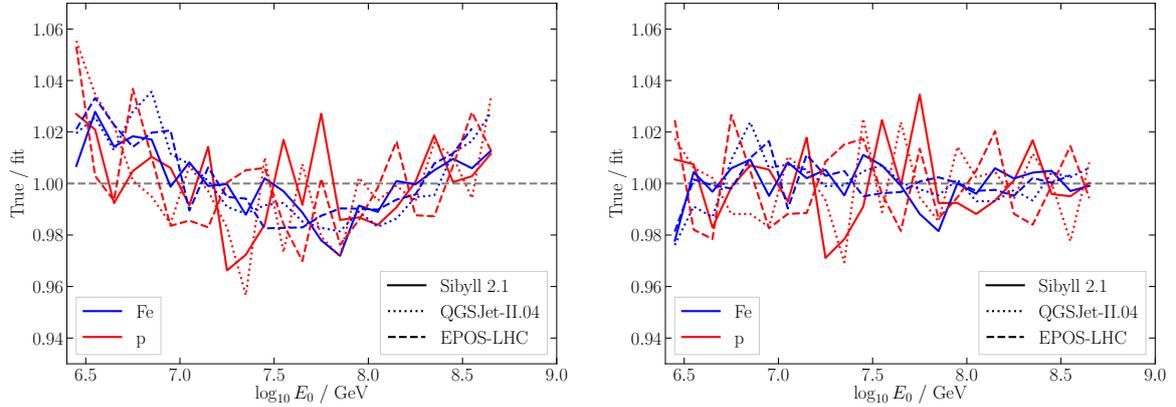
**Fig. 7.2:** Left: Distribution of the number of muons with energies higher than 500 GeV in simulated proton and iron showers at three different energies with a zenith of  $11.9^\circ$ . Right: Average multiplicity of muons above 500 GeV as a function of primary energy. Lines are fits to values obtained from simulations using three different hadronic interaction models, given by Eq. (7.5) and Table 7.1.

$\langle\theta\rangle = 11.0^\circ$  in the dataset after the event selection (Fig. 7.4), but rather  $\theta = 11.9^\circ$ , which corresponds to the angle equivalent to  $\langle 1/\cos\theta \rangle$  obtained from the dataset. The reason for this is that the muon multiplicity is expected to be proportional to  $1/\cos\theta$ , as in Eq. (2.23). For each primary type and energy, 500 showers were simulated. An example of the distributions of the number of muons above 500 GeV obtained using Sibyll 2.1 is shown in the left panel of Fig. 7.2. At each energy, the mean of the distribution is calculated with its error. The relation between the obtained  $\langle N_\mu > 500 \text{ GeV} \rangle$  and the primary energies is fitted with a function quadratic in the logarithms of the variables,

$$\log_{10}\langle N_\mu \rangle = p_0 + p_1 \log_{10} E_0 + p_2 \log_{10}^2 E_0, \quad (7.5)$$

where  $p_0$ ,  $p_1$ , and  $p_2$  are the fit parameters. This process is repeated for simulations using QGSJet-II.04 and EPOS-LHC.

The results are shown in the right panel of Fig. 7.2. While the relation looks very linear, as one could expect from the Heitler-Matthews model (Eq. (2.22)), it was found that the quadratic function significantly improves the fit, as illustrated in Fig. 7.3. The obtained values of the fit parameters are given in Table 7.1. The uncertainty on the fits is sufficiently small that it can be safely ignored in what follows. The fits will be used for various purposes throughout the analysis, such as for calculating  $z$ -values. The goal of this analysis is effectively to add values obtained from experimental data to the plot showing the expectations from MC, the right panel of Fig. 7.2.



**Fig. 7.3:** Ratio between a fit to and the actual values of  $\langle N_\mu > 500 \text{ GeV} \rangle$  obtained from simulation, using a linear (left) and quadratic (right) function. The quadratic function clearly improves the description of the behavior.

**Table 7.1:** Parameter values describing the expectations for the average multiplicity of muons above 500 GeV obtained from fits of Eq. (7.5) to simulations.

	Proton			Iron		
	$p_0$	$p_1$	$p_2$	$p_0$	$p_1$	$p_2$
Sibyll 2.1	-3.469	0.699	0.008	-2.531	0.596	0.011
QGSJet-II.04	-3.008	0.575	0.017	-2.212	0.514	0.017
EPOS-LHC	-3.256	0.632	0.013	-2.283	0.518	0.017

It is interesting to remark that these simulations indicate a  $\beta \approx 0.8$  in the Heitler-Matthews formula Eq. (2.22) for the TeV muons, compared to the  $\beta \approx 0.9$  that is found for the total number of muons in a shower.

## 7.4 Neural-network reconstruction

To estimate the primary energy  $E_0$  and muon multiplicity  $N_\mu$  ( $> 500$  GeV), various existing reconstructions will be used as inputs to neural networks. In this section we discuss the simulated datasets used for training and testing, several approaches for the neural-network reconstructions, and their performance.

### 7.4.1 MC dataset

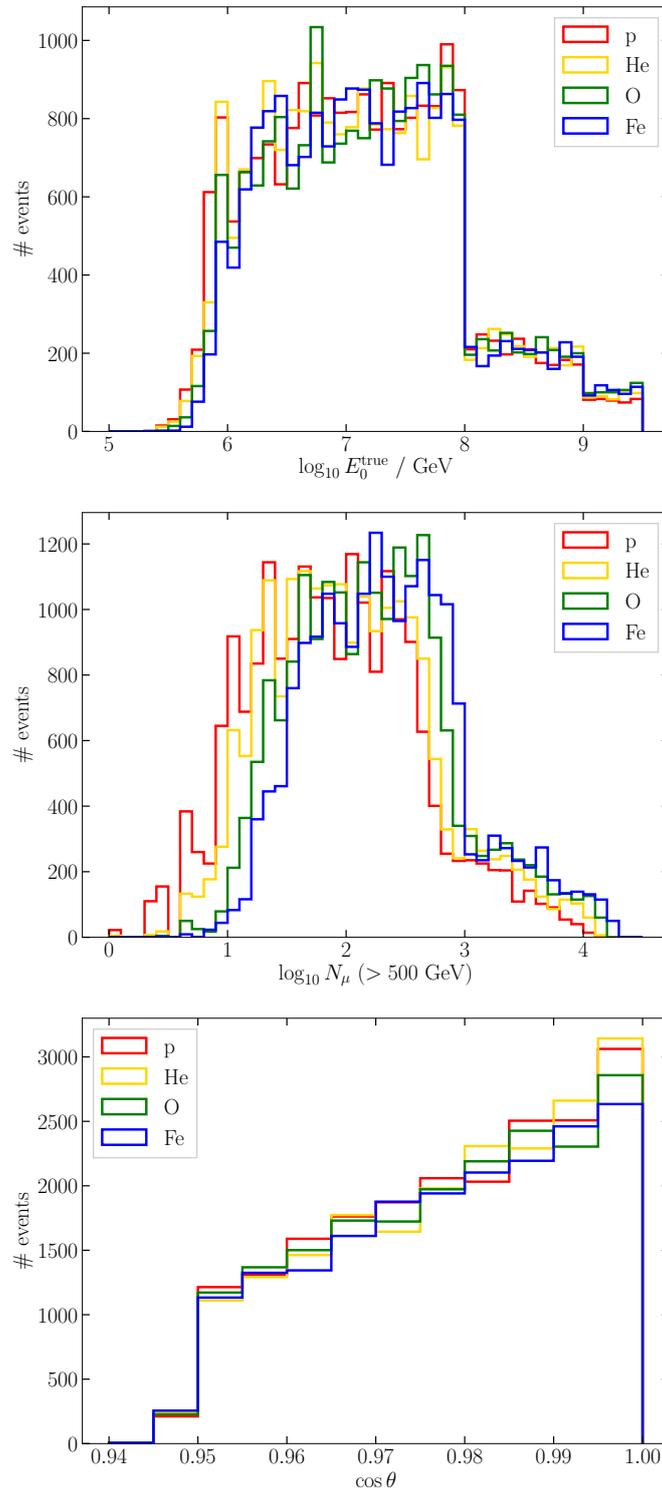
For training and testing of the neural network, the baseline IC86.2012 simulation sets based on Sibyll 2.1 listed in Table 4.3 are used. This includes four primary types (p, He, O, Fe) and covers an energy range  $10^5$ – $10^{9.6}$  GeV. The distribution of true energy, muon multiplicity, and zenith angle available in simulation after the selections of Section 7.2 are shown in Fig. 7.4. Events below  $10^{5.6}$  GeV are cut out during training, as they are not important for the further analysis and only add noise to the training process. The MC events will not be weighted during training. As such, the networks will see approximately equal amounts of the four primary types during training, and will train mostly on the 1 to 100 PeV range. In this way, we hope to obtain models with similar performance regardless of the primary mass over a large energy range.

### 7.4.2 Input

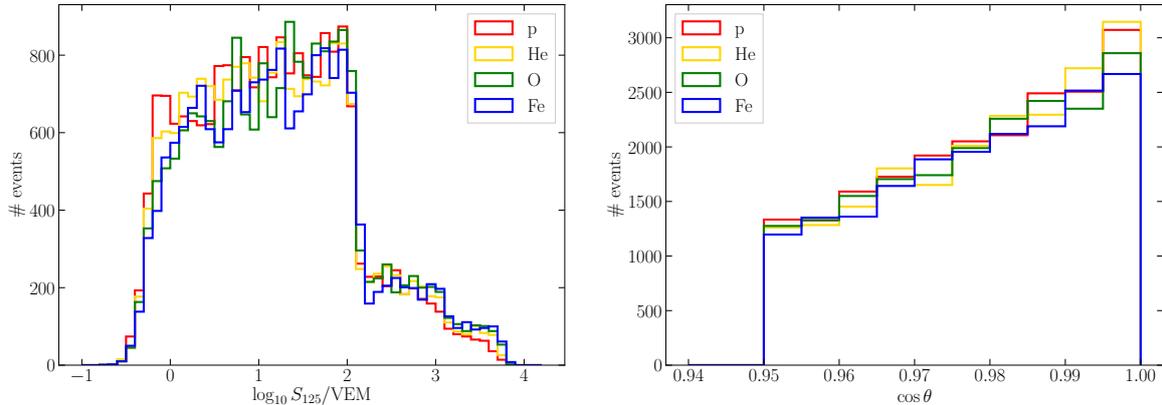
To reconstruct the primary energy and muon multiplicity of an event, neural networks will be trained using relevant input from IceTop and IceCube. The inputs are high-level variables taken from the existing air-shower and energy-loss reconstructions discussed in Chapter 4, and are summarized below. How exactly they are combined in different neural network models is discussed in Section 7.4.3.

#### IceTop input

The most important IceTop variable we use is the shower size  $S_{125}$ , obtained from the fit to the lateral distribution of charges observed in IceTop during the air-shower reconstruction, as described in Section 4.2.1.  $S_{125}$  will mainly be important for the energy reconstruction,



**Fig. 7.4:** True primary energy (top), muon multiplicity  $> 500$  GeV (middle), and cosine zenith (bottom) distributions in the Sibyll 2.1 datasets (not weighted) after the event selection.



**Fig. 7.5:** Distribution of the reconstructed shower size  $S_{125}$  (left) and zenith angle  $\theta$  (right) in the datasets used for training the neural networks.

as it is strongly correlated to the primary cosmic-ray energy with only a small dependence on the type of nucleus, as could be seen in Fig. 6.3. The distribution of  $S_{125}$  values in the MC datasets is shown in the left panel of Fig. 7.5.

Secondly, the reconstructed zenith angle  $\theta$  from the air-shower reconstruction will be used. This may provide extra information to the neural network for a more precise energy reconstruction, as well as for the muon multiplicity reconstruction, as it relates to the amount of energy muons lose before reaching the IceCube detector. The distribution of  $\theta$  in the MC datasets is shown in the right panel of Fig. 7.5.

### IceCube input

The input used from the IceCube detector is based on the energy loss reconstruction that is performed on the in-ice signals that are caused by the penetrating high-energy muon bundle, as described in Section 4.2.2. The Millipede algorithm performs a reconstruction of the energy loss in segments of 20 m along the seed track that is obtained from the air-shower reconstruction in IceTop. Starting from this reconstruction, we remove all segments that are outside the detector, which we define as a cylinder of height and radius 500 m around the center of the detector. The remaining segments in the detector, which can together be considered a 1-D vector, are padded with zeros to obtain a vector of equal length for each event. This was done initially to have the flexibility to use this energy-loss vector as an input to 1-D convolutional neural networks, which require an input of fixed size. The padding is done based on the geometry of the event and will as a result also carry some information on the zenith angle of the shower, which may

be useful for networks trained only on the IceCube input without the IceTop input (see next section).

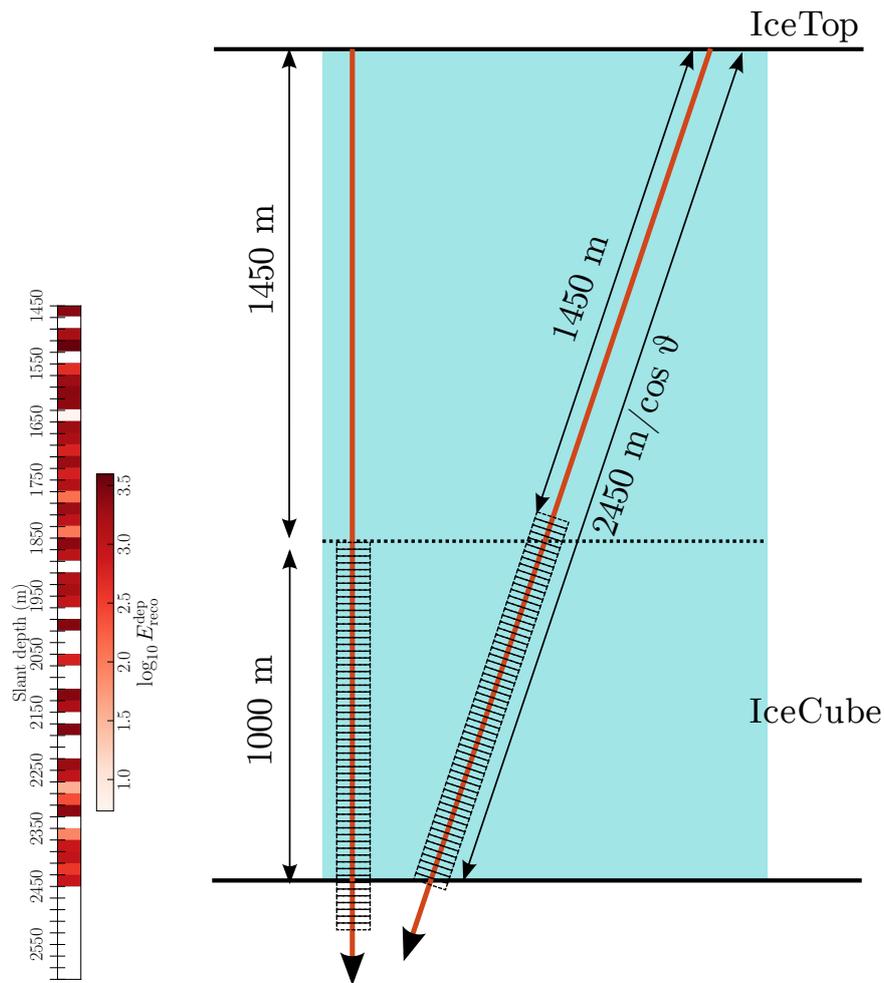
The vectors are constructed in such a way that the start of the first segment corresponds to a travelled distance of 1450 m in ice for the muons. This corresponds to the top of the detector for vertical showers. The last segment is chosen to correspond to the travelled distance of muons at the bottom of the detector for maximally inclined showers, which is in this  $2450 \text{ m} / \cos 0.95$ . To span this range of distances with segments of 20 m, a vector of length 57 is needed. For vertical showers, the segments at the end of the vector lie below the detector. They are given the value of zero. For more inclined showers, the first segments lie above the top of the detector, so the corresponding energy-loss vectors will be zero-valued at the start. This procedure is illustrated in Fig. 7.6. As already mentioned in Section 7.2, events whose resulting energy-loss vectors have less than three non-zero values are cut out from the analysis.

Some more minor manipulations are performed on the vectors. First, the logarithm of the energy losses is taken, as they span several orders of magnitude, which is difficult to deal with for neural networks. The distribution of vector entries in the MC dataset, i.e. 57 reconstructed energy loss values per event, is shown in the left panel of Fig. 7.7. As can be seen in the plot, there is a tail of energy losses extending toward extremely small values. As these values are not very physically meaningful, a constant value of -1 is given to all segments with a lower value. Finally, all vector entries are decreased by 2.8 and divided by 0.8. This results in a distribution that has a mean of approximately zero, and a standard deviation of approximately 1, which may help in training the neural network. The resulting distribution is shown in the right panel of Fig. 7.7. In hindsight, the manipulations bar the logarithm are probably not necessary, as the neural networks actually train rather effortlessly. However, even if it does not help, it will not hurt either.

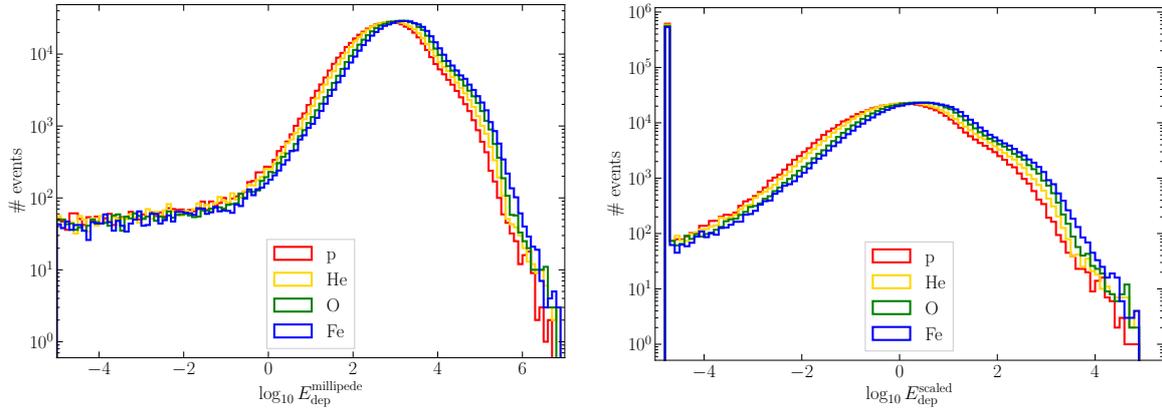
This energy-loss vector will mainly be important to determine the number of muons above the chosen threshold of 500 GeV in the shower. In Fig. 4.8, it was already shown how one simple variable obtained from performing a fit to the energy-loss profile,  $dE/dX_{1500}$ , is correlated to  $N_\mu$ . Using neural network techniques and the information of the full energy-loss profile in the detector, the aim is to improve on this.

### 7.4.3 Method

Combining the inputs described above with neural networks, we try to obtain an accurate estimate of the primary energy and muon multiplicity of the coincident events. The first steps toward this were taken by trying to obtain a reconstruction of the muon multiplicity based on only the energy-loss vector, as we already have a good estimator of



**Fig. 7.6:** Illustration of the construction of the energy-loss vector used as an input to the neural network. The vector has a fixed length defined by the maximal zenith angle in the dataset, and is padded with zeros at the top and/or bottom based on the reconstructed zenith of the event. Also shown is an example energy-loss vector for a vertical shower.



**Fig. 7.7:** Distribution of the values in the energy-loss vectors in the datasets used for training the neural networks. Each value of the vector enters the distribution, i.e. the number of entries is 57 times the number of events in the dataset.

the primary energy with  $S_{125}$ . Two main ways of using the energy-loss vector were tried: convolutional neural networks (CNNs) with 1-dimensional (1-D) convolution operations along the vector, and recurrent neural networks (RNNs). These are typical options for time series or more general sequence data, like the energy-loss profile of particles through IceCube. A detailed description of how these models work under the hood is outside the scope of this thesis; an excellent resource is Ref. [199], which gives an introduction to deep learning concepts and the Keras library [200], which is the library used for training the neural networks discussed here. The CNN approach for reconstructing  $N_\mu$  works well, but needs a lot of hyperparameter tuning (number of layers, convolution kernel size, etc.) to get optimal performance. It was found that a relatively simple RNN architecture, on the other hand, immediately performs equally well or even slightly better than the optimized CNN, with little variation when changing hyperparameters. One drawback of RNN models is that they train slower, but due to the limited size of the training set and the relative simplicity of the input (a 1-D vector of length 57), this never became an issue. Therefore, it was chosen to only continue with RNN models for further analysis. The exact RNN layer used is a bidirectional gated recurrent unit (GRU) with 64 units.

With that knowledge as a baseline, several approaches were explored to obtain a primary energy and muon multiplicity reconstruction. We introduce the different approaches and the nomenclature that will be used throughout this chapter below. A Keras visualization of all architectures is included in Appendix B.1.

**$N_\mu$  with Neural Network: “NNN”**

The NNN approach refers to a neural network model with as input the IceCube energy-loss vector and as output the multiplicity of muons above 500 GeV  $N_\mu$ . In this model, the energy-loss vector is an input to a recurrent layer as described above, the outputs of which go to a single node without activation function which gives a value for  $\log_{10} N_\mu$ .

 **$E_0$  with Neural Network: “ENN”**

The ENN approach refers to a neural network model with as input the shower size  $S_{125}$  and the reconstructed zenith angle  $\theta$  and as output the energy of the primary cosmic ray  $E_0$ . This is done by feeding the two inputs through two successive dense layers with 32 nodes, followed by a single node which predicts  $\log_{10} E_0$ .

 **$E_0$  with Conversion Function: “ECF”**

A different approach to obtain an estimate of the primary energy is through an approach that has already been discussed in Chapters 4 and 6, namely through a simple conversion function from  $S_{125}$ . Such a conversion function is of the form of Eq. (4.6), and examples derived from pure proton and iron simulation were shown in Fig. 6.3. In this case, we derive the conversion function from the four-component simulation weighted to H4a, similar to what was shown in Fig. 4.8. We will refer to this way of estimating  $E_0$  as the ECF approach.

 **$N_\mu$  and  $E_0$  with Neural Network: “NENN”**

The NENN approach refers to a neural network with as input both the IceTop inputs  $S_{125}$  and  $\theta$ , and the IceCube energy-loss vector, and as output both  $E_0$  and  $N_\mu$ . In such a combined model, the energy-loss vector is again an input to the same recurrent layer as before. Its outputs are now concatenated with the values of  $S_{125}$  and  $\theta$ , which collectively go into a dense layer with 128 units. This layer is followed by two nodes which return values for  $\log_{10} N_\mu$  and  $\log_{10} E_0$ .

**7.4.4 Training**

For the training of the neural networks, the MC dataset characterized by the distributions from Figs. 7.4, 7.5 and 7.7, is split into two parts: 60% is used for training the model, the remaining 40% is used to test the performance of the final model in an unbiased way and for further analysis. The split is done randomly on a per-primary-type basis, but it

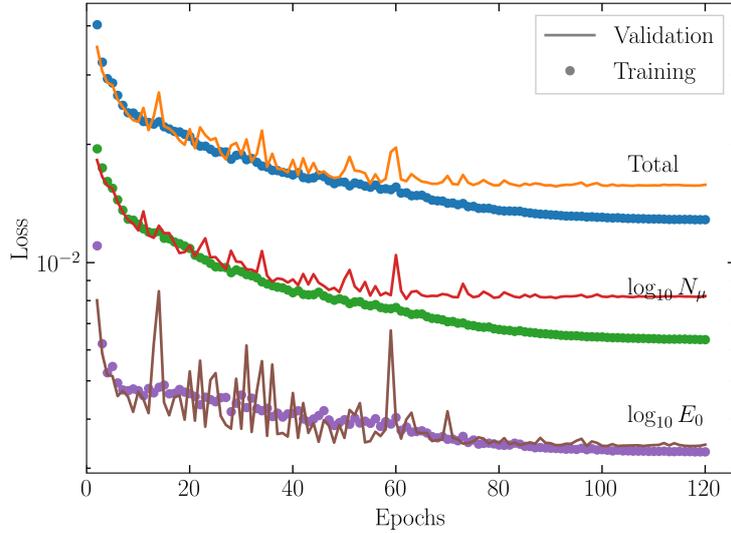
is checked explicitly that in each bin of  $\log_{10} N_\mu$  with width 0.1, the test set contains a fraction of the total events in that bin not outside the range of 20% to 60%, to ensure a somewhat homogeneous split that allows for sufficient events to test the performance over considered  $N_\mu$  range. The training set is split similarly into a part of 90%, which is used to actually train the model, and a validation set of 10%, which is used to monitor the training process and select at what point during the training the optimal model is obtained. The same train and test sets are used for every model we consider.

The problem we are dealing with here is clearly a regression task, where a neural network is used to predict a real-valued, continuous number. As neural networks generally do not handle regression targets that span several orders of magnitudes in value well, they are trained to predict the  $\log_{10}$  of  $E_0$  and/or  $N_\mu$ . For all models, a mean squared error (MSE) loss function is used. In case of the NENN model, the MSE loss function of  $E_0$  and  $N_\mu$  are combined with weights of 1 and 1.5 respectively, which was chosen based on the fact that  $N_\mu$  is the hardest quantity of the two to reconstruct. As optimizer, the Adam algorithm [201], a variation of stochastic gradient descent, is used with a learning rate of  $10^{-3}$ . An example of the evolution of the loss calculated on the training and validation sets is shown in Fig. 7.8. Around the epoch<sup>3</sup> where the loss starts to flatten, the learning rate is slowly decreased in an exponential way, with the aim of approaching the optimal neural-network configuration more closely. Throughout the training, the neural-network configuration is periodically saved. When we stop training the neural network after a large number of epochs, the different configurations from the region where the validation loss approximates a constant value are inspected in more detail. This is done by checking the performance on the validation set, namely the relative error in the prediction of  $N_\mu$  and/or  $E_0$ . The configuration which optimizes having minimal bias and a good resolution is then saved as the final model. In the case of Fig. 7.8, this was found to be the configuration obtained after epoch 92.

### 7.4.5 Performance

To examine the performance of the models that were obtained, they are applied to the test set. As the models have never seen the test set during training, this provides an unbiased measure of their accuracy. The accuracy will be demonstrated through different types of figures. For convenience, we will do this by separately considering the energy and muon multiplicity reconstructions for the different approaches. The emphasis here will

<sup>3</sup>Epoch is a term used for a full pass through the training dataset.



**Fig. 7.8:** Evolution of the loss functions during the training of the NENN model. Shown are the individual losses for the  $E_0$  and  $N_\mu$  reconstructions, together with the combined loss where the  $N_\mu$  loss has a weight of 1.5.

be on the muon multiplicity reconstruction, as this is the hardest quantity to reconstruct, and is where the main improvement of this work compared to previous work lies.

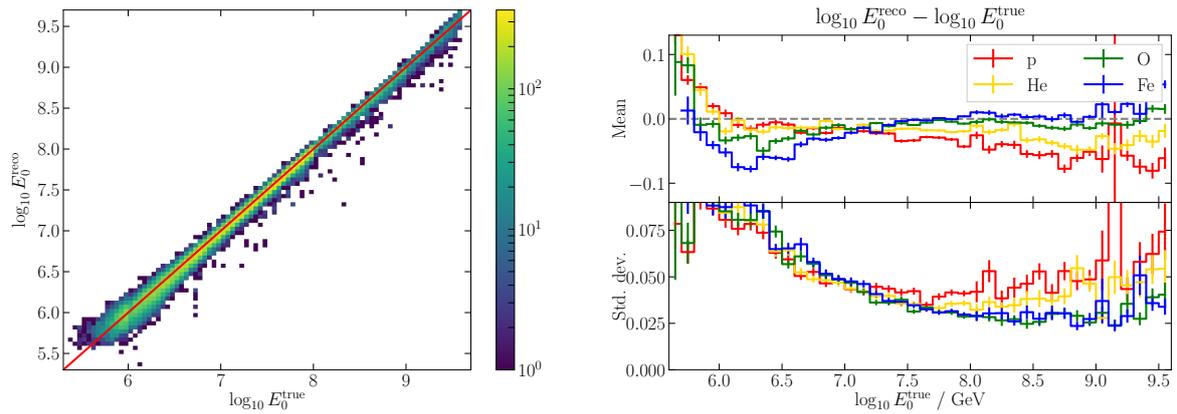
### Energy reconstruction

The performance of the energy reconstruction of the ECF, ENN, and NENN methods is shown in Figs. 7.9, 7.10 and 7.11 respectively. The left panel shows a correlation plot, where the neural-network prediction is plotted against the truth. This is done for the total test set combined, i.e. four types of primary, without any weighting. The red diagonal line shows the case of a perfect reconstruction. In the right panel, the bias and the resolution of the reconstruction in  $\log_{10} E_0$  is plotted. These are defined as the mean and standard deviation of the difference between the logarithm of the true and reconstructed values, in bins of true energy. They are derived for the four primaries separately to show possible mass-dependent effects. In this case, the simulation was also weighted to a realistic spectrum, the H4a all-particle spectrum (Section 1.5), although this has little impact on the actual values shown in the plots.

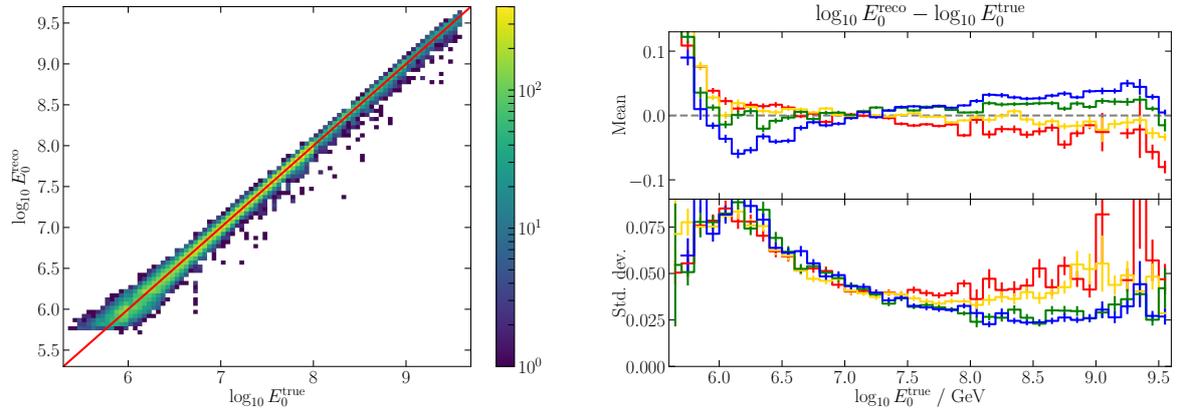
The correlation plots show that the energy reconstruction performs well over the entire energy range considered, without indications of problematic outliers. A more detailed understanding of the accuracy of the reconstructions can be obtained through the bias and resolution plots. All three methods show similar resolutions, decreasing

with primary energy until they flatten and sometimes slightly increase again from around 30 PeV. The resolution is slightly worse for lighter elements at higher energy. The ECF method, which uses a conversion function to convert  $S_{125}$  to  $E_0$ , and the ENN method, which uses a simple neural network based in  $S_{125}$  and  $\theta$ , also show similar behavior for the bias. It has a clear dependence on the primary type, with the difference between proton and iron becoming nearly 0.1 at its maximum. The ECF systematically underestimates the energy, so that it has a bias that is negative in most cases, while the bias of the ENN is more symmetric around zero. This could, however, be solved by simply shifting the conversion function up or down until one finds the desired behavior.

