Cosmic Ray Observations at the TeV Scale with the HAWC Observatory

By

Zigfried Hampel-Arias

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Physics)

at the
UNIVERSITY OF WISCONSIN–MADISON
2017

Date of final oral examination: April 18th 2017

The dissertation was reviewed by the following members of the Final Oral Committee:

Paolo Desiati, Senior Scientist, Physics
Michael DuVernois, Senior Scientist, Physics
Francis Halzen, Professor, Physics
Kael Hanson, Professor, Physics
Stefan Westerhoff, Professor, Physics
Ellen Zweibel, Professor, Astronomy

COSMIC RAY OBSERVATIONS AT THE TEV SCALE WITH THE HAWC OBSERVATORY

Zigfried Hampel-Arias

Under the supervision of Professor Stefan Westerhoff

At the University of Wisconsin-Madison

Over the past two decades, a more detailed understanding of TeV-scale cosmic rays has emerged which appears to deviate from the isotropic, single power law description of the cosmic ray flux. This may be the result of the distribution of sources within the Galaxy, changes in source spectra, effects from the propagation of cosmic rays from their sources to Earth, or a combination of the three. Supernova remnants are thought to be the most likely source of Galactic cosmic rays, providing a natural power law source spectrum with sufficient power to generate the observed cosmic ray energy density. Yet, recent results from balloon-borne experiments hint at a possible change in the spectral index between $20 - 50$ TeV. These direct detection apparatuses provide the most precise measurements of the cosmic ray flux up to ~ 30 TeV, beyond which they are limited by the combined effects of their physical dimensions, runtime durations, and a rapidly decreasing flux. Above ~ 100 TeV, the spectrum has been measured by ground based air shower arrays, with typical systematic uncertainties of order 10%. Despite having the combined measurements from various experimental techniques, their different energy scales and systematics imply that identifying finer structure between $10 - 100$ TeV requires a single experimental method to span the entire range. Furthermore, as the nearest potential source is hundreds of parsecs away and the Larmor radius of TeV scale charged cosmic rays in the Galaxy is of order 10^{-3} parsecs, the previously observed anisotropy in arrival directions of cosmic rays is unexpected. In order to attain the statistical power necessary to observe TeV cosmic ray anisotropy at the 10^{-3} level and below, the long data taking periods required are only attainable by air shower arrays. This thesis presents a measurement of the cosmic ray energy spectrum and the energy dependence of the anisotropy on small scales $O(10^{\circ})$ using data from the High Altitude Water Cherenkov (HAWC) Observatory, an air-shower array located near Puebla, Mexico that is sensitive to gamma rays and cosmic rays at TeV energies. The analyses in this work comprise data taking periods of order 1 yr containing $\sim 10^{10}$ events. An analysis of the cosmic ray Moon shadow is first presented as a verification of the angular resolution and energy scale of the detector. Next, a measurement of the all-particle cosmic ray energy spectrum from $10 - 500$ TeV is shown, with an indication of structure deviating from a single power law. The final results presented in this work show an improved spectral measurement of a particular region of cosmic ray excess at the 10^{-4} level, previously observed both in HAWC and in other experiments.

ACKNOWLEDGMENTS

The work presented in this thesis could not have been possible without the support of a great number of mentors, friends, and loved ones. Their never-ending encouragement provided me the foundation for success. First and foremost I am thankful to my family for all their support since I can remember. To my mother who challenged me from an early age to do no less than my best. It is because of you that my world feels without limits. To my father who first taught me how to tread the Earth beneath my feet so that I could explore the stars in the sky. My curiosity for the natural world would have never blossomed without your guidance. And to my much cooler brother, Jenz, who is and always will be a continuing source of inspiration. Your uncanny ability to bring anyone you meet to laugh and smile is unmatched, and your age betrays the infinite wisdom and kindness in your heart. I am immensely proud of who you have become, am honored to call you my brother, and look forward to what the future holds for you.

I am indebted to all the teachers and mentors who challenged me in my studies and encouraged me in my struggles. As a result of their guidance, my education has taken me far and wide both intellectually as well as around the world. To Don Dry and Bill Inkret for trusting a teenager in a nuclear chemistry lab. My experience mixing solutions started me on what seems to be a random walk of increasingly interesting and difficult scientific problems up to my present state. To Paul Padley, my undergraduate advisor, thanks for taking a chance on a severely inexperienced physics student to work inside the CMS cavern. That summer opened my eyes to the multicultural aspect of international collaboration, and laid the foundations for an insatiable love of physics. I extend my deep and sincere gratitude to my thesis advisor, Stefan Westerhoff. The UW HAWC group you formed provided a singular environment in which to learn and develop, which was integral to my success. Your keen guidance kept me on track of what otherwise seemed like an endless wild goose chase of fascinating physics problems, and your patience led me to produce work of which I am proud to present.

To the members of the HAWC collaboration, it has been a pleasure working with so many wonderful and talented scientists. I thank Megan Longo Proper and Katheryne Sparks Woodle for that chance meeting in Bariloche that led to my involvement with HAWC. Oh, and that includes Dan of course, didn't mean to forget you. To Jim Braun and Segev Benzvi, your limitless patience and politeness helped me to develop my programming skills, making me a more effective and competent researcher. In addition to having thoroughly

reviewed this manuscript, I am very greatful to Tom Weisgarber for sharing his expertise with me and challenging me in my work. Your exceptional ability to identify cracks in my non-linear thought process was invaluable to my progress, and I know you were root, Root, ROOT-ing for me the whole time. To my dear colleagues Dan Fiorino and Ian Wisher, thank you for all the wonderful times had in Chamberlin and abroad. In addition to righting me when lost and for supporting me when most needed, you two reminded me that the best part of our working together was laughter. The halls of Chamberlin are less bright without your presense, and you will forever be demasiado cool.

To Aimee and Renee Lefkow, the entire UW Physics department staff, Nils Irland, Kim Kreiger, Ben Stock, and the members of WIPAC. A special thank you for your patience with my requests, however big or small.

I am overwhelmed by the many friendships fostered during my tenure in Madison, without whose presence I would have faltered long ago. To Antonia Hubbard, there are too many fond memories to enumerate here. Whether singing at home, dancing in the office, or laughing in the streets, our time together got me through the ups and downs of graduate school. To James Duff, the cajun-sconnie who's down for whatever, thanks for your ever-beaming smile. Who's awesome? Jakob deOscarMaeyerWienerschnitzel's awesome. To Matt M., Mateo, and Goat, you have unwittingly saved my life many times over, and I will always consult you regarding adventures in the Midwest. You can always count on me, forever your hermano in Sánchez. I am grateful for the camaraderie of the Hippo Water Polo team, who helped me find myself again in difficult times. A Carlos, de Buenos Aires a Madison y más allá, hemos tenido aventuras inolvidables. Gracias por los grandes recuerdos y tu fuerte amistad sin límites. Eres mi hermanito del alma. Margarita, gracias por el tiempo que pasamos juntos cocinando, cantando y bailando. Siempre serás hermana mía, y se que nos vamos a encontrar en varias partes distintas del mundo.

I also extend my deep fondness for the movie nighters, whose affection and antics extend well beyond the affinity for food, wine, and films. For those who were there week after week (or every other week), Anna Castelnovo, Armand, Esteban (Es-Steve), Jaakko, Lampros, Mélanie, Merve, Will, and sometimes Jordi, among others, I will forever be amazed at the expanse of connections that cross continents and oceans as a result of our time together. Thanks to Kostas for being the one true $\mu\alpha\lambda\alpha\kappa\alpha$ of the three burritos. I hope you can one day enlighten me as to the historical circumstances regarding the presence of two popes. And a special thank you to Larissa for thinking of me one fateful evening that changed my life.

Et enfin merci à Olivia, dont l'amour et le soutien ne connaît pas de limites, et qui m'inspire toujours d'être mieux. Ces dernières années magiques m'ont apporté une joie immense, et j'ai hâte de continuer notre voyage ensemble.

DISCARD THIS PAGE

TABLE OF CONTENTS

Page

DISCARD THIS PAGE

LIST OF TABLES

DISCARD THIS PAGE

LIST OF FIGURES

Figure Page

Figure Page

Figure Page

Chapter 1

Cosmic Rays and Extensive Air Showers

In this thesis I present several cosmic ray analyses using data from the High Altitude Water Cherenkov Observatory (HAWC). These include measurements of the all-particle cosmic ray energy spectrum, the energy dependence of the cosmic ray Moon shadow, and continued work on the cosmic ray anisotropy with improved energy estimation. The main motivation for these studies is related to understanding the sources and propagation of TeV-scale cosmic rays, the details of which will be further elaborated in the present chapter. Following this is a description of the HAWC detector in Chapter 2, and the event reconstruction techniques in Chapters 3 and 4.

The analysis is divided into two main parts: first an initial verification of the energy scale with the cosmic ray Moon shadow, then measurements of the cosmic ray energy spectrum and the energy dependence of the cosmic ray anisotropy. The map making techniques used for the Moon shadow and anisotropy are first presented in Chapter 5, followed by the experimental results of the Moon shadow's evolution with energy in Chapter 6. Next is a description of the statistical method used for the "unfolding" of the all-particle spectrum in Chapter 7, with experimental results in Chapter 8. Finally, Chapter 9 presents results of the energy dependence of the cosmic ray anisotropy. The current chapter provides brief historical background regarding cosmic ray science, and concludes with an overview of the current status of TeV-scale observations relevant to this thesis.

1.1 Historical Background

The study of charged particles of cosmic origin has its roots on the ground at the turn of the 20th century. With the discovery of radioactivity, it was observed that the electroscopes designed to measure ionizing radiation were discharging in the absence of known sources. At first, this was thought to be due to natural radiation in the environment, primarily gamma rays which were known to be produced by radon present in the air. However, in the year 1912 the Austrian scientist Victor Hess made seven hot air balloon flights [2] that demonstrated the source of the ionizing radiation was not of terrestrial origin, hence the term cosmic rays. Hess made measurements with electroscopes at altitudes surpassing 5 km, observing that the rate of ionization discharges increased with increasing altitude (figure 1.1).

Subsequent flights by Kolhörster [3] to higher altitudes confirmed this, but eventually unmanned balloon observations up to 16 km demonstrated a drop in the ionization rate, leading Robert Millikan and others to conclude that the radiation being measured was in fact secondary, being produced by primary radiation higher in the atmosphere. With the advent of the Geiger-Müller tube in 1928, the nature of these secondary ionizing *corpuscles* could be examined using the coincidence detector technique, revealing energies in excess of 10^9 eV [4].

It was also realized that the secondary radiation consisted of charged particles, as they were deflected in magnetic fields. Indeed, using similar coincident techniques, Blackett and Occhialini [5] developed a self-triggering mechanism to take photographs of particle deflections due to an applied magnetic field in cloud chambers (figure 1.2). This supported observations by Anderson [6] and experimentally confirmed the electron/positron theory of Dirac. The same method was used in the discovery of the muon in 1936 [7]. For their contributions in discovering cosmic radiation and the positron, respectively, Hess and Anderson shared the Nobel prize in 1936.

Rossi, followed by Pierre Auger and his colleagues, began separating coincidence detectors over increasingly larger areas [8], and measured the physical extent to which these particles are correlated. They discovered extensive air showers implying primary particle energies of at least 10^{15} eV [9]. Around the same time, balloon-borne nuclear emulsions were deployed higher into the atmosphere to measure the primary particles directly. It was discovered that the primary spectrum is composed mostly of protons [10], $\sim 10\%$ helium nuclei and a remaining 1% contribution from heavier nuclei [11].

Through their pioneering work, these experiments laid the foundations for the methods still used today to study primaries with energies from $\sim 10^9$ eV extending up to 10^{20} eV. Yet despite subsequent decades of further research, the precise origin of these particles remains a mystery. Understanding the mechanisms that accelerate particles to the highest observed energies and the environments they encounter traveling to Earth remain the most fundamental questions in cosmic ray physics a century later.

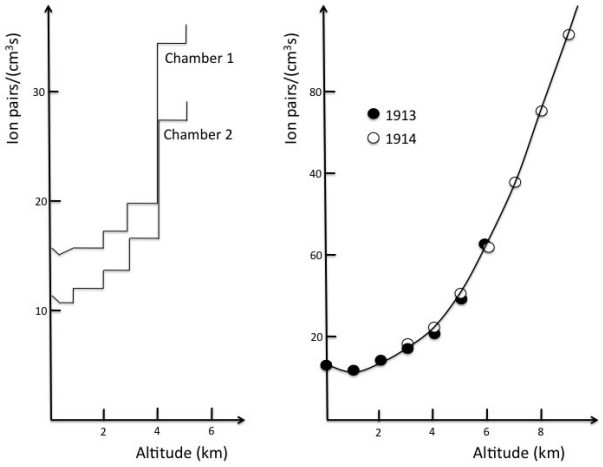

Figure 1.1: Measurements of the ionization rate with altitude from Hess [2] in 1912 and Kolhörster [3] from 1913–1914. The rise in rate with increasing altitude is attributed to the cosmic origin of the penetrating radiation.

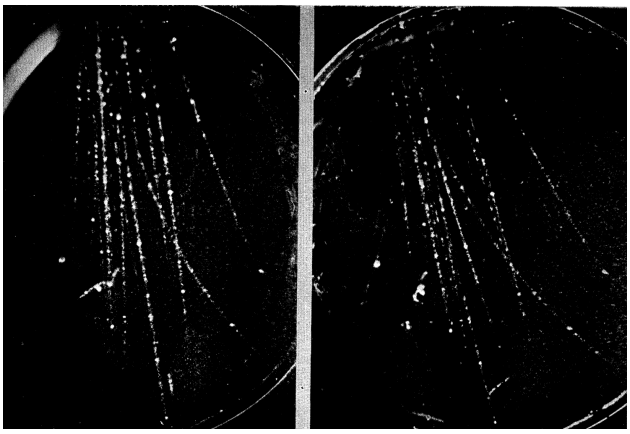

Figure 1.2: Photograph from Blackett and Occhialini's triggered cloud chamber device [5], showing creation of electron-positron pairs deflecting in opposite directions. Two particle counters were placed such that their simultaneous discharge initiated the taking of the photograph. The chamber was placed in a water cooled solenoid with a field strength of 0.3 T.

1.2 Extensive Air Showers

Primary cosmic ray particles with energies $\gtrsim$ 100 GeV incident at the top of the Earth's atmosphere produce extensive air showers, characterized by a laterally extended but thin disk of secondary particles (see figure 1.3). The nature of the primary particle determines the evolution of the shower with regard to the particle content and subsequently the shape and energy distribution as the shower develops through the atmosphere. This section introduces extensive air showers by first discussing purely electromagnetic cascades, and then addresses the more complicated processes involved in hadronic initiated particle showers.

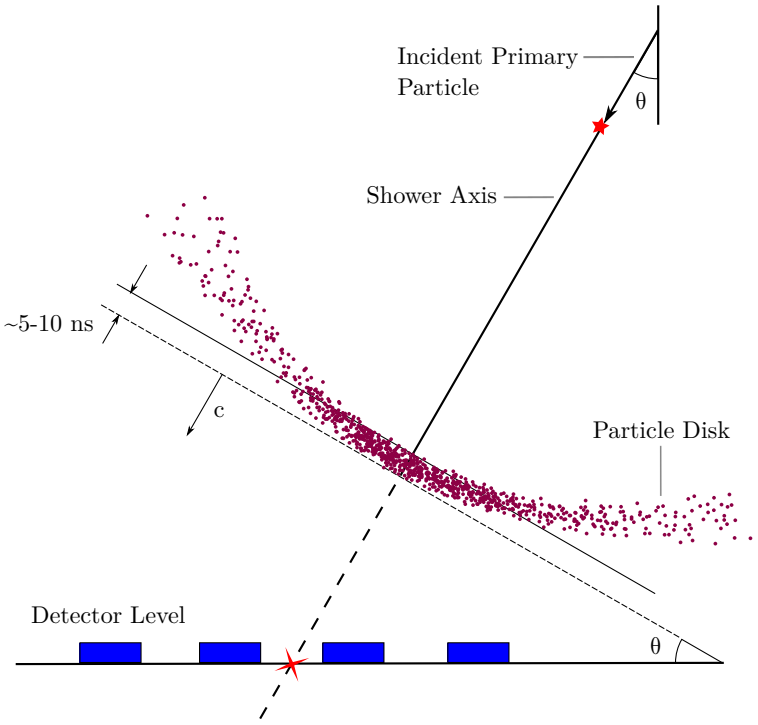

Figure 1.3: Diagram of an extensive air shower particle front. The particle disk moves forward at nearly the speed of light, and has a characteristic width and curvature. The shower axis direction is defined by the zenith angle θ, which projected onto the ground determines the core location, represented by the thin four-pointed star. Near the core the shower front is compact, smearing and thinning out laterally. Air shower reconstruction is done using the relative timing and integrated signals of detector elements on the ground. The five-pointed star represents the first interaction point of the primary particle.