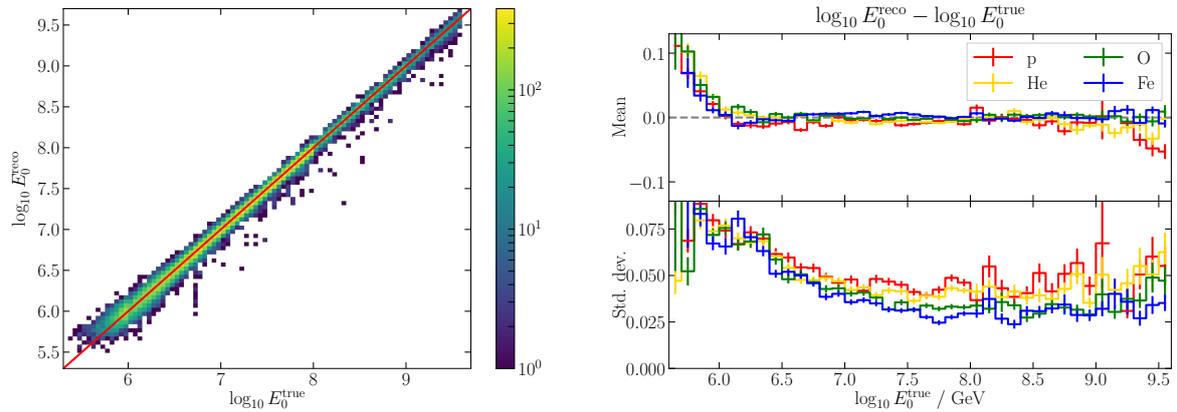
The NENN model improves somewhat upon the previous two. It outperforms them with respect to resolution at lower energies, and remarkably, the mass-dependent bias that was visible earlier is almost completely absent. This indicates that the combined neural network can use the extra information about the primary mass from the muon bundle signal to improve its energy reconstruction.



**Fig. 7.9:** Performance of the ECF primary energy reconstruction. Left: Two-dimensional histogram of the reconstructed and true values. The red line indicates the 1:1 correlation case. Right: Bias and resolution, defined as the mean and standard deviation of the difference between the reconstructed and true values in logarithmic scale, as a function of true energy.



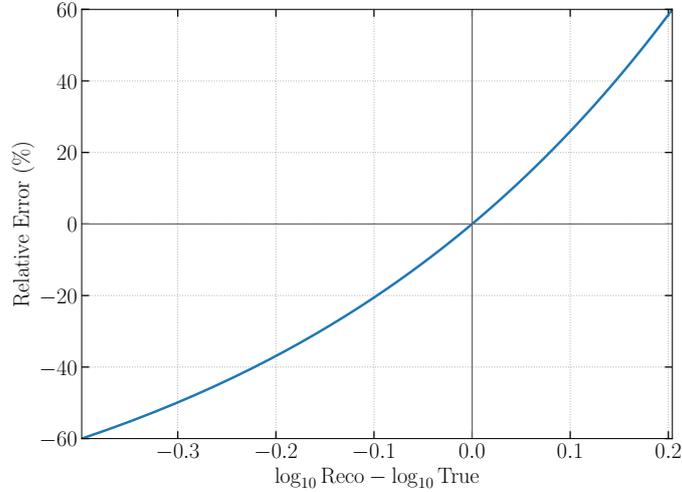
**Fig. 7.10:** Performance of the ENN primary energy reconstruction. Left: Two-dimensional histogram of the reconstructed and true values. The red line indicates the 1:1 correlation case. Right: Bias and resolution, defined as the mean and standard deviation of the difference between the reconstructed and true values in logarithmic scale, as a function of true energy.



**Fig. 7.11:** Performance of the NENN primary energy reconstruction. Left: Two-dimensional histogram of the reconstructed and true values. The red line indicates the 1:1 correlation case. Right: Bias and resolution, defined as the mean and standard deviation of the difference between the reconstructed and true values in logarithmic scale, as a function of true energy.

### Muon multiplicity reconstruction

For the reconstruction of the muon multiplicity, the NNN and NENN approaches are considered. We will similarly as before show correlation plots and bias and resolution plots in  $\log_{10} N_{\mu}$ . In addition, as for what follows  $N_{\mu}$  will be used more than  $\log_{10} N_{\mu}$ ,



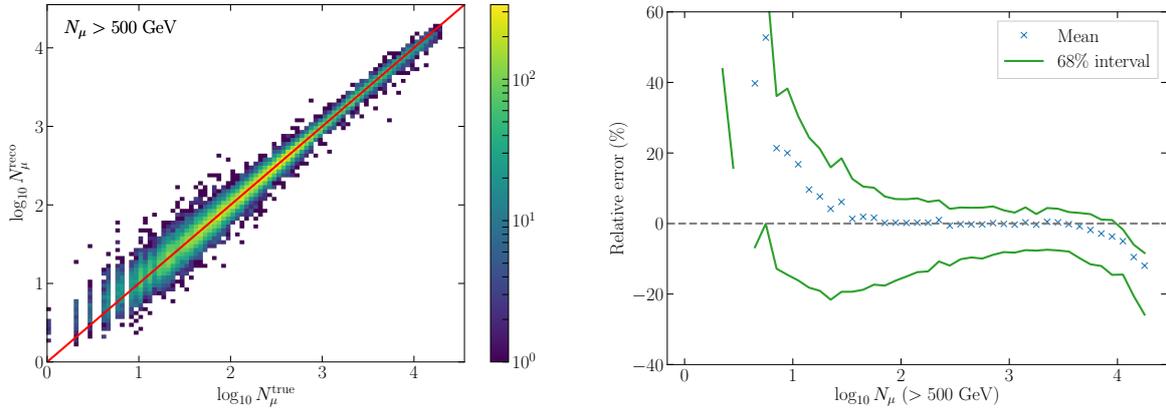
**Fig. 7.12:** Relation between the relative error in the reconstructed  $N_\mu$  compared to the true value, and the log difference between the two.

we include a plot of the relative error (RE) in the reconstruction of  $N_\mu$ , defined as

$$\text{RE} = \frac{N_\mu^{\text{reco}} - N_\mu^{\text{true}}}{N_\mu^{\text{true}}}. \quad (7.6)$$

In these plots, we will show the mean value of RE together with bands that include 68% of the events around it, derived from the entire test set. The RE is related to the logarithmic difference  $\Delta = \log_{10} N_\mu^{\text{reco}} - \log_{10} N_\mu^{\text{true}}$  as  $\Delta = \log_{10}(1 + \text{RE})$ , as illustrated in Fig. 7.12.

The correlation and relative-error plots for the NNN model, which predicts  $N_\mu$  based on the energy-loss profile of the bundle, are shown in Fig. 7.13. Starting from around 10 muons, the model manages to predict  $N_\mu$  quite accurately, with a spread that decreases toward only  $\pm 10\%$  with increasing number of muons. A common effect is visible at the edges of the dataset, where the lower end tends to get overestimated and the higher end tends to get underestimated. These events are, however, mainly outside the range of interest for this analysis. In Fig. 7.14, the bias and resolution of the reconstruction are shown per primary type. As was clear from the previous plots, the resolution improves with increasing number of muons, and flattens somewhat around 1000 muons. It is nearly independent of the primary mass. At low multiplicities, there is some mass-dependent bias, but this nearly disappears above 100 muons. This shows that in this case the muon bundle light deposit can be translated accurately into a number of muons, regardless



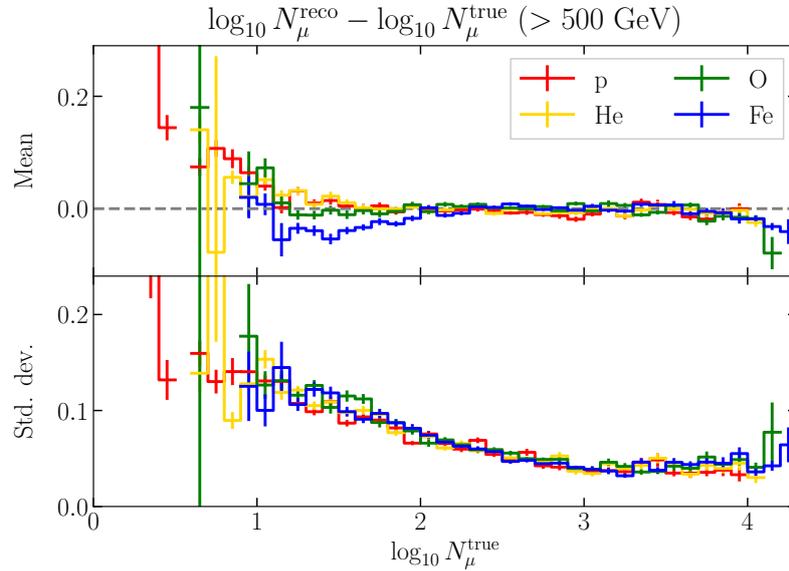
**Fig. 7.13:** Performance of the NNN multiplicity reconstruction. Left: Two-dimensional histogram of the reconstructed and true values. The red line indicates the 1:1 correlation case. Right: Mean and 68%-inclusion band of the relative error in the reconstructed value, defined by Eq. (7.6).

of the mass of the primary. This agrees with the naive expectation for muons of equal energies and with a constant energy loss.

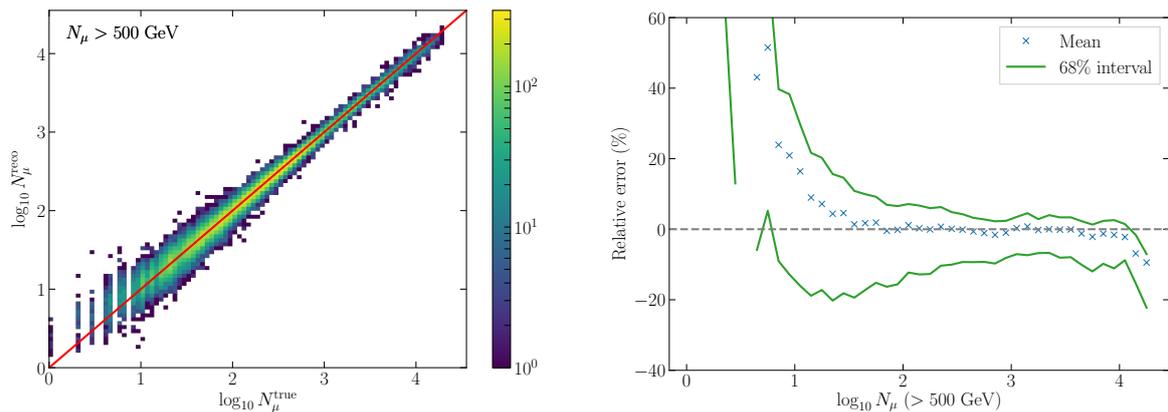
The accuracy of the  $N_\mu$  reconstruction that is obtained from the NENN model, is shown similarly in Figs. 7.15 and 7.16. Compared to the NNN model, there are fewer events that lie further away from the perfect reconstruction line. Similarly, the resolutions show a very slight improvement over the NNN model. However, there is a much larger mass-dependent bias visible in the reconstruction, and this for the entire range of multiplicities considered. Likely, this is related to the fact that the neural network also has information on the primary energy available. The primary energy is easier to predict well than the muon multiplicity. As a result, it may be beneficial for the network to learn the average relation between  $N_\mu$  and  $E_0$  to perform well for events that have a hard-to-reconstruct  $N_\mu$ . This average relation, however, causes a tendency for the number of muons in proton showers to get overestimated, while that in iron showers gets underestimated. A method for dealing with the impact of such biases in further steps in the analysis is discussed in the next section.

## 7.5 Calibration

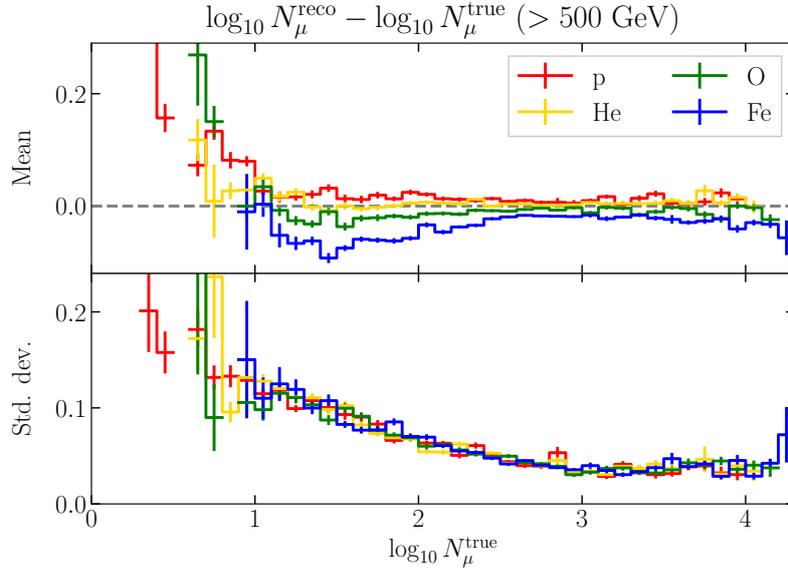
Given the results from the previous section, we will consider in what follows both the combination of the ENN and NNN models, as well as the combined NENN model. With these different methods to estimate  $E_0$  and  $N_\mu$  in hand, we can determine the average



**Fig. 7.14:** Bias and resolution of the NNN muon multiplicity reconstruction, defined as the mean and standard deviation of the difference between the reconstructed and true values in logarithmic scale, as a function of true multiplicity.



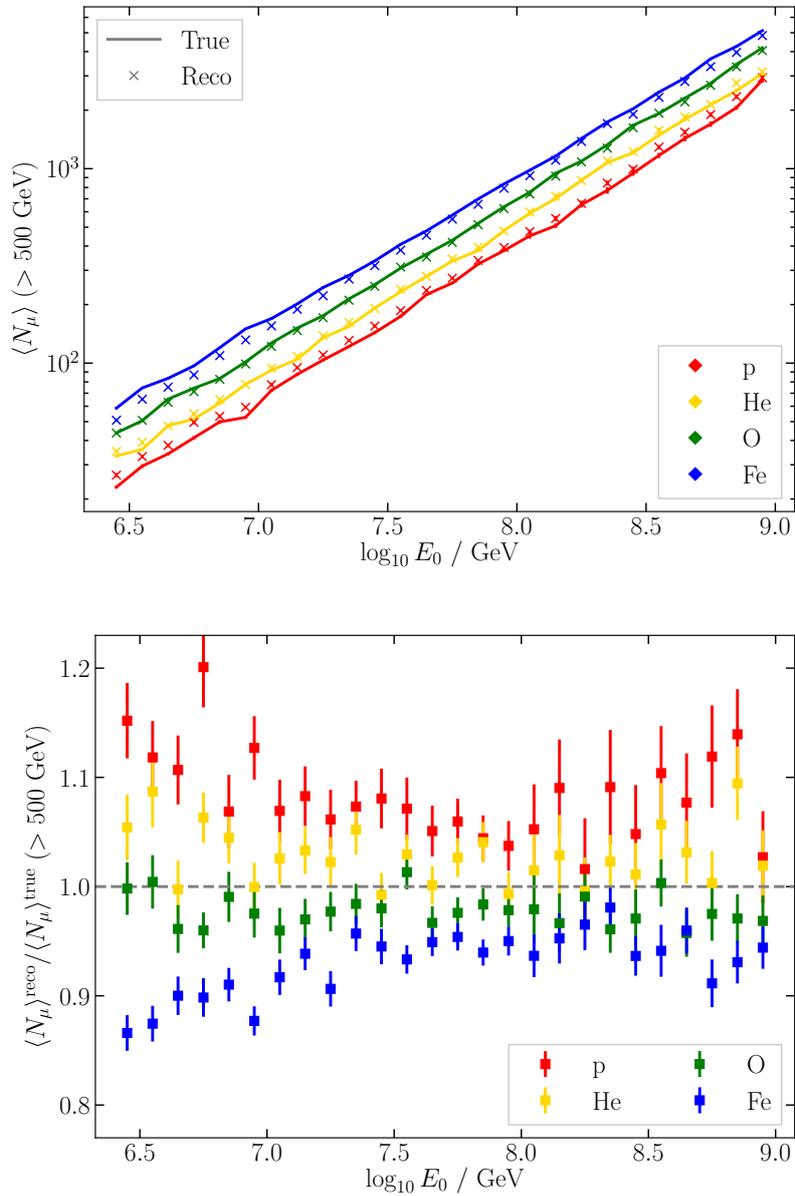
**Fig. 7.15:** Performance of the NENN multiplicity reconstruction. Left: Two-dimensional histogram of the reconstructed and true values. The red line indicates the 1:1 correlation case. Right: Mean and 68%-inclusion band of the relative error in the reconstructed value, defined by Eq. (7.6).



**Fig. 7.16:** Bias and resolution of the NENN muon multiplicity reconstruction, defined as the mean and standard deviation of the difference between the reconstructed and true values in logarithmic scale, as a function of true multiplicity.

multiplicity  $\langle N_\mu \rangle$  as a function energy. To do this, events will be binned in bins of reconstructed  $\log_{10} E_0$  with a width of 0.1, and in each bin the average reconstructed  $N_\mu$  will be calculated with its standard error. To assess how well this works, this can be done on simulation, and the results can be compared to the true values in the simulation, obtained by calculating the average of the true multiplicity in bins of true energy. This is shown for the case of the NENN model in the top panel of Fig. 7.17. The values obtained from the reconstruction capture the true behavior quite well, and the separation between the different elements is clear.

However, some systematic biases can be observed. This is illustrated more clearly in the bottom panel of the figure, which shows the ratio of the reconstructed and true values from the top plot. We see that for this model, the multiplicity of proton and iron showers respectively gets over- and underestimated by close to 10% on average. This bias is a result of imperfections in the neural-network reconstructions of  $N_\mu$ , as well as the  $E_0$  reconstruction as a result of bin migrations. This bias is an important effect that needs to be taken into account when applying the methods to experimental data. Below, we explain how this can be done.



**Fig. 7.17:** Comparison of obtaining  $\langle N_\mu \rangle$  as a function of energy when using the NENN predictions of  $N_\mu$  and  $E_0$  to the true values in simulation. The bottom plot shows the ratio of the values obtained from the reconstruction to the true ones.

### 7.5.1 Correction factors

Deriving  $\langle N_\mu \rangle$  in simulation gives us a handle on the bias that can be expected in the result due to imperfections in the neural networks. This knowledge can be used to derive correction factors which can later be applied to data. As was shown in Fig. 7.17, the ratio between the true and the reconstructed  $\langle N_\mu \rangle$  for each primary type varies smoothly with energy. To capture this behavior, we perform a fit to it, and the resulting curve will later on be used as a multiplicative correction factor. For the fits, a quadratic function is used of the form

$$\mathcal{C}(\log_{10} E_0) = p_0 + p_1 \log_{10} E_0 + p_2 \log_{10}^2 E_0. \quad (7.7)$$

The fit parameters define the correction factor  $\mathcal{C}$  necessary for a certain element. The statistical uncertainty on the correction factor is calculated as

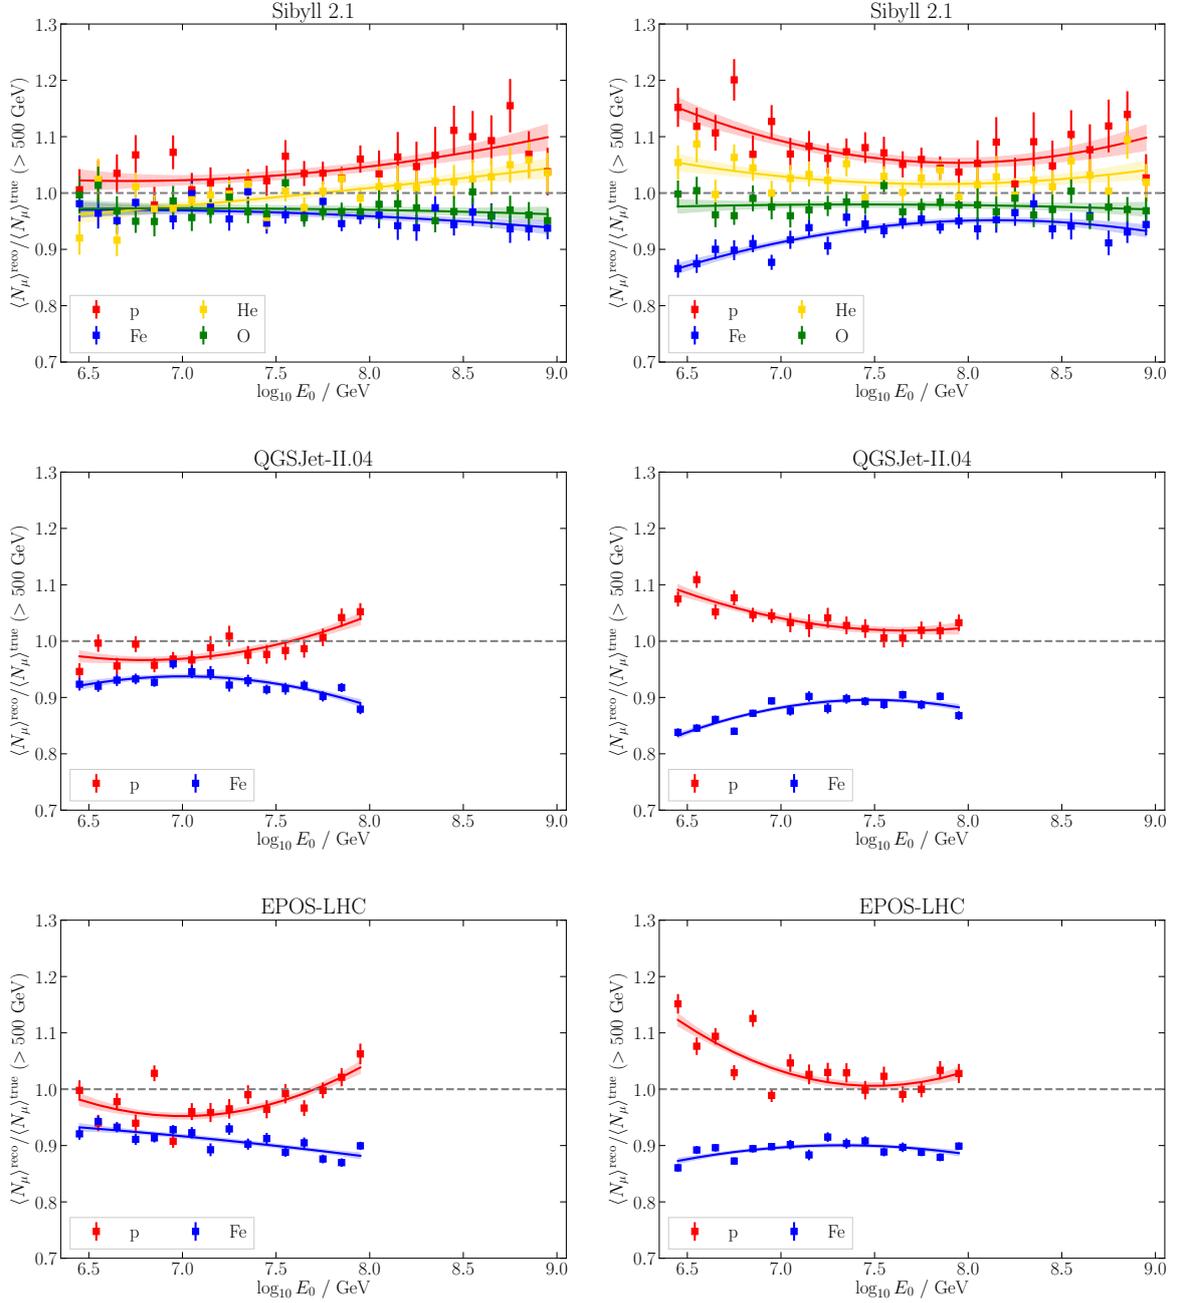
$$\delta\mathcal{C}^2 = \left( \frac{\delta\mathcal{C}}{\delta\vec{p}} \right)^T \mathbf{U} \left( \frac{\delta\mathcal{C}}{\delta\vec{p}} \right), \quad (7.8)$$

where  $\mathbf{U}$  is the covariance matrix of the fit parameters  $\vec{p} = (p_0, p_1, p_2)^T$ .

These correction factors are derived from simulation and will clearly make the final analysis result dependent on the hadronic interaction model used in the simulation. While so far Sibyll 2.1 was used for everything, the neural-network reconstructions can also be applied to the datasets simulated using QGSJet-II.04 and EPOS-LHC, listed in Table 4.3. This will allow us to later obtain results from experimental data under the assumption that different models are “correct”, even though the same neural network model is applied. The correction factors obtained from simulation using these different models and when using both the NNN+ENN and the NENN approach are shown in Fig. 7.18. Note that due to limitations in the dataset, correction factors for QGSJet-II.04 and EPOS-LHC are limited to an energy of 100 PeV and are only derived from proton and iron.

For all hadronic interaction models, the neural-network bias shows less mass dependence for the NNN+ENN method compared to the NENN method; the difference between the proton and iron correction factors is smaller over the entire energy range. For Sibyll 2.1, the correction factors are rather symmetric around 1, with proton having a higher value and iron a lower value. For QGSJet-II.04 and EPOS-LHC, the correction factors are similar in shape, but are shifted away from 1. This is a result of the fact that the neural network was trained on Sibyll 2.1 simulation.

One obvious problem that remains is that the correction that needs to be applied to data is clearly dependent on the mass composition, which is unknown. Two ways to



**Fig. 7.18:** Correction factors obtained for the NNN+ENN (left column) and NENN (right column) approach using three different hadronic interaction models. Quadratic functions are fit to the ratios between true and reconstructed values obtained from simulation; the bands give statistical uncertainties on the resulting correction factor.

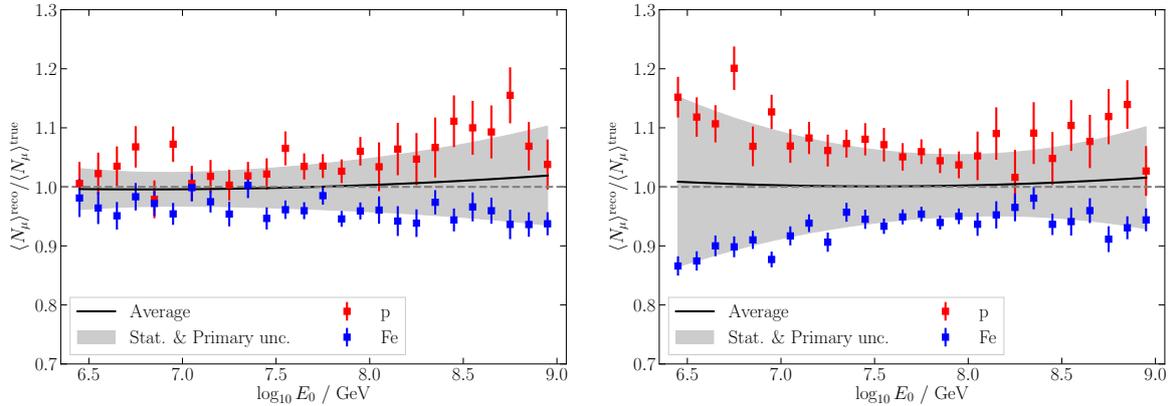
use the correction factors in practice will be discussed below. The first one is rather unrefined but can be used with all correction factors and can be used to get a rough estimate of the results. The second way will present a clear improvement over the first one, but demands a specific behavior for the correction factors, which is why we will end up using the NENN model over the others for the final  $\langle N_\mu \rangle$  measurement, even though the biases it produces are larger.

### 7.5.2 Simple correction procedure

Inspired by the method of the IceTop GeV muon density analysis, summarized in Section 4.3.3, one could choose to derive an average correction factor to apply to the result from data. This is done by taking the average of the proton and iron correction factors and assigning a systematic uncertainty that covers the spread of the correction factors to account for the unknown composition. This implicitly assumes that correction factors for intermediate elements lie in between proton and iron, which from Fig. 7.18 seems to hold approximately.

Examples of such average correction factors with their uncertainty are shown in Fig. 7.19. Here, the gray band covers the total uncertainty of the resulting correction factor. This includes a statistical uncertainty, which is the uncertainties on the proton and iron correction factors added in quadrature, and the composition uncertainty, which is simply the average deviation of the proton and iron factors from their average. The two are then added in quadrature. When applying the correction factor, its total uncertainty can be used as a systematic uncertainty on the final result.

We will show an example of how this would be applied using simulation of the four primary groups (p, He, O, Fe) weighted according to the H4a model. In Fig. 7.20, the true  $\langle N_\mu \rangle$  obtained from the simulation is shown together with what we obtain from applying the neural network models. For clarity, they are plotted as  $z$ -values as in Eq. (7.1), where the expectations for proton and iron are obtained from the fits of Section 7.3. While the reconstructions get close to the true values, systematic deviations are again visible. In Fig. 7.21, the average correction factor gets applied to these results, i.e. the result are divided by the correction factor. The total uncertainty of the correction factor gets assigned to the end result as a systematic uncertainty, indicated in the plots by the brackets. While the agreement between the center values and the truth does not necessarily improve by applying the average correction, the systematic brackets now cover an envelope in which the actual values are expected. The true values can be seen to be covered well by the brackets, except for the lowest energy points when using



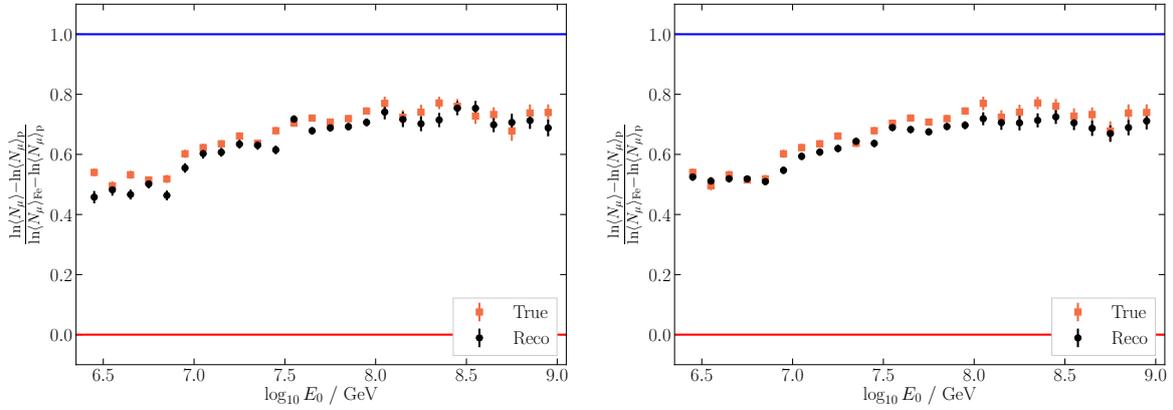
**Fig. 7.19:** Definition of an average correction factor based on Sibyll 2.1. The black line gives the average between the proton and iron correction factors obtained in Fig. 7.18, the gray band describes the statistical uncertainty and the spread between proton and iron.

the NNN+ENN method, which may be related to the behavior of the corresponding correction factor at these energies, which does not show a consistent mass-dependence.