1.2.1 Electromagnetic Cascades

Primary gamma rays generate cascades containing mostly electromagnetic particles, namely photons and electrons ($\gamma, e^{\pm}$). The most relevant processes for these particles are $e^{\pm}$-pair production by γs and bremsstrahlung by the charged $e^{\pm}$ in the vicinity of atmospheric nuclei, as illustrated in figure 1.4. In addition to pair production and bremsstrahlung, many other particle interactions can contribute to shower development, yet the simple Heitler model [12] illustrates the general features of electromagnetic cascades. By considering identical shower particles, the Heitler model assumes the generation of new secondaries occurs at fixed intervals defined by the interaction length, λ. The particle content is doubled after each λ, with the energy from the parent particles being distributed equally to the created particle pair. Given an incident primary with energy E_0, after n radiation lengths, the number of particles is given by $N = 2^n$, each with energy $E = E_0/2^n$. The shower continues to grow until particle energies reach a critical energy, E_c, which is a quantity dependent on the medium. At this point, energy is lost equally to both bremsstrahlung and through ionization. For electrons in air, $E_c \approx 84$ MeV. At this point the shower has reached its maximum, given by $\lambda \log_2 E_0/E_c$ interaction lengths, and the number of particles is $N_{\max} = E_0/E_c$. This simple model describes shower development only up to shower maximum, yet it qualitatively relates the primary's initial energy to the number of shower particles at a given stage of development.

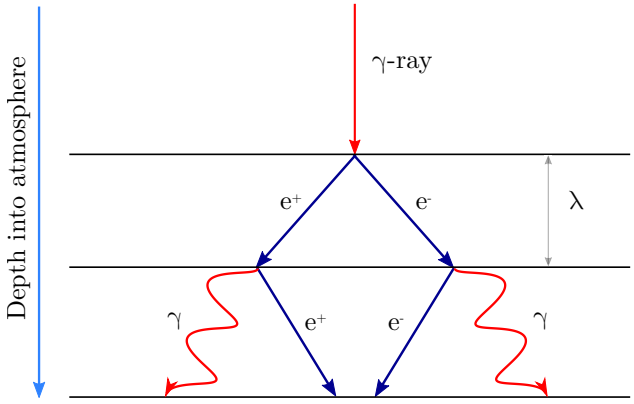

Figure 1.4: Diagram of a purely electromagnetic shower. The number of secondary particles increases exponentially as the shower develops deeper into the atmosphere. The primary gamma ray is represented by a straight ray to differentiate it from the secondary photons generated by bremsstrahlung. Each stage indicated by the horizontal lines represents an interaction length, λ, where the number of particles increases by a factor of two. This occurs by pair production ($\gamma \to e^+ + e^-$) and bremsstrahlung radiation ($e^{\pm} + A \to e^{\pm} + \gamma + A$). Eventually, the energy of the particles drops below a critical energy and the shower ceases to grow. This is the point of shower maximum.

The theory of the lateral spread of the electromagnetic particles was developed by Greisen [13] and Katamata & Nishimura [14]. At each interaction point, pair production and bremsstrahlung processes emit particles that are not aligned precisely along the primary's direction, and further Coulomb scattering of the $e^{\pm}$ results in a lateral smearing of the particle front. The approximation for the density of electrons at a distance r from the shower axis is given by the NKG function:

$$\rho_e(r) = \frac{N_e}{2\pi\ R_m^2}\ C(s)\ \left(\frac{r}{R_m}\right)^{s-2}\ \left(1+\frac{r}{R_m}\right)^{s-9/2}\ . \tag{1.1}$$

Here, N_e is the total number of electrons in the shower, and $C(s) = 0.366\ s^2\ (2.07-s)^{1.25}$ is a normalizing expression such that $\int_0^\infty 2\pi r\rho(r)dr\ =\ N_e$. The parameter s is referred to as the *age*, which defines development relative to the shower maximum. By definition, showers are initiated at $s = 0$, reach maximum at $s = 1$, and diminish to $O(1)$ particle at $s = 2$. The Molière length (R_m) incorporates the effect of Coulomb scattering. At sea level $R_m \simeq 79$ m and at the altitude of HAWC $R_m \simeq 120$ m. The radial form of the NKG function means that electromagnetic cascades are densely collimated, with particle density decreasing rapidly with increasing lateral distance from the shower axis.

1.2.2 Hadronic Cascades

Primaries of hadronic origin induce showers that develop in a manner qualitatively similar to electromagnetic cascades, but due to the nuclear interactions involved, they exhibit some important differences. As depicted in figure 1.5, hadronic showers develop from an indicent cosmic ray colliding with a nucleus A in the atmosphere, and are composed of an electromagnetic component ($\gamma, e^{\pm}$), as well as hadronic constituents such as charged and neutral pions ($\pi^{\pm}, \pi^0$). The charged pions interact with atmospheric nuclei, inducing additional meson production, until decaying to muons via

$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$
$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu\ .$$

To a much lesser extent (branching ratio 0.000123) the pions decay to $e^{\pm}$ and electron neutrinos. There are approximately twice as many charged pions as neutral pions which decay almost immediately to two photons, $\pi^0 \rightarrow 2\gamma$. These generate electromangetic sub-showers as described previously. Both the hadronic and electromagnetic sub-cascades continue to develop until the particles reach respective energy thresholds, after which particle production stops. Despite the presence of nucleons and mesons, their total number in the shower is significantly smaller than for $\mu^{\pm}$, $e^{\pm}$ and γ.

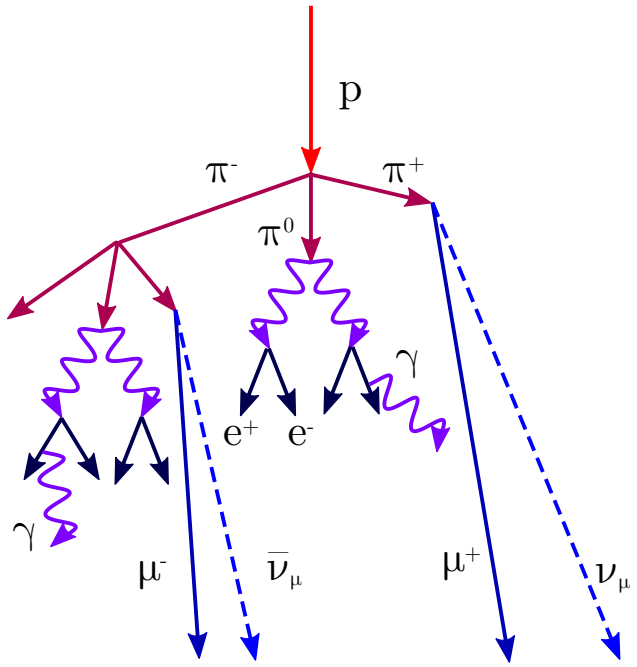

Figure 1.5: Diagram of a shower initiated by a proton, i.e. a hadronic shower. Charged pions ($\pi^{\pm}$) generated through nuclear interactions can carry energy laterally outward, resulting in decays to penetrating muons far from the central shower core. An electromagnetic component primarily arises from neutral pion decay ($\pi^0 \rightarrow 2\gamma$) and accounts for about 90% of the particle content of the shower.

Due to uncertainty in the atmospheric depth of the first interaction, as well as the inherent fluctuations during hadronic cascade development, the total numbers of these particles can vary significantly. Yet, the identity of the primary and its energy can still be estimated by measuring the respective muon and electron component populations. For example, Matthews [15] approaches hadronic shower development in a manner similar to Heitler, which leads to an estimate for N_μ provided the superposition approximation for heavier primary nuclei holds. This states that the shower from a primary nuclei of mass number A and energy E_0 develops as would A proton showers of energy E_0/A. The muon number from heavier nuclei is related to that for protons via

$$N_\mu^A = A^{1-\beta} {N_\mu}^p \ ,$$

where $\beta = 0.85$. While this approximates the muon content, the electromagnetic component is complicated by the fact that heavier showers are more likely to interact higher in the atmosphere. Thus, electron and photon numbers on the ground will be smaller than that for proton initiated showers.

The capability to definitively identify muons in data to measure N_e/N_μ was not possible for this analysis. But considering that a decreased electron number is equivalent to observing a more developed shower of the

same energy, comparing measured s and energy values can be used to discriminate heavier nuclei from the dominant proton and helium components of the cosmic ray spectrum. This composition measurement will be the subject of a future analysis.

1.2.3 Air Shower Detection

Air shower particles propagate through the atmosphere at nearly the speed of light, forming a nearly spherically curved wavefront that is roughly symmetric with respect to the primary particle's direction vector (figure 1.3). As will be discussed in chapter 3, the differences in arrival times of the particles and the density of deposited energy in a detector are used to determine the shower axis and the energy-rich shower *core*. Provided an estimated core location, the lateral distribution of secondaries and their energy content provide a means to estimate the total energy of the primary. The energy estimation method used in this work is presented in chapter 4.

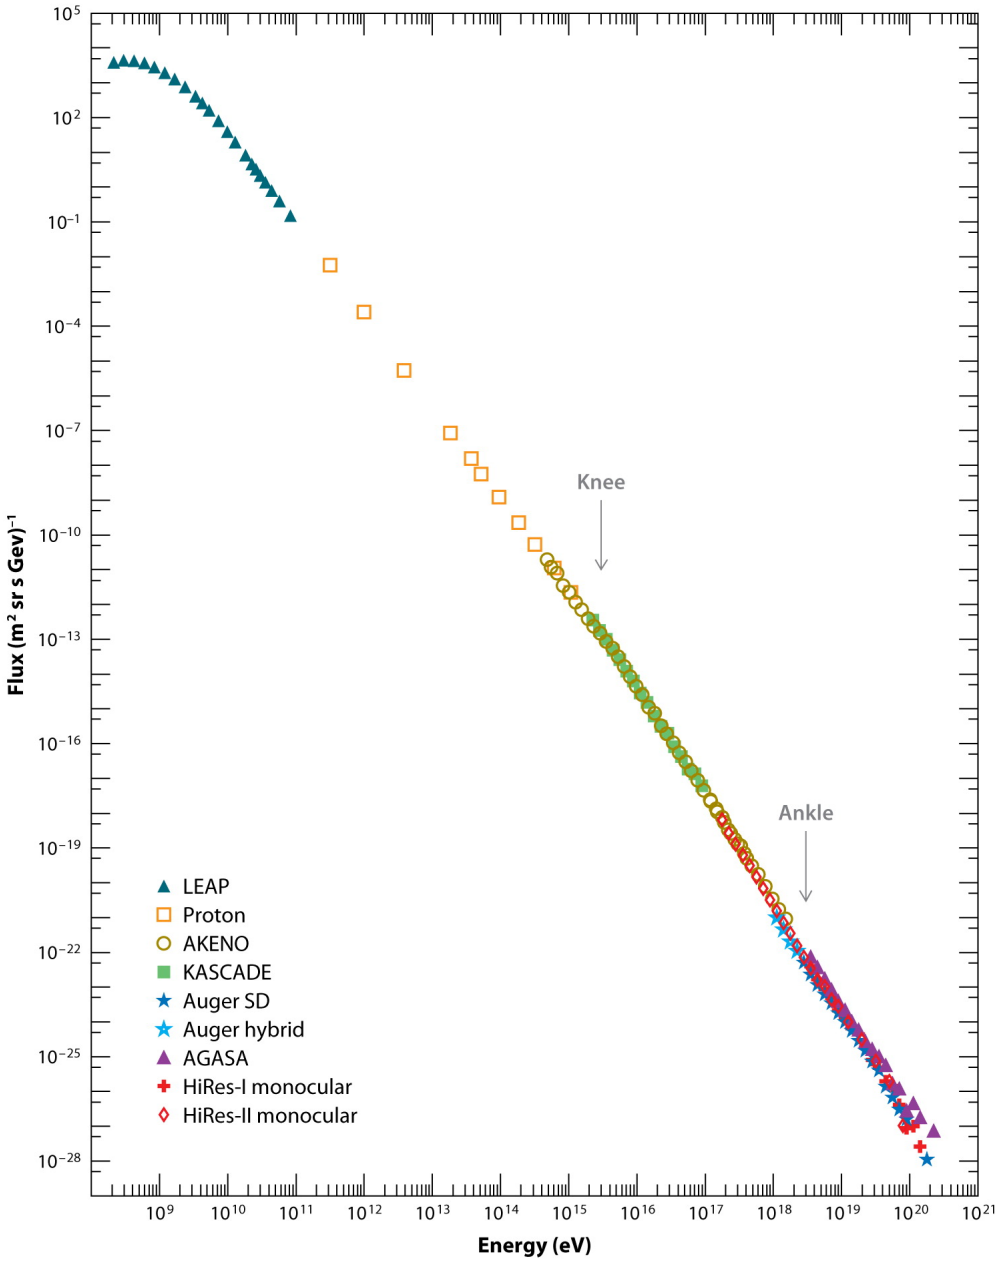

Figure 1.6: The all-particle cosmic ray differential energy spectrum as measured by a variety of experiments [16]. The flux follows a nearly constant power law with index $\gamma = 2.7$ with breaks at the knee and ankle, around 3×10^{15} eV and 4×10^{18} eV, respectively. At the lowest energies, below $\sim 10^{10}$ eV, solar activity influences the spectrum. Up to $O(100\,\text{TeV})$, direct measurements are provided by balloon-borne and satellite experiments (represented here by LEAP and Proton). Ground based detectors operate from TeV scales to the highest energies.

1.3 Energy Spectrum & Composition

The primary cosmic ray spectrum spans a remarkable range covering over ten decades in energy and thirty decades in flux [17, 18]. The spectrum is typically expressed in terms of flux, $\mathscr{F}$, defined by the number of particles of energy E arriving per unit area, solid angle, time and energy:

$$\mathscr{F}(E) = \frac{dN(E)}{dA\, d\Omega\, dt\, dE} \ . \tag{1.2}$$

The observed cosmic ray flux is well described by a nearly single inverse power law of index $\gamma \approx 2.7$, with two breaks at $\sim 3 \times 10^{15}$ eV and $\sim 4 \times 10^{18}$ eV. These are referred to as the *knee* and *ankle*, respectively. The structure of the observed all-particle energy flux is shown in figure 1.6. The non-thermal nature of the spectrum, including the breaks, carries information regarding the dynamics of the environments in which cosmic rays are accelerated and those that they traverse.

Since charged particles are deflected in magnetic fields, there is an immediate implication concerning cosmic ray sources beyond 10^{15} eV, or PeV energies. For a charged particle with charge Z and energy E traversing space in the presence of a magnetic field with strength B, the Larmor radius r_g in pc is given by

$$\begin{aligned} r_g &= 1.1 \times 10^{-3}\,\mathrm{pc} \left(\frac{E}{\mathrm{TeV}}\right) \frac{1}{|Z|} \left(\frac{B}{\mu\mathrm{G}}\right)^{-1} \\ &= 1.1 \times 10^{-3}\,\mathrm{pc} \left(\frac{R}{\mathrm{TV}}\right) \left(\frac{B}{\mu\mathrm{G}}\right)^{-1} \ , \end{aligned} \tag{1.3}$$

where $R = E/Z$ is the rigitidy measured in volts and represents a charged particle's resistance to the influence of magnetic fields. Within the inner Milky Way disk (~ 300 pc thick), the Galaxy has an ordered magnetic field of $\sim 1\,\mu$G along with a turbulent component of similar strength. Particles at PeV energies near the knee have a Larmor radius of $r_g(E = \mathrm{PeV}) \simeq 0.3\,\mathrm{pc}$, and thus are effectively confined to randomized, diffusive motion within the galaxy. At higher energies, $r_g(E = 10^3\,\mathrm{PeV}) \simeq 300\,\mathrm{pc}$, these particles can begin to escape the galaxy, which indicates a transition of sources from Galactic to extragalactic origin.

Of course, r_g is also inversely proportional to the particle charge so features in the spectrum should vary for different elements. For the knee, this suggests that Galactic cosmic rays are expected to become heavier (higher Z) in this energy region. Figure 1.7 shows direct flux measurements for proton up to iron nuclei that appear to maintain nearly constant proportions just prior to the knee. At the lowest energies, however, there appear to be shifts in the locations of the solar bumps. In fact, the spectra would indeed exhibit constant ratios as a function of *rigidity*, which suggests that cosmic ray sources are effectively composition blind. This is called the universality of acceleration spectra.

Yet recent results from the PAMELA [20] satellite experiment demonstrate a decrease in the proton to helium flux ratio up to TV rigidities or $\sim$ TeV energies (figure 1.8). Similarly, the CREAM balloon-borne detector [21, 22, 23] reports further hardening of the helium energy spectrum, surpassing the proton flux at

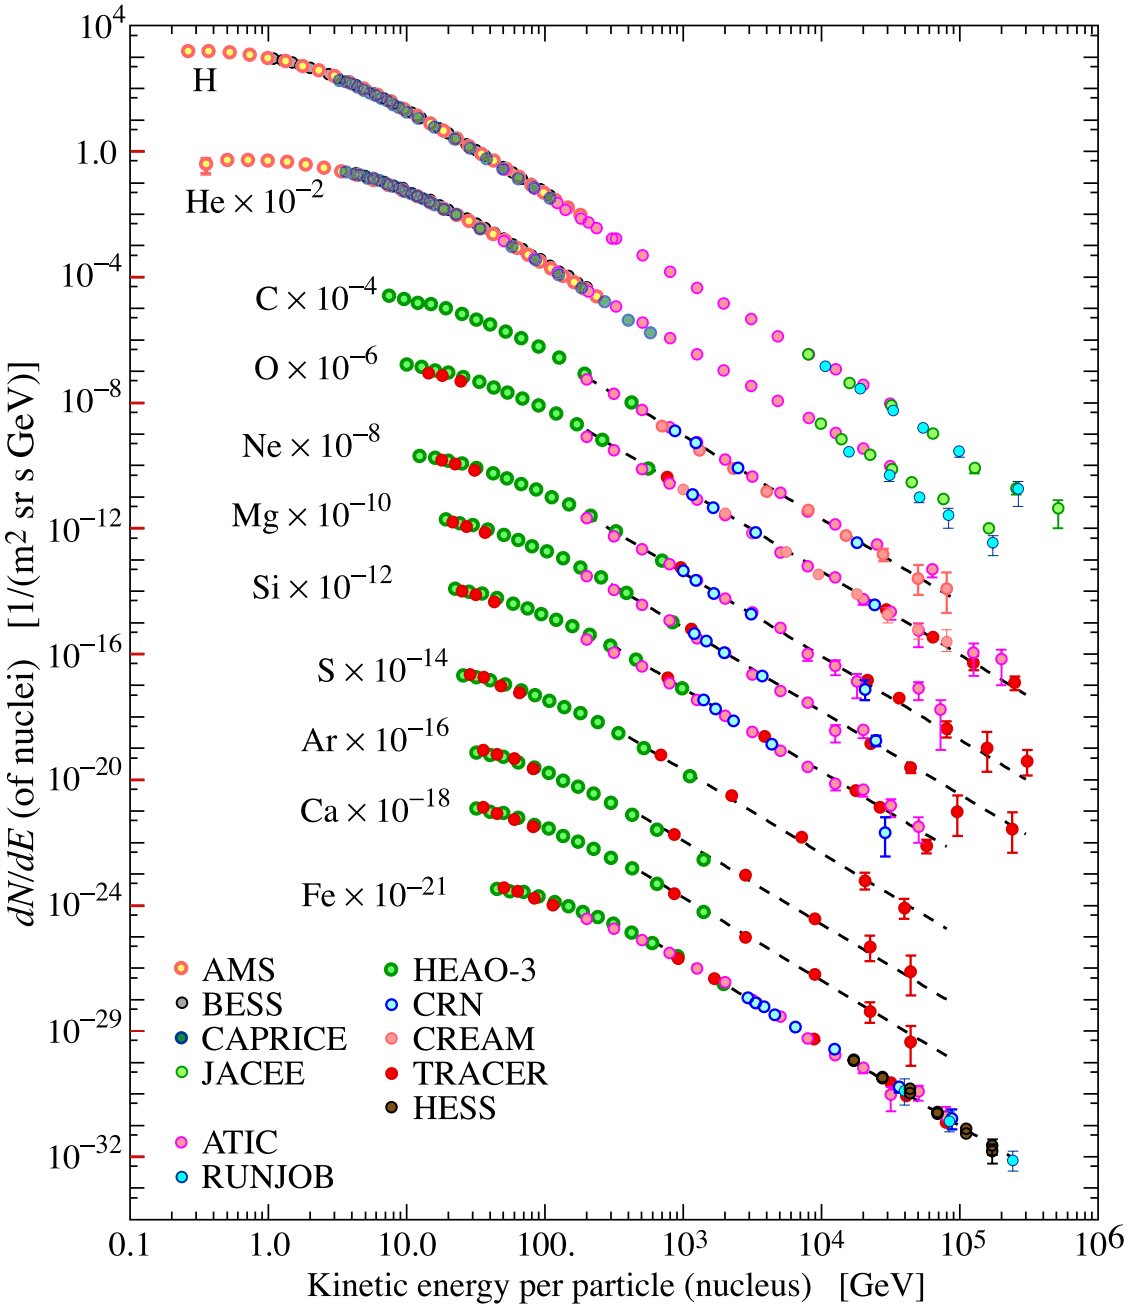

Figure 1.7: The differential energy spectra for various species as a function of kinetic energy [19]. Each species is scaled by an indicated factor for clarity. The data are primarily from direct detection experiments, providing precision measurements up to TeV energies, where detector exposure rapidly diminishes.

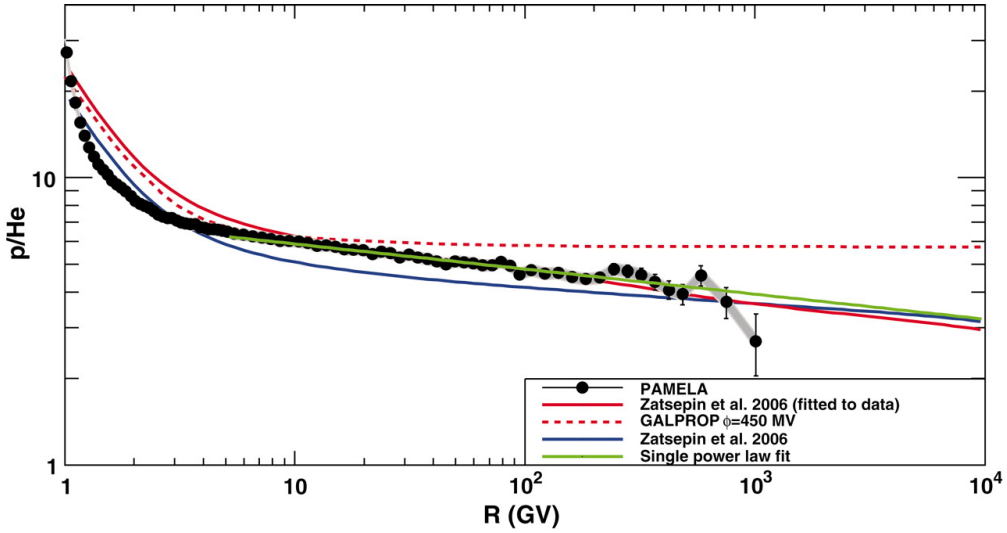

Figure 1.8: The proton to helium flux ratio up to 1 TV with data taken from PAMELA [20]. The gray shaded band represents the systematic uncertainty, while the curves denote various model fits.

approximately 10 TeV (see figure 1.9). The flattening of the helium flux relative to protons also has been reported by the ATIC experiment [24]. This type of structure could be an indication of different source populations [25], or a nearby source that is proton rich up to TeV energies [26]. It has also been suggested that the spectral hardening could be attributed to anomalous diffusion [27].

Probing the cosmic ray spectrum via direct detection becomes a challenge in the 10–100 TeV range and beyond due to limited detector exposures fueled by a rapidly diminishing flux. Ground-based air shower arrays are not as limited by their collection area but are not as sensitive to the identity of the primary particle. Yet, as demonstrated by the recent measurements of the proton and helium spectrum by ARGO-YBJ [28, 29] and composition studies by GRAPES-III [30], ground arrays carry significant potential for extending measurements past the TeV scale. Results from CREAM and ARGO-YBJ are shown in figure 1.9 for comparison.

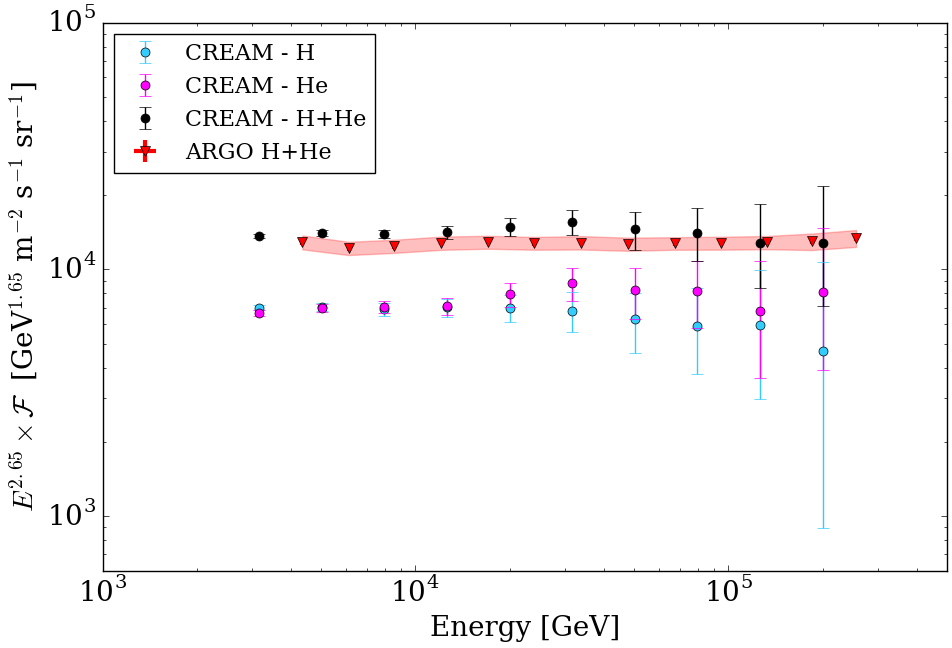

Figure 1.9: Proton and helium flux data from CREAM direct measurements [23], and the combined H+He spectrum from the ARGO-YBJ air shower array [29]. The red band around the ARGO data represents the systematic error. There is an indication of the helium spectrum hardening, or becoming less steep with energy, surpassing the proton flux around 10 TeV.

1.4 Production

It is generally believed that supernova remnants (SRNs) are responsible for the acceleration of cosmic rays at least to 10^{14} eV and possibly up to about 10^{18} eV [31]. Supernova explosions provide a mechanism for particle acceleration that naturally produces power law spectra. As the shock front encounters the surrounding environment, particle energies are boosted by repeatedly crossing the front. This process was initially proposed by Fermi [32] in the context of charged particle interactions with turbulent interstellar magnetic clouds, as shown in figure 1.10. In this scenario (adapted from [33]), relativistic particles ($E_0 \simeq p_0 c$) scatter off of many clouds, accelerating with each interaction in a manner proportional to its energy. In the reference frame of a cloud moving with speed $\beta_{\rm cl} = v_{\rm cl}/c$, the energy of the particle is transformed to

$$E' = \gamma_{\rm cl}(E_0 + \beta_{\rm cl} p_0) \ .$$

Assuming the particle interacts in an elastic manner, its velocity is reversed in the cloud frame, i.e. a simple reflection. Transforming back, the particle's new energy is

$$E_1 = \gamma_{\rm cl}(E' + \beta_{\rm cl} p') = E_0 \gamma_{\rm cl}^2 (1 + \beta_{\rm cl})^2 \ ,$$

representing a relative gain in energy $\Delta E/E_0$ given by

$$\frac{\Delta E}{E_0} = \frac{E_1 - E_0}{E_0} = \gamma_{\rm cl}^2 (1 + \beta_{\rm cl})^2 - 1 = \xi \ . \tag{1.4}$$

After n interactions with similar clouds, the particle energy is $E_n = E_0(1+\xi)^n$. Thus the number of interactions to reach an energy E_n is

$$n = \frac{\log\left(E_n/E_0\right)}{\log\left(1+\xi\right)} \ . \tag{1.5}$$

If the particle has a probability $p_{\rm esc}$ of escaping the cloud environment, then the probability of remaining after n interactions is $(1 - p_{\rm esc})^n$. For an ensemble of particles, the number that survive exceeding E_n is given by

$$N(\geq E_n) \propto \sum_n^\infty (1 - p_{\rm esc})^n = (1 - p_{\rm esc})^n \sum_{m=0}^\infty (1 - p_{\rm esc})^m = \frac{(1 - p_{\rm esc})^n}{p_{\rm esc}} \ . \tag{1.6}$$

Using the expression for n in equation 1.5 and the energy gain in equation 1.4 leads to an integral power law spectrum:

$$N(\geq E_n) \propto \left(\frac{E_n}{E_0}\right)^{-\alpha} \ , \tag{1.7}$$

where $\alpha \simeq p_{\rm esc}/\xi$ is the power law index. Since $\xi \propto \beta_{\rm cl}^2$, this process is referred to as second-order Fermi acceleration.

This magnetic mirror scenario has some strong underlying assumptions. For example, the cloud's motion is assumed not to be influenced by the interaction, i.e. a pure particle reflection. In addition, the velocities

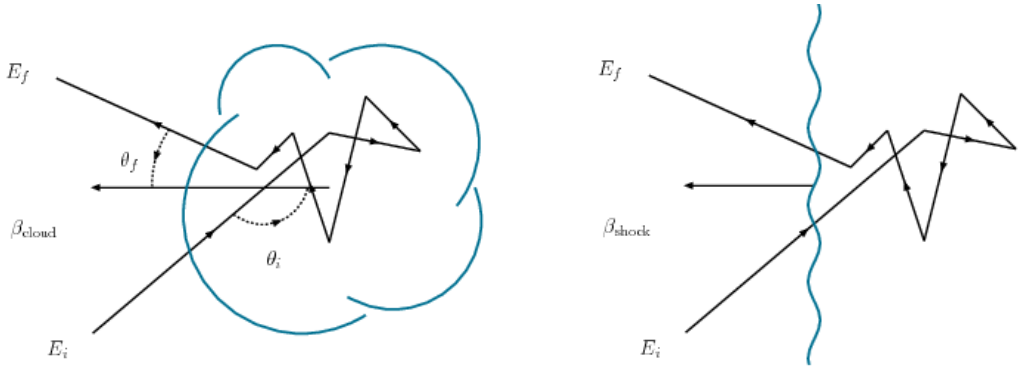

Figure 1.10: Diagrams of Fermi acceleration. The left panel depicts the magnetic cloud scenario, or 2nd order Fermi acceleration. The right panel shows shockfront acceleration.

of the cloud and particle are parallel (and anti-parallel), as other configurations influence the value of ξ. Yet, even with this simple argument, Fermi illustrated an acceleration process that manifests a natural power law.

In supernovae, magnetic inhomogeneties on both sides of the shock front conceptually take the place of the magnetic clouds, so the particle similarly gains energy upon each crossing of the shock boundary. The power law index, however, is related to the dynamics of the interstellar environment, specifically the Mach number of the shock and the adiabatic index of the interstellar gas. Nevertheless, particle acceleration at shock boundaries results in power law energy spectra with $\alpha \sim 1$. Since equation 1.7 represents the integral spectrum, taking the derivative with respect to E_n reveals a differential spectrum with index $\gamma = \alpha + 1 \sim 2$, harder but qualitatively near the observed cosmic ray index of 2.7.

While unequivocal evidence for this scenario is needed, the rate of supernova explosions and the power released is sufficient to fuel the Galactic cosmic ray flux density. Integrating the observed spectrum in figure 1.6, the local cosmic ray energy density is $\rho_{\text{E}} \sim 1$ eV cm^{-3}. Following the arguments of Stanev [33], we assume this is a constant throughout the galactic disc of thickness 300 pc and radius 15 kpc. We also suppose that particle diffusion leads to cosmic ray residence times $t_{\text{GD}} \sim 10^7$ yrs, so we can estimate the total power required to accelerate these particles via

$$L_{\text{CR}} = \frac{V_{\text{GD}} \rho_E}{\tau_{\text{GD}}} \simeq 3 \times 10^{40} \text{ergs/s}\,, \tag{1.8}$$

where the volume of the galactic disk $V_{\text{GD}} = \pi\, 300\,\text{pc}\,(\text{kpc})^2 \sim 10^{67}$ cm^3. The observed rate of supernovae is approximately 3 per century. Assuming progenitor stars of mass $10\,\text{M}_\odot$ expanding at 5×10^8 cm/s, the resulting luminosity is 3×10^{42} ergs/s. Thus, an acceleration process with an efficiency of merely 1% is sufficient to supply the cosmic ray energy density in a power law spectrum.

In both acceleration scenarios above, many particle interactions are required to reach the observed energies. Furthermore, the fact that we observe them means they attain energies such that escape from acceleration environments is possible. At this point, the particles no longer accelerate, and have reached the maximal attainable energy, which by rearranging equation 1.3 and scaling by the speed of the front, β_s, is

$$E_{\max}\,[\mathrm{eV}] = 10^{18}\,\beta_s\,Z\left(\frac{d}{\mathrm{kpc}}\right)\left(\frac{B}{\mu\mathrm{G}}\right)\ . \tag{1.9}$$

In this context, d is the size of the accelerator; thus as soon as a particle's energy results in $r_g \gtrsim d$, the particle is able to escape. This geometric condition is known as the Hillas criterion [34], and is useful in searching for potential acceleration sites at all cosmic ray energy scales, as depicted in figure 1.11.