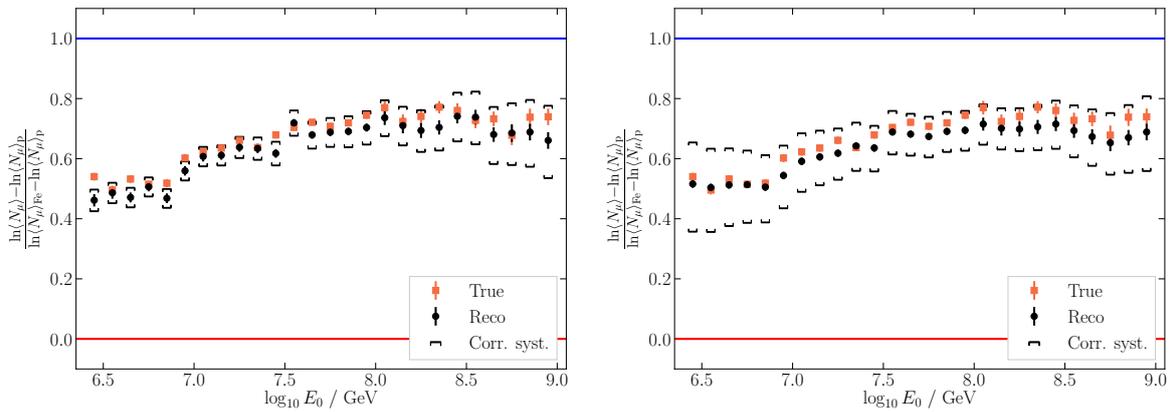
This method provides a quick and easy way to obtain an envelope in which the actual  $\langle N_{\mu} \rangle$  values should be contained. However, it is not very refined and comes with large uncertainties. This is a result of the relatively large differences observed between the correction factors that are derived for a pure proton and pure iron case. This differs from the GeV muon density analysis, where the proton and iron correction factors agree rather well with each other. Alternatively, one could also assume a composition model and derive a correction factor from this. However, given that there is a lot of uncertainty in our knowledge of the composition, this also isn't a great option. Clearly, it is desirable to implement a better method which allows us to reconstruct the muon multiplicity with greater precision, no matter the composition. Such a method is discussed below.

### 7.5.3 Iterative correction procedure

A better method to correct the result obtained for  $\langle N_{\mu} \rangle$  biases resulting from the neural-network reconstructions has to deal with the difficulty that the exact correction that needs to be applied depends on the mass composition, as was shown in Fig. 7.18. To this end, we can recall that the multiplicity of muons in a shower itself is related to the composition, as can be understood through the Heitler-Matthews model of Section 2.2.1. This implies that we can use the first estimate for  $\langle N_{\mu} \rangle$  to estimate the mass composition,



**Fig. 7.20:** Average muon multiplicity in simulation weighted according to the H4a flux model expressed in  $z$ -scale. Shown are the true values and those obtained when using the neural-network reconstructions for  $E_0$  and  $N_\mu$ . Left: NNN+ENN model. Right: NENN model.



**Fig. 7.21:** Average muon multiplicity in simulation weighted according to the H4a flux model. A proton-iron averaged correction factor (Fig. 7.19) has been applied to the reconstructed values, and the brackets indicate the corresponding uncertainty. Left: NNN+ENN model. Right: NENN model.

from which a correction factor can be defined to obtain a better estimate of  $\langle N_\mu \rangle$ , and then iterate this process.

First, we need a way in which to derive an effective correction factor  $\mathcal{C}_{\text{eff}}$  given a specific mass composition. A simple way could be by combining the correction factors of the different mass groups given the respective fractions of the total flux  $f_i$  they constitute as

$$\mathcal{C}_{\text{eff}} = f_{\text{p}}\mathcal{C}_{\text{p}} + f_{\text{He}}\mathcal{C}_{\text{He}} + f_{\text{O}}\mathcal{C}_{\text{O}} + f_{\text{Fe}}\mathcal{C}_{\text{Fe}}, \quad (7.9)$$

where  $\mathcal{C}_i$  are the correction factors derived for the different groups  $i = \text{p, He, O, Fe}$ . The issue here is that the muon measurement gives an estimate of the average composition  $\langle \ln A \rangle$ , but this does not define a unique way to combine the four components. It does, however, define a way to combine two components to obtain the average composition, as we have two equations to solve for two fractions. Choosing proton and iron as the two components, these equations read

$$f_{\text{p}} \ln A_{\text{p}} + f_{\text{Fe}} \ln A_{\text{Fe}} = \langle \ln A \rangle, \quad (7.10)$$

$$f_{\text{p}} + f_{\text{Fe}} = 1. \quad (7.11)$$

Given that  $A_{\text{p}} = 1$ , the fractions can be written as

$$f_{\text{Fe}} = \frac{\langle \ln A \rangle}{\ln 56}, \quad f_{\text{p}} = 1 - f_{\text{Fe}}. \quad (7.12)$$

The average composition is simply estimated from the muon measurement expressed in  $z$ -scale, inspired by the Heitler-Matthews model,

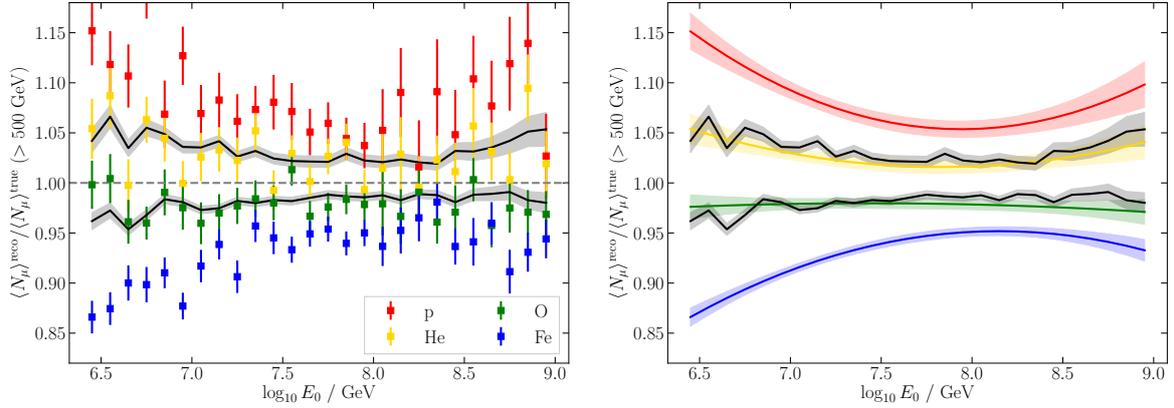
$$z = \frac{\ln \langle N_\mu \rangle - \ln \langle N_\mu \rangle_{\text{p}}}{\ln \langle N_\mu \rangle_{\text{Fe}} - \ln \langle N_\mu \rangle_{\text{p}}} \approx \frac{\langle \ln A \rangle}{\ln 56}. \quad (7.13)$$

With this, we can finally calculate the effective correction factor based on a measurement of  $\langle N_\mu \rangle$  as

$$\mathcal{C}_{\text{eff}} = f_{\text{p}}\mathcal{C}_{\text{p}} + f_{\text{Fe}}\mathcal{C}_{\text{Fe}}, \quad (7.14)$$

where the fractions  $f_{\text{p}}$  and  $f_{\text{Fe}}$  are given by Eq. (7.12) and describe the average composition implied by  $\langle N_\mu \rangle$  through Eq. (7.13). Note that this approach is equivalent to linearly interpolating the correction factors in  $\ln A$  as

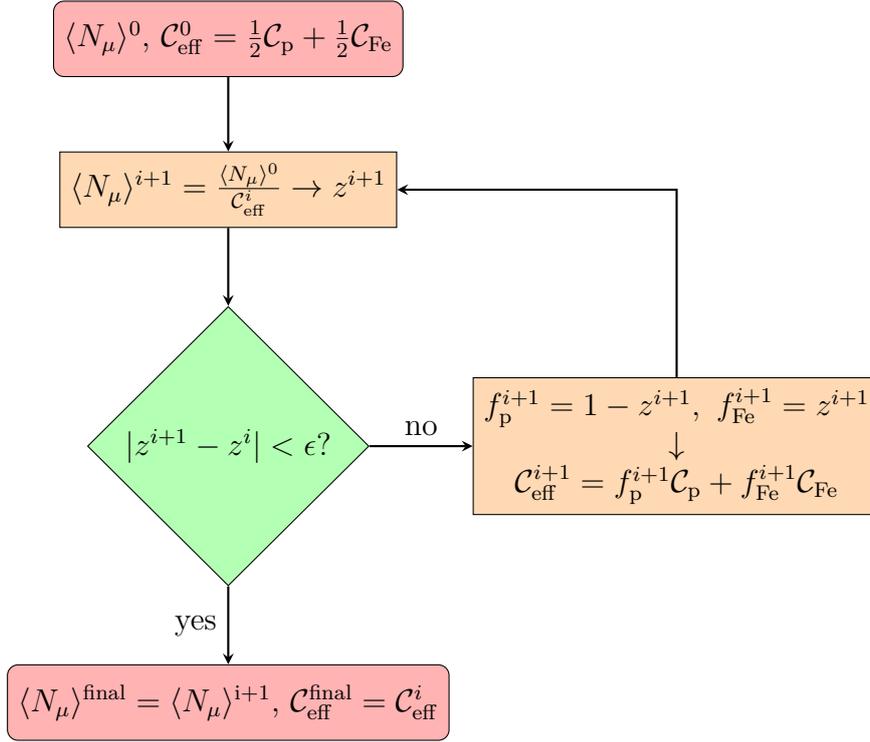
$$\mathcal{C}_{\text{eff}}(\ln A) = \mathcal{C}_{\text{p}} + \frac{\mathcal{C}_{\text{Fe}} - \mathcal{C}_{\text{p}}}{\ln 56} \ln A. \quad (7.15)$$



**Fig. 7.22:** Test of the assumption that correction factors can be obtained by combining the proton and iron correction factors based on the muon multiplicity observed in a dataset. The gray bands show the correction factor obtained from the true muon multiplicity in helium and oxygen simulation. They show good agreement with the correction factors fitted to the simulation. Left plot compares to the actual points obtained from simulation, right plot to the fits.

This approach assumes that the correction factors show a linear dependence on  $\ln A$ , so that the correction factors for intermediate elements can simply be obtained through interpolation. From the Sibyll 2.1 plots in Fig. 7.18, we can see by eye that this will be approximately true for the correction factors derived when working with the NENN model, but not for those from the NNN+ENN model, where the oxygen correction factor nearly overlaps with the one for iron. This is why we continue to work only with NENN for the final  $\langle N_{\mu} \rangle$  measurement that will be presented in the next sections. A more explicit check of the validity of this assumption for the NENN model is shown in Fig. 7.22. Here, we calculate an effective correction factor using Eq. (7.14) based on the muon multiplicity in helium and oxygen simulation, and compare it to the helium and oxygen correction factors we derived earlier in Fig. 7.18. It can be seen that there is good agreement between the two, so that we can confidently use the method for any composition.

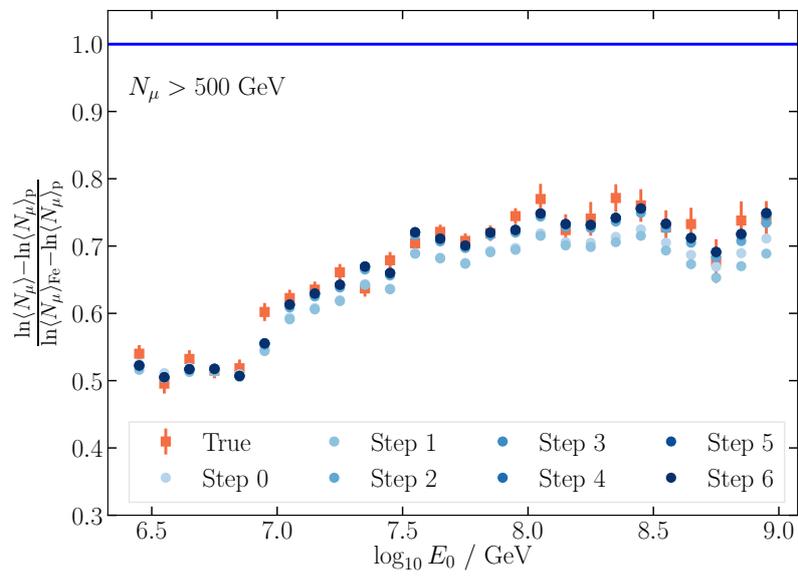
Having defined how we can obtain a correction factor based on an estimate of the mass, an iterative method can be defined for the correction. It starts by simply applying the neural-network reconstruction to obtain  $N_{\mu}$  and  $E_0$  for each event, from which we calculate a first estimate of  $\langle N_{\mu} \rangle$  in bins of  $\log_{10} E_0$ . Using this, an effective correction factor can be calculated as described above. Dividing the obtained values for the multiplicity by this correction factor gives us an updated estimate of  $\langle N_{\mu} \rangle$ , which in turn allows us to calculate a new correction factor. The updated correction factor gets applied again to



**Fig. 7.23:** Flowchart describing the iterative correction procedure. The process begins with the ‘raw’ estimate of  $\langle N_\mu \rangle$  in bins of  $\log_{10} E_0$  from the NENN model and a first guess correction factor. The iteration starts by calculating a corrected value of  $\langle N_\mu \rangle$ . From this, a new composition estimate and a new correction factor can be obtained for a next step. If the difference between the  $z$ -value obtained from the newly corrected  $\langle N_\mu \rangle$  and that from the previous step is smaller than  $\epsilon$ , the process stops.

the original values obtained for  $\langle N_\mu \rangle$ . This process is repeated for several iterations, until the difference in  $z$  calculated from  $\langle N_\mu \rangle$  in subsequent steps becomes smaller than some threshold value, chosen to be  $10^{-5}$ . This process is schematically drawn in Fig. 7.23. An example, again based on the H4a model, is shown in Fig. 7.24. It can be seen how the offset which is present between the first obtained  $\langle N_\mu \rangle$  values and the truth disappears almost entirely after several iteration steps, ending up with a result that is compatible with the truth. The improvement of this method over the one discussed above is clear. More similar tests of the method will be discussed in the next section.

The final result is obtained by dividing the initial  $\langle N_\mu \rangle$  estimate before any correction by the correction factor obtained in the final step. As such, the statistical uncertainty on the final result is divided by the same factor, so that the relative statistical uncertainty remains unchanged. The correction factor of course comes with a statistical uncertainty,



**Fig. 7.24:** Example of the iterative correction procedure applied to simulation weighted according to the H4a model. The true muon multiplicity values are shown in orange, those obtained from the neural-network reconstructions in the lightest shade of blue. After several steps of applying a correction factor, the multiplicities converge to values compatible with the true values.

which is calculated as

$$\sigma_{C_{\text{eff}}^{\text{final}}} = \sqrt{\left(f_{\text{p}}^{\text{final}}\sigma_{C_{\text{p}}}\right)^2 + \left(f_{\text{Fe}}^{\text{final}}\sigma_{C_{\text{Fe}}}\right)^2} \quad (7.16)$$

This correction factor uncertainty is taken into account in the final  $\langle N_{\mu} \rangle$  result as a systematic uncertainty inherent to the method.

## 7.6 Simulation tests

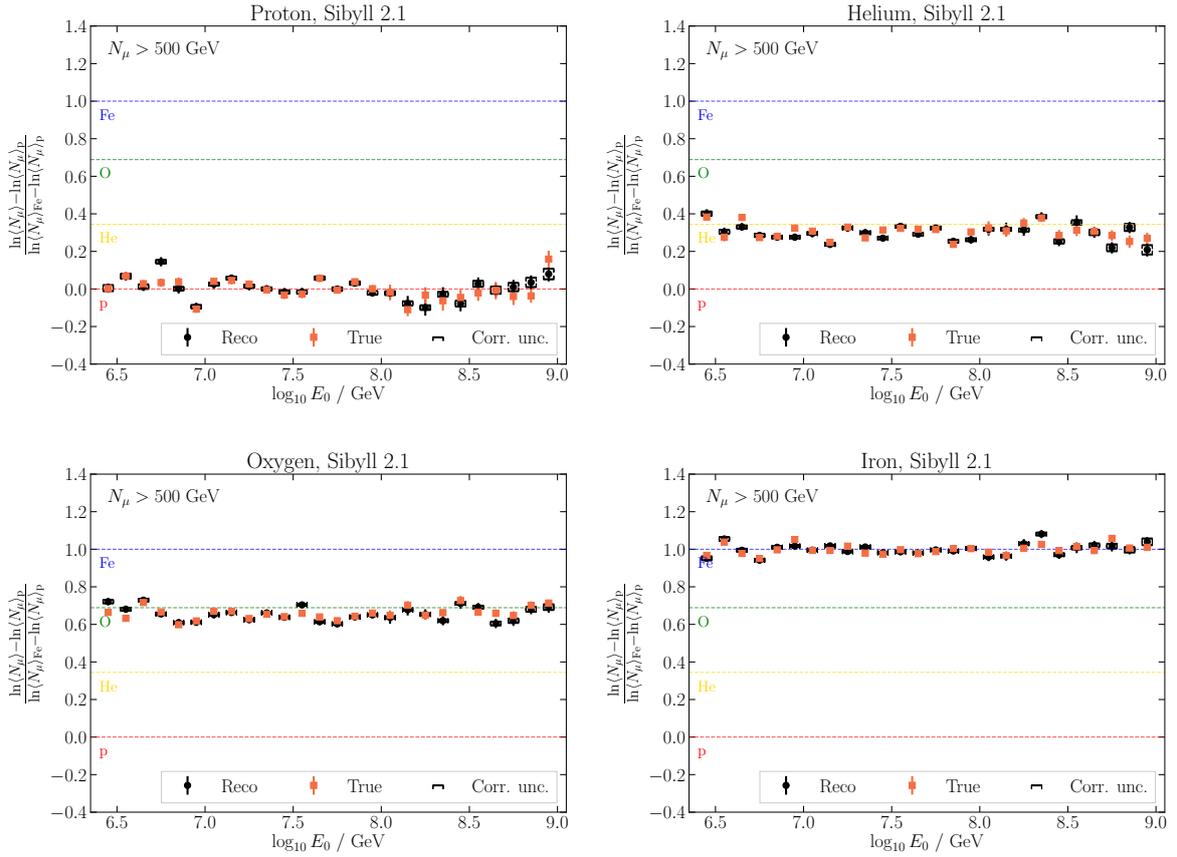
In the previous sections, neural network models were trained to determine the primary energy  $E_0$  and multiplicity of muons above 500 GeV  $N_{\mu}$  for an event, and a method to correct for biases when using these estimates to determine the average multiplicity  $\langle N_{\mu} \rangle$  in bins of  $\log_{10} E_0$  was presented. A first demonstration of the performance of these two analysis steps was shown in Fig. 7.24. To gain confidence in the analysis, it is useful to perform more tests on simulation to test if it works generally. As the iterative correction procedure only works with the NENN model, we will only consider this one going forward. Furthermore, before applying these methods to data, it needs to be assured that there is good agreement between data and the simulation used to develop the analysis.

### 7.6.1 Analysis performance tests

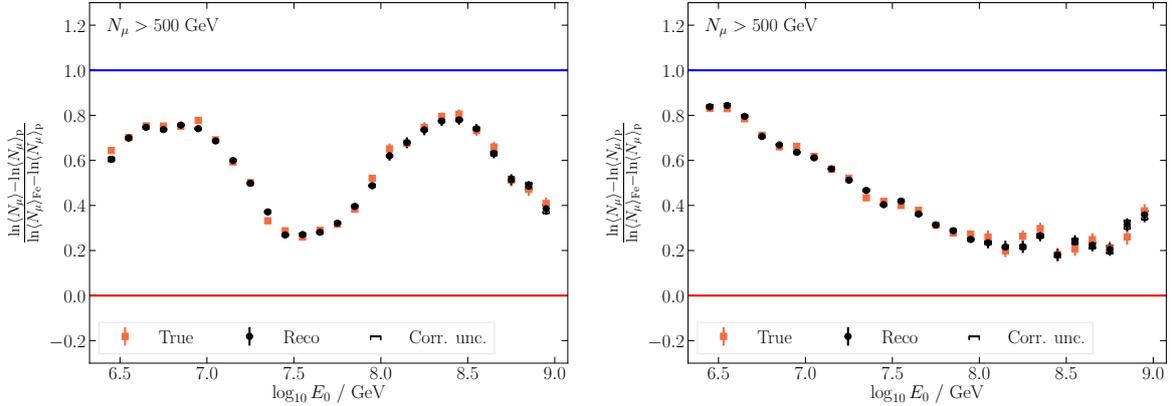
We test the performance of the neural network and subsequent correction in deriving  $\langle N_{\mu} \rangle$  on a number of different composition assumptions, or equivalently, true multiplicity assumptions.

Fig. 7.25 shows the case of a pure composition. This is obtained by taking simulation of one specific primary, and weighting it according to the all-particle H4a spectrum. This therefore corresponds to the assumption that the entire cosmic-ray flux is caused by only a single type of primary. The plots show that the method obtains good agreement with the true values for all considered primaries: p, He, O, and Fe. Note that from Fig. 7.17, we know that pure proton and iron get over/underestimated by about 10% when just using the neural network outputs. The iterative correction clearly gets rid of this bias.

Additionally, tests are performed on different artificial mixed composition cases. These are obtained by weighting simulation of four primary types to a realistic flux, and then modifying the event weights in a smooth energy dependent way that is different for each type, and combining them. Two arbitrary cases are shown in Fig. 7.26. Once again, good agreement is obtained.



**Fig. 7.25:** Average muon multiplicity (expressed in  $z$ -values) obtained using the neural-network reconstructions followed by the iterative correction for the case of pure proton, helium, oxygen, and iron simulation. The results are compatible with the true values. The horizontal lines for He and O are drawn at  $\ln 4 / \ln 56$  and  $\ln 16 / \ln 56$  for reference.



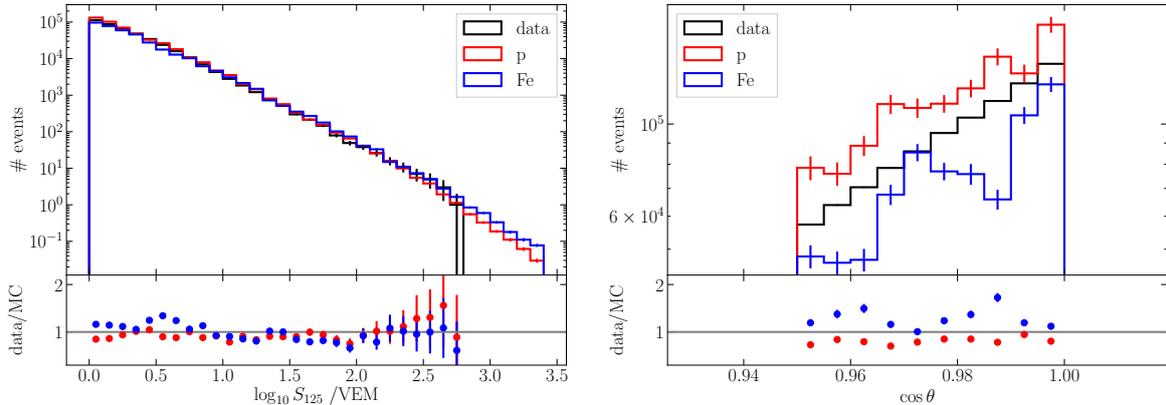
**Fig. 7.26:** Average muon multiplicity (expressed in  $z$ -values) obtained using the neural-network reconstructions followed by the iterative correction for two artificial composition cases. The results are compatible with the true values.

More tests on artificial compositions, together with plots demonstrating more clearly that the reconstructed  $\langle N_\mu \rangle$  are compatible with the true values, are shown in Appendix B.2. The tests on the pure composition cases, a realistic composition case like H4a, and several artificial compositions confirm that the analysis works well generally, independent of the actual composition.

### 7.6.2 Data-MC comparison

In the previous section, we have established that the analysis method works well on simulation. Before applying it to experimental data, it is important to make sure that the simulations provide a good description of this data. This is especially important when using machine learning methods that are trained on simulation. We will here specifically compare the high-level variables that are used as input to the neural networks in data and MC, as well as the distributions of neural network outputs.

To do this, some assumptions need to be made for weighting the simulation in a meaningful way. As the composition in data is not known precisely, it isn't ideal to weight the simulation to a specific model and compare data to it, as any disagreements may stem from errors in the composition model. We do expect, however, that any physically reasonable composition will on average not be lighter than pure proton and not heavier than pure iron. Therefore, we will use proton and iron simulation and weight it to describe the all-particle energy spectrum, which has much smaller uncertainties than the mass composition, and compare data to these two extreme cases. If simulation and data agree well, data should be in between of the proton and iron distributions. If



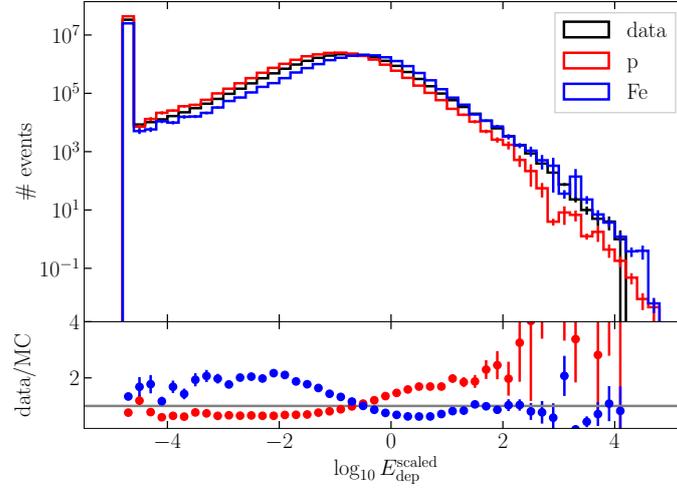
**Fig. 7.27:** Comparison of data and MC for the neural network IceTop input variables: shower size  $S_{125}$  (left) and reconstructed zenith  $\theta$  (right).

any considerable deviations from this show up, this indicates a possible issue with the simulation. In the plots, the event rates obtained for simulation after weighting it to the all-particle H4a spectrum will be increased by another 25%, a factor which was found in extensive low-level data-MC comparisons [202]. The data used is 10% of IC86.2012 data, corresponding to a total livetime of 32.7 days.

A comparison for the IceTop input variables,  $S_{125}$  and the reconstructed  $\theta$ , is shown in Fig. 7.27. The distribution of  $\theta$  for data lies nicely in between p and Fe. The  $S_{125}$  spectrum shows acceptable agreement, with some small deviation between 1 and 2 in  $\log_{10} S_{125}$ . This is likely attributable to disagreements between the H4a and the true energy spectrum, and no cause for concern.

The energy-loss vector used as a neural-network input from IceCube is not a single variable that can be compared. In Fig. 7.28, the distribution of all values in the vectors is shown (the distributions have  $\# \text{ events} \times \text{vector length}$  entries). A good agreement between data and simulation is found. As the energy-loss vector and the evolution of values throughout the detector are analyzed in detail by the neural network, it is desirable to perform some more detailed tests. In Appendix B.3, we include plots comparing various properties derived in different parts of the vector. No concerning disagreements are observed.

Finally, as a simple test of the generalization of the neural network from data to simulation, we compare the distributions of the neural network output variables obtained in simulation to this in data. This is shown in Fig. 7.29. For the  $E_0$  predictions, the data/MC ratio looks very similar to the  $S_{125}$  input distribution, which is to be expected



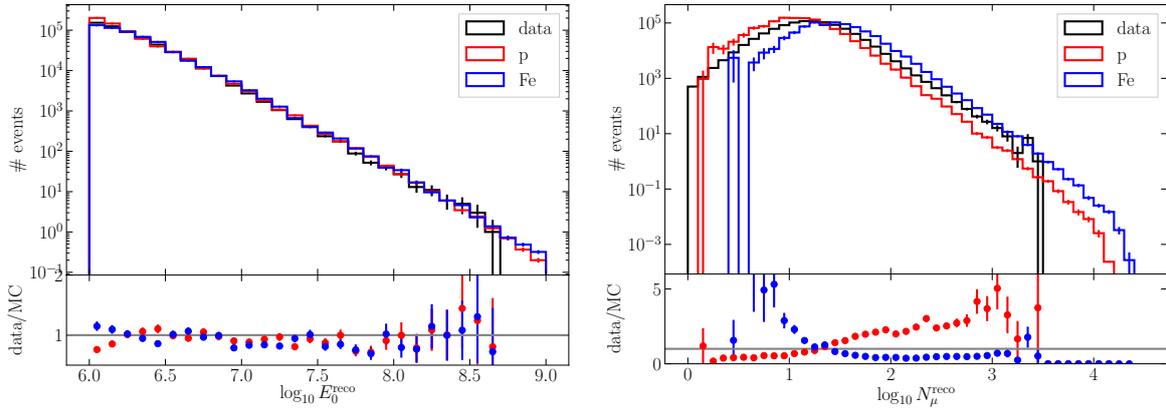
**Fig. 7.28:** Comparison of data and MC for the neural network IceCube input energy-loss vector. The distribution includes all 57 vector entries for each event.

as they are both good estimators for the actual energy. For  $N_\mu$ , the distribution in data is nicely contained between proton and iron.

The distributions shown in this section and in Appendix B.3 give confidence that for the purposes of this analysis, the simulation gives a good description of experimental data and the neural network can be used reliably.

## 7.7 Results

In the previous section, the method of using the NENN neural network model to obtain  $\langle N_\mu \rangle$  followed by the application of a correction was shown to work well on simulation regardless of the mass composition. The good agreement between simulation and experimental data for the different inputs used in the analysis, as well as for the neural network outputs, supports that the analysis can be applied to experimental data to reliably determine the average muon multiplicity as a function of energy, albeit in a model-dependent way as a result of the correction factors obtained from simulation. In this section, the analysis is applied to data, different systematic uncertainties are discussed, and the final results are presented.



**Fig. 7.29:** Comparison of data and MC for the neural-network predictions: primary energy  $E_0$  and muon multiplicity  $N_\mu$ .

### 7.7.1 Application to data

The analysis methods are applied to 10% of the IC86.2012 data, obtained by taking all 8-hour runs with a run number ending in 0, corresponding to a livetime of about 32.7 days<sup>45</sup>. After applying the L3 processing and all selections described in Section 7.2, the NENN neural network from Section 7.4.3 is applied to all events to obtain an estimate for  $E_0$  and  $N_\mu$ . Next,  $\langle N_\mu \rangle$  is determined in bins of  $E_0$ , after which the iterative correction procedure from Section 7.5 is applied. This is done with the correction factors obtained from the hadronic interaction models Sibyll 2.1, QGSJet-II.04, and EPOS-LHC, so that finally the average muon multiplicity result is obtained under the assumption of these three different models. The obtained results, and the impact of the correction factor, are shown in Fig. 7.30. The results derived using a certain hadronic model is always compared to the predictions from proton and iron simulation using the same model, as derived in Section 7.3. We show both the actual  $\langle N_\mu \rangle$  result, and the  $z$ -scale representation (Eq. (7.1)). For QGSJet-II.04 and EPOS-LHC, we limit the results to 100 PeV, as there is no high-energy simulation available to determine the correction factor.