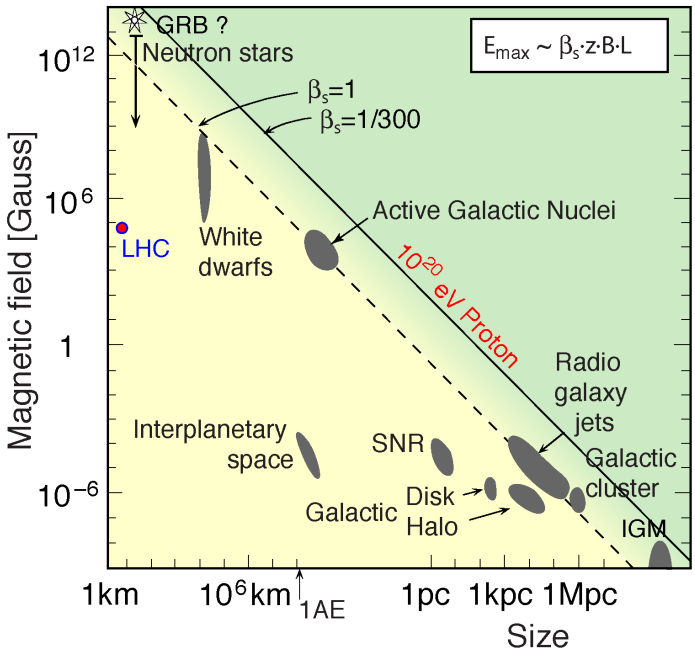

Figure 1.11: Hillas diagram indicating astrophysical objects capable of accelerating cosmic rays. The location of the LHC is included for reference. The energies attainable for each source depends on its dimensions and the strength of its magnetic field. Used with permission from J. Aguilar.

1.5 Propagation

Upon being ejected from their sources, charged particles are subject to a number of other processes in interstellar space. The propagation of cosmic rays within the Galaxy is depicted in figure 1.12 and it is governed by the diffusion-loss equation [35]:

$$\frac{\partial N_i}{\partial t} = D\,\nabla^2 N_i + \frac{\partial}{\partial E}\left[b(E)\,N_i\right] + Q_i - \frac{N_i}{\tau_i} + \sum_{j<i} \frac{P_{ji}}{\tau_j} N_j \tag{1.10}$$

which is composed of terms accounting for various sources of gains and losses to the number density N_i of species i. Here, Q represents the cosmic ray production rate by sources, D is the diffusion coefficient, $b(E)$ parameterizes the rate of energy loss, N_i/τ_i describes the rate of escape from the Galaxy with lifetime τ_i, and the last term incorporates nuclear spallation to other species j with probability P_{ji}. Thus equation 1.10 represents the time evolution of the integral spectrum of species i. Taking the approach of the previous section, some simple assumptions are made to extract general behavior.

As a basic approximation, we ignore the diffusion, spallation and energy loss terms. Furthermore we assume that the observed flux is at equilibrium, i.e. independent of time. This leaves us with the following for a single species:

$$\frac{\partial N(E)}{\partial t} = -\frac{1}{\tau_{\rm esc}} N(E) + Q(E) = 0 \tag{1.11}$$

where we have emphasized the energy dependence and that τ is the escape time scale. Rearranging, we see that $N(E) = \tau_{\rm esc}\,Q(E)$, i.e. cosmic ray production and loss rates are equal. Since we observe $N(E) \sim E^{-2.7}$ and from the previous section for supernovae $Q(E) \sim E^{-2}$, we can conclude that the escape time is energy dependent: $\tau_{\rm esc}(E) \sim E^{-0.7}$.

As a simple numerical estimate of $\tau_{\rm esc}$, we return to equation 1.10 considering only the diffusion term. If we assume that the Galactic magnetic field is responsible for the diffusive process, then the mean free path length between interactions λ is of the same order as the gyroradius, r_g. For general three-dimensional diffusion, the r.m.s. distance of a particle from its source after time τ is

$$\left\langle r^2 \right\rangle \simeq 2\,D\,\tau\ . \tag{1.12}$$

And for relativistic charged particles in a magnetic field, the theory of Bohm gives the diffusion coefficient

$$D = \frac{1}{3}\,c\,r_g\ . \tag{1.13}$$

If we consider a PeV proton in the Galactic disk of thickness $z_{\rm GD} = 300$ pc and equate $\langle r^2 \rangle \rightarrow z_{\rm GD}^2$, then the escape time can be estimated:

$$\begin{aligned} \tau_{\rm esc} &= \frac{z_{\rm GD}^2}{2\,D} = \frac{3}{2}\frac{z_{\rm GD}^2}{c\,r_g} \\ &= \frac{3}{2}\frac{(300{\rm pc})^2}{(0.307\,{\rm pc/yr})\,(0.3\,{\rm pc})} \\ &\simeq 10^6\,{\rm yr}\ . \end{aligned} \tag{1.14}$$

Of course, an immediate consequence of these assumptions is that $\tau_{\rm esc} \sim E^{-1}$, not quite the fractional index as found above. Yet, we do gain some insight into the form of $\tau_{\rm esc}$ as this energy dependence originates from the gyroradius, r_g. This really implies $\tau_{\rm esc} \propto (E/|Z|)^{-\delta} \sim R^{-\delta}$, where $R = E\,/\,|Z|$ is the particle rigidity. While not surprising, this does indicate that the difference in δ due to the simple approximation must be composition dependent. It also means that escape times increase with increasing Z (or equivalently mass), so one expects the contribution of the heavier elements to the all-particle flux to increase with energy.

The real story is complicated by nuclear fragmentation during propagation. Spallation increases particle diffusion and alters species abundances via interactions with the insterstellar gas. By choosing a nuclide with an appropriate half-life, it is possible to investigate the origin of this δ discrepancy.

The canonical example is the isotope beryllium-10. Its nuclear interactions (MeV–GeV) are well understood, but the abundance of ^{10}Be is rare in ordinary matter, and nearly all ^{10}Be cosmic rays are formed by the fragmentation of heavier nuclei inelastically interacting with the interstellar medium. It has a radioactive half-life of $t_{1/2} = 1.39 \times 10^6$ yr, so the relative absence of ^{10}Be in the cosmic ray spectrum at GeV energies would indicate propagation lifetimes much larger than $t_{1/2}$. Indeed, there is an observed reduction of $\sim 0.2 - 0.4$ in ^{10}Be abundance to what is expected [36], meaning that $\tau'_{\rm esc} = 5\,\tau_{\rm esc} \sim 10^7$ yr (hence the value of $t_{\rm GD}$ in equation 1.8). If we just assume that δ absorbs all differences between the magnetic and spallative diffusion, then we simply solve for $\tau_{\rm esc}$ and $\tau^{'}_{\rm esc}$ at the same rigidity to get:

$$\begin{aligned} \tau_{\rm esc} &\propto R^{-1} \\ \tau^{'}_{\rm esc} &\propto R^{-\delta} \\ \frac{\tau^{'\delta}_{\rm esc}}{\tau_{\rm esc}} &= \frac{(5\,\tau_{\rm esc})^{\delta}}{\tau_{\rm esc}} = 1\ . \end{aligned}$$

Solving with $\tau_{\rm esc} = t_{1/2} \sim 10^6$ yr, we obtain $\delta \sim 0.9$, which is nearer to 0.7, but again, we have assumed that the difference in magnetic and spallative diffusion constants is negligible. Actual measurements estimate the total diffusion constant to obtain more accurate estimates of δ. For example, the TRACER detector [37] recently measured the boron to carbon ratio, obtaining $\delta = 0.53 \pm 0.06$ up to about TeV per nucleon energies. While these are very specific examples, it is clear that by the time cosmic rays reach Earth, their original source spectra are influenced by both magnetic diffusion and spallation processes.

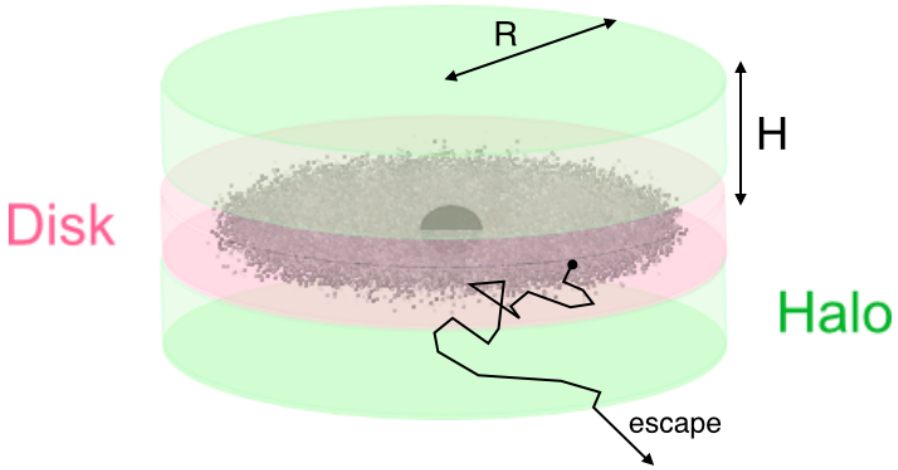

Figure 1.12: Diagram of a simplified galaxy, with radius r, disk height d, and halo height H. In the various approximations above, cosmic rays diffuse around the Galactic disk, interacting with the Milky Way's magnetic field and the interstellar medium. Particles can escape the disk if their escape times are long enough. Used with permission from J. Aguilar.

1.6 Anisotropy

The prior discussion assumed that the Galaxy is filled with randomly distributed sources in a leaky Galaxy, producing a steady, isotropic particle stream. And since the Larmor radius of Galactic TeV cosmic rays is about five times the size of the Solar System (~ 0.001 pc), orders of magnitude below the distance to the nearest SNRs, $O(100\text{pc})$, and dimensions of the Galaxy, such particles retain no memory of their origin upon arrival at Earth. But of course, any interpretation of measured arrival distribution, energy spectra and escape times must include the Earth's placement within the Galaxy and relation to possible sources that are both spatially and temporally discrete. For example, the value of τ_{esc} measured at Earth is in fact an average, for particles that originated nearer the disk edge can escape faster than those that are born towards the denser Galactic center, where there is more diffusive action and also more sources. At most a weak dipole should be present [38, 39], aligned with the source-dense Galactic center.

The study of the cosmic ray anisotropy is a powerful probe of acceleration sites and propagation, and is complementary to spectral and composition studies. Indeed, over the last decade many experiments have confirmed the observation of TeV anisotropy around the $10^{-4} - 10^{-3}$ level in relative intensity. In the northern hemisphere, these include measurements from Tibet-ASγ [40], Milagro [41], ARGO-YBJ [42] and a recent measururement by the HAWC Observatory [1]. In the southern hemisphere, IceCube [43] and IceTop [44] report qualitatively similar results. These experiments are ground based detectors, which is no coincidence, as large instrumented areas operating for long data taking periods are needed to reach per mille sensitivity at TeV–PeV energies.

The observed anisotropy has features on two angular scales: a large-scale dipole with sub-components from a few higher multipoles, and a small-scale structure with several excesses and deficits of about $10^\circ - 20^\circ$. Several measurements reported by IceCube and IceTop are presented in figure 1.13, showing energy dependence of the large-scale anisotropy. The noticeable flip in phase of the dipole structure could be the result of a change in the local source distribution, including sources being subject to internal energy cutoffs. The small-scale anisotropy measured with HAWC at 2 TeV is shown in figure 1.14. Several areas of excess are clearly visible in the map. The most significant excess is Region A, near right ascension $\alpha = 59.2^\circ$ and declination $\delta = -7.2^\circ$. Region A's evolution with energy is shown in figure 1.15, where it appears to steepen to $\sim$ 14 TeV before dropping. These fluctuations on top of the dipole anisotropy could be the result of young, nearby sources [39], an indication of non-diffusive propagation in the turbulent interstellar magnetic field [45, 46], or potentially heliospheric effects that could explain the spectral hardening [47]. Thus it is important to study the small-scale structure as a function of energy.

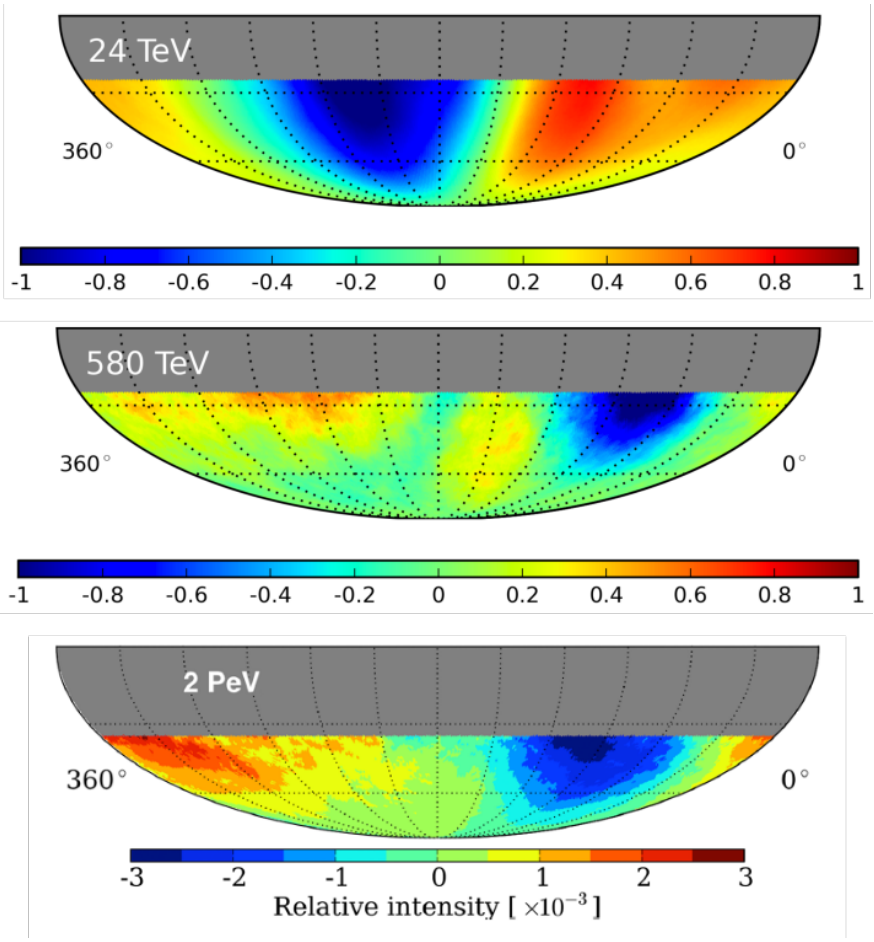

Figure 1.13: The large-scale structure measured with relative intensity as observed with IceCube [48] at 24 TeV and 580 TeV median energy, and IceTop [44] at 2 PeV. The phase of the dipole structure clearly flips with increasing energy. The relative intensity is scaled by the value 10^{-3} for each map as indicated by the label.

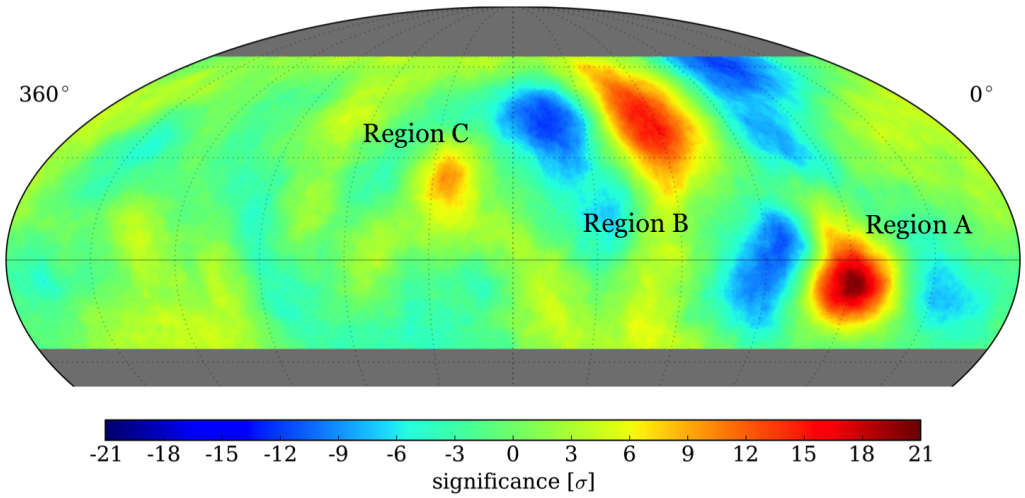

Figure 1.14: Significance map of small-scale anisotropy at a median energy of 2 TeV measured with HAWC [1] Three regions of excess are labelled, with Region A having the largest significance.