It can be seen that for all models, the obtained result is bracketed by the expectations of proton and iron. The impact of the correction factor on the Sibyll 2.1 result is small. This can be understood through the fact that the obtained muon multiplicity is close to being in the middle of proton and iron (in  $z$  representation). From the correction factors

<sup>4</sup>Due to time constraints, the analysis as presented here was approved by the IceCube collaboration using only 10% of a year of data, and presented in this way at the European Cosmic Ray Symposium 2022. The 10% allows us to already cover energies up to about 250 PeV with good statistics. Updates to the analysis using more data are expected to follow soon.

<sup>5</sup>Note that this analysis, like the analysis of Chapter 6, was not performed blinded.

shown in Fig. 7.18, which are roughly symmetric around 1, this shows that the correction that will ultimately be applied in this case will also be close to 1. For QGSJet-II.04 and EPOS-LHC, the correction factors are not centered around 1, a result of the neural network being trained on Sibyll 2.1. Therefore, a larger correction factor will be applied to the neural network output when assuming these models.

## 7.7.2 Systematic uncertainties

With the nominal  $\langle N_\mu \rangle$  results and its statistical uncertainties obtained in the previous section, it is time to consider any systematic uncertainties. Two sources are considered: an uncertainty related to the correction factor, and detector systematics.

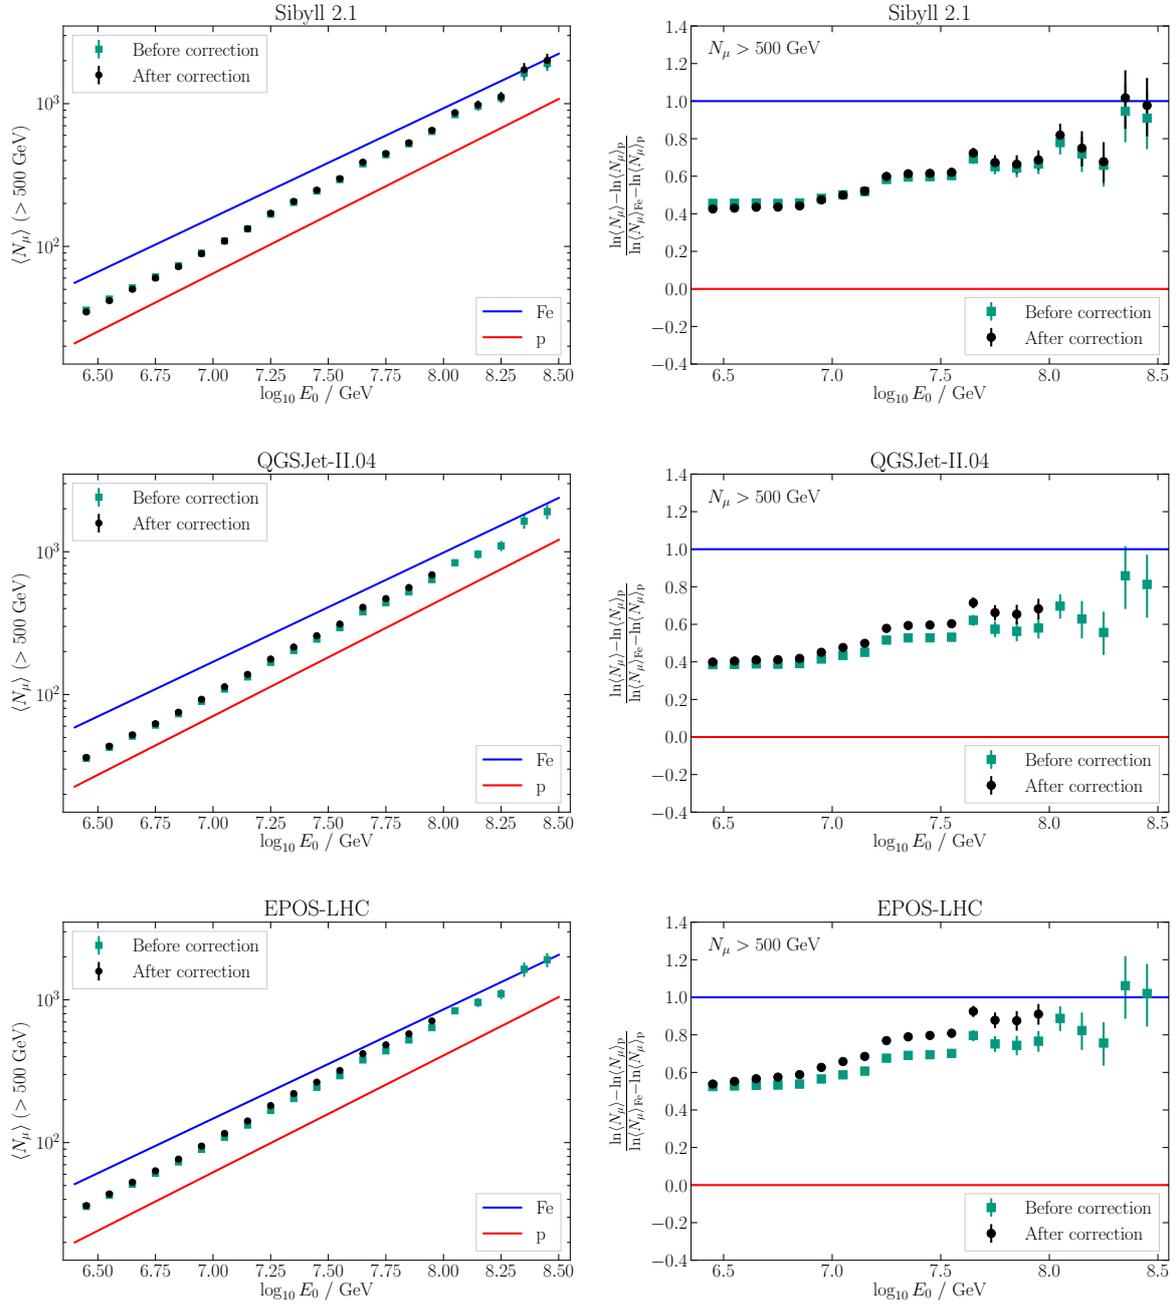
### Correction uncertainty

The statistical uncertainty on the correction factor is taken into account as an uncertainty on the final result. In Section 7.5.3, this was already explained: the statistical uncertainty on the final correction factor obtained in the last iteration step is calculated from the uncertainties on the proton and iron correction factors through Eq. (7.16). A systematic uncertainty of the same relative magnitude is then assigned to the final muon multiplicity result. This is a small uncertainty of only  $\sim 1\%$ .

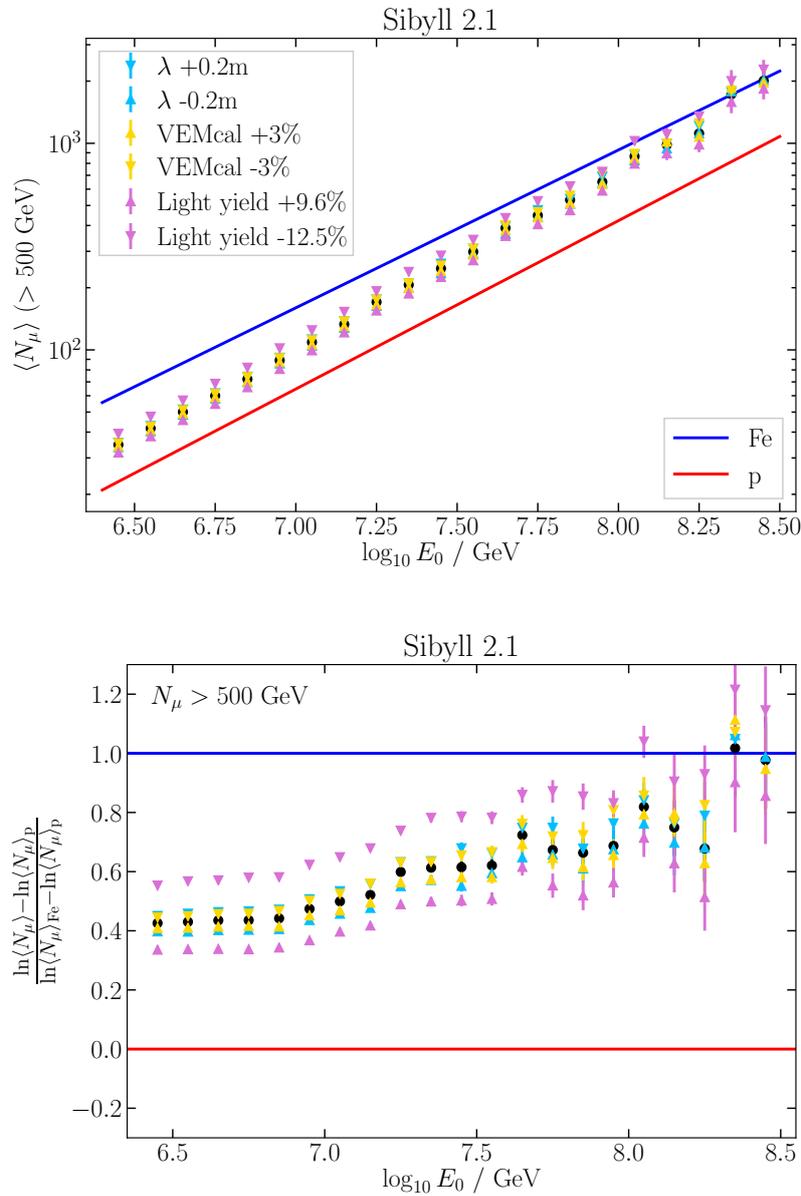
### Detector uncertainties

There are several detector uncertainties that could impact the result. Due to the strong similarity in reconstructions used, we will consider the same ones as in Chapter 6, which are also those considered in the composition analysis summarized in Section 4.3.1. They are uncertainties in the snow attenuation factor in the air shower reconstruction, the definition of the VEM unit, and the in-ice light yield. We discuss their origin and their effect on the final result, shown in Fig. 7.31, below. All of them will be added in quadrature to obtain a total systematic uncertainty for the final result.

The snow accumulation on top of IceTop impacts the signal expected in the tanks as it absorbs (mainly EM) shower particles. This is taken into account during the air shower reconstruction described in Section 4.2.1 by modifying the charge expectation in the likelihood according to Eq. (4.3), which utilizes an effective snow attenuation length  $\lambda_{\text{eff}}$ . From studies on data comparing the  $S_{125}$  spectrum in parts of the detector with less and more snow, the value of  $\lambda$  for IC86.2012 was found to be 2.1 m. An uncertainty of  $\pm 0.2$  m was assigned to it due to slight variations with energy and zenith angle [161]. This uncertainty may impact the muon multiplicity analysis, mainly by changing the



**Fig. 7.30:** Results for the average muon multiplicity ( $> 500 \text{ GeV}$ ) obtained in experimental data, before and after application of the iterative correction based on correction factors obtained from three different hadronic interaction models. Results are compared to the predictions for proton and iron of the corresponding hadronic model.



**Fig. 7.31:** Impact of the systematic detector uncertainties on the muon multiplicity result obtained from data. The difference to the nominal values is used as systematic uncertainty for the final result after adding all contributions in quadrature.

$S_{125}$  value and related to this the energy estimate, but also by changing the shower-axis reconstruction and as a result the in-ice energy loss reconstruction. To gauge the impact on the analysis, the L3 processing on data was redone with modified values of  $\lambda_{\text{eff}}$ , and the full analysis was repeated. The differences with the nominal result are used as the systematic uncertainty due to snow. The uncertainty is of the order  $\pm 3\%$  for  $\lambda_{\text{eff}} \pm 2$  m.

The charge unit used to express the IceTop signals is the VEM, the typical signal produced in a tank by a muon crossing it vertically. As explained in Section 3.4.1, the definition of the VEM in data is determined in regular intervals through the VEMCal procedure. The absolute calibration of IceTop is accomplished through a simulation of the VEMCal process using low-energy showers. Simulated charge signals are calibrated in VEM units such that the simulated muon peak is at the same location as in experimental data. This defines the energy scale of IceTop. An uncertainty of  $\pm 3\%$  was found in the simulation study presented in Ref. [194] as a result of atmospheric changes, hadronic interaction model, snow accumulation, energy spectrum, and others. This may influence the analysis in a similar way as the snow attenuation length uncertainty. To determine the impact on the analysis results, all IceTop charges in data were shifted by  $\pm 3\%$ , and the entire analysis was repeated. The differences with the nominal result are used as systematic uncertainty due to the VEM definition or absolute energy scale. The effect is of the order  $\mp 2\%$  for the  $\pm 3\%$  shifts.

Finally, there are the uncertainties affecting mainly the IceCube part of the analysis. There is an uncertainty related to the ice model, describing the scattering and absorption of photons in the ice. Furthermore, there is also the effective scattering length in the hole ice (see Section 3.2.2). For the bulk ice, three points on the  $1\sigma$  error ellipse around the nominal values [140] are considered: +10% scattering coefficient, +10% absorption coefficient, and -7.1% absorption and scattering coefficients. For the hole ice model, scattering lengths of 30 cm and 100 cm were used. A dedicated study on CORSIKA simulations in the zenith range  $0^\circ$ – $30^\circ$  determined the impact on the observed charge distributions. The magnitude of the shifts is given in Table 7.2. Next to the ice model, there is also an uncertainty on the photon detection efficiency of the DOMs, related to a shadow effect of the DOM cable, the PMT quantum efficiency, the optical gel, etc. Typically, an uncertainty of  $\pm 10\%$  is included on the normalization of the absolute DOM efficiency [137], although recent work suggests values around  $\pm 3\%$  [203]. This is also what was cited in the composition analysis, Ref. [19]. All of these uncertainties can be combined in quadrature into a total uncertainty on the in-ice light yield of (-12.5%, +9.6%). The impact on the analysis is determined by modifying the DOM efficiency parameter in the Millipede energy loss reconstruction by these percentages. The difference

**Table 7.2:** Systematic uncertainties for the in-ice light yield. Taken from Ref. [19].

Effect	Light yield shift
+10% scattering	+3.6%
+10% absorption	-11.8%
-7.1% scattering and absorption	+7%
30 cm hole ice scattering	+4.5%
100 cm hole ice scattering	-2.9%
DOM efficiency	$\pm 3\%$
Total light yield effect	+9.6%, -12.5%

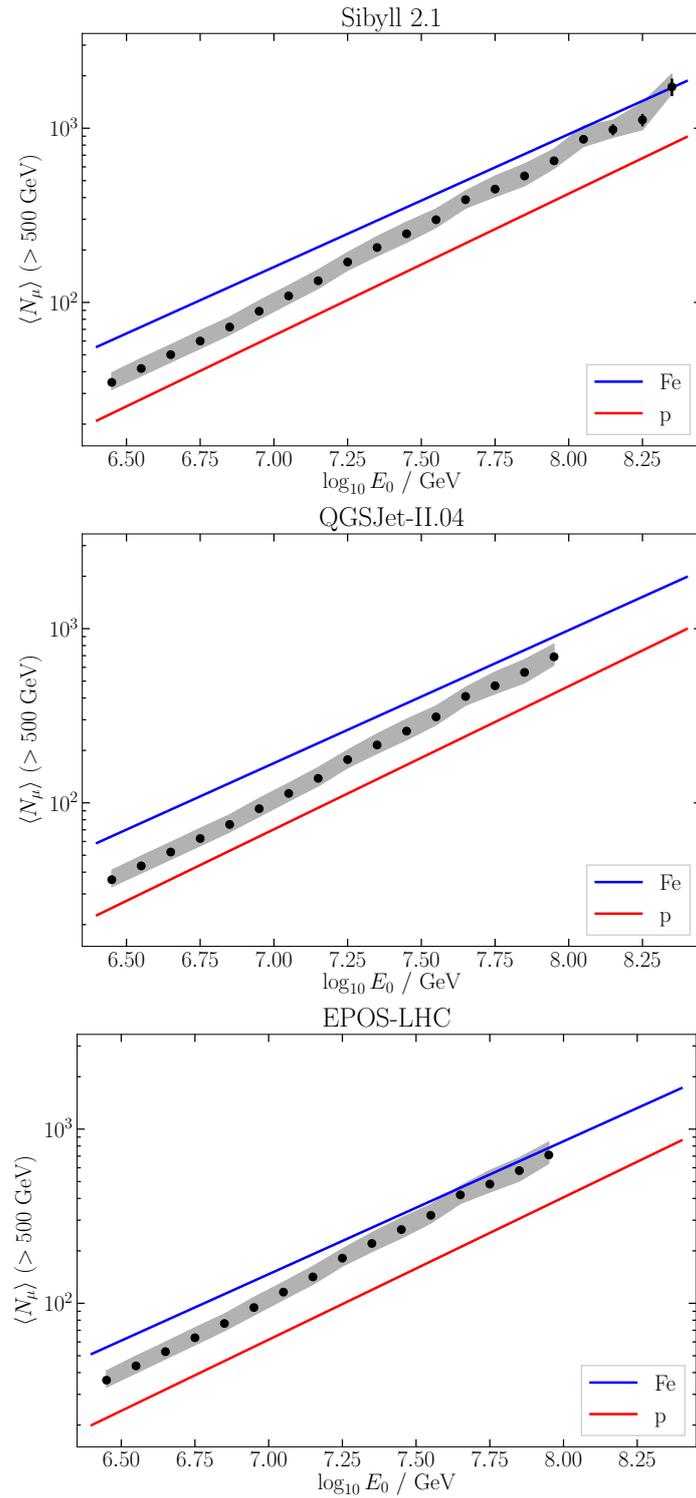
with the nominal result is used as total light yield systematic uncertainty. The effect on  $\langle N_\mu \rangle$  is of the order (+14%, -9%) for the light yield shifts of (-12.5%, +9.6%). The IceCube uncertainty therefore has the largest impact on the analysis.

We remark again that these systematic uncertainties have been adopted from previous works, and that with dedicated studies some of them could likely be reduced. For example, more modern ice models are available than to the one used in the simulation used here [141]. The uncertainty on  $\lambda_{\text{eff}}$  could likely be reduced because of the very limited zenith range that is considered here. Improved techniques to deal with snow are also being developed [162]. A new method for the propagation of systematic uncertainties in IceCube known as the SnowStorm method [204] also seems promising for an analysis like the one presented here. In contrast to the possible improvements in systematics listed here, there are indications of tension between the IceTop energy scale and that of other experiments [50]. This is something that should be studied in more detail in the future.

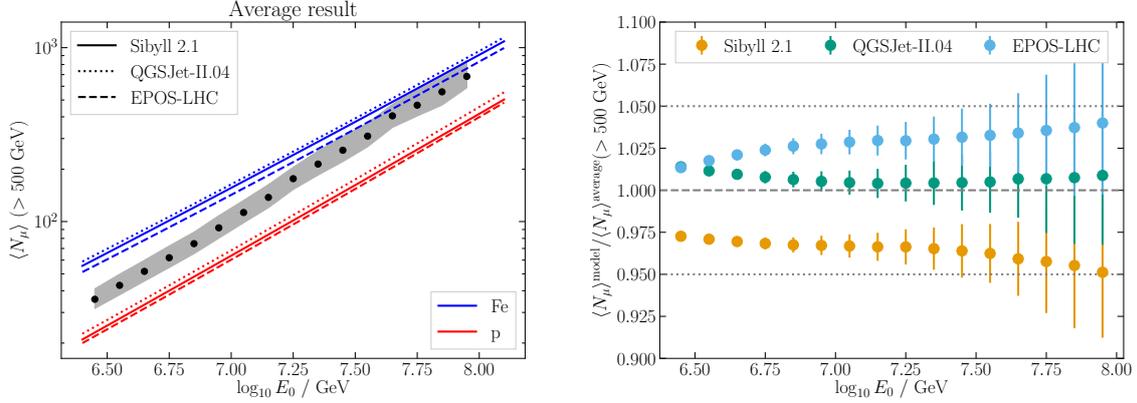
### 7.7.3 Final results

The nominal results obtained in data are combined with the total systematic uncertainty to obtain our final results on the average multiplicity of muons above 500 GeV in air showers, in an energy range starting from 2.5 PeV up to 100 PeV for QGSJet-II.04 and EPOS-LHC, and 250 PeV for Sibyll 2.1. The results, compared to the corresponding expectations from simulation, are shown in Fig. 7.32.

As the logarithmic scale makes it difficult to see differences between the results from individual models, we calculate also an average result and see how the individual results compare to it. The left panel of Fig. 7.32 shows a simple average of the three  $\langle N_\mu \rangle$  results compared to the expectations of the three models, where also the statistical have been averaged, as they are nearly identical for the individual results. The gray band



**Fig. 7.32:** Average multiplicity of muons with energy larger than 500 GeV in experimental data as a function of primary energy, obtained under the assumption of different hadronic interaction models. Error bars indicate the statistical uncertainty, the band indicates the systematic uncertainty. Shown for comparison are the multiplicities from proton and iron simulation using the corresponding model.



**Fig. 7.33:** Left: Simple average of the muon multiplicity results obtained assuming different hadronic interaction models, as shown in Fig. 7.32. The band now covers the entire range of values covered by the systematic uncertainties of the individual results. Shown for comparison are the predictions for proton and iron using the three models. Right: Ratio of the individual results obtained assuming a hadronic interaction model to the average result shown on the left.

now covers the entire envelope covered by the bands of the individual models. In the right panel, the ratio of the individual models to the average is shown. It can be seen that the differences between the results are small, varying only about  $\pm 3\%$  around the average, with the result assuming Sibyll 2.1 giving the lowest number of muons, and that assuming EPOS-LHC giving the most.

The results expressed in  $z$  are shown in Fig. 7.34. For comparison, also expectations from different composition models, discussed in Section 1.5, are included. These expectations are calculated by taking the average of the muon multiplicity expected for the different mass groups included in the model, weighted by the contribution to the total flux of the group. For example,

$$\langle N_\mu \rangle_{\text{model}} \approx f_p \langle N_\mu \rangle_p + f_{\text{He}} \langle N_\mu \rangle_{\text{He}} + f_{\text{CNO}} \langle N_\mu \rangle_{\text{CNO}} + \dots, \quad (7.17)$$

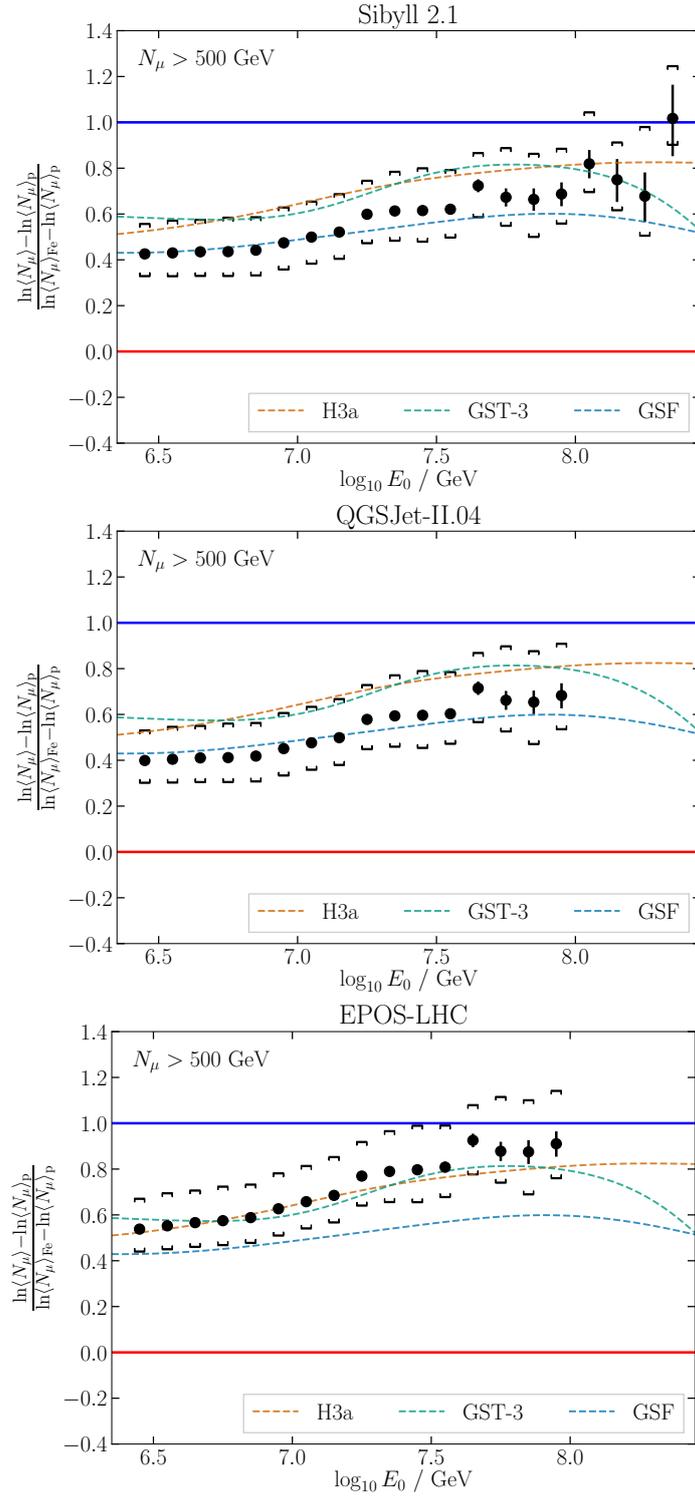
where the fraction  $f_i$  for a mass group  $i$  has been calculated as the flux of this group divided by the total flux using the CRFLUX python package [205, 206]. The muon multiplicities  $\langle N_\mu \rangle_p$  and  $\langle N_\mu \rangle_{\text{Fe}}$  are those given in Section 7.3 for a certain hadronic interaction model. For the intermediate mass groups, the expectation for  $\langle N_\mu \rangle$  is interpolated in  $\ln A$  as

$$\langle N_\mu \rangle_A \approx \langle N_\mu \rangle_p \times A^{\frac{\ln \langle N_\mu \rangle_{\text{Fe}} - \ln \langle N_\mu \rangle_p}{\ln 56}}. \quad (7.18)$$

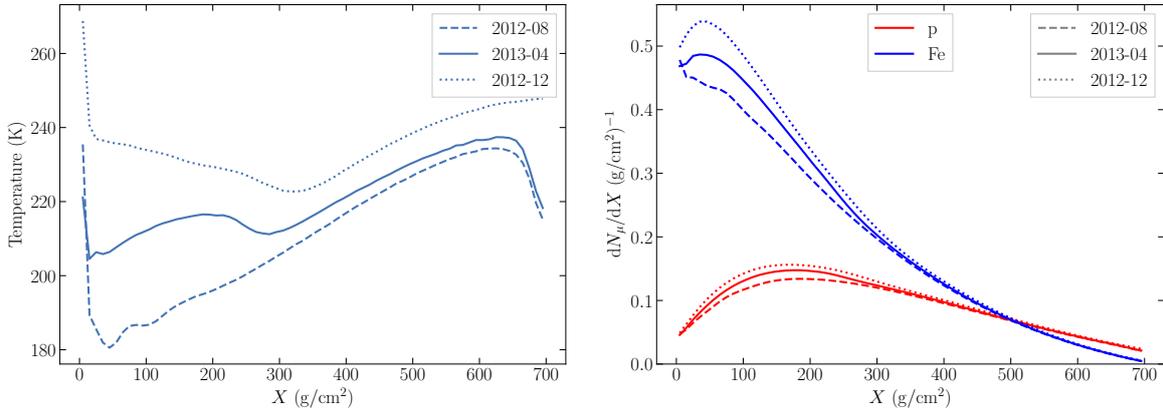
All results are bracketed by proton and iron and indicate a mixed composition that becomes heavier with increasing primary energy. The results for Sibyll 2.1 and QGSJet-II.04 are very similar in  $z$  and agree well with the expectations of the different composition models, especially GSF. The EPOS-LHC result indicates a slightly heavier composition, but also shows reasonable agreement with the composition models<sup>6</sup>.

---

<sup>6</sup>As the covariance matrix of the GSF model is not yet published, no uncertainty band was included. The other models do not include uncertainties.



**Fig. 7.34:** Average multiplicity of muons above 500 GeV in data obtained using different hadronic interaction models. The multiplicities are scaled according to expectations from proton and iron simulation according to Eq. (7.1), such that  $\ln\langle N_\mu \rangle$  is 0 for proton and 1 for iron. Shown for comparison are the expected muon multiplicity according to three cosmic-ray flux models (Section 1.5).



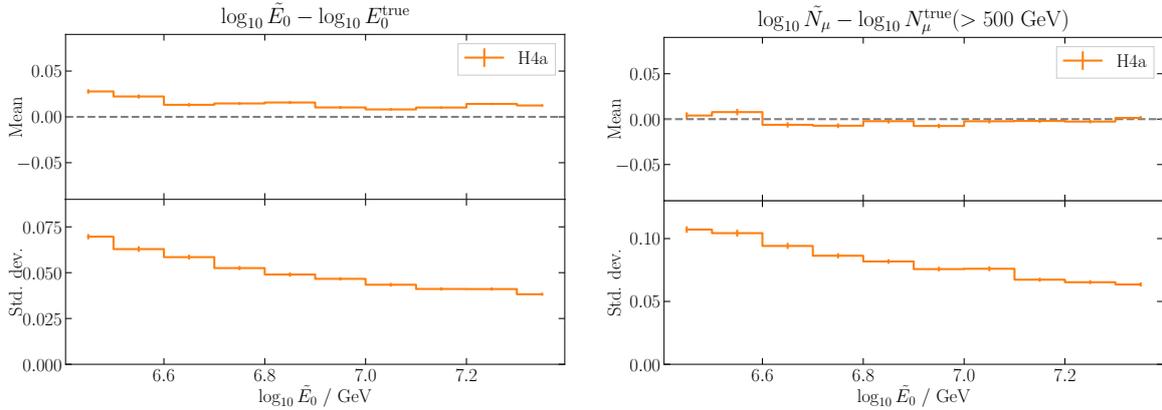
**Fig. 7.35:** Left: Average temperature profile of the atmosphere at the South Pole obtained from the NASA AIRS dataset [51] for three different months in the IC86.2012 season. Right: Muon production profiles describing the average production of muons with an energy higher than 500 GeV in 10 PeV proton and iron showers, obtained using the parameterization of Chapter 5 for the different atmospheric temperature profiles of the left plot.

## 7.8 First look at seasonal variations

In Chapter 5, the dependence of the production of muons in air showers on the atmospheric density was discussed, and a way to estimate the resulting seasonal variations in the number of high-energy muons was presented. In Fig. 7.35, average temperature profiles for different months are shown with how they are expected to change the muon production profile. This seasonal variation in the muon multiplicity can be expected to be seen using some of the tools presented in this chapter. We will present here a basic first inspection of these effects.

The number of high-energy muons in an air shower reaching the IceCube detector is expected to vary with atmospheric changes. Evidence for seasonal variations in the muon bundle signal, more specifically the  $dE/dX_{1500}$  parameter, was already presented in Fig. 5.16. On the other hand, only limited seasonal effects are expected for the shower size observable  $S_{125}$  [160], and we will neglect them. Due to the fact that only the in-ice signal shows a significant seasonal variation, we will use in this section the ENN+NNN approach, where the estimate of  $E_0$  and  $N_\mu$  are decoupled, in contrast to the NENN model which also utilizes the average relation between the two and can therefore be expected to capture seasonal effects less accurately.