The presence of this anisotropy clearly implies that cosmic ray production and propagation are not fully understood. As we shall see, the energy response of the full HAWC Observatory to hadronic air showers allows for an improved measurement of the small-scale anisotropy's energy dependence, as well as a detailed measurement of the all-particle cosmic ray spectrum above 10 TeV. This is due in part to the large field of view and long observation times required, the comparable extent of multi-TeV showers and the containment area of HAWC, as well as the array's proximity to shower maximum at these energies, where shower fluctuations are at a minimum. This places the HAWC Observatory in a unique position to bridge cosmic ray measurements between direct detection apparatuses and larger PeV-scale air shower array experiments.

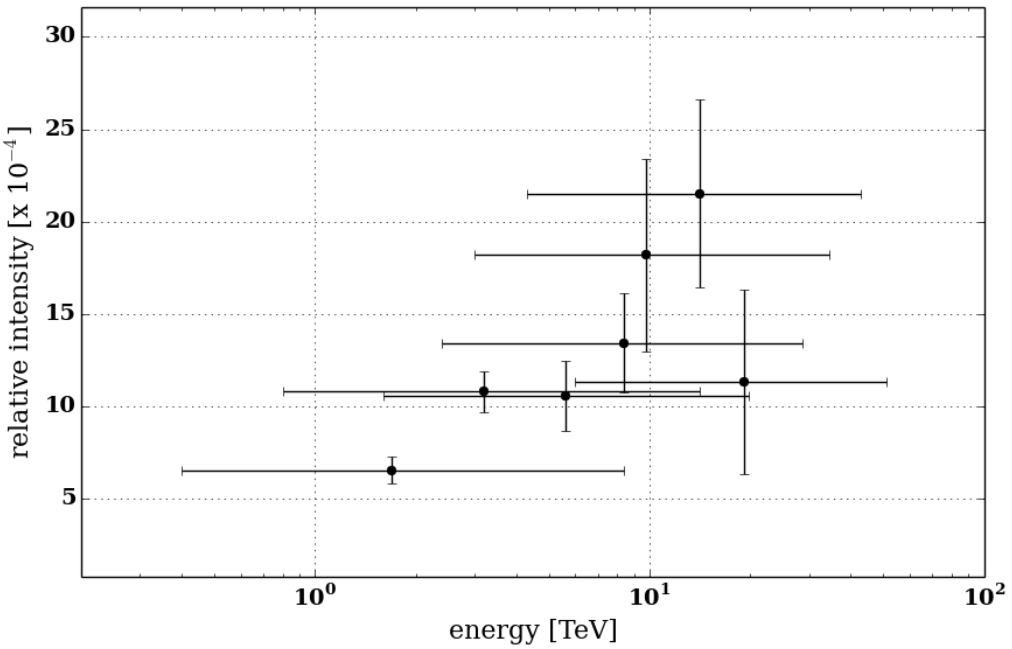

Figure 1.15: Relative intensity as a function of energy for Region A [1]. Uncertainties in both the energy and relative intensity represent the 68% central containment region.

Chapter 2

The HAWC Detector

The HAWC Observatory is an air shower array located at 4100 m a.s.l. on the slopes of Volcán Sierra Negra in the state of Puebla, México. The detector comprises a large 22,000 m^2 array of 300 close-packed water Cherenkov detectors (WCDs), of which 294 have been instrumented. HAWC is designed to detect air showers from primary gamma rays in the 100 GeV to 100 TeV energy range, but its altitude and physical dimensions permit measurements of primary hadronic cosmic rays up to multi-PeV energies. A photograph of the observatory during the final construction phase is shown in figure 2.1. The white building in the center of the array is the Counting House, home to the data acquisition system, online computing farm, and laser calibration system. Buried coaxial cables run from here to each tank, providing high voltage to and receiving signals from the WCDs. The large white rectangle at the top right of the image is the HAWC Utility Building (HUB), which houses the water filtration facility.

Figure 2.1: Photograph of the HAWC Observatory as viewed on the 19th of August, 2014 from the neighboring peak, Sierra Negra. This was taken during the final stages of construction; as one can see the last row of tanks is outlined on the bottom left corner of the image. Used with permission from Jordan Goodman.

2.1 Water Cherenkov Detectors (WCDs)

The Cherenkov detector unit in HAWC is a large, 5 meter tall, 7.3 meter wide water tank. Each WCD consists of a cylindrical corrugated steel wall protecting an inner light-tight plastic liner, or bladder (figure 2.3), which holds ∼188,000 liters of purified water. This corresponds to a water height of 4.5 m, which was chosen such that 100 MeV photons centrally entering the tank produce Cherenkov light that uniformly illuminates the tank bottom. On top sits a protective canvas dome supported by a steel frame to prevent rain or snow accumulation. The modular design and choice of materials meant that each tank could be assembled and instrumented one at a time. This allowed the experiment to operate in data-taking mode and commissioning of the detector elements to occur during the various phases of construction. As shown in the left panel of figure 2.2, the 300 WCDs form an array covering an area of 22,000 m^2, while the total instrumented area amounts to 12,500 m^2.

The secondary air shower particles primarily consist of γ and $e^{\pm}$, and some highly energetic muons. When these particles pass through the water in the tank, the Cherenkov light emitted propagates and scatters in the water, and reflects diffusively off of the liner walls. This continues until the light is either absorbed in the water or the liner plastic, or lands on one of the four photomultiplier tubes (PMTs) mounted on the bottom of the tank. The secondary γ do not emit Cherenkov radiation themselves, however, upon entering the tank, produce $e^{\pm}$ pairs which do. A schematic diagram of a single WCD is shown in the right panel of figure 2.2.

2.2 Cherenkov Radiation

Cherenkov radiation is light produced when a relativistic charged particle travels through a refractive material at a speed faster than the speed of light in that medium. For a generic material with an index of refraction n, the speed of light in that material is given by $c_n = c/n$ where c is the speed of light *in vacuo*. Hence, Cherenkov radiation is emitted if the particle's velocity $v > c/n$. This radiation is the result of the **medium's** response to the charged particle, and can be conceptually likened to macro-scale phenomena such as shock-waves. As a simple example, a duck swimming faster than water waves can travel produces a cone-like wavefront in still water.

The geometric representation of this process is depicted in figure 2.4. Because the particle moves superluminally through the medium, the constructive interference of the electromagnetic radiation creates a conical wave behind the particle. The coherent wavefront propagates with speed c_n at a fixed angle $\cos\theta = \frac{1}{n}$ with respect to the velocity vector of the particle. For a relativistic charged particle ($\beta = v/c \sim 1$) passing through water ($n \sim 1.33$), the Cherenkov light cone propagates at an angle of $\theta \sim 41.2°$.

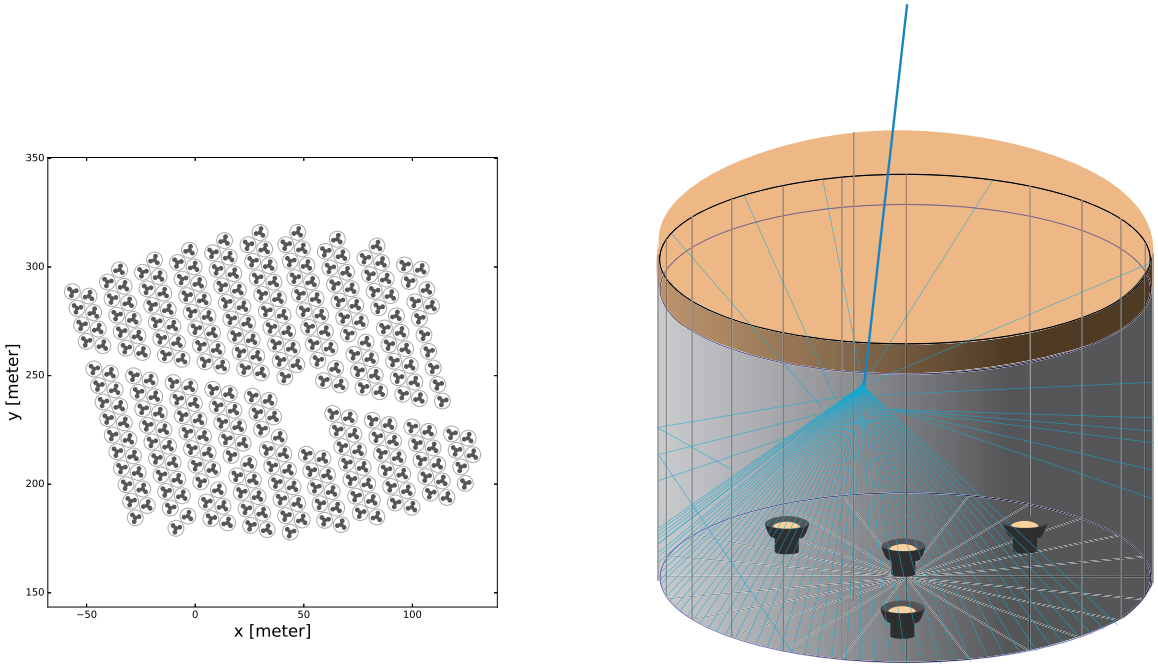

Figure 2.2: The left panel shows the layout of the entire HAWC detector. Each WCD is indicated by a large circle encompassing the smaller, darker circles which identify the PMT locations. The right panel depicts the representation of a single WCD including the steel tank, the protective roof, and the four PMTs. The penetrating dark blue line represents a high-energy secondary air shower particle, which emits Cherenkov radiation indicated by the cyan rays inside the WCD volume.

Figure 2.3: A HAWC bladder being inflated for testing. The exterior of the bladder is white, while the interior face is black and absorptive. Used with permission from Jordan Goodman.

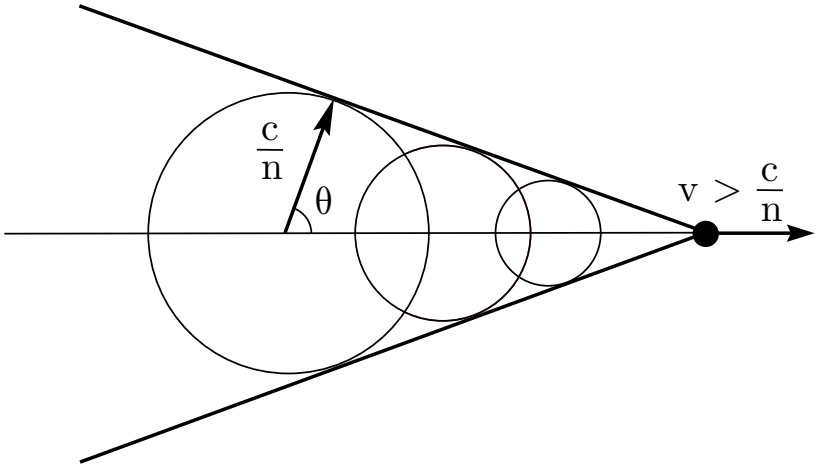

Figure 2.4: Huygens' wavefront construction illustrating Cherenkov radiation. The light travels with speed $c_n = \frac{c}{n}$ at an angle θ from the particle velocity vector, which traverses the medium at a speed $v > c_n$. The radiation forms a coherent wavefront in the shape of a cone with apex at the particle location.

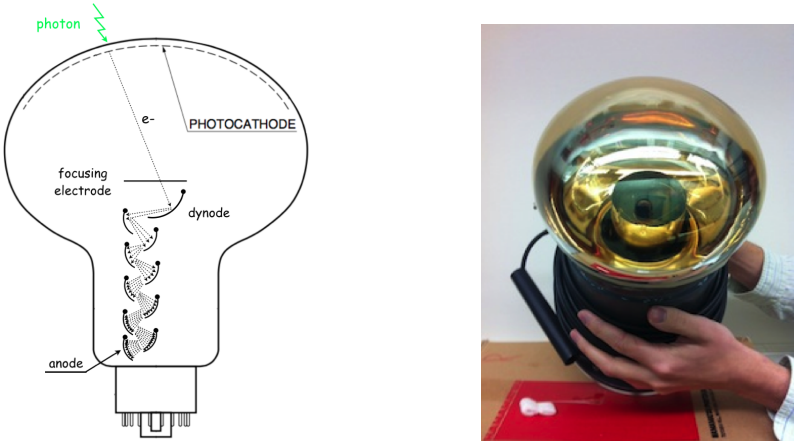

Figure 2.5: Left: diagram of the inner workings of a PMT. An incident photon releases an electron from the photocathode, which induces a cascade in the dynode chain. Used with permission from Joshua Wood. Right: a Hamamatsu R5912 8" photomultiplier tube used in HAWC.

Since high energy cosmic rays generate cascades of relativistic secondary particles such as $e^{\pm}$ and $\pi^{\pm}$, detectors can be built with appropriate dimensions and materials optimized to measure the Cherenkov light produced as the cascade particles pass through the detector. The detector elements deployed in HAWC operate under these principles.

2.3 Photomultiplier Tubes (PMTs)

Photomultiplier tubes are highly sensitive photon counting devices, capable of detecting single photons. The four upward-facing PMTs used in HAWC WCDs are anchored to the liner floor, ensuring their surveyed locations are known precisely, which is of prime importance for air-shower event reconstruction. Three 8" Hamamatsu R5912 PMTs passed on from the Milagro experiment are arranged in an equilateral triangle, each spaced 1.8 m from a central 10" high quantum efficiency Hamamatsu R7081-02 PMT. The right panel of figure 2.5 shows a photograph of an 8" HAWC PMT, and the left panel shows a schematic diagram of the inner workings of a general PMT.

When a Cherenkov photon is incident upon the PMT cathode surface, an electron can be ejected via the photoelectric effect with probability represented by the quantum efficiency (QE). The QE is dependent on the wavelength of the impinging photon and the location on the surface of the PMT cathode, and it typically includes the probability for the ejected electron to land on the first dynode. The manufacturer-quoted peak QE for the R5912 PMT is $\sim 25\%$, and $\sim 40\%$ for the R7081-02 PMT. Upon leaving the cathode surface, the electron accelerates onto the first dynode where it ejects more electrons. These in turn get accelerated by subsequent dynodes, all of which are held at increasing high voltage to amplify the developing electron

cascade at each step. The electrons finally are collected on the anode where the resulting signal is output for measurement. The ratio of the mean integrated output from a single incident photon to the fundamental electron charge defines the PMT gain. This value depends on the number of dynodes and the total voltage at which the PMT is operated. The PMTs in each WCD are gain matched to ensure uniform response in the data acquisition electronics.

2.4 Data Acquisition System

HAWC employs two data acquisition (DAQ) systems [49]: a main DAQ which records individual air shower events, and a scaler system which records single PMT rates. The scaler DAQ is intended for monitoring transient gamma ray phenomena such as GRBs and flaring AGNs, as well as fluctuations in the cosmic ray rate from Solar activity. This analysis uses events reconstructed through the main DAQ.

The PMT signals are transferred via RG59 coaxial cables to a central counting house, where they are amplified, shaped, and discriminated on custom front-end boards (FEBs) using two voltage thresholds: one at 20 mV and the other at 50 mV. The time stamps when signals cross these thresholds are recorded by CAEN VX1190A time-to-digital converters (TDCs) with a precision of 100 ps. The resulting time-over-threshold (ToT) is proportional to the logarithm of the pulse's total charge and is further detailed in section 3.1. The TDCs then send the data to a computer farm for further processing. A simple multiplicity trigger condition is enforced to identify candidate air shower events, ensuring that a minimum of N_{thresh} PMTs record signals within a defined time window. The value used in this analysis is $N_{\text{thresh}} > 28$ within 150 ns.

2.5 Calibration System

Accurate reconstruction of air-shower events requires that the measured PMT signals also must be measured accurately. Though the PMTs were relatively calibrated prior to deployment, other effects can influence measured signals. For example, the water quality within WCDs influences the attenuation and scattering length of emitted Cherenkov light, changing the number of photons incident on PMT cathodes. Similarly, the cable length from the PMTs to the FEBs affects the signal propagation and can smear out PMT pulses. A dedicated calibration system is installed at the HAWC site, performing an *in situ* measurement of PMT signals provided by a known input light source, effectively taking these potential issues into account.

A laser held in the counting house sends pulses of light that travel through optical fibers to each WCD tank. A set of neutral density filters establishes the light intensity, covering six orders of magnitude in dynamic range. Entering through the top, the light is scattered with an optical diffuser, which illuminates the surfaces of all four PMTs. Since the amount of light entering the tank is well known, the measured PMT signal times and charges are calibrated. The signals registered during calibration are flagged in the DAQ so

as to not include pulses from the laser light into the air shower data taking chain. The light source is itself monitored closely by a set of photodiodes installed both before and after the filter.

2.6 Science Goals

As a detector capable of observing TeV scale gamma rays and cosmic rays, the HAWC Observatory operates within the scope of multimessenger astrophysics. Hence the science program of HAWC is diverse and falls into four main, yet interconnected categories: Galactic, extragalactic, cosmic ray, and beyond standard model physics.