To have sufficient statistics for studying the variations, the data will be binned by month. For this preliminary study, no corrections as performed for the  $\langle N_\mu \rangle$  analysis



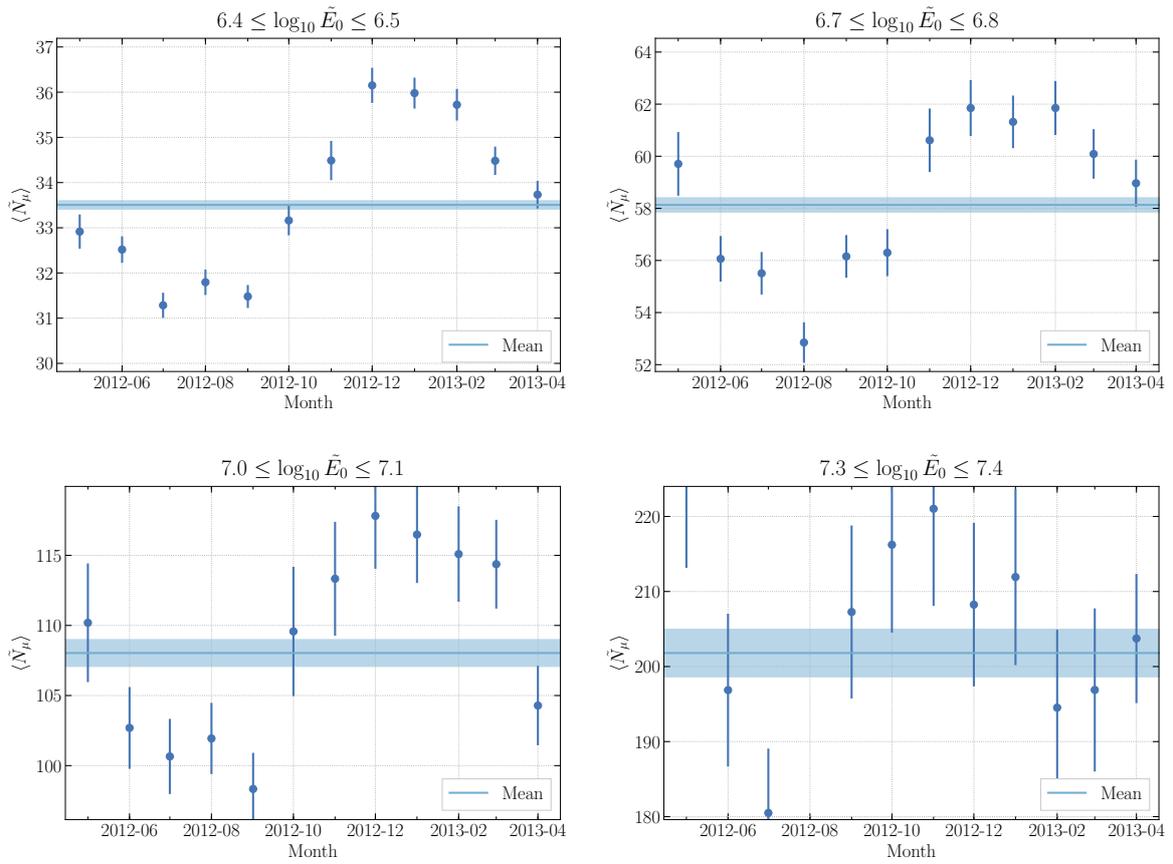
**Fig. 7.36:** Performance plots for the NNN+ENN approach on simulation weighted according to H4a. Shown are the bias and resolution of the primary energy (left) and muon multiplicity (right) prediction in bins of reconstructed energy.

discussed before will be included. Instead, we will examine the variations in the neural network outputs for  $E_0$  and  $N_\mu$ . To make clear that they are neural-network predictions rather than unfolded or calibrated results, they will be marked with a tilde, i.e.  $\tilde{E}_0$  and  $\tilde{N}_\mu$ . For a more detailed study, it would be desirable to have simulation based on the atmosphere of different months to test the reconstruction performance in these cases, rather than just the yearly average that has been used so far.

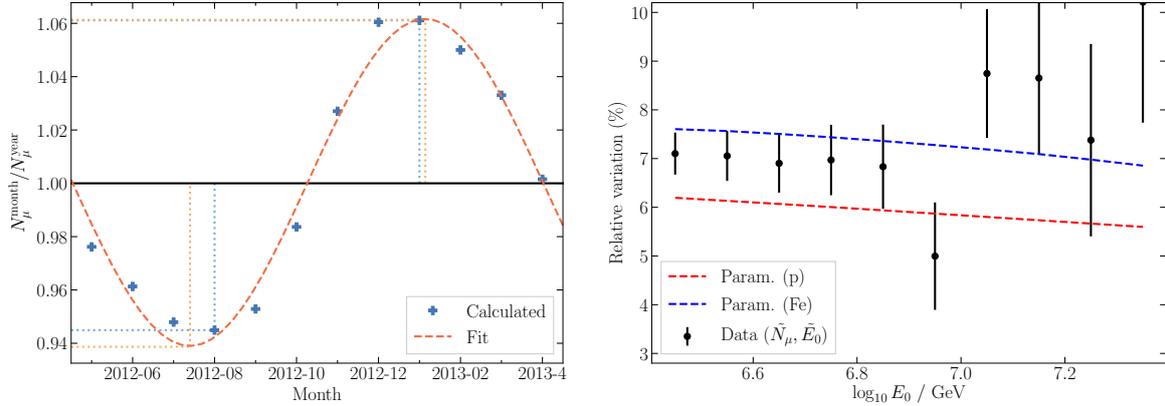
As the magnitude of the seasonal variations is not expected to be constant with primary energy, the data will also be binned in the energy estimator  $\tilde{E}_0$ . To gain insight into the accuracy that can be expected for the energy and muon multiplicity reconstruction in bins of  $\tilde{E}_0$  for experimental data, simulation is weighted according to H4a, and the performance plots derived therefrom are shown in Fig. 7.36. Due to the limited dataset and the monthly binning, the energy range is limited to  $6.4 \leq \log_{10} \tilde{E}_0 \leq 7.4$ . The energy reconstruction has a good resolution, but overestimates on average. The estimate of the muon multiplicity in each bin turns out to be rather trustworthy.

The simplest way to visualize possible seasonal variations in each  $\tilde{E}_0$  bin is to calculate the average  $\tilde{N}_\mu$  for each month. This is what's shown in Fig. 7.37. In the lowest energy bins, a variation of  $\langle \tilde{N}_\mu \rangle^{\text{month}}$  is clearly visible. A summer maximum and winter minimum are observed, and the values for April agree well with the yearly average. Toward higher energies, where the number of events is low, the statistical fluctuations become too large to observe seasonal variations.

It is tempting to compare the values shown in Fig. 7.37 to calculations made based on the parameterization from Chapter 5. While doing this shows a similar seasonal pattern,



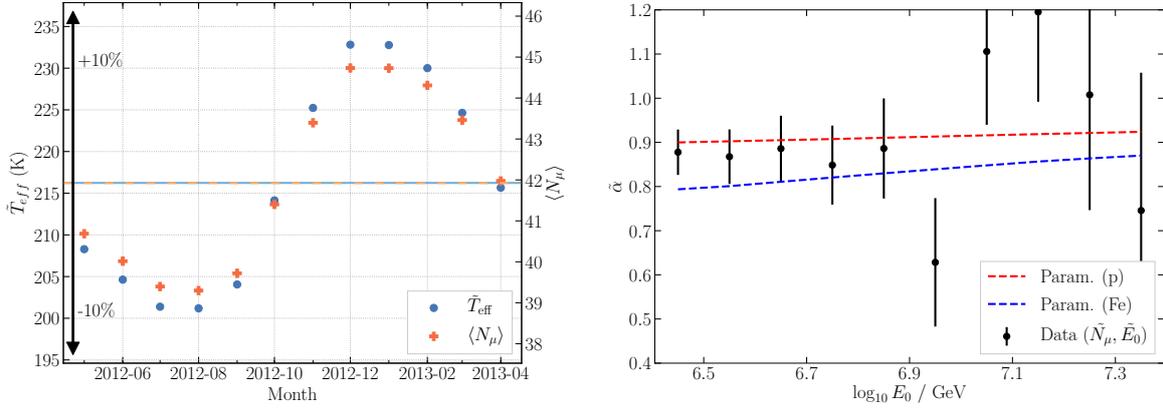
**Fig. 7.37:** Monthly average values of the neural-network reconstructed muon multiplicity  $\tilde{N}_\mu$  in different bins of neural-network reconstructed primary energy  $\tilde{E}_0$ . Seasonal variations are clearly visible in the bins with sufficient statistics.



**Fig. 7.38:** Left: Example of the average muon multiplicity over the yearly average as predicted for proton showers using the parameterization of Chapter 5. Fitted is a sine function, as a way to define the amplitude of the variation. Right: Amplitude of the relative variation in  $\tilde{N}_\mu$  in bins of  $\tilde{E}_0$  compared to predictions for proton and iron from the parameterization.

there is always an offset between the two that prevents a direct comparison. This is to be expected for several reasons. First, it is known from the discussions in Chapter 5 that the parameterization is not perfect, and that while it tends to describe the shape of a muon production profile well, the normalization is often off by a few percent. Secondly, one needs to assume a composition to calculate the expected  $\langle N_\mu \rangle$ , which we know will also not be a perfect description of reality. Finally, the  $\tilde{N}_\mu$  values obtained from data may be biased somewhat due to neural network imperfections, and it is also not immediately clear at what energy the parameterization should be evaluated to accurately describe the  $\tilde{E}_0$  bin. However, while there may be an offset in the exact number of muons, the relative variations can be expected to be similar, as they are not expected to depend strongly on energy, and are not influenced by constant offsets in normalization or neural-network biases that are assumed to not show extreme variations between months. We assess the relative variations through their amplitude. For this, the monthly values are divided by the yearly average, and the resulting values are fit with a sine function, as in the left panel of Fig. 7.38. The results obtained from data are shown in the right panel of the Figure, together with expectations from the parameterization for proton and iron. The magnitude of the relative variations are observed to be consistent between the two, around 7% in our energy range.

A second way of describing the relative variations is by looking at the correlation with an effective temperature  $\tilde{T}_{\text{eff}}$ , as the one that was already defined in Eq. (5.11). This effective temperature represents the atmospheric temperature profile as a single



**Fig. 7.39:** Left: Example of the monthly variation in muon multiplicity and effective temperature calculated using the parameterization. Right: Correlation coefficient  $\tilde{\alpha}$  fitted to the relative variations in  $\tilde{N}_\mu$  versus those in  $\tilde{T}_{\text{eff}}$ . Shown for comparison are predictions calculated from the parameterization.

number, by multiplying it with a muon production profile and integrating. The effective temperature will show similar variations as the muon multiplicity, but the amplitude of the muon multiplicity variations are typically smaller. An example of multiplicity and effective temperature calculated using the parameterization and based on the H4a composition model is shown in the left panel of Fig. 7.39. A correlation coefficient  $\tilde{\alpha}$  can be defined as

$$\frac{\Delta N_\mu}{\langle N_\mu \rangle} = \tilde{\alpha} \frac{\Delta \tilde{T}_{\text{eff}}}{\langle \tilde{T}_{\text{eff}} \rangle}, \quad (7.19)$$

where  $\langle \tilde{T}_{\text{eff}} \rangle$  is the yearly average and  $\Delta \tilde{T}_{\text{eff}} = \tilde{T}_{\text{eff}}^{\text{month}} - \langle \tilde{T}_{\text{eff}} \rangle$  and similar for the muon multiplicity. By fitting the relative differences in both quantities with a straight line one obtains  $\tilde{\alpha}$ , similar to Fig. 5.16. The results obtained in data using again  $\tilde{N}_\mu$  in bins of  $\tilde{E}_0$ , compared to expectations calculated using the parameterization, are shown in the right panel of Fig. 7.39. The correlation coefficient obtained from the neural-network reconstructions agrees well with the expectations and is around 0.87 in the considered energy range.

We conclude that a seasonal variation effect is clearly visible in the raw neural network outputs of  $\tilde{N}_\mu$ . The effect agrees qualitatively with the expectations based on the parameterization of muon production profiles presented in Chapter 7, with the multiplicity being maximal in summer and minimal in winter. The yearly average agrees well with the values obtained in October and April. Looking into the magnitude of the effect through the amplitude of the relative variations and correlation with an effective

temperature, these first results seem to agree with expectations, showing an amplitude around 7% and a correlation around 0.87.

In the future, a more detailed analysis of the seasonal variations should be performed, including some type of unfolding or other correction for neural network misreconstructions and possibly other systematic effects. Increasing the statistics and energy range would permit studying the energy dependence of the effect, and may provide insights into the physics of atmospheric muon production, such as relative contributions from pions and kaons.

## 7.9 Discussion & outlook

In this chapter, a first analysis to determine the average multiplicity of TeV muons in air showers detected with IceTop and IceCube was presented. The study of atmospheric muons is motivated by their relation to the cosmic-ray mass composition, and by the discrepancies observed between air-shower data and simulation in the muon component, known as the Muon Puzzle. In Chapter 6, inconsistencies were observed between a proxy variable for the high-energy muons observed in IceTop, and measurements related to low-energy muons in IceTop. Therefore, we chose to pursue a more detailed measurement of the high-energy muon component, complementary to the existing analysis of the density of GeV muons in IceTop.

The analysis focuses on near-vertical showers and borrows heavily from previous event selections and reconstructions. The shower size and zenith angle from the air-shower reconstruction with IceTop, together with an energy loss reconstruction performed on the signal of the muon bundle propagating through IceCube, are used as inputs to neural networks which are trained to find the primary energy and multiplicity of muons with an energy higher than 500 GeV in the shower. Afterwards, corrections derived from MC are applied to the values obtained from data, leading to results for the average muon multiplicity in air showers as a function of primary energy, under the assumption of three different hadronic interaction models: Sibyll 2.1, QGSJet-II.04, and EPOS-LHC.

The muon multiplicity results were shown in Fig. 7.32, compared to expectations for proton and iron simulation. For all three models, the results are nicely bracketed by proton and iron, which means there is no excess or deficit of TeV muons in the considered primary energy range; 2.5–250 PeV for Sibyll 2.1, 2.5–100 PeV for the two post-LHC models. Plotting the results in the  $z$ -scale (Eq. (7.1)) allows to more easily interpret the results in terms of the average mass composition. From Fig. 7.34, it can be seen that for all hadronic models, the results indicate a composition that becomes heavier between the

first and the second knee, in agreement with the KASCADE result that was discussed in Chapter 1. The results for Sibyll 2.1 and QGSJet-II.04 are very similar in  $z$ , and are in very good agreement with the expectations for the GSF composition model, and are also quite close to predictions from the H4a and GST models, also introduced in Chapter 1. The conclusions are similar for EPOS-LHC, but here a somewhat heavier composition interpretation is found.

The  $z$ -scale results can also easily be compared to other measurements, as they should point to the same underlying mass composition. Of special interest here is the comparison to the density of GeV muons measured in near-vertical showers with IceTop. The results of this analysis were shown in Fig. 4.11. It can be seen that the GeV and TeV muon results obtained for Sibyll 2.1 lead to a consistent composition interpretation, close to expectations from recent flux models. However, inconsistencies are observed between the GeV and TeV muons for the post-LHC models QGSJet-II.04 and EPOS-LHC, qualitatively consistent with those found in Chapter 6; the GeV muon results indicate a lighter mass composition than the TeV muons, even compatible with pure proton. This has important implications. It could imply that there are issues with the lateral distribution of muons in simulations, so that the measurement of the GeV muons performed at the large lateral distances of 600 and 800 m, do not agree well with data. If this is not the case, and if the muons at large lateral distance are a good proxy for the total number of muons in a shower (dominated by GeV muons), then the observed inconsistency could be an indication that the shape of the muon energy spectrum in simulations does not match that in experimental data. In that case, the mismatch seen for the low-energy muons is not just a matter of normalization or increasing the total number of muons in the shower, but changes to processes influencing also the shape of the muon spectrum should be considered for future hadronic models (see e.g. Fig. 2.11). Further interpretation in this direction will be left to the modelers. The mismatch between the GeV and TeV muon results furthermore highlight again the large uncertainties that come with indirect mass-composition measurements, even with modern hadronic interaction models.

As mentioned earlier in this thesis, several other experiments combining an air-shower array with an underground detector were operated in the past. It is worth briefly mentioning some related results. One setup consisted the predecessors of IceTop and IceCube, namely the SPASE-2 scintillator array and the AMANDA detector. AMANDA sits at a similar depth in ice as IceCube, and thus sees muons with a similar energy. No muon multiplicity analysis was published, but a composition analysis based on the energy loss of the muon bundle, conceptually similar to the IceCube composition analysis (Section 4.3.1), was performed in the 500 TeV to 5 PeV primary energy range [84]. While

there is a small overlap in energy range, a direct comparison to the results of this chapter does not seem meaningful due to the fact that a very old hadronic interaction model was used, and that the absolute composition was calibrated to direct cosmic-ray measurements. Other experiments are the EAS-TOP air-shower array with the MACRO underground detector at the National Gran Sasso Laboratories, and the Andyrchy air-shower array with the Baksan Underground Scintillation Telescope (BUST) of the Baksan Neutrino Observatory. Both the MACRO and BUST detectors are significantly smaller than IceCube or AMANDA, typically not containing the entire high-energy muon bundle in an event. These experiments have, on the other hand, an excellent resolution for the muon multiplicity, often able to accurately reconstruct the individual muon tracks in low-multiplicity events. The MACRO detector has a muon energy threshold of about 1.3 TeV. Together with EAS-TOP, which sits at a zenith angle of about  $30^\circ$  relative to MACRO, the average number of muons detected was determined for primary energies in the 1-10 PeV region, and the primary energy spectrum for a light and heavy component were derived [82], leading to similar conclusions about the average composition as in this work. Interestingly, the paper explicitly comments on the good agreement with measurements of KASCADE and EAS-TOP alone based on low-energy muons. These analyses, however, use the old hadronic interaction model QGSJET [207]. The Andyrchy-BUST experiment, with an energy threshold of 230 GeV for muons, performed a muon multiplicity measurement in the knee region, and similarly finds a mixed composition that becomes increasingly heavy with energy [208] (using the QGSJet-II.03 model [209]). A re-analysis of the data of these experiments with modern hadronic interaction models would, if possible, be worthwhile. A combined study of muon measurements, as discussed in Section 2.4.3, but now including a dependence on the large range of muon energy thresholds spanned by different experiments, could be especially interesting to determine at which muon energies, or in which kinematic regime, issues arise.

In the future, the work presented in this chapter could be improved in several ways. The production of more simulation to open the higher-energy range for QGSJet-II.04 and EPOS-LHC is one, the production of simulation based on the more modern Sibyll 2.3d is another. Including more experimental data, for example combining different seasons, is also necessary to explore higher primary energies. In that case, possible systematic effects due to the snow accumulation should be studied, especially when training neural networks on simulation representing only a single point in time. Another point of improvement is the systematic uncertainty, especially the in-ice light yield uncertainty, which should be in relatively close reach with recent improvements in ice models. It would also be interesting to repeat the analysis for a larger zenith range, or split up in different zenith bins, as

this probes muon production in different parts of the atmosphere. A first study in this direction has been performed in Ref. [210], still using events that are contained inside IceTop. Loosening this requirement, which should be possible with the RockBottom reconstruction that is being developed [156], would significantly increase the possibilities for measurements in more inclined showers. Furthermore, a more detailed study of the seasonal variations of TeV muons than the first look presented in Section 7.8, should be performed. Lastly, dedicated combined studies of the GeV and TeV muons could be performed, as their correlations may contain important information on the shower development. Such combined measurements of the GeV and TeV muon component are expected to make a significant contribution toward the resolution of the Muon Puzzle.

Muon measurements at the IceCube Neutrino Observatory will also benefit from various planned extensions to the existing instrumentation. The so-called IceCube Upgrade will consist of seven new strings, instrumenting a volume near the bottom center of the array even more densely than DeepCore, with horizontal and vertical spacings of about 20 m and 3 m respectively, and with improved DOM designs [211]. Such a dense array may prove useful for performing detailed measurements of different muon bundle properties. The Upgrade will also come with a variety of calibration devices, which will improve our understanding of the detector medium and the hole ice (Section 3.2.2), which will allow re-analysis of existing data with reduced systematic uncertainties. An enhancement of the surface instrumentation is also planned, consisting of an array of scintillators and radio antennas inside the footprint of IceTop [101, 102]. Stations consisting of 8 scintillator panels and three radio antennas, elevated above the snow, will be placed among the IceTop tanks. The scintillator panels will reduce the threshold for air-shower detection to about 0.5 PeV, increasing the accessible energy range for coincident measurements with IceCube. They may also be used in combination with the existing IceTop tanks to calibrate the impact of snow on the measurements, and to improve separation of EM particles and GeV muons. The radio array will bring a calorimetric measurements of the shower energy, which may reduce energy scale uncertainties, and a measurement of the depth of shower maximum (especially at high energies of 100 PeV and above), relevant to test hadronic interaction models. Longer term, the IceCube-Gen2 detector is proposed [212], consisting of a larger instrumented in-ice volume of about 8 km<sup>2</sup>, albeit with a lower density of strings ( $\sim 240$  m horizontal spacing), and a surface array consisting of the scintillator-radio antenna stations of the IceTop enhancement, following the same grid as the in-ice strings [213]. The larger surface array will allow the study of primary cosmic rays beyond EeV energies, and the combination with the larger in-ice detector greatly improves the zenith angle range for coincident events.



# Chapter 8

## Summary and discussion

In this thesis, various studies related to muons in cosmic-ray induced extensive air showers were performed. Muons in air showers are of special interest, as they predominantly originate from the decay of charged mesons, and are therefore the fingerprints of the hadronic cascade. Consequently, the muon component of an air shower contains information about the mass of the primary nucleus, and it can be used to study the mass composition of the cosmic-ray flux, a probe of the sources of cosmic rays. They also allow the indirect study of high-energy hadronic interactions, which reach beyond the phase space accessible with human-made accelerators. While air showers are therefore of interest for particle physics, this also comes with complications. As hadronic interaction models used in air-shower simulations need to extrapolate from accelerator measurements, this induces large uncertainties in the interpretation of air-shower observations. This has been clearly established by several observed discrepancies between muons in experimental data and expectations from simulations. This is the context in which the work presented here has been performed.

Studies have been performed based on observations obtained at the South Pole with the IceCube Neutrino Observatory. The Observatory consists out of a surface air-shower detector array, IceTop, and a cubic kilometer in-ice Cherenkov detector 1.5 km below it. A specific class of cosmic ray events has been the focus of this work: near-vertical air-shower events in the PeV-EeV range whose core lands inside the IceTop array, with a bundle of high-energy muons penetrating all the way into the IceCube detector. Such events are especially interesting for hadronic-interaction studies for several reasons. The large number of electromagnetic particles reaching IceTop allows for a good determination of the primary cosmic-ray energy. Furthermore, also muons can be seen in IceTop, especially at large lateral distance, with energies typically in the GeV range. The muon bundle signal in IceCube can be used to study muons with energies above several hundred GeV.

The muon measurements in two energy ranges probe the muon energy spectrum in air showers and the early and late stages of the air-shower development in a way that is unique to this experiment, and can be used to perform important tests of hadronic interaction models. It is the high-energy muon component, referred to as the TeV muons, that has been the main subject of study in this work.

In Chapter 5, air-shower simulations were performed to gain insight into the longitudinal production of high-energy muons, and its dependence on primary energy, mass, zenith angle, and the muon threshold energy. The results of the simulations were fitted with a physically motivated function, leading to a parameterization of muon production profiles. Of special interest in this parameterization is the explicit dependence on the atmospheric temperature. As such, the expected muon production profile can be obtained given a temperature profile, and possible seasonal variations in the muon yield at ground can be estimated. It provides a simple and quick alternative to a full simulation or to solving cascade equations to obtain expectations for a certain experiment or to analyze seasonal variations in experimental data, especially when percent-level precision is not necessary. While it is already useful in its current form, it could be improved in certain ways, such as in generalizing to different energy regimes and quantifying the uncertainties resulting from the parameterization.

In Chapter 6, the ability of IceTop and IceCube to observe different shower components was put to use to test the internal consistency of hadronic interaction models. Using the coincident events described above, different composition-sensitive observables were defined: the slope of the IceTop lateral charge distribution  $\beta$ , the density of GeV muons  $\rho_\mu$  in IceTop at a lateral distance of 600 and 800 m, and a single variable obtained from the muon bundle energy loss in IceCube  $dE/dX_{1500}$ . In a single dataset, the average value of these different observables was obtained as a function of the primary energy proxy  $S_{125}$ , the IceTop shower size. By comparing the results to expectations from simulations of proton- and iron-induced air showers, the distributions can be interpreted in terms of the average mass composition. As the mass composition is a property of the primary cosmic-ray flux, all observables should give a consistent interpretation. However, inconsistencies were found using simulations based on Sibyll 2.1 and two state-of-the-art hadronic interaction models QGSJet-II.04 and EPOS-LHC. In Sibyll 2.1, good agreement was found between  $\rho_\mu$  and  $dE/dX_{1500}$ , but  $\beta$  was inconsistent. Inconsistencies were found between all three observables for QGSJet-II.04. For EPOS-LHC,  $dE/dX_{1500}$  was consistent with  $\beta$ , but showed a larger inconsistency with  $\rho_\mu$  than the other two models. While all models and observables indicate a composition that becomes heavier between roughly 2.5 and 80 PeV, the large shifts between the observables make it difficult

---

to make any other meaningful conclusions. This highlights the difficulty of indirect composition measurements through air-shower observations, and indicates that neither of the considered models gives an adequate description of experimental data.

Motivated by the observed inconsistencies, a first analysis to obtain the actual multiplicity of high-energy muons in air showers was developed in Chapter 7. The analysis uses again the near-vertical, coincident events, and starts from various existing, high-level reconstructions of the air shower in IceTop and the muon bundle energy loss in IceCube. Several resulting variables from both IceTop and IceCube were subsequently used as input to neural networks, with as aim to obtain on an event-by-event basis an estimate of the primary cosmic-ray energy and the multiplicity of muons with an energy higher than the chosen threshold of 500 GeV. These estimates were used to determine the average multiplicity in bins of energy. Correction factors were derived from Monte Carlo to take into account biases on the results that originate from imperfect reconstructions of both the multiplicity and primary energy. Results for the average multiplicity were obtained under the assumption of Sibyll 2.1 in the energy range of 2.5–250 PeV, and assuming QGSJet-II.04 and EPOS-LHC up to 100 PeV, limited by the availability of simulation. The individual results were found to be close to each other, and to agree well with the expectations based on recent composition models. No indications of excesses or deficits in the TeV muon component were found. Furthermore, the neural networks were used to perform a first study of seasonal variations in the high-energy muon multiplicity. The preliminary results showed qualitative agreement with calculations performed using the methods of Chapter 5.

The TeV muon results obtained in Chapter 7 can be compared to the dedicated measurement of the density of GeV muons at large lateral distance in IceTop [168], similarly to what was done in Chapter 6. The composition interpretation obtained from the GeV and TeV muons is found to be in agreement when using the Sibyll 2.1 model. However, for the post-LHC models QGSJet-II.04 and EPOS-LHC, the GeV muons indicate a much lighter composition. This may indicate that these models produce too many GeV muons in the primary energy range of IceTop, in contrast to what is found in measurements of ultra-high-energy cosmic rays at, for example, the Pierre Auger Observatory [119, 120]. This disagreement between GeV and TeV muons explicitly shows the presence of a mismodeling in QGSJet-II.04 and EPOS-LHC, and point toward possible issues in the muon energy spectrum. The obtained results in this work, and future improvements toward measurements of both the GeV and TeV muon component with IceTop and IceCube, may therefore serve as important constraints toward muon production models in future iterations of hadronic interaction models, as was discussed in

Chapter 2. Improved hadronic interaction models will, in turn, be indispensable toward a better understanding of the high-energy cosmic rays: the transition between Galactic and extragalactic cosmic rays, the sources of the highest energy particles in the Universe, and various related measurements in neutrino- and gamma ray astronomy.

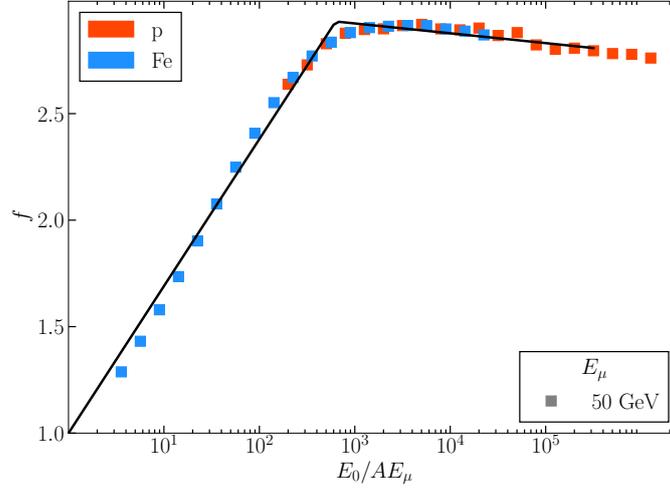
# Appendix A

## Low-energy muon profile parameterization

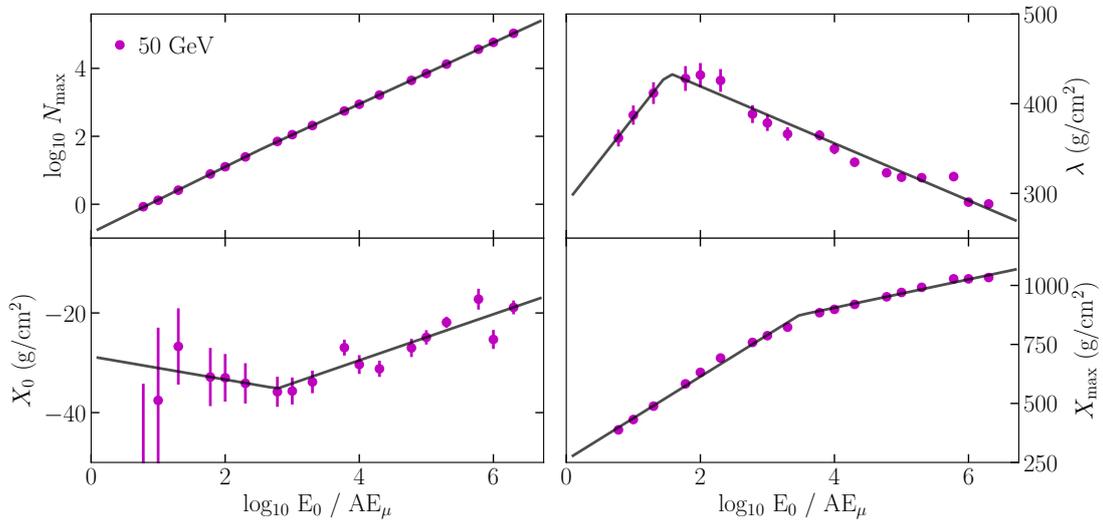
As discussed in Chapter 5, the parameterization does not scale perfectly with the ratio of the energy per nucleon  $E_0/A$  and the muon energy threshold  $E_\mu$ , and there is a remaining dependence of the profile shape on  $E_\mu$ . As a result, it is recommended to re-optimize the parameters when performing calculations with a very different  $E_\mu$  from the original parameterization of Table 5.1. In Ref. [182], we present calculations relevant for the NOvA Near Detector [187] with a muon energy threshold around 50 GeV. A second set of parameters was derived especially for this case. They are included here in Table A.1, and the fits are shown in Figs. A.1 and A.2. Both sets of parameters are included in the script provided on GitHub<sup>1</sup>.

---

<sup>1</sup><https://github.com/verpoest/muon-profile-parameterization>



**Fig. A.1:**  $f$ -factor derived from simulations for a muon energy threshold of 50 GeV. See Section 5.2.2 and Fig. 5.2 for definition.



**Fig. A.2:** Optimal fit values obtained from fitting Eq. (5.2) to vertical proton showers with  $E_\mu > 50$  GeV. Black lines show the fits to these values as given in Table A.1.

**Table A.1:** Parameter values for Eq. (5.9) for  $> 50$  GeV muons. To be combined with Eqs. (5.9) and (5.2).

	$i$	$c_i$	$p_i$	$q$
$N_{\max}$	1	0.144	0.972	2.557
	2	0.213	0.905	
	$i$	$a_i$ (g/cm <sup>2</sup> )	$b_i$ (g/cm <sup>2</sup> )	$q$
$X_{\max}$	1	260.9	176.4	3.476
	2	665.1	60.1	
$\lambda$	1	289.4	95.0	1.526
	2	483.0	-31.8	
$X_0$	1	-28.7	-2.3	2.778
	2	-48.0	4.6	
$f$	1	1	0.69	2.80
	2	3.06	-0.05	



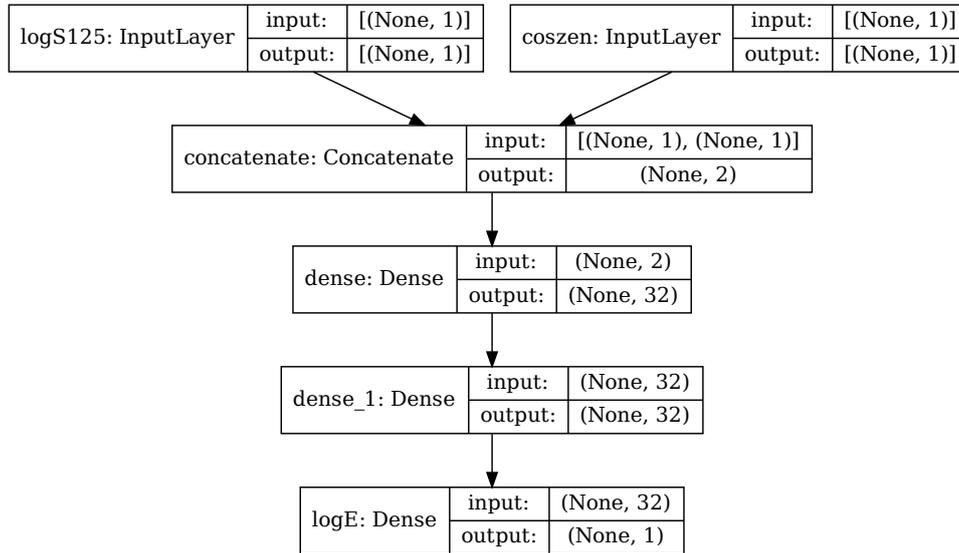
# Appendix B

## TeV muon multiplicity supplementary material

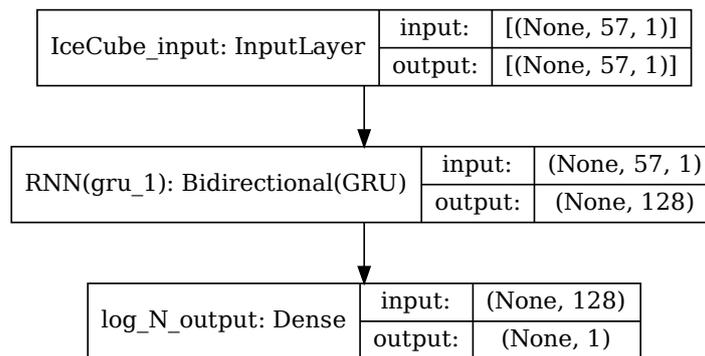
This Appendix contains some content that was not included in Chapter 7 so as not to bloat the text.

### B.1 Neural network architectures

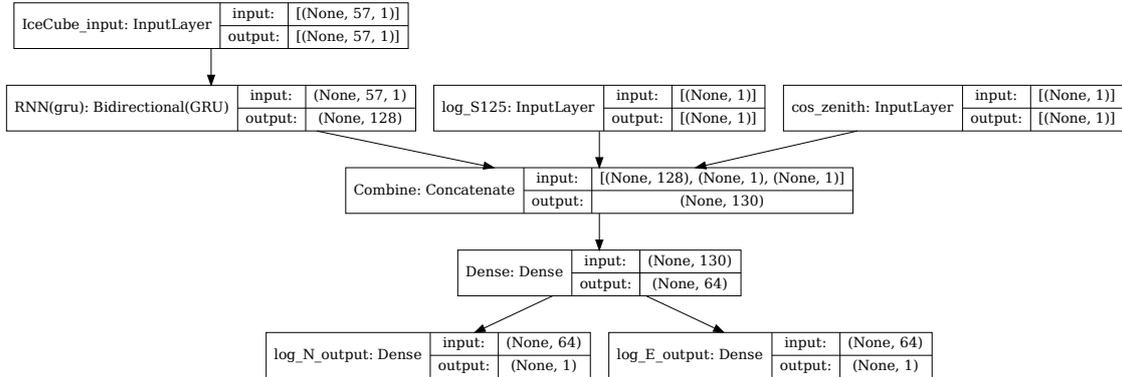
A visualization of the neural network architectures, produced with the Keras library [199], is shown in Figs. B.1, B.2 and B.3 for ENN, NNN, and NENN respectively (see Section 7.4.3 for description).



**Fig. B.1:** Architecture of the ENN neural network.



**Fig. B.2:** Architecture of the NNN neural network.



**Fig. B.3:** Architecture of the NENN neural network.

## B.2 MC tests

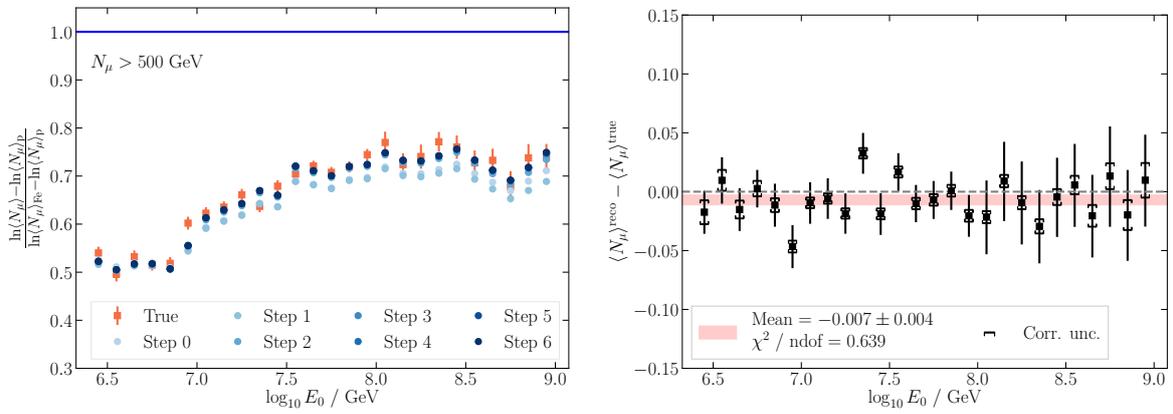
We include more plots testing the accuracy of the analysis method in simulation weighted to different composition cases, as in Section 7.6. The analysis consists out of applying the NENN neural network of Section 7.4.3, followed by the iterative correction of Section 7.5.3.

We will include plots comparing the true and reconstructed average multiplicity of muons above 500 GeV, expressed in  $z$  (Eq. (7.1)), as well as plots showing that the difference between true and reconstructed  $\langle N_\mu \rangle$  is consistent with zero.

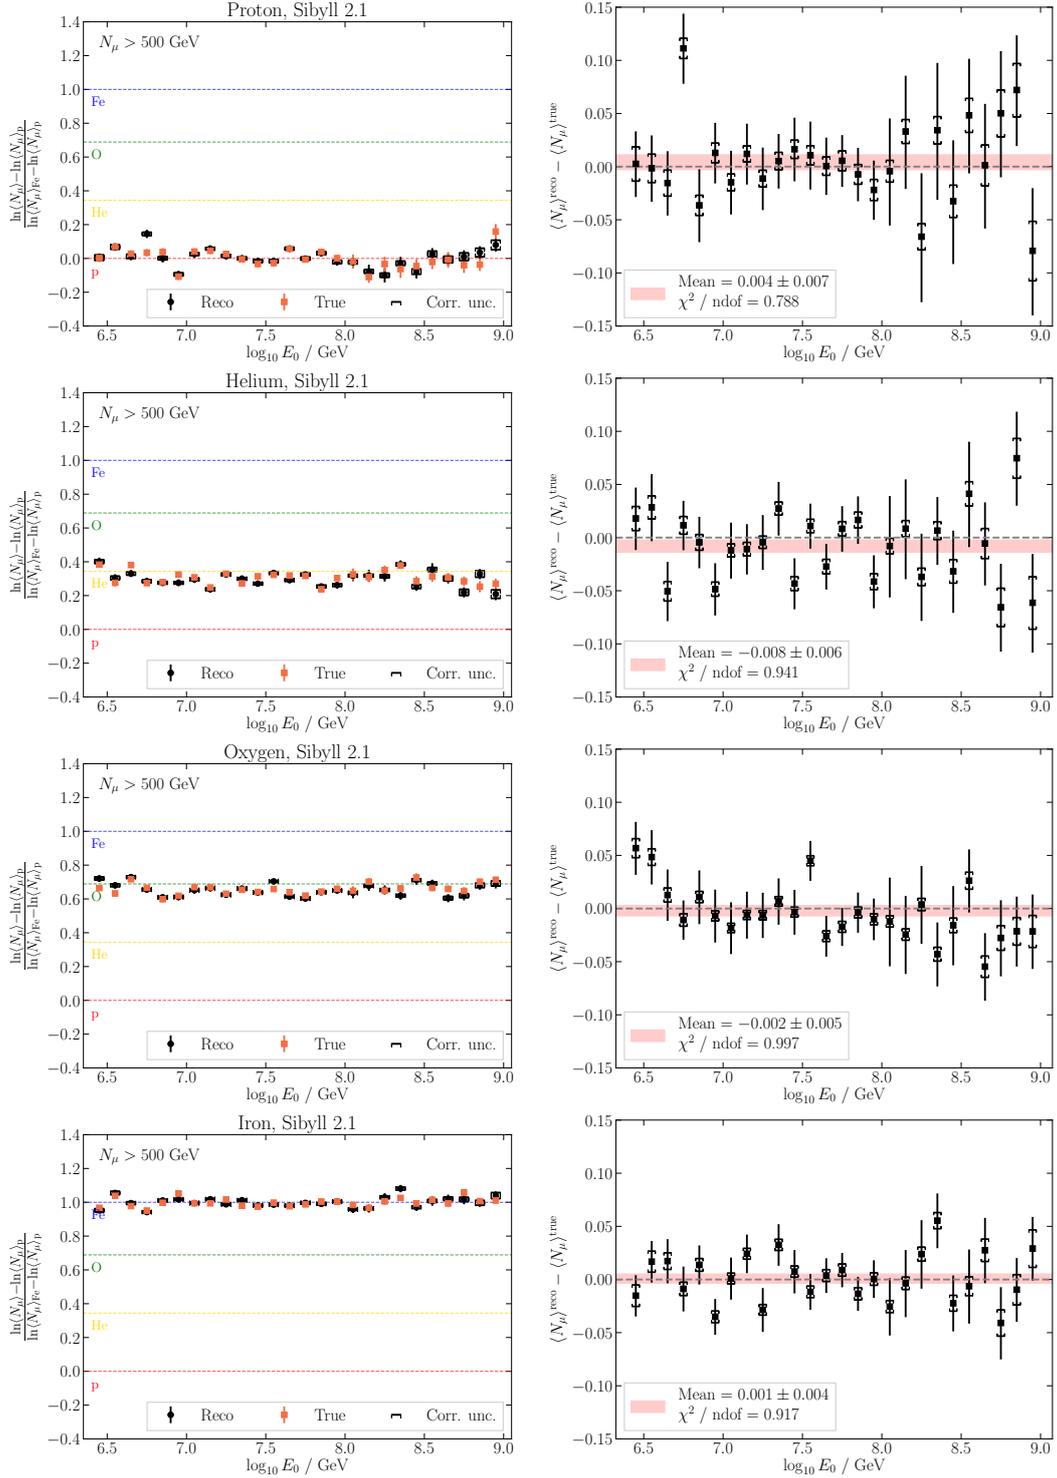
Fig. B.4 repeats the plot shown in Fig. 7.24, together with a plot showing the difference in true and reconstructed  $z$ -values. The reconstructed values come with statistical uncertainties and an uncertainty due to the statistical uncertainty of the correction factor that is applied to the results, shown by the brackets. The mean of the differences over all energy bins is derived, taking into account the total uncertainty obtained by adding the statistical and correction uncertainties in quadrature. The mean of the difference is very small and within two standard deviations of zero. There are no energy dependent trends, as confirmed by the reduced  $\chi^2$  values.

The plots for the case of a pure composition of a certain primary type causing the entire cosmic-ray flux measured on Earth is shown in Fig. B.5. Good agreement is obtained between the true and reconstructed  $\langle N_\mu \rangle$  for all primaries: p, He, O, and Fe.

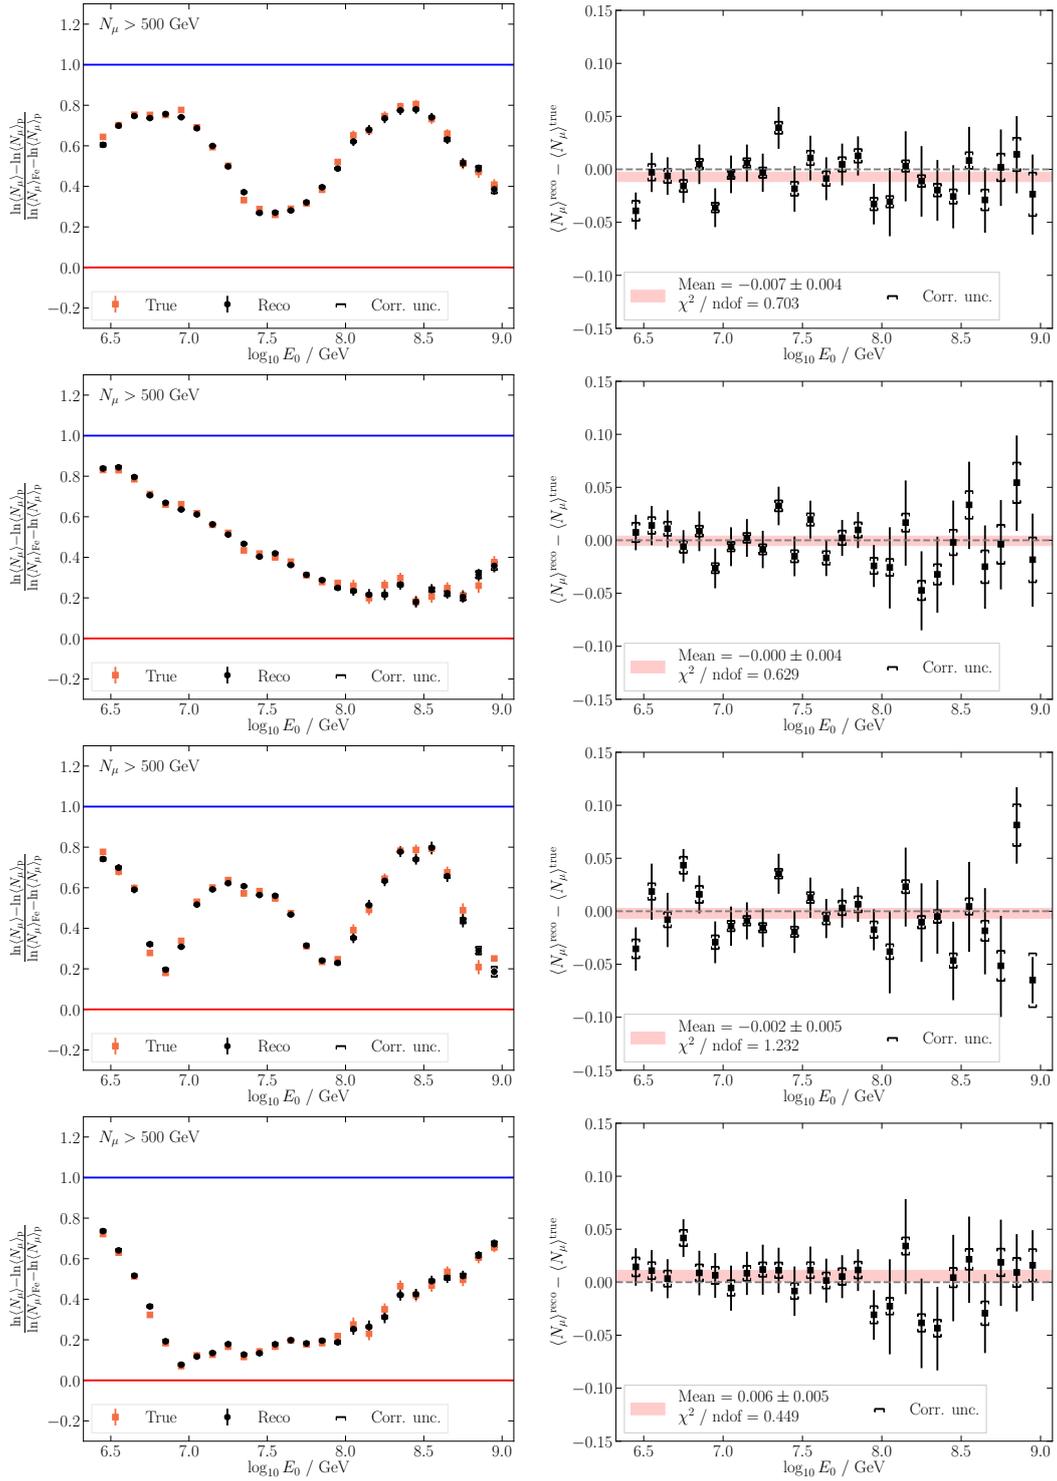
Finally, some more tests using random compositions obtained by combining the four primary groups in different ways are shown in Fig. B.6. Again, good agreement is obtained no matter the composition, confirming that the method works well universally.



**Fig. B.4:** Comparison of the reconstructed  $\langle N_\mu \rangle$  to true values in MC weighted according to H4a. Left: Comparison of the  $z$ -values. Right: Difference between the true and reconstructed values. Error bars represent the statistical uncertainties, brackets the uncertainty from the correction factor. The mean with its uncertainty is also shown.



**Fig. B.5:** Comparison of the reconstructed  $\langle N_\mu \rangle$  to true values in MC pure proton, helium, oxygen, and iron MC weighted according to the H4a all-particle spectrum. Left: Comparison of the z-values. Right: Difference between the true and reconstructed values. Error bars represent the statistical uncertainties, brackets the uncertainty from the correction factor. The mean with its uncertainty is also shown.



**Fig. B.6:** Comparison of the reconstructed  $\langle N_\mu \rangle$  to true values in MC weighted according to four artificial composition cases. Left: Comparison of the  $z$ -values. Right: Difference between the true and reconstructed values. Error bars represent the statistical uncertainties, brackets the uncertainty from the correction factor. The mean with its uncertainty is also shown.

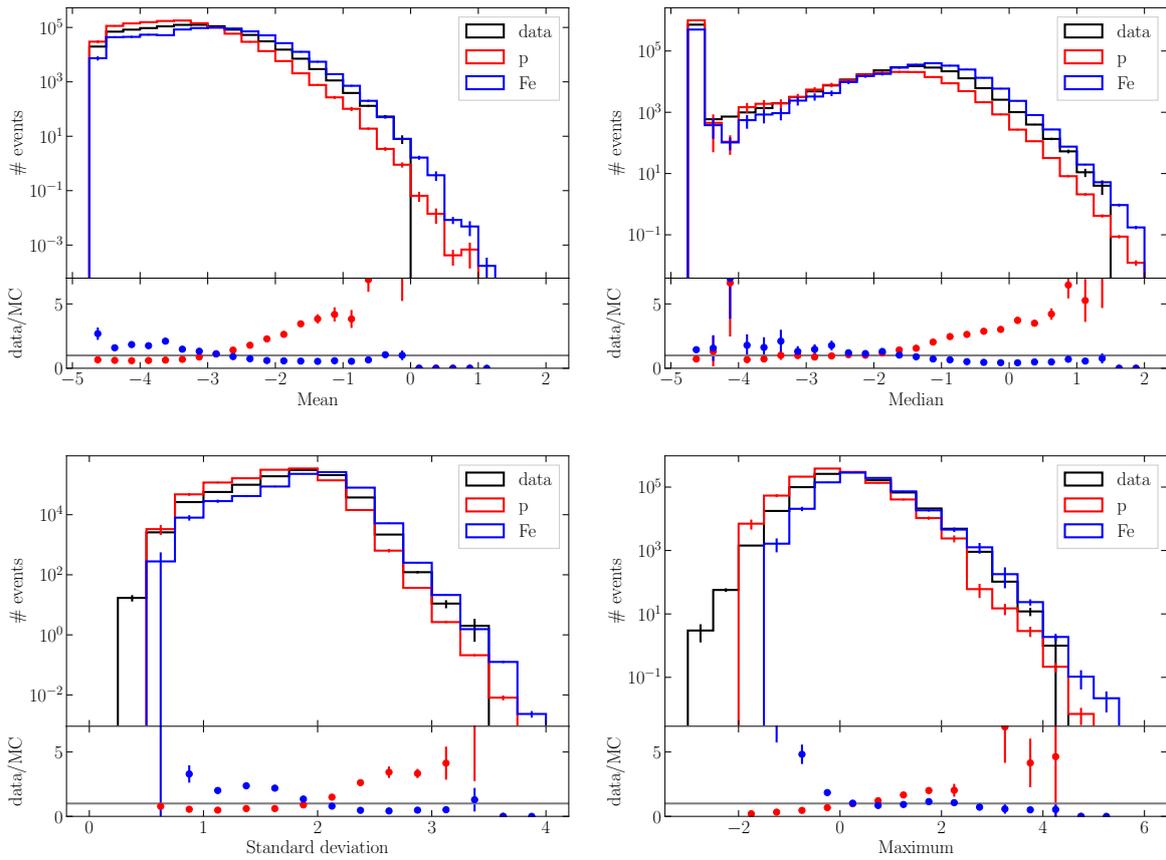
## B.3 Data-MC comparison

In this section, a more detailed data-MC comparison is given for the energy loss vector that is used as an input to the neural networks for reconstructing  $N_\mu$ . As these vectors are input to a recurrent neural network layer, it is important that the energy loss profiles are well described by simulation throughout the detector. Therefore, different variables probing the properties of different parts of the vectors will be checked, complementary to the distributions that were already shown in Fig. 7.28.

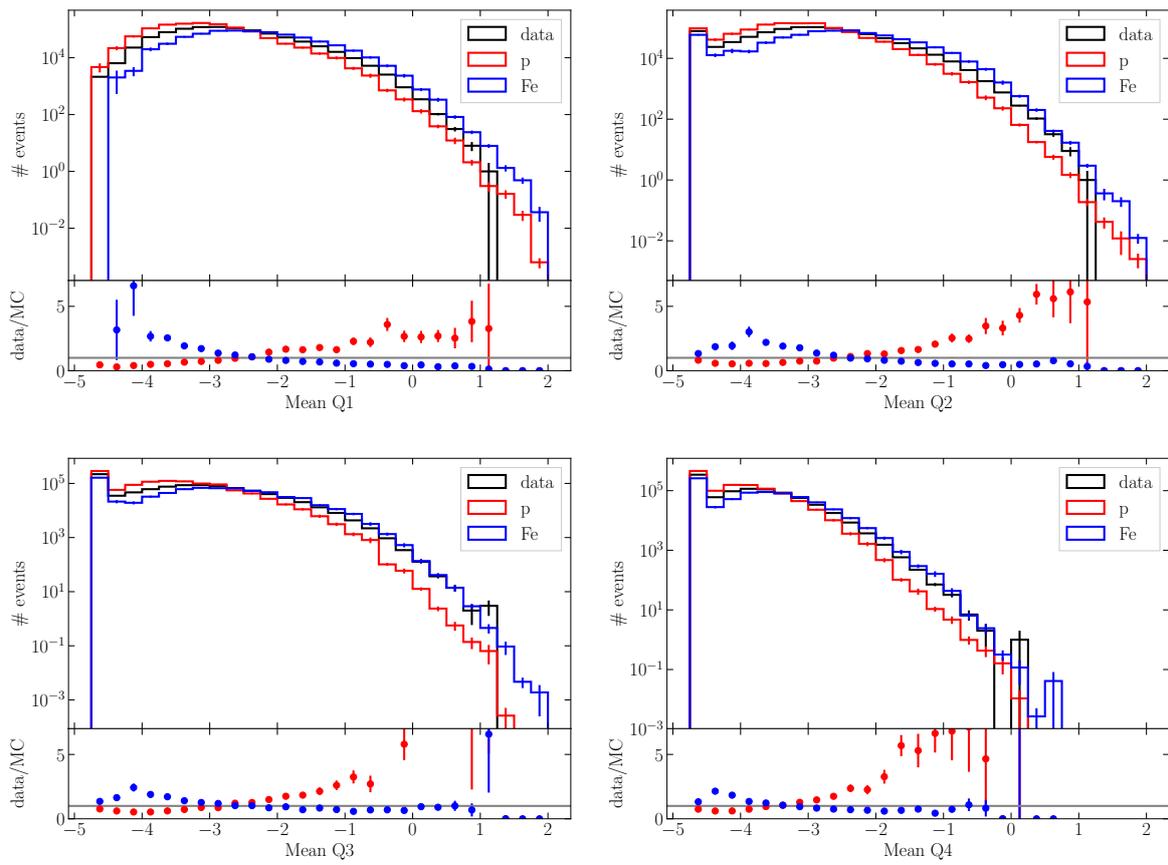
First, the distributions of simple quantities of the vectors are compared between data and MC. As described in Section 7.6.2, data is compared to pure proton and iron MC weighted to describe the entire cosmic-ray flux. In Fig. B.7, a comparison is shown for the mean, median, standard deviation, and maximum value of the vectors. For all four, the distribution in data is nicely bracketed by proton and iron.

As the shape or the evolution of the energy loss profile is important for the neural network, we check some quantities in subparts of the vector. The vector is divided into four sections of length 14, ignoring the last bin. In Fig. B.8, the distribution of the mean value in each of the section is compared between data and MC. Good agreement is again observed. For Fig. B.8, the difference between the maximal and minimal value is calculated, and the resulting distributions again do not show reasons for concern.

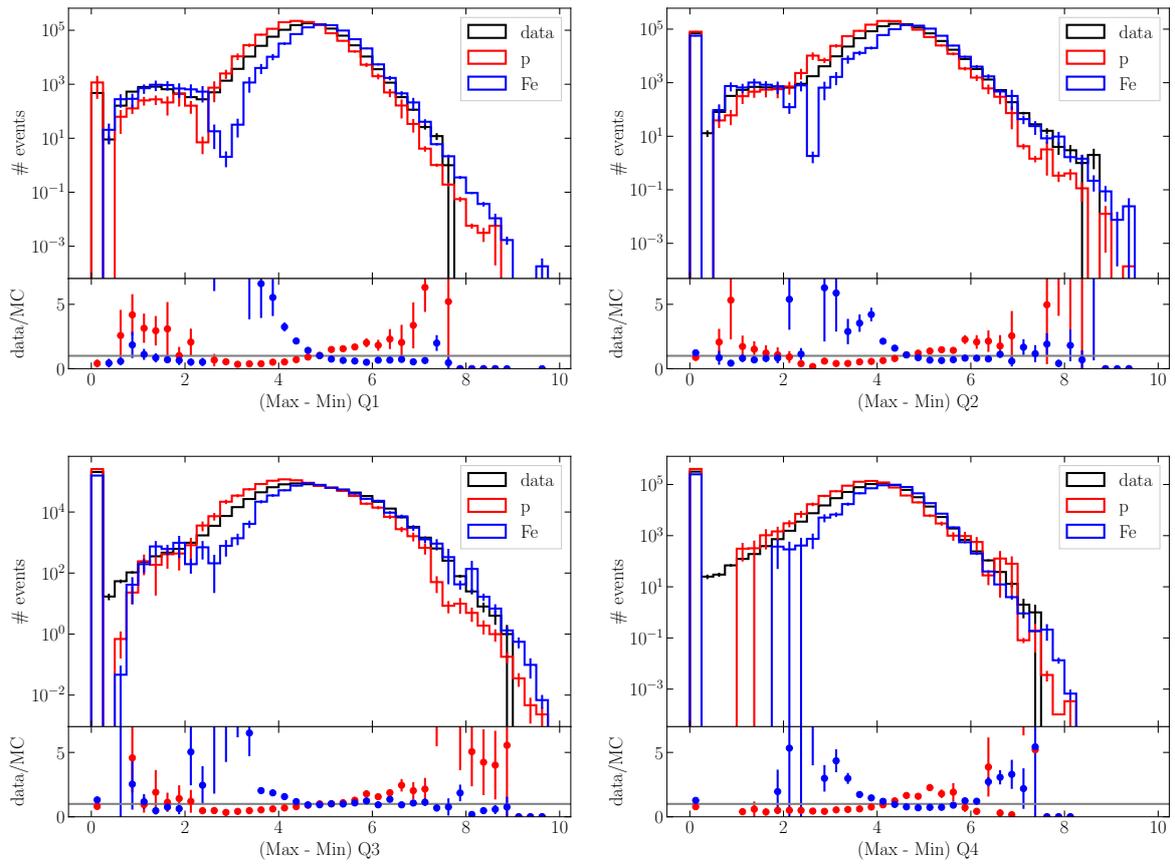
From these comparisons, we conclude that the simulation seems to give a good description of the muon bundle energy loss in IceCube, so that we can confidently use it in neural networks.



**Fig. B.7:** Comparison of the distribution of different properties of the energy loss vector used as input to the neural network reconstructions. Top left shows the mean of the vector, top right the median, bottom left the standard deviation, and bottom right the maximum.



**Fig. B.8:** Comparison of the data and MC distributions of the mean value in four sections of the energy loss vector.



**Fig. B.9:** Comparison of the data and MC distributions of the difference between the maximum and minimum value in four sections of the energy loss vector.