The Galactic science program involves understanding Galactic TeV gamma ray sources, such as supernova remnants (SNRs) and extended objects like the Fermi Bubbles [50]. Characterizing the shape of these sources' gamma ray energy spectra provides clues regarding whether source environments are dominated by leptonic processes or hadronic particle interactions. For point sources, it is expected that gamma rays produced by electrons undergoing synchrotron self-Compton emission should have an energy spectrum that cuts off earlier than if produced by a purely hadronic mechanism. For extended emission, cosmic rays interacting with interstellar gas produce pions that further decay into high energy gamma rays. The measurement of the diffuse flux can thus aid in the search for regions of cosmic ray sources within the Galaxy.

HAWC can also probe sources outside our Galaxy, where the most violent environments in the universe reside. Due to its wide field of view and nearly constant duty cycle, HAWC can search for transient phenomena such as gamma-ray bursts (GRBs) and flaring sources such as active galactic nuclei (AGNs). These observations are part of a real-time alert system involving instruments that operate at other wavelengths from radio up to TeV gamma rays. In turn, HAWC can be alerted to flaring sources by other observatories, providing complementary observations at the highest energies, where for example space-based detectors can not measure due to limiting detector acceptances. Also within this context is the study of the extragalactic environment between us and sources. Specifically, the extragalactic background light (EBL) and the intergalactic magnetic field (IGMF) influence gamma rays on their path to Earth [51, 52]. By observing gamma ray spectra from blazars (AGNs with their jets pointing in our direction), these background fields can be probed to further understand the universe's cosmological history.

So far, the discussion has focused on the observation of gamma rays to probe high energy environments, yet HAWC has a healthy TeV-scale cosmic ray science program as well. As will be discussed in this thesis, HAWC can measure the all-particle cosmic ray energy spectrum, whose shape is important for understanding the nature of nearby sources as well as the environment in which cosmic rays propagate. Another measurement which takes advantage of HAWC's wide field of view and continuous operation is the observation of the cosmic ray anisotropy to the 10^{-3} level, which provides an additional lens with which to probe our local cosmic ray environment.

Finally, HAWC contributes to new physics topics by putting limits on exotic phenomena. This includes constraining Lorentz invariance violation with transient sources, and searching for signals from dark matter and massive relic particles like SUSY Q-Balls and primordial black holes [53]. These special searches are implemented within the normal observational capacity of the other science topics.

2.7 The HAWC Collaboration

The HAWC Collaboration consists of more than 100 scientists from institutions in Costa Rica, Germany, Mexico, Poland, and the United States. It formed in 2007 as a joint project originally between Mexico and the US to build upon the success of the Milagro experiment [54]. Milagro was a first-generation, wide field of view gamma ray experiment located at 2,630 m a.s.l. in the Jemez Mountains near Los Alamos, NM. With seven years of operation, Milagro pioneered the water Cherenkov technique for TeV-scale observations, leading to the first unbiased deep survey of the gamma ray sky [55], identification of new gamma ray sources [56], and discovery of the small-scale cosmic ray anisotropy [56].

The HAWC Observatory's design improves upon Milagro by moving to higher altitude, increasing the detector area five-fold, and isolating detector elements. This has significant impact on the sensitivity to gamma rays between 100 GeV–100 TeV, and extends cosmic ray observation capabilities up to the cosmic ray knee. Construction began with a prototype, VAMOS, starting in June 2010. The VAMOS array operated from June 2011–May 2012, and served as a test for various operations necessary to build the full HAWC array, such as detector construction, water delivery to the site, and electronics testing. The first elements of the main array were installed mid-2012, and HAWC was built in stages of 30, 111, 250, and finally 300 tanks. The inauguration of the full detector took place in March of 2015. Development of the site and detector deployment are shown in figure 2.6.

Figure 2.6: The HAWC site in various stages of construction. The top panel shows the initial 7-tank test array, VAMOS, followed by the leveling of the site in preparation for the main HAWC array. The 30-tank array was in operation by August of 2012, and 111 WCDs were completed by mid-2013. The final panel shows the completed HAWC detector at the end of 2014. The large peak in the background is Pico de Orizaba.

Chapter 3

Reconstruction of Air Showers with HAWC

The generalized data object of the HAWC detector is an air shower event, which consists of a variety of measurable and observable quantities. The simplest observable is the number of PMTs that record a signal within a predefined time window, N_{hit}. If a minimum number of PMTs in the array do not have a signal, the threshold multiplicity condition is not met, and the potential air shower event is rejected. For the analyses in this thesis, the most important quantities are composite, requiring sophisticated algorithms using the information from the collected charge and relative timing of PMT signals. These include an event's directionality, core location, and the estimated energy of the primary particle. The energy reconstruction is discussed in chapter 4. This chapter begins by describing the fundamental event reconstruction in HAWC, the simulation used to verify the reconstruction procedures, and finally the data reduction required to conduct the analyses in this thesis.

3.1 PMT Hits

A PMT "hit" is the basic component required for air shower reconstruction. It is defined by the position, time, and integrated signal or charge from light detected by a single PMT. Typical PMT signals from Cherenkov light are characterized by a steep rise followed by a relatively long tail. In HAWC, entire analog voltage waveforms are not recorded but are instead digitized to time stamps by the FEBs, providing a continuous data stream of leading and trailing edges.

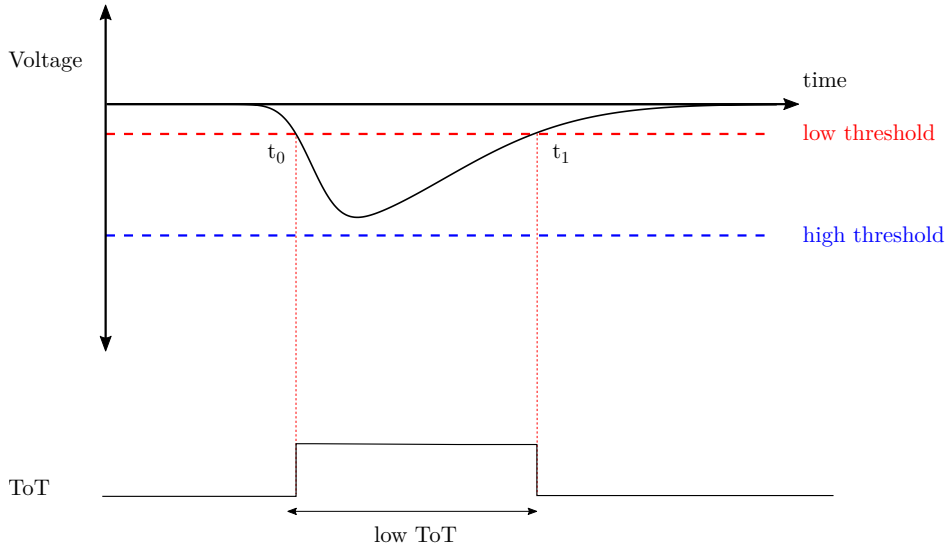

Figure 3.1: Diagram of a two-edge hit. The PMT analog signal is represented by the black curve. Only the lower voltage threshold is crossed, resulting in the digital low ToT edges shown below.

Each edge corresponds to the time the PMT signal crosses one of two thresholds: a low threshold of 20 mV and a high threshold corresponding to 50 mV. A signal which only exceeds the low threshold produces a 2-edge hit, since it also crosses as the signal decays, while 4-edge hits comes from larger signals that pass both signal thresholds. The time difference between the respective thresholds crossings is referred to as the time-over-threshold (ToT), hence there is a low ToT and a high ToT. The calibration system described in section 2.5 uses the laser light source to define the conversion factor between the ToT and the integrated signal, ensuring that PMT charges in air shower events are measured accurately. Figures 3.1 and 3.2 illustrate this digitization scheme for producing ToT hits. Occasional "bad" hits are recorded due to combinations of very small pulses or prompt afterpulsing in PMTs. These result in very long 2-edge and 6-edge hits (figure 3.3), respectively; however, the hit finding algorithm is programmed to identify and remove these undesireable hits from the data stream.

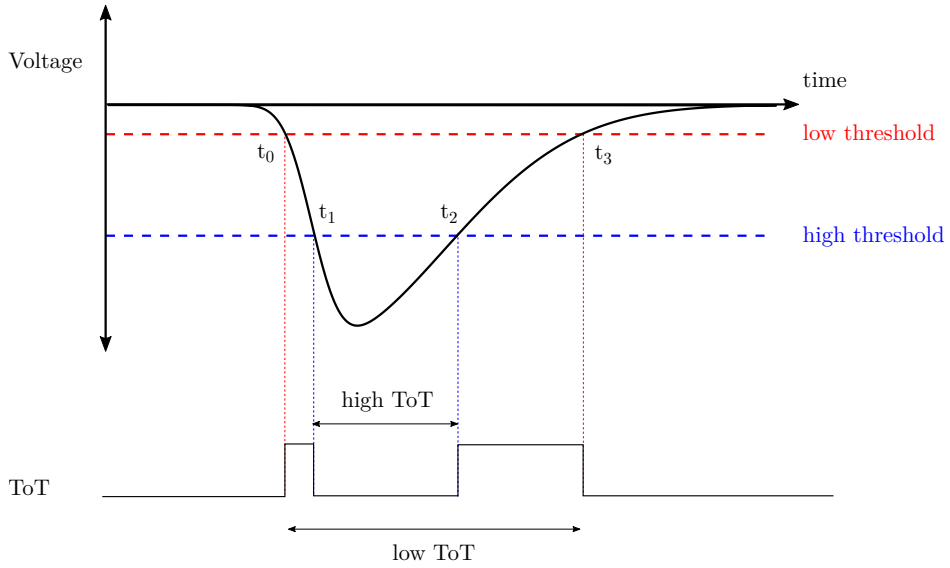

Figure 3.2: Diagram of a four-edge hit. The PMT analog signal is represented by the black curve. Both the low and high voltage thresholds are crossed, and the pulse gives a four edge digital ToT signal.

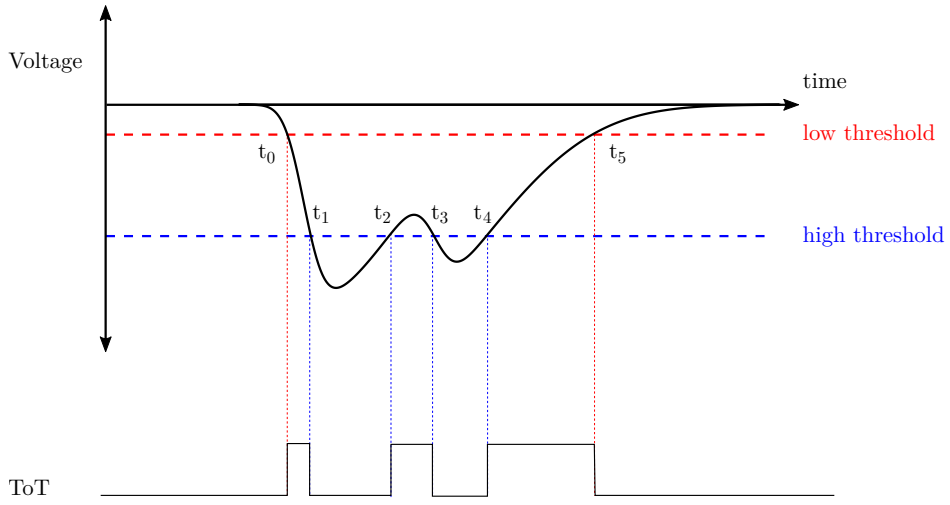

Figure 3.3: Diagram of a six-edge hit. Two crossings of the higher threshold occurred. This type of hit is flagged by the edge-finding software and removed from the data stream.

The recorded ToT hits are sorted into an ordered time series, and the collection of hits within a predefined time window defines a candidate event. But due to various sources of noise such as PMT dark noise and small fragments of distant showers, random hits enter the data stream, so a minimum hit threshold condition is implemented to identify reconstructable showers. During the data taking period for this analysis, the requirement $N_{\mathrm{hit}} \geq 28$ within a time window of 150 ns defines an air shower event. The resulting trigger rate is ~ 25 kHz.

3.2 Air Shower Reconstruction

3.2.1 Core Location

As discussed in section 1.2, extensive air showers propagate along the axis defined by the incident primary particle. The shower front is densest along this axis, and its projection onto the ground is called the core. In an air shower detector such as HAWC, the core is the location of maximal energy desposition. Figure 3.4 shows the charge distribution measured in HAWC for an example air shower event. Farther from the core, the shower front becomes less populated and wider in time. For the electromagnetic component of the shower, the expected particle density a certain lateral distance from the core is well approximated by the NKG function of equation 1.1. This is an estimate to the electromagnetic *particle* density in the shower front; however, HAWC WCDs measure the *energy* density, so a more appropriate function is NKG$\times 1/r$.

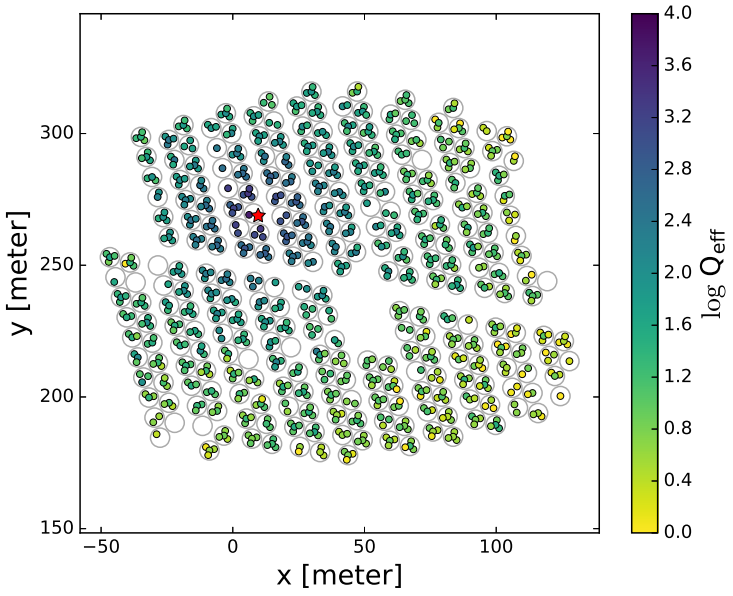

Figure 3.4: Example event from data showing charge distribution on the array. The SFCF core fit indicated by the red star is within the region of greatest energy density deposition, and is located at $\mathbf{X}_{\mathrm{SFCF}} = (9.6\ \mathrm{m}, 268.8\ \mathrm{m})$.

Due to performance limitations, we instead implement a simplification that includes a Gaussian component, which we call the super-fast core fitter (SFCF). The SFCF takes the following form:

$$S(A, \mathbf{x}, \mathbf{x}_i) = A \cdot \left(\frac{1}{2\pi\sigma^2} e^{\frac{-|\mathbf{x}_i - \mathbf{x}|^2}{2\sigma^2}} + \frac{N}{(0.5 + \frac{|\mathbf{x}_i - \mathbf{x}|}{R_m})^3} \right), \tag{3.1}$$

where A is the amplitude, $\mathbf{x}$ is the core location, $\mathbf{x}_i$ is the location of a PMT signal measurement, R_m is the Molière radius of the atmosphere (124.2 m at HAWC altitude), σ is the width of the Gaussian, and N is the normalization of the NKG tail. The values of σ and N are fixed to 10 m and 5×10^{-5}, respectively. The SFCF form is shown in figure 3.5 along with comparisons to the NKG-type functions, as well as a pure Gaussian function. Like the NKG, the SFCF function has radial symmetry, and despite including a less cuspy Gaussian component near the core, it still incorporates an extended exponential decay characteristic of the NKG tail.

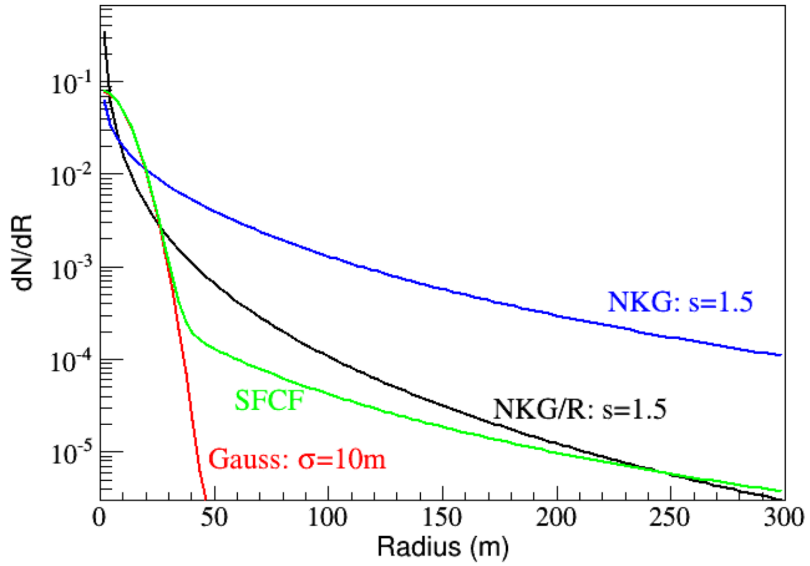

Figure 3.5: Comparison of several radial distribution functions. The rapidly rising NKG-type functions near $r = 0$ prove difficult to fit without having to iterate extensively as compared to the SFCF, which has a less cuspy-like form closer to the core location. The Gaussian function, while well-behaved near the core, falls off too fast to describe air-shower tails.