# References

- [1] V. F. Hess, “Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten”, *Phys. Z.* **13** (1912) 1084–1091.
- [2] P. Auger et al., “Extensive cosmic ray showers”, *Rev. Mod. Phys.* **11** (1939) 288–291, [doi:10.1103/RevModPhys.11.288](https://doi.org/10.1103/RevModPhys.11.288).
- [3] D. J. Bird et al., “Detection of a cosmic ray with measured energy well beyond the expected spectral cutoff due to cosmic microwave radiation”, *Astrophys. J.* **441** (1995) 144–150, [doi:10.1086/175344](https://doi.org/10.1086/175344), [arXiv:astro-ph/9410067](https://arxiv.org/abs/astro-ph/9410067).
- [4] C. D. Anderson, “The Positive Electron”, *Phys. Rev.* **43** (1933) 491–494, [doi:10.1103/PhysRev.43.491](https://doi.org/10.1103/PhysRev.43.491).
- [5] C. D. Anderson and S. H. Neddermeyer, “Cloud Chamber Observations of Cosmic Rays at 4300 Meters Elevation and Near Sea-Level”, *Phys. Rev.* **50** (1936) 263–271, [doi:10.1103/PhysRev.50.263](https://doi.org/10.1103/PhysRev.50.263).
- [6] C. M. G. Lattes, H. Muirhead, G. P. S. Occhialini, and C. F. Powell, “PROCESSES INVOLVING CHARGED MESONS”, *Nature* **159** (1947) 694–697, [doi:10.1038/159694a0](https://doi.org/10.1038/159694a0).
- [7] G. D. Rochester and C. C. Butler, “Evidence for the Existence of New Unstable Elementary Particles”, *Nature* **160** (1947) 855–857, [doi:10.1038/160855a0](https://doi.org/10.1038/160855a0).
- [8] L. Evans and P. Bryant, “LHC Machine”, *JINST* **3** (2008) S08001, [doi:10.1088/1748-0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001).
- [9] J. Albrecht et al., “The Muon Puzzle in cosmic-ray induced air showers and its connection to the Large Hadron Collider”, *Astrophys. Space Sci.* **367** (2022), no. 3, 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).
- [10] AMS Collaboration, M. Aguilar et al., “Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station”, *Phys. Rev. Lett.* **114** (2015) 171103, [doi:10.1103/PhysRevLett.114.171103](https://doi.org/10.1103/PhysRevLett.114.171103).
- [11] H. S. Ahn et al., “Discrepant hardening observed in cosmic-ray elemental spectra”, *Astrophys. J. Lett.* **714** (2010) L89–L93, [doi:10.1088/2041-8205/714/1/L89](https://doi.org/10.1088/2041-8205/714/1/L89), [arXiv:1004.1123](https://arxiv.org/abs/1004.1123).

- [12] AMS Collaboration, M. Aguilar et al., “The Alpha Magnetic Spectrometer (AMS) on the international space station: Part II — Results from the first seven years”, *Phys. Rept.* **894** (2021) 1–116, [doi:10.1016/j.physrep.2020.09.003](https://doi.org/10.1016/j.physrep.2020.09.003).
- [13] Y. S. Yoon et al., “Proton and Helium Spectra from the CREAM-III Flight”, *Astrophys. J.* **839** (2017), no. 1, 5, [doi:10.3847/1538-4357/aa68e4](https://doi.org/10.3847/1538-4357/aa68e4), [arXiv:1704.02512](https://arxiv.org/abs/1704.02512).
- [14] HAWC Collaboration, R. Alfaro et al., “All-particle cosmic ray energy spectrum measured by the HAWC experiment from 10 to 500 TeV”, *Phys. Rev. D* **96** (2017), no. 12, 122001, [doi:10.1103/PhysRevD.96.122001](https://doi.org/10.1103/PhysRevD.96.122001), [arXiv:1710.00890](https://arxiv.org/abs/1710.00890).
- [15] N. Gorbunov et al., “Energy spectra of abundant cosmic-ray nuclei in the NUCLEON experiment”, *Adv. Space Res.* **64** (2019), no. 12, 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).
- [16] TIBET III Collaboration, M. Amenomori et al., “The All-particle spectrum of primary cosmic rays in the wide energy range from  $10^{14}$  eV to  $10^{17}$  eV observed with the Tibet-III air-shower array”, *Astrophys. J.* **678** (2008) 1165–1179, [doi:10.1086/529514](https://doi.org/10.1086/529514), [arXiv:0801.1803](https://arxiv.org/abs/0801.1803).
- [17] V. V. Prosin et al., “Tunka-133: Results of 3 year operation”, *Nucl. Instrum. Meth. A* **756** (2014) 94–101, [doi:10.1016/j.nima.2013.09.018](https://doi.org/10.1016/j.nima.2013.09.018).
- [18] KASCADE-Grande Collaboration, S. Schöo et al., “The energy spectrum of cosmic rays in the range from  $10^{14}$  to  $10^{18}$  eV”, *PoS ICRC2015* (2016) 263, [doi:10.22323/1.236.0263](https://doi.org/10.22323/1.236.0263).
- [19] IceCube Collaboration, M. G. Aartsen et al., “Cosmic ray spectrum and composition from PeV to EeV using 3 years of data from IceTop and IceCube”, *Phys. Rev. D* **100** (2019), no. 8, 082002, [doi:10.1103/PhysRevD.100.082002](https://doi.org/10.1103/PhysRevD.100.082002), [arXiv:1906.04317](https://arxiv.org/abs/1906.04317).
- [20] D. Ivanov, “TA Spectrum Summary”, *PoS ICRC2015* (2016) 349, [doi:10.22323/1.236.0349](https://doi.org/10.22323/1.236.0349).
- [21] Pierre Auger Collaboration, P. Abreu et al., “The energy spectrum of cosmic rays beyond the turn-down around  $10^{17}$  eV as measured with the surface detector of the Pierre Auger Observatory”, *Eur. Phys. J. C* **81** (2021), no. 11, 966, [doi:10.1140/epjc/s10052-021-09700-w](https://doi.org/10.1140/epjc/s10052-021-09700-w), [arXiv:2109.13400](https://arxiv.org/abs/2109.13400).
- [22] Pierre Auger Collaboration, A. Aab et al., “Measurement of the cosmic-ray energy spectrum above  $2.5 \times 10^{18}$  eV using the Pierre Auger Observatory”, *Phys. Rev. D* **102** (2020), no. 6, 062005, [doi:10.1103/PhysRevD.102.062005](https://doi.org/10.1103/PhysRevD.102.062005), [arXiv:2008.06486](https://arxiv.org/abs/2008.06486).
- [23] D. Maurin, F. Melot, and R. Taillet, “A database of charged cosmic rays”, *Astron. Astrophys.* **569** (2014) A32, [doi:10.1051/0004-6361/201321344](https://doi.org/10.1051/0004-6361/201321344), [arXiv:1302.5525](https://arxiv.org/abs/1302.5525).

- [24] D. Maurin et al., “Cosmic-Ray Database Update: Ultra-High Energy, Ultra-Heavy, and Antinuclei Cosmic-Ray Data (CRDB v4.0)”, *Universe* **6** (2020), no. 8, 102, [doi:10.3390/universe6080102](https://doi.org/10.3390/universe6080102), [arXiv:2005.14663](https://arxiv.org/abs/2005.14663).
- [25] HiRes Collaboration, R. U. Abbasi et al., “First observation of the Greisen-Zatsepin-Kuzmin suppression”, *Phys. Rev. Lett.* **100** (2008) 101101, [doi:10.1103/PhysRevLett.100.101101](https://doi.org/10.1103/PhysRevLett.100.101101), [arXiv:astro-ph/0703099](https://arxiv.org/abs/astro-ph/0703099).
- [26] Telescope Array Collaboration, D. Ivanov, “Energy Spectrum Measured by the Telescope Array”, *PoS ICRC2019* (2019) 298, [doi:10.22323/1.358.0298](https://doi.org/10.22323/1.358.0298).
- [27] Pierre Auger Collaboration, A. Aab et al., “Features of the Energy Spectrum of Cosmic Rays above  $2.5 \times 10^{18}$  eV Using the Pierre Auger Observatory”, *Phys. Rev. Lett.* **125** (2020), no. 12, 121106, [doi:10.1103/PhysRevLett.125.121106](https://doi.org/10.1103/PhysRevLett.125.121106), [arXiv:2008.06488](https://arxiv.org/abs/2008.06488).
- [28] Particle Data Group Collaboration, R. L. Workman and Others, “Review of Particle Physics”, *PTEP* **2022** (2022) 083C01, [doi:10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [29] V. Ptuskin, “Propagation of galactic cosmic rays”, *Astroparticle Physics* **39-40** (2012) 44–51, [doi:https://doi.org/10.1016/j.astropartphys.2011.11.004](https://doi.org/10.1016/j.astropartphys.2011.11.004).
- [30] T. K. Gaisser, R. Engel, and E. Resconi, “Cosmic Rays and Particle Physics”. Cambridge University Press, 2016. ISBN 9780521016469.
- [31] E. Fermi, “On the origin of the cosmic radiation”, *Phys. Rev.* **75** (1949) 1169–1174, [doi:10.1103/PhysRev.75.1169](https://doi.org/10.1103/PhysRev.75.1169).
- [32] A. M. Hillas, “The Origin of Ultrahigh-Energy Cosmic Rays”, *Ann. Rev. Astron. Astrophys.* **22** (1984) 425–444, [doi:10.1146/annurev.aa.22.090184.002233](https://doi.org/10.1146/annurev.aa.22.090184.002233).
- [33] B. Peters, “Primary cosmic radiation and extensive air showers”, *Il Nuovo Cimento* **22** (1961) 800–819, [doi:10.1007/BF02783106](https://doi.org/10.1007/BF02783106).
- [34] IceCube Collaboration, M. G. Aartsen et al., “Astrophysical neutrinos and cosmic rays observed by IceCube”, *Adv. Space Res.* **62** (2018) 2902–2930, [doi:10.1016/j.asr.2017.05.030](https://doi.org/10.1016/j.asr.2017.05.030), [arXiv:1701.03731](https://arxiv.org/abs/1701.03731).
- [35] Fermi-LAT Collaboration, M. Ackermann et al., “Detection of the Characteristic Pion-Decay Signature in Supernova Remnants”, *Science* **339** (2013) 807, [doi:10.1126/science.1231160](https://doi.org/10.1126/science.1231160), [arXiv:1302.3307](https://arxiv.org/abs/1302.3307).
- [36] H.E.S.S. Collaboration, A. Abramowski et al., “Acceleration of petaelectronvolt protons in the Galactic Centre”, *Nature* **531** (2016) 476, [doi:10.1038/nature17147](https://doi.org/10.1038/nature17147), [arXiv:1603.07730](https://arxiv.org/abs/1603.07730).
- [37] Z. Cao et al., “Ultrahigh-energy photons up to 1.4 petaelectronvolts from 12  $\gamma$ -ray Galactic sources”, *Nature* **594** (2021) 33–36, [doi:10.1038/s41586-021-03498-z](https://doi.org/10.1038/s41586-021-03498-z).

- [38] Pierre Auger Collaboration, A. Aab et al., “Searches for Anisotropies in the Arrival Directions of the Highest Energy Cosmic Rays Detected by the Pierre Auger Observatory”, *Astrophys. J.* **804** (2015), no. 1, 15, [doi:10.1088/0004-637X/804/1/15](https://doi.org/10.1088/0004-637X/804/1/15), [arXiv:1411.6111](https://arxiv.org/abs/1411.6111).
- [39] IceCube, Fermi-LAT, MAGIC, AGILE, ASAS-SN, HAWC, H.E.S.S., INTEGRAL, Kanata, Kiso, Kapteyn, Liverpool Telescope, Subaru, Swift NuSTAR, VERITAS, VLA/17B-403 Collaboration, M. G. Aartsen et al., “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”, *Science* **361** (2018), no. 6398, eaat1378, [doi:10.1126/science.aat1378](https://doi.org/10.1126/science.aat1378), [arXiv:1807.08816](https://arxiv.org/abs/1807.08816).
- [40] K. Greisen, “End to the cosmic ray spectrum?”, *Phys. Rev. Lett.* **16** (1966) 748–750, [doi:10.1103/PhysRevLett.16.748](https://doi.org/10.1103/PhysRevLett.16.748).
- [41] G. T. Zatsepin and V. A. Kuzmin, “Upper limit of the spectrum of cosmic rays”, *JETP Lett.* **4** (1966) 78–80.
- [42] J. L. Puget, F. W. Stecker, and J. H. Bredekamp, “Photonuclear Interactions of Ultrahigh-Energy Cosmic Rays and their Astrophysical Consequences”, *Astrophys. J.* **205** (1976) 638–654, [doi:10.1086/154321](https://doi.org/10.1086/154321).
- [43] KASCADE Collaboration, T. Antoni et al., “KASCADE measurements of energy spectra for elemental groups of cosmic rays: Results and open problems”, *Astropart. Phys.* **24** (2005) 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).
- [44] KASCADE Grande Collaboration, W. D. Apel et al., “Kneelike structure in the spectrum of the heavy component of cosmic rays observed with KASCADE-Grande”, *Phys. Rev. Lett.* **107** (2011) 171104, [doi:10.1103/PhysRevLett.107.171104](https://doi.org/10.1103/PhysRevLett.107.171104), [arXiv:1107.5885](https://arxiv.org/abs/1107.5885).
- [45] K.-H. Kampert and M. Unger, “Measurements of the Cosmic Ray Composition with Air Shower Experiments”, *Astropart. Phys.* **35** (2012) 660–678, [doi:10.1016/j.astropartphys.2012.02.004](https://doi.org/10.1016/j.astropartphys.2012.02.004), [arXiv:1201.0018](https://arxiv.org/abs/1201.0018).
- [46] A. M. Hillas, “Can diffusive shock acceleration in supernova remnants account for high-energy galactic cosmic rays?”, *J. Phys. G* **31** (2005) R95–R131, [doi:10.1088/0954-3899/31/5/R02](https://doi.org/10.1088/0954-3899/31/5/R02).
- [47] T. K. Gaisser, “Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio”, *Astropart. Phys.* **35** (2012) 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).
- [48] Pierre Auger Collaboration, A. Aab et al., “Evidence for a mixed mass composition at the ankle in the cosmic-ray spectrum”, *Phys. Lett. B* **762** (2016) 288–295, [doi:10.1016/j.physletb.2016.09.039](https://doi.org/10.1016/j.physletb.2016.09.039), [arXiv:1609.08567](https://arxiv.org/abs/1609.08567).
- [49] T. K. Gaisser, T. Stanev, and S. Tilav, “Cosmic Ray Energy Spectrum from Measurements of Air Showers”, *Front. Phys. (Beijing)* **8** (2013) 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).

- [50] H. P. Dembinski et al., “Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to  $10^{11}$  GeV”, *PoS ICRC2017* (2018) 533, [doi:10.22323/1.301.0533](https://doi.org/10.22323/1.301.0533), [arXiv:1711.11432](https://arxiv.org/abs/1711.11432).
- [51] NASA, “Atmospheric Infrared Sounder (AIRS)”. <https://airs.jpl.nasa.gov/>.
- [52] P. Gondolo, G. Ingelman, and M. Thunman, “Charm production and high-energy atmospheric muon and neutrino fluxes”, *Astropart. Phys.* **5** (1996) 309–332, [doi:10.1016/0927-6505\(96\)00033-3](https://doi.org/10.1016/0927-6505(96)00033-3), [arXiv:hep-ph/9505417](https://arxiv.org/abs/hep-ph/9505417).
- [53] “AMRC, SSEC, UW-Madison”. <http://amrc.ssec.wisc.edu/>.
- [54] T. Kuwabara. Private Communication, 2019.
- [55] R. M. Wagner, “Measurement of Very High Energy Gamma-Ray Emission from Four Blazars Using the MAGIC Telescope and a Comparative Blazar Study”. PhD thesis, Tech. U. Munich, Germany, 2006.
- [56] Z. Hampel-Arias, “Cosmic Ray Observations at the TeV Scale with the HAWC Observatory”. PhD thesis, University of Wisconsin-Madison, United States, August, 2017.
- [57] A. Fedynitch, “Cascade equations and hadronic interactions at very high energies”. PhD thesis, KIT, Karlsruhe, Dept. Phys., 11, 2015. [doi:10.5445/IR/1000055433](https://doi.org/10.5445/IR/1000055433).
- [58] A. Fedynitch et al., “Hadronic interaction model sibyll 2.3c and inclusive lepton fluxes”, *Phys. Rev. D* **100** (2019), no. 10, 103018, [doi:10.1103/PhysRevD.100.103018](https://doi.org/10.1103/PhysRevD.100.103018), [arXiv:1806.04140](https://arxiv.org/abs/1806.04140).
- [59] J. Matthews, “A Heitler model of extensive air showers”, *Astropart. Phys.* **22** (2005) 387–397, [doi:10.1016/j.astropartphys.2004.09.003](https://doi.org/10.1016/j.astropartphys.2004.09.003).
- [60] W. Heitler, “The quantum theory of radiation”, volume 5 of *International Series of Monographs on Physics*. Oxford University Press, Oxford, 1936.
- [61] C. Meurer et al., “Muon production in extensive air showers and its relation to hadronic interactions”, *Czech. J. Phys.* **56** (2006) A211, [doi:10.1007/s10582-006-0156-9](https://doi.org/10.1007/s10582-006-0156-9), [arXiv:astro-ph/0512536](https://arxiv.org/abs/astro-ph/0512536).
- [62] J. M. C. Montanus, “An extended Heitler-Matthews model for the full hadronic cascade in cosmic air showers”, *Astropart. Phys.* **59** (2014) 4–11, [doi:10.1016/j.astropartphys.2014.03.010](https://doi.org/10.1016/j.astropartphys.2014.03.010).
- [63] S. Grimm, R. Engel, and D. Veberic, “Heitler-Matthews model with leading-particle effect”, *PoS ICRC2017* (2018) 299, [doi:10.22323/1.301.0299](https://doi.org/10.22323/1.301.0299).
- [64] J. Alvarez-Muñiz et al., “Hybrid simulations of extensive air showers”, *Phys. Rev. D* **66** (2002) 033011, [doi:10.1103/PhysRevD.66.033011](https://doi.org/10.1103/PhysRevD.66.033011).
- [65] R. Engel, D. Heck, and T. Pierog, “Extensive air showers and hadronic interactions at high energy”, *Ann. Rev. Nucl. Part. Sci.* **61** (2011) 467–489, [doi:10.1146/annurev.nucl.012809.104544](https://doi.org/10.1146/annurev.nucl.012809.104544).

- [66] UA1 Collaboration, C. Albajar et al., “A Study of the General Characteristics of  $p\bar{p}$  Collisions at  $\sqrt{s} = 0.2\text{-TeV}$  to  $0.9\text{-TeV}$ ”, *Nucl. Phys. B* **335** (1990) 261–287, [doi:10.1016/0550-3213\(90\)90493-W](https://doi.org/10.1016/0550-3213(90)90493-W).
- [67] G. Moliere, “Theory of the scattering of fast charged particles. 2. Repeated and multiple scattering”, *Z. Naturforsch. A* **3** (1948) 78–97.
- [68] L. Cazon, R. Conceição, and F. Riehn, “Probing the energy spectrum of hadrons in proton air interactions at ultrahigh energies through the fluctuations of the muon content of extensive air showers”, *Phys. Lett. B* **784** (2018) 68–76, [doi:10.1016/j.physletb.2018.07.026](https://doi.org/10.1016/j.physletb.2018.07.026), [arXiv:1803.05699](https://arxiv.org/abs/1803.05699).
- [69] J. W. Elbert, “Multiple muons produced by cosmic ray interactions.”, in *Proceedings of the DUMAND Summer Workshop*, pp. 101–121. Scripps Institution of Oceanography, La Jolla CA, 1979.
- [70] IceCube Collaboration, M. G. Aartsen et al., “Characterization of the Atmospheric Muon Flux in IceCube”, *Astropart. Phys.* **78** (2016) 1–27, [doi:10.1016/j.astropartphys.2016.01.006](https://doi.org/10.1016/j.astropartphys.2016.01.006), [arXiv:1506.07981](https://arxiv.org/abs/1506.07981).
- [71] N. Chiba et al., “Akeno giant air shower array (AGASA) covering 100 km<sup>2</sup> area”, *Nucl. Instrum. Meth. A* **311** (1992) 338–349, [doi:10.1016/0168-9002\(92\)90882-5](https://doi.org/10.1016/0168-9002(92)90882-5).
- [72] A. Borione et al., “A Large air shower array to search for astrophysical sources emitting  $\gamma$ -rays with energies  $\geq 10^{14}$  eV”, *Nucl. Instrum. Meth. A* **346** (1994) 329–352, [doi:10.1016/0168-9002\(94\)90722-6](https://doi.org/10.1016/0168-9002(94)90722-6).
- [73] EAS-TOP Collaboration, M. Aglietta et al., “The cosmic ray primary composition in the “knee” region through the EAS electromagnetic and muon measurements at EAS-TOP”, *Astropart. Phys.* **21** (2004) 583–596, [doi:10.1016/j.astropartphys.2004.04.005](https://doi.org/10.1016/j.astropartphys.2004.04.005).
- [74] GRAPES-3 Collaboration, S. K. Gupta et al., “GRAPES-3: A high-density air shower array for studies on the structure in the cosmic-ray energy spectrum near the knee”, *Nucl. Instrum. Meth. A* **540** (2005) 311–323, [doi:10.1016/j.nima.2004.11.025](https://doi.org/10.1016/j.nima.2004.11.025).
- [75] KASCADE Collaboration, T. Antoni et al., “The Cosmic ray experiment KASCADE”, *Nucl. Instrum. Meth. A* **513** (2003) 490–510, [doi:10.1016/S0168-9002\(03\)02076-X](https://doi.org/10.1016/S0168-9002(03)02076-X).
- [76] W. D. Apel et al., “The KASCADE-Grande experiment”, *Nucl. Instrum. Meth. A* **620** (2010) 202–216, [doi:10.1016/j.nima.2010.03.147](https://doi.org/10.1016/j.nima.2010.03.147).
- [77] A. P. Garyaka et al., “Rigidity-dependent cosmic ray energy spectra in the knee region obtained with the GAMMA experiment”, *Astropart. Phys.* **28** (2007) 169–181, [doi:10.1016/j.astropartphys.2007.04.004](https://doi.org/10.1016/j.astropartphys.2007.04.004), [arXiv:0704.3200](https://arxiv.org/abs/0704.3200).
- [78] T. V. Danilova et al., “The ANI experiment: On the investigation of interactions from hadrons and nuclei in the energy range  $10^3$  TeV to  $10^5$  TeV”, *Nucl. Instrum. Meth. A* **323** (1992) 104–107, [doi:10.1016/0168-9002\(92\)90276-A](https://doi.org/10.1016/0168-9002(92)90276-A).

- [79] A. V. Glushkov et al., “Measurements of Energy Spectrum of Primary Cosmic Rays in the Energy Range Above  $10^{17}$  eV”, *Izv. Akad. Nauk Ser. Fiz.* **40** (1976) 1023–1025.
- [80] M. T. Dova et al., “The mass composition of cosmic rays near  $10^{18}$  eV as deduced from measurements made at Volcano Ranch”, *Astropart. Phys.* **21** (2004) 597–607, doi:[10.1016/j.astropartphys.2004.04.009](https://doi.org/10.1016/j.astropartphys.2004.04.009), arXiv:[astro-ph/0312463](https://arxiv.org/abs/astro-ph/0312463).
- [81] R. Walker and A. A. Watson, “Measurement of the fluctuations in the depth of maximum of showers produced by primary particles of energy greater than  $1.5 \times 10^{17}$  eV”, *Journal of Physics G Nuclear Physics* **8** (1982) 1131–1140, doi:[10.1088/0305-4616/8/8/016](https://doi.org/10.1088/0305-4616/8/8/016).
- [82] MACRO, EAS-TOP Collaboration, M. Aglietta et al., “The Primary cosmic ray composition between  $10^{15}$  and  $10^{16}$  eV from extensive air showers electromagnetic and TeV muon data”, *Astropart. Phys.* **20** (2004) 641–652, doi:[10.1016/j.astropartphys.2003.10.004](https://doi.org/10.1016/j.astropartphys.2003.10.004), arXiv:[astro-ph/0305325](https://arxiv.org/abs/astro-ph/0305325).
- [83] V. N. Bakatanov, Y. F. Novoseltsev, and R. V. Novoseltseva, “Observation of ‘the knee’ in cosmic ray energy spectrum with underground muons and primary mass composition in the range  $10^{15}$  eV to  $10^{17}$  eV”, *Astropart. Phys.* **12** (1999) 19–24, doi:[10.1016/S0927-6505\(99\)00074-2](https://doi.org/10.1016/S0927-6505(99)00074-2).
- [84] AMANDA, SPASE Collaboration, J. Ahrens et al., “Measurement of the cosmic ray composition at the knee with the SPASE-2/AMANDA-B10 detectors”, *Astropart. Phys.* **21** (2004) 565–581, doi:[10.1016/j.astropartphys.2004.04.007](https://doi.org/10.1016/j.astropartphys.2004.04.007).
- [85] SPASE, AMANDA Collaboration, J. Ahrens et al., “Calibration and survey of AMANDA with the SPASE detectors”, *Nucl. Instrum. Meth. A* **522** (2004) 347–359, doi:[10.1016/j.nima.2003.12.007](https://doi.org/10.1016/j.nima.2003.12.007).
- [86] T. Gaisser and A. Hillas, “Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers”, *Proceedings, 15th International Cosmic Ray Conference (ICRC1977): Ploudiv, Bulgaria* **8** (1977) 353–357.
- [87] G. L. Cassiday et al., “Measurements of cosmic-ray air shower development at energies above  $10^{17}$  eV”, *Astrophys. J.* **356** (1990) 669, doi:[10.1086/168873](https://doi.org/10.1086/168873).
- [88] HiRes Collaboration, R. U. Abbasi et al., “A Study of the composition of ultrahigh energy cosmic rays using the High Resolution Fly’s Eye”, *Astrophys. J.* **622** (2005) 910–926, doi:[10.1086/427931](https://doi.org/10.1086/427931), arXiv:[astro-ph/0407622](https://arxiv.org/abs/astro-ph/0407622).
- [89] Pierre Auger Collaboration, J. Abraham et al., “The Fluorescence Detector of the Pierre Auger Observatory”, *Nucl. Instrum. Meth. A* **620** (2010) 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).
- [90] Telescope Array Collaboration, J. N. Matthews et al., “The Telescope Array Experiment”, in *32nd International Cosmic Ray Conference*, volume 2, p. 273. 8, 2011. doi:[10.7529/ICRC2011/V02/1306](https://doi.org/10.7529/ICRC2011/V02/1306).

- [91] A. Karle, “Design and performance of the angle integrating Cherenkov array AIROBICC”, *Astropart. Phys.* **3** (1995) 321–347, [doi:10.1016/0927-6505\(95\)00009-6](https://doi.org/10.1016/0927-6505(95)00009-6).
- [92] EAS-TOP, MACRO Collaboration, M. Aglietta et al., “The cosmic ray proton, helium and CNO fluxes in the 100-TeV energy region from TeV muons and EAS atmospheric Cherenkov light observations of MACRO and EAS-TOP”, *Astropart. Phys.* **21** (2004) 223–240, [doi:10.1016/j.astropartphys.2004.01.005](https://doi.org/10.1016/j.astropartphys.2004.01.005).
- [93] J. W. Fowler et al., “A Measurement of the cosmic ray spectrum and composition at the knee”, *Astropart. Phys.* **15** (2001) 49–64, [doi:10.1016/S0927-6505\(00\)00139-0](https://doi.org/10.1016/S0927-6505(00)00139-0), [arXiv:astro-ph/0003190](https://arxiv.org/abs/astro-ph/0003190).
- [94] B. A. Antokhonov et al., “A new 1 km<sup>2</sup> EAS Cherenkov array in the Tunka valley”, *Nucl. Instrum. Meth. A* **639** (2011) 42–45, [doi:10.1016/j.nima.2010.09.142](https://doi.org/10.1016/j.nima.2010.09.142).
- [95] S. P. Knurenko and A. Sabourov, “Spectrum and mass composition of cosmic rays in the energy range 10<sup>15</sup>-10<sup>18</sup> eV derived from the Yakutsk array data array”, in *32nd International Cosmic Ray Conference*, volume 1, pp. 189–192. 12, 2011. [arXiv:1112.2430](https://arxiv.org/abs/1112.2430). [doi:10.7529/ICRC2011/V01/0182](https://doi.org/10.7529/ICRC2011/V01/0182).
- [96] F. G. Schröder, “Radio detection of Cosmic-Ray Air Showers and High-Energy Neutrinos”, *Prog. Part. Nucl. Phys.* **93** (2017) 1–68, [doi:10.1016/j.pnpnp.2016.12.002](https://doi.org/10.1016/j.pnpnp.2016.12.002), [arXiv:1607.08781](https://arxiv.org/abs/1607.08781).
- [97] G. A. Askaryan, “Excess negative charge of an electron-photon shower and its coherent radio emission”, *Zh. Eksp. Teor. Fiz.* **41** (1961) 616–618.
- [98] P. A. Bezyazeev et al., “Reconstruction of cosmic ray air showers with Tunka-Rex data using template fitting of radio pulses”, *Phys. Rev. D* **97** (2018), no. 12, 122004, [doi:10.1103/PhysRevD.97.122004](https://doi.org/10.1103/PhysRevD.97.122004), [arXiv:1803.06862](https://arxiv.org/abs/1803.06862).
- [99] S. Buitink et al., “Method for high precision reconstruction of air shower  $X_{max}$  using two-dimensional radio intensity profiles”, *Phys. Rev. D* **90** (2014), no. 8, 082003, [doi:10.1103/PhysRevD.90.082003](https://doi.org/10.1103/PhysRevD.90.082003), [arXiv:1408.7001](https://arxiv.org/abs/1408.7001).
- [100] Pierre Auger Collaboration, A. Aab et al., “Energy Estimation of Cosmic Rays with the Engineering Radio Array of the Pierre Auger Observatory”, *Phys. Rev. D* **93** (2016), no. 12, 122005, [doi:10.1103/PhysRevD.93.122005](https://doi.org/10.1103/PhysRevD.93.122005), [arXiv:1508.04267](https://arxiv.org/abs/1508.04267).
- [101] IceCube Collaboration, A. Haungs, “A Scintillator and Radio Enhancement of the IceCube Surface Detector Array”, *EPJ Web Conf.* **210** (2019) 06009, [doi:10.1051/epjconf/201921006009](https://doi.org/10.1051/epjconf/201921006009), [arXiv:1903.04117](https://arxiv.org/abs/1903.04117).
- [102] IceCube Collaboration, F. G. Schröder, “Science Case of a Scintillator and Radio Surface Array at IceCube”, *PoS ICRC2019* (2020) 418, [doi:10.22323/1.358.0418](https://doi.org/10.22323/1.358.0418), [arXiv:1908.11469](https://arxiv.org/abs/1908.11469).
- [103] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw, “CORSIKA: A Monte Carlo code to simulate extensive air showers”, *Wissenschaftliche Berichte, Forschungszentrum Karlsruhe* (1998).