In order to estimate the core location with a general lateral function $Q(\mathbf{x}_i)$, we maximize the quantity

$$\mathscr{X} = \sum_{i=1}^{N_{\mathrm{PMT}}} \frac{(Z_i - Q(\mathbf{x}_i))^2}{Q(\mathbf{x}_i) + \sigma_i^2} \ , \tag{3.2}$$

where the index i runs over all operational PMTs, N_{PMT}, each with charge measurement Z_i and estimated uncertainty in the charge measurement, σ_i, which is derived from the calibration. Photomultiplier tubes measure charge in number of PEs so the expected signal measurement, $Q(\mathbf{x}_i)$, has Poissonian uncertainties, i.e. $\sqrt{Q(\mathbf{x}_i)}$. The denominator accounts for these two uncertainties with their quadrature sum. Maximization of equation 3.2 can not be solved analytically for either the SFCF or NKG functions, so we use the Levenberg-Marquardt [57] iterative procedure.

Regardless of the form of a shower profile function, fitting the shower core is CPU-intensive because all core-to-PMT distances (including operational PMTs that do not record a signal) must be evaluated at each iteration of the fit. For the full HAWC array, this amounts to $\sim 10^3$ PMT-core calculations for each core position hypothesis. Fits to the the SFCF function converge faster than for the NKG because it is simpler and its derivatives can be computed analytically. It also ameliorates excessive iterations due to the NKG function's rapidly increasing amplitude near $r = 0$, as well as expensive calls to fitting the exponential terms involving the age parameter, s.

3.2.2 Shower Age

Despite preferring the SFCF for the core fitting procedure, the age parameter s must still be estimated using the NKG hypothesis. However, the SFCF estimate fixes the core location and thus sets the PMT-core lateral distances. Furthermore, fitting the NKG age exponentials is mitigated by considering the logarithm of equation 1.1 scaled by $1/r$:

$$\log\left(\rho(r)/r\right) = \log A + s \cdot \left[\log\frac{r}{R_m} + \log\left(1+\frac{r}{R_m}\right)\right] - 3\cdot\log\frac{r}{R_m} - 4.5\cdot\log\left(1+\frac{r}{R_m}\right), \tag{3.3}$$

where r is the distance from the fit shower core, and the amplitude A absorbs various constants and exponential factors. For each event, the best fits to the amplitude A and age s are found by minimizing the χ^2 of equation 3.3.

3.2.3 Arrival Direction

To reconstruct the shower's arrival direction, the timing differences amongst the hit PMTs are used. As a first approximation, the front of shower particles can be represented by a plane traveling at the speed of light with some nominal thickness of ∼5 ns. This is true near the shower axis, where the particles have maintained the primary's directionality, but particles farther from the core have undergone multiple scatterings causing the particle disk to widen. Thus the shower front is curved, and fitting it to a plane must take into account the time delays for hits in the shower tail. This is illustrated in figure 3.6 using the same example event from figure 3.4. The typical curvature correction is of order 0.15 ns per meter from the shower core. A further complication comes from the fact that PMT hits with larger numbers of PEs are detected earlier than smaller hits. This sampling effect means that on average, detected hits from the shower tail will be further delayed by a few ns. Hence an additional sampling time correction is performed based on a PMT's measured signal.

Accounting for these timing effects, the corrected PMT hit times effectively project the shower front onto a planar front. A χ^2-minimization fit to an arrival plane is performed, weighting each PMT by its measured charge. The resulting θ, ϕ define the estimated local direction vector of the primary particle.

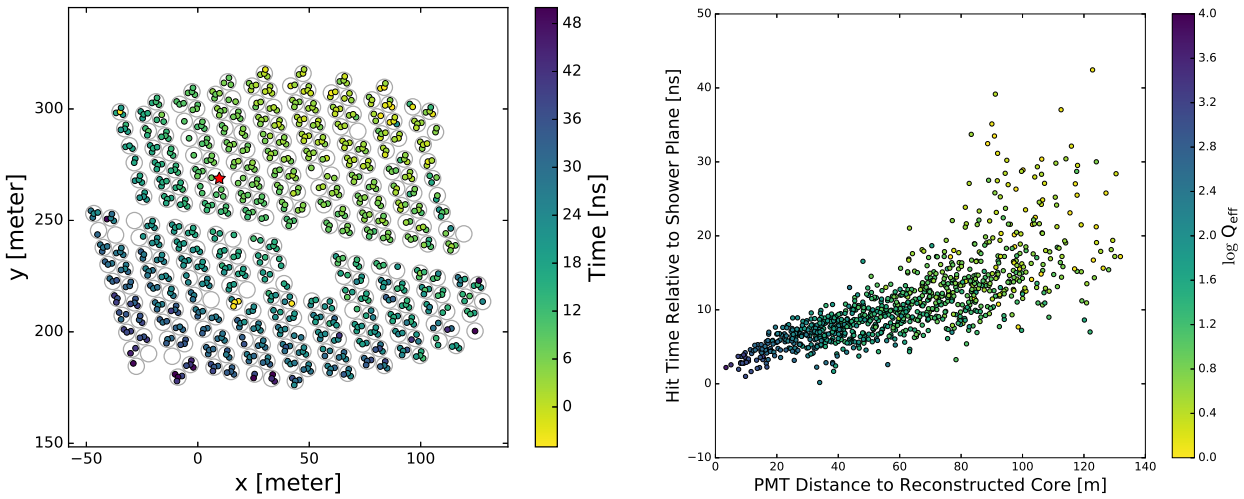

Figure 3.6: Timing distributions for the example event from figure 3.4. Left panel: arrival times on array, relative to the first recorded hit. The shower front comes from the upper right, moving across the detector towards the lower left. The SFCF fit core location is indicated by the red star. Right panel: time delays from planar front as a function of distance from the reconstructed core. The shower curvature with increasing lateral distance is evident. The reconstructed zenith and azimuthal angles are $\theta = 4.0°$ and $\phi = 46.7°$, respectively.

3.2.4 Iterative Core and Angle Fits

In practice, the core and angle fitting routines are performed iteratively, as an accurate core location is essential to making the curvature corrections to the planar front, and only in the shower plane is the particle distribution radially symmetric. The iterative core-angle fitting procedure starts with simple center center-of-mass core ($\mathbf{X}_{\mathrm{COM}}$) estimate, with lateral distances calculated in the plane of the detector, i.e.

$$\mathbf{X}_{\mathrm{COM}} = \frac{1}{Q} \sum_{i=0}^{N_{\mathrm{hit}}} q_i \, \mathbf{x}_i \, ,$$

where $\mathbf{x}_i, q_i$ respectively are the positions and charges of each of the N_{hit} PMTs in the event, and $Q = \sum_{i=0}^{N_{\mathrm{hit}}} q_i$ is the shower's total measured charge. This serves as a seed to a first fit with the SFCF function, again done in the detector plane. Then an initial angle fit is performed with this SFCF estimate and accounting for the

curvature correction. In an attempt to remove influence from very early or late hit times, a selection of hits within ±50 ns of the resulting fit plane is made which gets passed to another fit with the SFCF. In this final core fit, the PMT-to-core distances are calculated more appropriately in the shower plane. Using the same hit selection and the last SFCF core result, the angle fit is done once more, including the curvature correction factors. Resulting lateral distributions of charge in the shower plane are shown for two reconstructed events in figure 3.7. The likely gamma shower (left panel), characterized by mostly electromagnetic secondaries, is much smoother than the hadronic shower on the right.

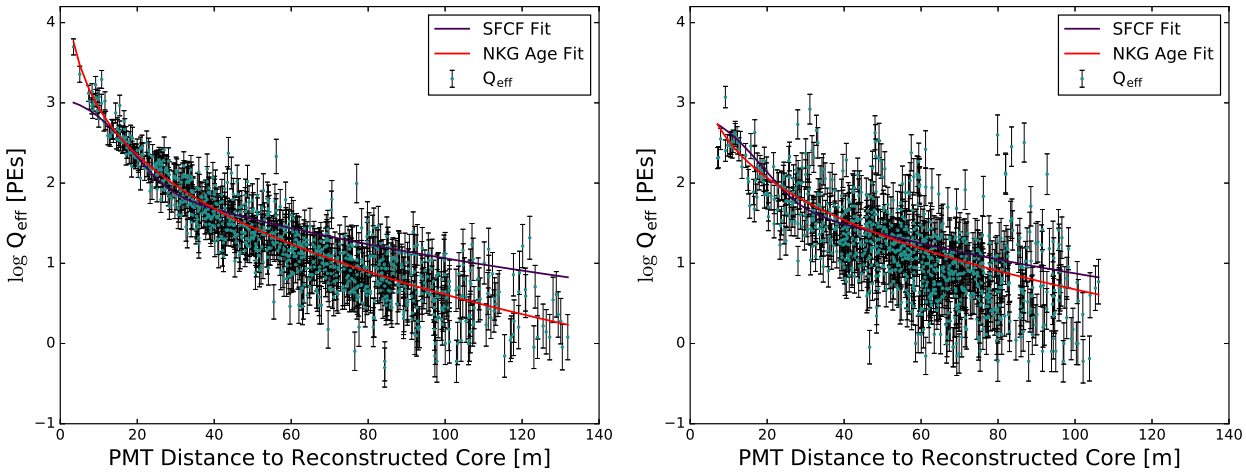

Figure 3.7: Lateral distribution of charge for a likely gamma ray event (left) and a cosmic ray event (right). Note the smoother nature of the gamma ray, as its development is primarily electromagnetic in nature. The SFCF fit used to find the core is represented by the purple curve, while the NKG function used to estimate the age is shown by the red curve.

3.3 HAWC Simulation

The core and angle fits provide the fundamental components to conduct TeV-scale cosmic ray analysis with HAWC. Building upon these quantities we can develop more sophisticated variables that relate to other aspects of the primary particle, such as energy estimators and particle discriminators, for example. Of course, these provide event-by-event estimates, but the detector resolution to these quantities or any biases in the reconstruction process can only be assessed by looking at a large collection of events that are well understood. Thus we characterize the response of HAWC via Monte Carlo (MC) simulation. This is divided into to independent components: the simulation of shower development, and the simulation of the detector to those showers.

3.3.1 Extensive Air Showers

For a primary particle incident on the Earth's atmosphere, the resulting air shower is simulated using the CORSIKA [58] package (v740), a standard toolkit in the cosmic ray community. CORSIKA tracks shower development using models based on the most up-to-date particle physics data from accelerator experiments. The implementation defining the standard simulation for this study was set up with FLUKA [59, 60] and QGSJET-II-03 [61] as the low- and high-energy particle physics interaction models, respectively. Smaller simulation sets were generated for hadronic interaction systematic studies using the EPOS [62] and SIBYLL [63] high energy models. The resulting systematic effects from these models on the spectral analysis are discussed in chapter 8.

Primaries of the eight species measured by the CREAM flights [21, 23] (H, ^{4}He, ^{12}C, ^{16}O, ^{20}Ne, ^{24}Mg, ^{28}Si, ^{56}Fe) were generated on an E^{-2} differential energy spectrum in energy ranges dependent on the primary. The lower limit was chosen to include showers that do not trigger the detector due to insufficient secondaries reaching ground level. The respective limits along with the total number of simulated showers are shown in the first data column of table 3.1. The simulated zenith angle range was $0^{\circ} - 70^{\circ}$, azimuthally symmetric, and weighted to a $\sin\theta\cos\theta$ arrival distribution.

	N_{showers}	E_{min} [GeV]	E_{max} [GeV]
H	10^5	5.0	2×10^3
He	10^5	10.0	2×10^3
C	5×10^4	50.0	2×10^3
O	5×10^4	50.0	2×10^3
Ne	5×10^4	100.0	2×10^3
Mg	5×10^4	100.0	2×10^3
Si	5×10^4	100.0	2×10^3
Fe	5×10^4	200.0	2×10^3

Table 3.1: Shower numbers and energy limits in CORSIKA simulation. The spectrum for all primaries is a power law with index $\gamma = -2$.

3.3.2 HAWC Events

Simulated HAWC events are created by subjecting the CORSIKA showers to the entire data analysis chain. Monte Carlo events thus have identical reconstructed data objects as real events. In this thesis, reconstructed variables from simulation are labeled either with "MC" or "true" subscripts, unless the risk of confusion is absent. The secondary charged particles interacting with the HAWC detector were simulated with the GEANT4 [64] package. A detailed model of the WCDs is used, including particle interactions, Cherenkov photon propagation and tracking, PMT response, and full electronics simulation. Effects from sources of noise and signal smearing are simulated as well, and treated as sources of systematic uncertainty from the nominal simulation. The CORSIKA showers are distributed over a 1 km radius circular area centered on the HAWC array to sample shower tails that can trigger the detector. Dedicated HAWC software was used to reconstruct simulated events using the same reconstruction algorithms applied to data.

3.3.3 Composition

The cosmic ray composition assumptions play an important role in understanding the detector response from simulation. Individual primaries were weighted by a double power law spectrum,

$$\frac{dN}{dE}(E) \propto \begin{cases} A \cdot (E/E_0)^{\gamma_1} & E < E_{\rm br} \\ A \cdot (E_{\rm br}/E_0)^{\gamma_2-\gamma_1} \cdot (E/E_0)^{\gamma_2} & E \geq E_{\rm br}, \end{cases} \tag{3.4}$$

where A is the normalization at E_0 with spectral index γ_1, and a second index γ_2 starting at energy $E_{\rm br}$. For all species' fits, $E_0 = 100$ GeV. The nominal composition used in simulation (referred to as the CREAM model) is the best fit of equation 3.4 to direct measurement data provided by AMS [65, 66], CREAM [21, 23], and PAMELA [20]. Due to the dominance of the lighter elements (H, ^{4}He) in the all-particle flux, the higher mass elements' ($Z > 2$) fluxes were varied together. The direct measurement data with corresponding fits are shown in figure 3.8. The abundances of the lighter elements were allowed to vary independently. The fit parameters with associated errors are provided in table 3.2. During the analysis, different mass compositions were used by varying the nominal composition by the estimated errors of the best fit parameters, and reweighting the simulated showers accordingly. The Hörandel [67] (also referred to as Polygonato), H4a [68], and Gaisser-Stanev-Tilav (GST4-gen) models also were implemented to further estimate the sensitivity of the results to the composition assumption.

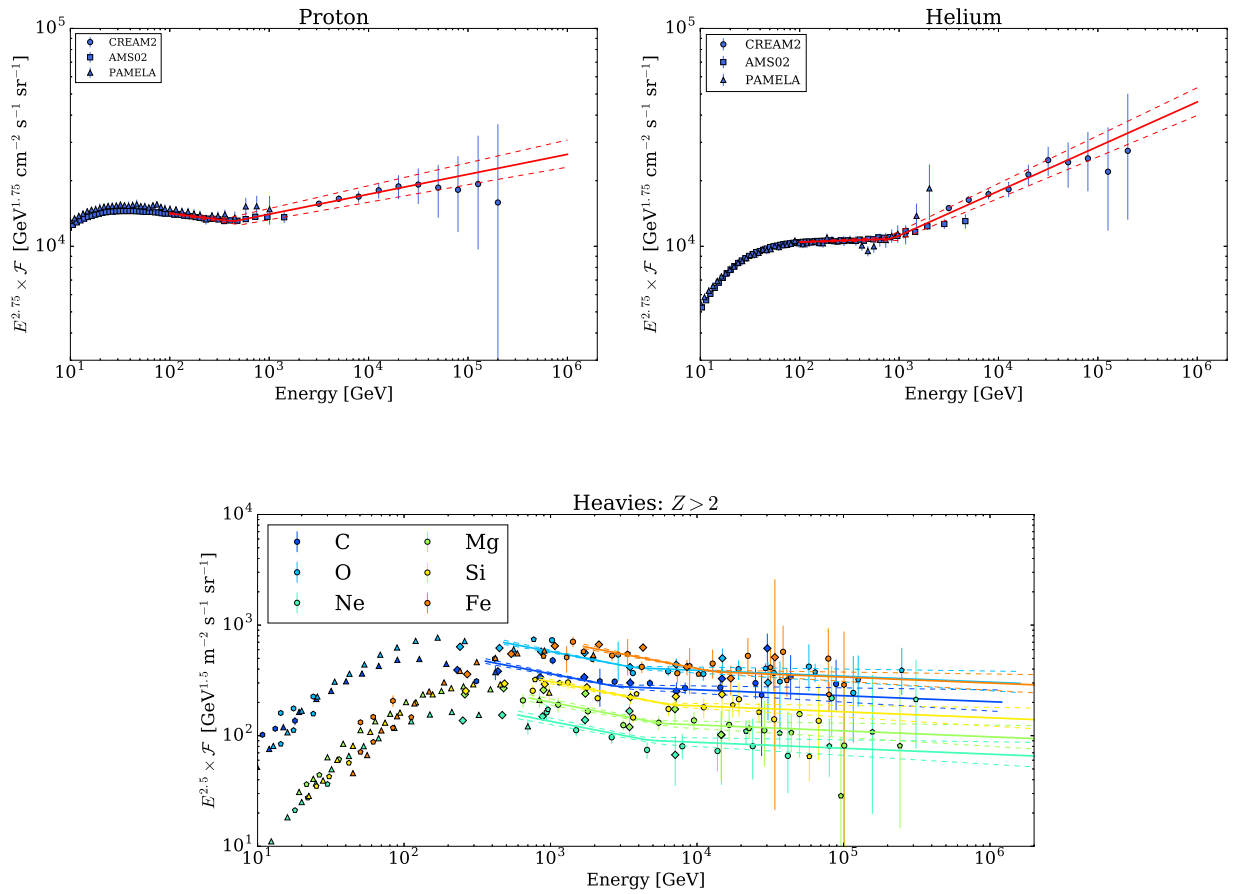

Figure 3.8: Fits to direct detection data from AMS [65, 66], CREAM [21, 23], and PAMELA [20] experiments. The flux is scaled to show the break energies, and includes the region bounded by the fit uncertainties. The fit parameters are presented in table 3.2, and define the nominal cosmic ray composition model used in this study.