- [104] T. Huege, M. Ludwig, and C. W. James, “Simulating radio emission from air showers with CoREAS”, *AIP Conf. Proc.* **1535** (2013), no. 1, 128, [doi:10.1063/1.4807534](https://doi.org/10.1063/1.4807534), [arXiv:1301.2132](https://arxiv.org/abs/1301.2132).
- [105] A. M. Hillas, “Shower simulation: Lessons from MOCCA”, *Nucl. Phys. B Proc. Suppl.* **52** (1997) 29–42, [doi:10.1016/S0920-5632\(96\)00847-X](https://doi.org/10.1016/S0920-5632(96)00847-X).
- [106] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft, “FLUKA: A multi-particle transport code (Program version 2005)”, [doi:10.2172/877507](https://doi.org/10.2172/877507).
- [107] T. T. Böhlen et al., “The FLUKA Code: Developments and Challenges for High Energy and Medical Applications”, *Nucl. Data Sheets* **120** (2014) 211–214, [doi:10.1016/j.nds.2014.07.049](https://doi.org/10.1016/j.nds.2014.07.049).
- [108] S. A. Bass et al., “Microscopic models for ultrarelativistic heavy ion collisions”, *Prog. Part. Nucl. Phys.* **41** (1998) 255–369, [doi:10.1016/S0146-6410\(98\)00058-1](https://doi.org/10.1016/S0146-6410(98)00058-1), [arXiv:nuc1-th/9803035](https://arxiv.org/abs/nuc1-th/9803035).
- [109] M. Bleicher et al., “Relativistic hadron-hadron collisions in the ultrarelativistic quantum molecular dynamics model”, *J. Phys. G* **25** (1999) 1859–1896, [doi:10.1088/0954-3899/25/9/308](https://doi.org/10.1088/0954-3899/25/9/308), [arXiv:hep-ph/9909407](https://arxiv.org/abs/hep-ph/9909407).
- [110] H. Fesefeldt, “The simulation of hadronic showers: physics and applications”. PITHA.: Technische Hochschule Aachen. Physikal. Inst., 1985.
- [111] F. Riehn et al., “Hadronic interaction model Sibyll 2.3d and extensive air showers”, *Phys. Rev. D* **102** (2020), no. 6, 063002, [doi:10.1103/PhysRevD.102.063002](https://doi.org/10.1103/PhysRevD.102.063002), [arXiv:1912.03300](https://arxiv.org/abs/1912.03300).
- [112] S. Ostapchenko, “Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model”, *Phys. Rev. D* **83** (2011) 014018, [doi:10.1103/PhysRevD.83.014018](https://doi.org/10.1103/PhysRevD.83.014018), [arXiv:1010.1869](https://arxiv.org/abs/1010.1869).
- [113] S. Ostapchenko, “LHC data on inelastic diffraction and uncertainties in the predictions for longitudinal extensive air shower development”, *Phys. Rev. D* **89** (2014), no. 7, 074009, [doi:10.1103/PhysRevD.89.074009](https://doi.org/10.1103/PhysRevD.89.074009), [arXiv:1402.5084](https://arxiv.org/abs/1402.5084).
- [114] T. Pierog et al., “EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider”, *Phys. Rev. C* **92** (2015), no. 3, 034906, [doi:10.1103/PhysRevC.92.034906](https://doi.org/10.1103/PhysRevC.92.034906), [arXiv:1306.0121](https://arxiv.org/abs/1306.0121).
- [115] S. Roesler, R. Engel, and J. Ranft, “The Monte Carlo event generator DPMJET-III”, in *International Conference on Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications (MC 2000)*, pp. 1033–1038. 12, 2000. [arXiv:hep-ph/0012252](https://arxiv.org/abs/hep-ph/0012252). [doi:10.1007/978-3-642-18211-2\\_166](https://doi.org/10.1007/978-3-642-18211-2_166).
- [116] E.-J. Ahn et al., “Cosmic ray interaction event generator SIBYLL 2.1”, *Phys. Rev. D* **80** (2009) 094003, [doi:10.1103/PhysRevD.80.094003](https://doi.org/10.1103/PhysRevD.80.094003), [arXiv:0906.4113](https://arxiv.org/abs/0906.4113).

- [117] HiRes, MIA Collaboration, T. Abu-Zayyad et al., “Evidence for Changing of Cosmic Ray Composition between  $10^{17}$ -eV and  $10^{18}$ -eV from Multicomponent Measurements”, *Phys. Rev. Lett.* **84** (2000) 4276–4279, [doi:10.1103/PhysRevLett.84.4276](https://doi.org/10.1103/PhysRevLett.84.4276), [arXiv:astro-ph/9911144](https://arxiv.org/abs/astro-ph/9911144).
- [118] A. G. Bogdanov et al., “Investigation of the properties of the flux and interaction of ultrahigh-energy cosmic rays by the method of local-muon-density spectra”, *Phys. Atom. Nucl.* **73** (2010) 1852–1869, [doi:10.1134/S1063778810110074](https://doi.org/10.1134/S1063778810110074).
- [119] Pierre Auger Collaboration, A. Aab et al., “Muons in Air Showers at the Pierre Auger Observatory: Mean Number in Highly Inclined Events”, *Phys. Rev. D* **91** (2015), no. 3, 032003, [doi:10.1103/PhysRevD.91.032003](https://doi.org/10.1103/PhysRevD.91.032003), [arXiv:1408.1421](https://arxiv.org/abs/1408.1421). [Erratum: *Phys.Rev.D* 91, 059901 (2015)].
- [120] Pierre Auger Collaboration, A. Aab et al., “Measurement of the Fluctuations in the Number of Muons in Extensive Air Showers with the Pierre Auger Observatory”, *Phys. Rev. Lett.* **126** (2021), no. 15, 152002, [doi:10.1103/PhysRevLett.126.152002](https://doi.org/10.1103/PhysRevLett.126.152002), [arXiv:2102.07797](https://arxiv.org/abs/2102.07797).
- [121] Pierre Auger Collaboration, A. Aab et al., “Depth of maximum of air-shower profiles at the Pierre Auger Observatory. II. Composition implications”, *Phys. Rev. D* **90** (2014), no. 12, 122006, [doi:10.1103/PhysRevD.90.122006](https://doi.org/10.1103/PhysRevD.90.122006), [arXiv:1409.5083](https://arxiv.org/abs/1409.5083).
- [122] EAS-MSU, IceCube, KASCADE-Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array, Yakutsk EAS Array Collaboration, H. P. Dembinski et al., “Report on Tests and Measurements of Hadronic Interaction Properties with Air Showers”, *EPJ Web Conf.* **210** (2019) 02004, [doi:10.1051/epjconf/201921002004](https://doi.org/10.1051/epjconf/201921002004), [arXiv:1902.08124](https://arxiv.org/abs/1902.08124).
- [123] EAS-MSU, IceCube, KASCADE-Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array, Yakutsk EAS Array Collaboration, D. Soldin, “Update on the Combined Analysis of Muon Measurements from Nine Air Shower Experiments”, *PoS ICRC2021* (2021) 349, [doi:10.22323/1.395.0349](https://doi.org/10.22323/1.395.0349), [arXiv:2108.08341](https://arxiv.org/abs/2108.08341).
- [124] R. Ulrich, R. Engel, and M. Unger, “Hadronic Multiparticle Production at Ultra-High Energies and Extensive Air Showers”, *Phys. Rev. D* **83** (2011) 054026, [doi:10.1103/PhysRevD.83.054026](https://doi.org/10.1103/PhysRevD.83.054026), [arXiv:1010.4310](https://arxiv.org/abs/1010.4310).
- [125] S. Baur et al., “Core-corona effect in hadron collisions and muon production in air showers”, [arXiv:1902.09265](https://arxiv.org/abs/1902.09265).
- [126] ALICE Collaboration, J. Adam et al., “Enhanced production of multi-strange hadrons in high-multiplicity proton-proton collisions”, *Nature Phys.* **13** (2017) 535–539, [doi:10.1038/nphys4111](https://doi.org/10.1038/nphys4111), [arXiv:1606.07424](https://arxiv.org/abs/1606.07424).
- [127] LHCb Collaboration, R. Aaij et al., “LHCb Detector Performance”, *Int. J. Mod. Phys. A* **30** (2015), no. 07, 1530022, [doi:10.1142/S0217751X15300227](https://doi.org/10.1142/S0217751X15300227), [arXiv:1412.6352](https://arxiv.org/abs/1412.6352).

- [128] NA61 Collaboration, N. Abgrall et al., “NA61/SHINE facility at the CERN SPS: beams and detector system”, *JINST* **9** (2014) P06005, [doi:10.1088/1748-0221/9/06/P06005](https://doi.org/10.1088/1748-0221/9/06/P06005), [arXiv:1401.4699](https://arxiv.org/abs/1401.4699).
- [129] C. M. LHCb Collaboration, “LHCb SMOG Upgrade”, technical report, CERN, Geneva, May, 2019. [doi:10.17181/CERN.SAQC.EOWH](https://doi.org/10.17181/CERN.SAQC.EOWH).
- [130] L. Radel and C. Wiebusch, “Calculation of the Cherenkov light yield from low energetic secondary particles accompanying high-energy muons in ice and water with Geant 4 simulations”, *Astropart. Phys.* **38** (2012) 53–67, [doi:10.1016/j.astropartphys.2012.09.008](https://doi.org/10.1016/j.astropartphys.2012.09.008), [arXiv:1206.5530](https://arxiv.org/abs/1206.5530).
- [131] J. Jelley and U. K. A. E. Authority, “Čerenkov Radiation, and Its Applications”. United Kingdom Atomic Energy Authority, 1958. ISBN 9780080131276.
- [132] I. M. Frank and I. E. Tamm, “Coherent visible radiation of fast electrons passing through matter”, *Compt. Rend. Acad. Sci. URSS* **14** (1937), no. 3, 109–114, [doi:10.3367/UFNr.0093.196710o.0388](https://doi.org/10.3367/UFNr.0093.196710o.0388).
- [133] D. E. Groom and S. R. Klein, “Passage of particles through matter: in Review of Particle Physics (RPP 2000)”, *Eur. Phys. J. C* **15** (2000) 163–173, [doi:10.1007/BF02683419](https://doi.org/10.1007/BF02683419).
- [134] P. B. Price and K. Woschnagg, “Role of group and phase velocity in high-energy neutrino observatories”, *Astropart. Phys.* **15** (2001) 97–100, [doi:10.1016/S0927-6505\(00\)00142-0](https://doi.org/10.1016/S0927-6505(00)00142-0), [arXiv:hep-ex/0008001](https://arxiv.org/abs/hep-ex/0008001).
- [135] J. H. Koehne et al., “PROPOSAL: A tool for propagation of charged leptons”, *Comput. Phys. Commun.* **184** (2013) 2070–2090, [doi:10.1016/j.cpc.2013.04.001](https://doi.org/10.1016/j.cpc.2013.04.001).
- [136] J. Pretz, “Detection of atmospheric muon neutrinos with the IceCube 9-string detector”. PhD thesis, University of Maryland, College Park, United States, January, 2006.
- [137] IceCube Collaboration, M. G. Aartsen et al., “The IceCube Neutrino Observatory: Instrumentation and Online Systems”, *JINST* **12** (2017), no. 03, 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).
- [138] IceCube Collaboration, R. Abbasi et al., “IceTop: The surface component of IceCube”, *Nucl. Instrum. Meth. A* **700** (2013) 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).
- [139] IceCube Collaboration, R. Abbasi et al., “Calibration and Characterization of the IceCube Photomultiplier Tube”, *Nucl. Instrum. Meth. A* **618** (2010) 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).
- [140] IceCube Collaboration, M. G. Aartsen et al., “Measurement of South Pole ice transparency with the IceCube LED calibration system”, *Nucl. Instrum. Meth. A* **711** (2013) 73–89, [doi:10.1016/j.nima.2013.01.054](https://doi.org/10.1016/j.nima.2013.01.054), [arXiv:1301.5361](https://arxiv.org/abs/1301.5361).

- [141] IceCube Collaboration, M. Rongen and D. Chirkin, “Advances in IceCube ice modelling & what to expect from the Upgrade”, *JINST* **16** (2021), no. 09, C09014, doi:10.1088/1748-0221/16/09/C09014, arXiv:2108.03291.
- [142] M. Ackermann et al., “Optical properties of deep glacial ice at the South Pole”, *J. Geophys. Res.* **111** (2006), no. D13, D13203, doi:10.1029/2005JD006687.
- [143] ICECUBE Collaboration, D. Chirkin, “Evidence of optical anisotropy of the South Pole ice”, in *33rd International Cosmic Ray Conference*, p. 0580. 2013.
- [144] IceCube Collaboration, S. Tilav et al., “Atmospheric Variations as observed by IceCube”, in *Proc. 31st Int. Cosmic Ray Conference, Lodz, Poland*. 2010. arXiv:1001.0776.
- [145] IceCube Collaboration, M. G. Aartsen et al., “Energy Reconstruction Methods in the IceCube Neutrino Telescope”, *JINST* **9** (2014) P03009, doi:10.1088/1748-0221/9/03/P03009, arXiv:1311.4767.
- [146] IceCube Collaboration, M. G. Aartsen et al., “Search for non-relativistic Magnetic Monopoles with IceCube”, *Eur. Phys. J. C* **74** (2014), no. 7, 2938, doi:10.1140/epjc/s10052-014-2938-8, arXiv:1402.3460. [Erratum: *Eur.Phys.J.C* 79, 124 (2019)].
- [147] GEANT4 Collaboration, S. Agostinelli et al., “GEANT4—a simulation toolkit”, *Nucl. Instrum. Meth. A* **506** (2003) 250–303, doi:10.1016/S0168-9002(03)01368-8.
- [148] J. Allison et al., “Geant4 developments and applications”, *IEEE Trans. Nucl. Sci.* **53** (2006) 270, doi:10.1109/TNS.2006.869826.
- [149] F. Kislak, T. Melzig, A. V. Overloop, and T. Waldenmaier, “IceTop tank simulation using g4-tankresponse”. IceCube Internal Report, 2012. [https://internal-apps.icecube.wisc.edu/reports/data/icecube/2012/07/001/icecube\\_201207001\\_v1.pdf](https://internal-apps.icecube.wisc.edu/reports/data/icecube/2012/07/001/icecube_201207001_v1.pdf).
- [150] D. Chirkin, “dCORSIKA - CORSIKA for AMANDA”. <https://user-web.icecube.wisc.edu/~dima/work/BKP/DCS/REPORT8/paper.pdf>.
- [151] IceCube Collaboration, R. Abbasi et al., “Search for Point Sources of High Energy Neutrinos with Final Data from AMANDA-II”, *Phys. Rev. D* **79** (2009) 062001, doi:10.1103/PhysRevD.79.062001, arXiv:0809.1646.
- [152] D. Chirkin et al., “Photon propagation using gpus by the icecube neutrino observatory”, in *2019 15th International Conference on eScience (eScience)*, pp. 388–393. 2019. doi:10.1109/eScience.2019.00050.
- [153] C. Kopper, “clsim”. <https://github.com/claudiok/clsim>.
- [154] J. Lundberg et al., “Light tracking through ice and water—Scattering and absorption in heterogeneous media with PHOTONICS”, *Nuclear Instruments and Methods in Physics Research A* **581** (2007) 619–631, doi:10.1016/j.nima.2007.07.143, arXiv:astro-ph/0702108.

- [155] M. J. Larson, “Simulation and identification of non-Poissonian noise triggers in the IceCube neutrino detector”, Master’s thesis, Alabama U., United States, 2013.
- [156] IceCube Collaboration, X. Bai, E. Dvorak, J. Gonzalez, and D. Soldin, “A Three-dimensional Reconstruction of Cosmic Ray Events in IceCube”, *PoS ICRC2019* (2020) 244, [doi:10.22323/1.358.0244](https://doi.org/10.22323/1.358.0244), [arXiv:1908.07582](https://arxiv.org/abs/1908.07582).
- [157] S. De Ridder, “Sensitivity of IceCube cosmic ray measurements to the hadronic interaction models”. PhD thesis, Ghent University, Belgium, 2019.
- [158] T. Feusels, “Measurement of cosmic ray composition and energy spectrum between 1PeV and 1EeV with IceTop and IceCube”. PhD thesis, Ghent University, Belgium, 2013.
- [159] S. Klepser et al., “Lateral Distribution of Air Shower Signals and Initial Energy Spectrum above 1 PeV from IceTop”, in *International Cosmic Ray Conference*, volume 4 of *International Cosmic Ray Conference*, pp. 35–38. January, 2008.
- [160] S. Klepser, “Reconstruction of extensive air showers and measurement of the cosmic ray energy spectrum in the range of 1-PeV - 80-PeV at the South Pole”. PhD thesis, Humboldt-Universität zu Berlin, Germany, 6, 2008.
- [161] IceCube Collaboration, M. G. Aartsen et al., “Measurement of the cosmic ray energy spectrum with IceTop-73”, *Phys. Rev. D* **88** (2013), no. 4, 042004, [doi:10.1103/PhysRevD.88.042004](https://doi.org/10.1103/PhysRevD.88.042004), [arXiv:1307.3795](https://arxiv.org/abs/1307.3795).
- [162] K. Rawlins. Private Communication, 2022.
- [163] G. F. Knoll, “Radiation Detection and Measurement”. John Wiley and Sons, New York, 3rd edition, 2000. ISBN 978-0-471-07338-3, 978-0-471-07338-3.
- [164] N. Whitehorn, J. van Santen, and S. Lafebre, “Penalized Splines for Smooth Representation of High-dimensional Monte Carlo Datasets”, *Comput. Phys. Commun.* **184** (2013) 2214–2220, [doi:10.1016/j.cpc.2013.04.008](https://doi.org/10.1016/j.cpc.2013.04.008), [arXiv:1301.2184](https://arxiv.org/abs/1301.2184).
- [165] K. Erpels, “Influence of Snow on Cosmic Rays observed by IceTop”, Master’s thesis, Ghent University, 2015.
- [166] IceCube Collaboration, M. G. Aartsen et al., “Cosmic ray spectrum from 250 TeV to 10 PeV using IceTop”, *Phys. Rev. D* **102** (2020) 122001, [doi:10.1103/PhysRevD.102.122001](https://doi.org/10.1103/PhysRevD.102.122001), [arXiv:2006.05215](https://arxiv.org/abs/2006.05215).
- [167] IceCube Collaboration, S. De Ridder, E. Dvorak, and T. K. Gaisser, “Sensitivity of IceCube Cosmic-Ray measurements to the hadronic interaction models”, *PoS ICRC2017* (2018) 319, [doi:10.22323/1.301.0319](https://doi.org/10.22323/1.301.0319).
- [168] IceCube Collaboration, R. Abbasi et al., “Density of GeV muons in air showers measured with IceTop”, *Phys. Rev. D* **106** (2022), no. 3, 032010, [doi:10.1103/PhysRevD.106.032010](https://doi.org/10.1103/PhysRevD.106.032010), [arXiv:2201.12635](https://arxiv.org/abs/2201.12635).

- [169] P. H. Barrett et al., “Interpretation of Cosmic-Ray Measurements Far Underground”, *Rev. Mod. Phys.* **24** (1952), no. 3, 133–178, [doi:10.1103/RevModPhys.24.133](https://doi.org/10.1103/RevModPhys.24.133).
- [170] IceCube Collaboration, S. Tilav, T. K. Gaisser, D. Soldin, and P. Desiati, “Seasonal variation of atmospheric muons in IceCube”, *PoS ICRC2019* (2020) 894, [doi:10.22323/1.358.0894](https://doi.org/10.22323/1.358.0894), [arXiv:1909.01406](https://arxiv.org/abs/1909.01406).
- [171] T. Gaisser and S. Verpoest, “Seasonal variation of atmospheric muons”, *PoS ICRC2021* (2021) 1202, [doi:10.22323/1.395.1202](https://doi.org/10.22323/1.395.1202).
- [172] IceCube Collaboration, P. Desiati et al., “Seasonal Variations of High Energy Cosmic Ray Muons Observed by the IceCube Observatory as a Probe of Kaon/Pion Ratio”, in *32nd International Cosmic Ray Conference*, volume 1, pp. 78–81. 2011. [doi:10.7529/ICRC2011/V01/0662](https://doi.org/10.7529/ICRC2011/V01/0662).
- [173] IceCube Collaboration, P. Heix, S. Tilav, C. Wiebusch, and M. Zöcklein, “Seasonal Variation of Atmospheric Neutrinos in IceCube”, *PoS ICRC2019* (2020) 465, [doi:10.22323/1.358.0465](https://doi.org/10.22323/1.358.0465), [arXiv:1909.02036](https://arxiv.org/abs/1909.02036).
- [174] ICECUBE Collaboration, S. De Ridder and T. Feusels, “Seasonal variation of the muon multiplicity in cosmic rays at South Pole”, in *33rd International Cosmic Ray Conference*, p. 0763. 2013.
- [175] IceCube Collaboration, T. Fuchs, “Development of a Machine Learning Based Analysis Chain for the Measurement of Atmospheric Muon Spectra with IceCube”, in *25th European Cosmic Ray Symposium*. 1, 2017. [arXiv:1701.04067](https://arxiv.org/abs/1701.04067).
- [176] IceCube Collaboration, D. Soldin, “Atmospheric Muons Measured with IceCube”, *EPJ Web Conf.* **208** (2019) 08007, [doi:10.1051/epjconf/201920808007](https://doi.org/10.1051/epjconf/201920808007), [arXiv:1811.03651](https://arxiv.org/abs/1811.03651).
- [177] IceCube Collaboration, R. Abbasi et al., “Lateral Distribution of Muons in IceCube Cosmic Ray Events”, *Phys. Rev. D* **87** (2013), no. 1, 012005, [doi:10.1103/PhysRevD.87.012005](https://doi.org/10.1103/PhysRevD.87.012005), [arXiv:1208.2979](https://arxiv.org/abs/1208.2979).
- [178] D. Soldin, “Laterally Separated Muons from Cosmic Ray Air Showers Measured with the IceCube Neutrino Observatory”. PhD thesis, Wuppertal U., 2017.
- [179] P. Billoir, “A sampling procedure to regenerate particles in a ground detector from a “thinned” air shower simulation output”, *Astroparticle Physics* **30** (2008), no. 5, 270–285, [doi:https://doi.org/10.1016/j.astropartphys.2008.10.002](https://doi.org/10.1016/j.astropartphys.2008.10.002).
- [180] J. Madsen and A. Tamburro, “Simulating UHE cosmic-ray showers in IceTop”. IceCube Internal Report, 2009. [https://internal-apps.icecube.wisc.edu/reports/data/icecube/2009/11/001/icecube\\_200911001\\_v2.pdf](https://internal-apps.icecube.wisc.edu/reports/data/icecube/2009/11/001/icecube_200911001_v2.pdf).
- [181] J. Van Houtte, “Air shower measurements with IceTop: the effect of (un)thinning simulation techniques”, Master’s thesis, Ghent University, 2016.

- [182] T. K. Gaisser and S. Verpoest, “Profiles of energetic muons in the atmosphere”, *Astropart. Phys.* **133** (2021) 102630, [doi:10.1016/j.astropartphys.2021.102630](https://doi.org/10.1016/j.astropartphys.2021.102630), [arXiv:2106.12247](https://arxiv.org/abs/2106.12247).
- [183] S. Verpoest and T. K. Gaisser, “Parameterization of muon production profiles in the atmosphere”, *PoS ICRC2021* (2021) 434, [doi:10.22323/1.395.0434](https://doi.org/10.22323/1.395.0434), [arXiv:2107.11068](https://arxiv.org/abs/2107.11068).
- [184] T. Gaisser and S. Verpoest, “Seasonal variation of atmospheric muons”, *PoS ICRC2021* (2021) 1202, [doi:10.22323/1.395.1202](https://doi.org/10.22323/1.395.1202), [arXiv:2107.12913](https://arxiv.org/abs/2107.12913).
- [185] A. Fedynitch et al., “Calculation of conventional and prompt lepton fluxes at very high energy”, *EPJ Web Conf.* **99** (2015) 08001, [doi:10.1051/epjconf/20159908001](https://doi.org/10.1051/epjconf/20159908001), [arXiv:1503.00544](https://arxiv.org/abs/1503.00544).
- [186] C. Song, “Longitudinal profile of extensive air showers”, *Astropart. Phys.* **22** (2004) 151–158, [doi:10.1016/j.astropartphys.2004.05.011](https://doi.org/10.1016/j.astropartphys.2004.05.011).
- [187] NOvA Collaboration, M. A. Acero et al., “Observation of seasonal variation of atmospheric multiple-muon events in the NOvA Near Detector”, *Phys. Rev. D* **99** (2019), no. 12, 122004, [doi:10.1103/PhysRevD.99.122004](https://doi.org/10.1103/PhysRevD.99.122004), [arXiv:1904.12975](https://arxiv.org/abs/1904.12975).
- [188] B. Alper et al., “Large angle inclusive production of protons, anti-protons and kaons, and particle composition at the cern isr”, *Phys. Lett. B* **47** (1973) 275–280, [doi:10.1016/0370-2693\(73\)90730-2](https://doi.org/10.1016/0370-2693(73)90730-2).
- [189] P. Lipari and T. Stanev, “Propagation of multi - TeV muons”, *Phys. Rev. D* **44** (1991) 3543–3554, [doi:10.1103/PhysRevD.44.3543](https://doi.org/10.1103/PhysRevD.44.3543).
- [190] Pierre Auger Collaboration, P. Abreu et al., “The effect of the geomagnetic field on cosmic ray energy estimates and large scale anisotropy searches on data from the Pierre Auger Observatory”, *JCAP* **11** (2011) 022, [doi:10.1088/1475-7516/2011/11/022](https://doi.org/10.1088/1475-7516/2011/11/022), [arXiv:1111.7122](https://arxiv.org/abs/1111.7122).
- [191] C. Forti et al., “Simulation of atmospheric cascades and deep underground muons”, *Phys. Rev. D* **42** (1990) 3668–3689, [doi:10.1103/PhysRevD.42.3668](https://doi.org/10.1103/PhysRevD.42.3668).
- [192] IceCube Collaboration, S. Verpoest et al., “Testing Hadronic Interaction Models with Cosmic Ray Measurements at the IceCube Neutrino Observatory”, *PoS ICRC2021* (2021) 357, [doi:10.22323/1.395.0357](https://doi.org/10.22323/1.395.0357), [arXiv:2107.09387](https://arxiv.org/abs/2107.09387).
- [193] D. Soldin, “IC79 in-ice systematics”, 2014. <http://at-web.physik.uni-wuppertal.de/~soldin/systematics.html>.
- [194] IceCube Collaboration, A. Van Overloop, “Simulation of IceTop VEM calibration and the dependency on the snow layer”, in *32nd International Cosmic Ray Conference*, volume 1, pp. 97–101. 2011. [doi:10.7529/ICRC2011/V01/0899](https://doi.org/10.7529/ICRC2011/V01/0899).
- [195] IceCube Collaboration, R. Abbasi et al., “Density of GeV Muons Measured with IceTop”, *PoS ICRC2021* (2021) 342, [doi:10.22323/1.395.0342](https://doi.org/10.22323/1.395.0342), [arXiv:2107.09583](https://arxiv.org/abs/2107.09583).

- [196] IceCube Collaboration, M. Schauffel and K. Andeen, “IceAct, small Imaging Air Cherenkov Telescopes for IceCube”, *PoS ICRC2019* (2021) 179, [doi:10.22323/1.358.0179](https://doi.org/10.22323/1.358.0179), [arXiv:1908.11177](https://arxiv.org/abs/1908.11177).
- [197] H. P. Dembinski, “Computing mean logarithmic mass from muon counts in air shower experiments”, *Astropart. Phys.* **102** (2018) 89–94, [doi:10.1016/j.astropartphys.2018.05.008](https://doi.org/10.1016/j.astropartphys.2018.05.008), [arXiv:1711.05737](https://arxiv.org/abs/1711.05737).
- [198] F. Gesualdi et al., “On the muon scale of air showers and its application to the AGASA data”, *PoS ICRC2021* (2021) 473, [doi:10.22323/1.395.0473](https://doi.org/10.22323/1.395.0473), [arXiv:2108.04824](https://arxiv.org/abs/2108.04824).
- [199] F. Chollet et al., “Keras”. <https://keras.io>, 2015.
- [200] F. Chollet, “Deep Learning with Python”. Manning Publications Co., USA, 1st edition, 2017. ISBN 1617294438.
- [201] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization”, *arXiv e-prints* (2014) arXiv:1412.6980, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [202] S. Tilav. Private Communication, 2022.
- [203] N. Kulacz, “In situ measurement of the icecube dom efficiency factor using atmospheric minimum ionizing muons”, Master’s thesis, University of Alberta, 2019.
- [204] IceCube Collaboration, M. G. Aartsen et al., “Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube”, *JCAP* **10** (2019) 048, [doi:10.1088/1475-7516/2019/10/048](https://doi.org/10.1088/1475-7516/2019/10/048), [arXiv:1909.01530](https://arxiv.org/abs/1909.01530).
- [205] A. Fedynitch, J. Becker Tjus, and P. Desiati, “Influence of hadronic interaction models and the cosmic ray spectrum on the high energy atmospheric muon and neutrino flux”, *Phys. Rev. D* **86** (2012) 114024, [doi:10.1103/PhysRevD.86.114024](https://doi.org/10.1103/PhysRevD.86.114024).
- [206] A. Fedynitch and H. Dembinski, “crflux: Models/parametrizations of the cosmic ray flux at the top of the atmosphere”. <https://github.com/afedynitch/crflux>, 2015.
- [207] N. N. Kalmykov, S. S. Ostapchenko, and A. I. Pavlov, “Influence of the Landau-Pomeranchuk-Migdal effect on the features of extensive air showers”, *Phys. Atom. Nucl.* **58** (1995) 1728–1731.
- [208] V. B. Petkov, “The Andyrchy-BUST experiment: primary spectrum and composition around the knee”, *arXiv e-prints* (2009) arXiv:0911.5679, [arXiv:0911.5679](https://arxiv.org/abs/0911.5679).
- [209] S. Ostapchenko, “QGSJET-II: Towards reliable description of very high energy hadronic interactions”, *Nucl. Phys. B Proc. Suppl.* **151** (2006) 143–146, [doi:10.1016/j.nuclphysbps.2005.07.026](https://doi.org/10.1016/j.nuclphysbps.2005.07.026), [arXiv:hep-ph/0412332](https://arxiv.org/abs/hep-ph/0412332).

- 
- [210] T. D. Meulenaere, “Reconstructing the high-energy muon multiplicity in cosmic-ray induced air showers at icecube using neural networks”, Master’s thesis, Ghent University, 2022.
- [211] IceCube Collaboration, A. Ishihara, “The IceCube Upgrade - Design and Science Goals”, *PoS ICRC2019* (2021) 1031, [doi:10.22323/1.358.1031](https://doi.org/10.22323/1.358.1031), [arXiv:1908.09441](https://arxiv.org/abs/1908.09441).
- [212] IceCube-Gen2 Collaboration, M. G. Aartsen et al., “IceCube-Gen2: the window to the extreme Universe”, *J. Phys. G* **48** (2021), no. 6, 060501, [doi:10.1088/1361-6471/abbd48](https://doi.org/10.1088/1361-6471/abbd48), [arXiv:2008.04323](https://arxiv.org/abs/2008.04323).
- [213] IceCube-Gen2 Collaboration, F. G. Schroeder et al., “The Surface Array planned for IceCube-Gen2”, *PoS ICRC2021* (2021) 407, [doi:10.22323/1.395.0407](https://doi.org/10.22323/1.395.0407), [arXiv:2108.00364](https://arxiv.org/abs/2108.00364).