	A	$E_{\rm br}$	γ_1	γ_2
	$[\text{GeV s sr m}^2]^{-1}$	[GeV]		
H	$4.48 \pm 0.04 \times 10^{-2}$	$440.6^{+87.8}_{-62.7}$	-2.81 ± 0.01	-2.66 ± 0.01
He	$3.31 \pm 0.02 \times 10^{-2}$	$854.8^{+125.7}_{-105.5}$	-2.73 ± 0.01	-2.54 ± 0.01
C	$6.96 \pm 0.18 \times 10^{-6}$	$2882^{+904.4}_{-481.9}$	-2.76 ± 0.03	-2.55 ± 0.04
O	$5.00 \pm 0.09 \times 10^{-6}$	3843^{+1206}_{-643}	-2.76 ± 0.03	-2.55 ± 0.04
Ne	$6.31 \pm 0.35 \times 10^{-7}$	4803^{+1507}_{-803}	-2.76 ± 0.03	-2.55 ± 0.04
Mg	$5.70 \pm 0.26 \times 10^{-7}$	5764^{+1809}_{-964}	-2.76 ± 0.03	-2.55 ± 0.04
Si	$5.70 \pm 0.13 \times 10^{-7}$	6725^{+2110}_{-1124}	-2.76 ± 0.03	-2.55 ± 0.04
Fe	$2.00 \pm 0.04 \times 10^{-7}$	13450^{+4220}_{-2249}	-2.76 ± 0.03	-2.55 ± 0.04

Table 3.2: Fit results to equation 3.4 defining the standard cosmic ray composition model for this work. Parameters were obtained as best fits to CREAM, AMS, and PAMELA data. Uncertainties in the fits are included in estimating the systematic uncertainties in flux measurements due to composition assumptions.

3.4 Event Selection

The MC data set passes through all the same reconstruction steps as data events. Thus, we can characterize the quality of HAWC events and perform data reduction in a manner that balances detector response resolution and event rate. The simplest reconstruction condition is requiring that the core and angle fits succeed. Without an angle estimate, no directional information is available for the event, and a reconstructed core location is required to estimate the primary energy. These fits are accompanied by estimated *fit* uncertainties, but a simple method to identify higher quality fits is to set a level of minimum information content, i.e. the number of hit PMTs.

Apart from the trigger condition, an additional N_{hit} cut requiring $\geq 6\%$ operational PMTs hit in an event is made, which, in accordance with previous cosmic ray analyses [1] sets an initial mean angular resolution of a few degrees. The bare core resolution is on par with the array dimensions, but this is due primarily to the multitude of events landing just off of the array, effectively hiding the true core region. Instead of increasing the multiplicity condition, we can simply demand that a core location have at least some nearby hits, which is of course not required to obtain a general core fit. This is summarized by the N_{r40} variable, which measures the total number of hits within a 40 m radius of the reconstructed core. A set of example events showing the implementation of this cut is shown in figure 3.9. Implementing the cut $N_{r40} \geq 40$ improves the median angular resolution to $\sim 0.5^\circ$ and the core resolution to ~ 10 m, a drastic improvement and paramount to energy resolution. The core and angular resolutions upon successively applying these cuts on simulation are shown in figure 3.10, and the reconstructed core distribution is shown in figure 3.11.

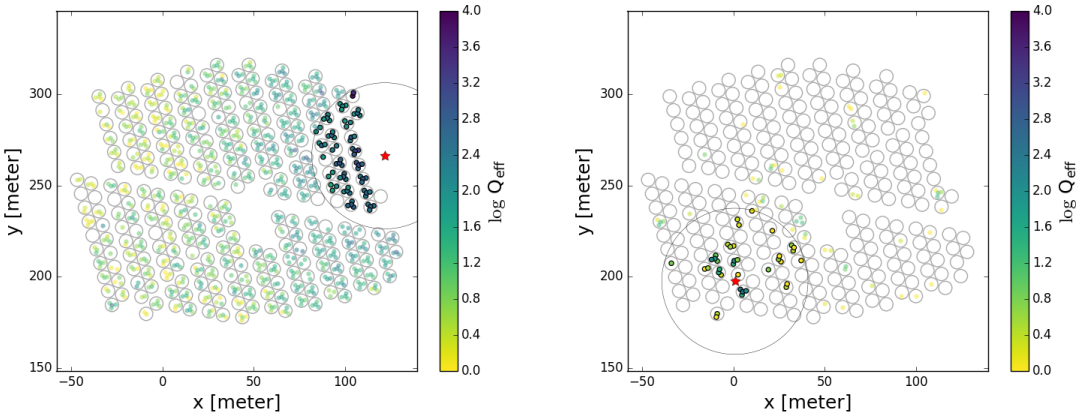

Figure 3.9: Example of the N_{r40} selection cut with two events. The 40 m circle is centered on the fit core location, outlining the set of highlighted PMTs. Here, the cut value is $N_{r40} \geq 40$, so the event on the left clearly passes this selection criterion, while the one on the right ($N_{r40} = 35$) is not accepted.

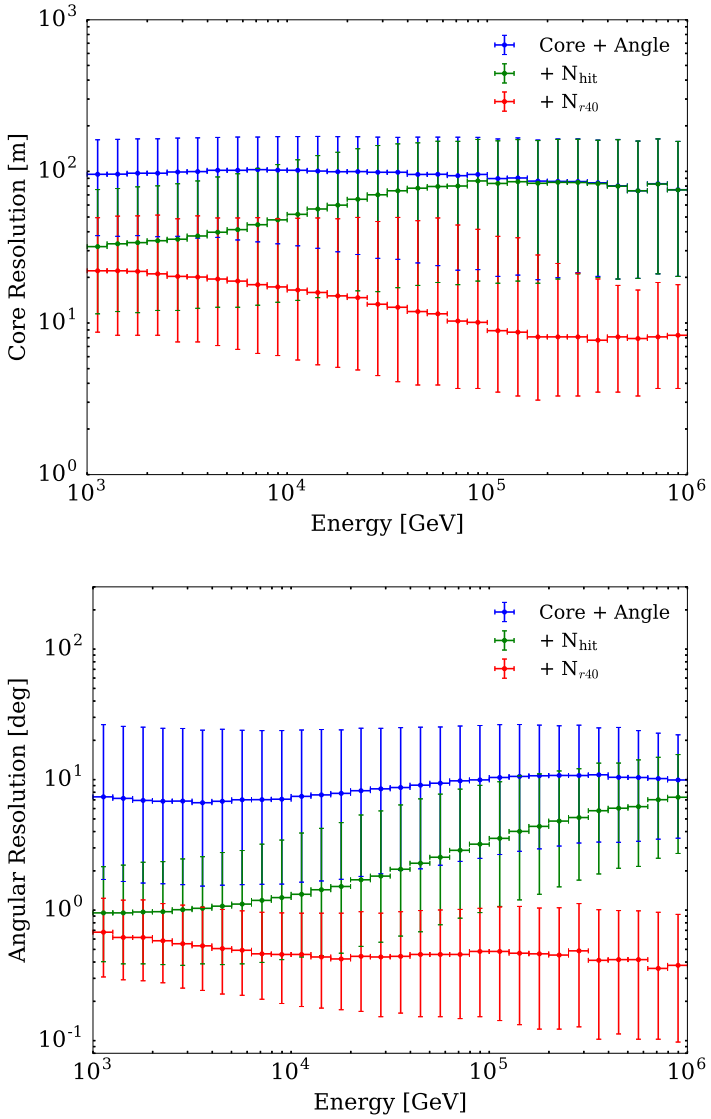

Figure 3.10: Core (upper) and angular (lower) resolutions as functions of Monte Carlo energy for the basic analysis cuts. The median values with 68 % containment are displayed. All events with $\theta_{\mathrm{MC}} \leq 60°$ were included, and cuts of $N_{\mathrm{hit}} \geq 75$ and $N_{r40} \geq 40$ were used. Each cut includes the previous as indicated by the + sign. Note the initial drop for lower energies with just the multiplicity cut, while at higher energies resolutions remain higher. This is due to large showers triggering the array but lacking suffient sampling for accurate reconstruction. The N_{r40} cut provides a simple geometric condition to drastically improve core and angular resolutions. The angular resolution is slightly less dependent on these cuts with increasing energy simply because the shower front is nearly planar. The core position, however, strongly depends on sampling nearby.

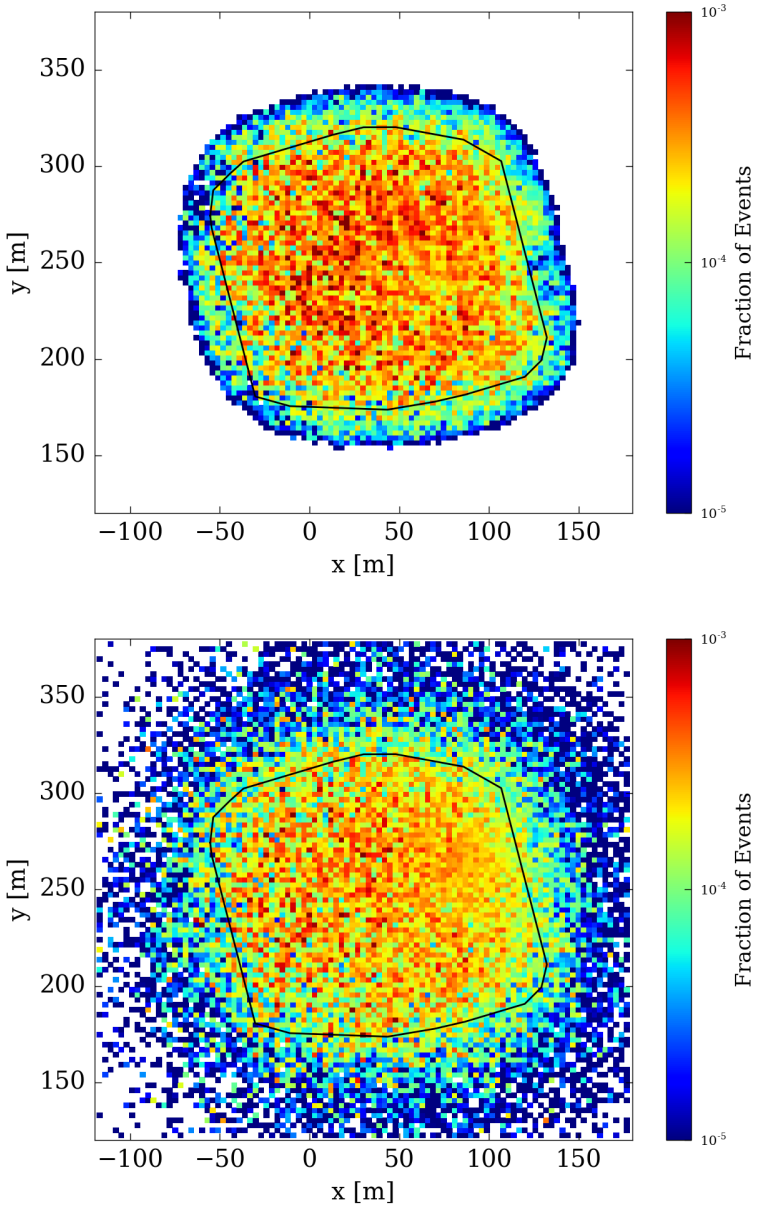

Figure 3.11: Reconstructed (top) and true (bottom) core distributions after applying analysis cuts for all simulated cosmic ray showers. The black outline indicates the edges of the array. By design the N_{r40} cut selects reconstructed cores within ∼15 m of the array, and the deviation from the true core location follows the core resolution from figure 3.10.

Of course, the number of showers passing the selection cuts depends on their energy. We quantify the *efficiency* to primary cosmic rays with the effective area, $A_{\text{eff}}(E)$. This quantity is estimated via simulation, and is paramount to making spectral measurements, since the flux is inversely proportional to the effective detection area (equation 1.2). Showers are simulated landing over an area $A_{\text{throw}} = \pi\, 10^6\,\text{m}^2$ about the detector, so we relate the number of events selected (N_{sel}) and the total number simulated (N_{sim}) to the effective area via

$$A_{\text{eff}}(E) = A_{\text{throw}} \frac{N_{\text{sel}}(E)}{N_{\text{sim}}(E)}\,. \tag{3.5}$$

The impact of the cuts discussed above on the effective area is shown in figure 3.12. We see that for $E \gtrsim 30$ TeV, A_{eff} approaches the detector area, A_{HAWC}, indicating that the cuts will select all showers of these energies that land on or very near to the array.

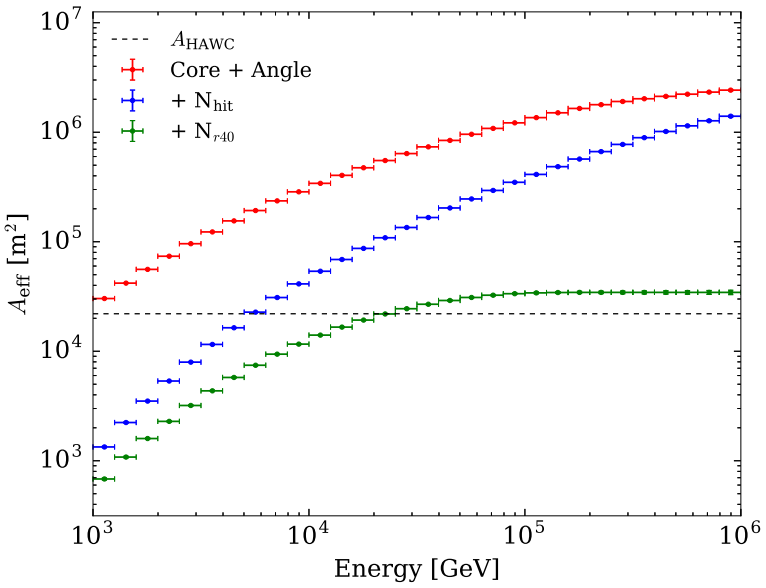

Figure 3.12: The effective area of HAWC to cosmic rays as a function of energy for the basic selection criteria. In addition to the listed cuts, particles within 60° of zenith were included. Including all cuts the A_{eff} assumes a value ∼50% larger than the instrumented detector area ($A_{\text{HAWC}} = 22{,}000\,\text{m}^2$) for energies above ∼30 TeV. This is expected as the N_{r40} cut does not perfectly contain cosmic ray showers to the array, allowing reconstructed cores to land within approximately 15 m of the edges.

Chapter 4

Energy Estimation

This study concerns estimating primary energies from hadronic primaries, which produce showers subject to significant fluctuations during development. Yet the number, type, and lateral distribution of secondary shower particles is intimately related to the primary's atomic mass and energy. Unfortunately, the nature of the WCDs in HAWC does not permit the absolute identification of individual secondary particles. However, since the PMTs detect Cherenkov light from them, each PMT's signal amplitude and distance from the measured core location provide information towards estimating primary particle energies. Of course, as illustrated with the Heitler and NKG approximations, any shower model is subject to assumptions and inherent uncertainties. With this in mind, we describe here a likelihood energy estimator based on simulation that uses the lateral distribution of measured PMT signals as a function of the primary energy. We begin by introducing the methods to build and evaluate the estimator. This includes the simulation for modeling shower development, the form of the model itself, and the maximum likelihood methods used to provide the energy estimate. At the end of this chapter, we summarize the performance of the technique in preparation for use with HAWC data.

4.1 Probability Tables

Recalling from section 1.2, extensive air showers are subject to fluctuations during development. While clearly influencing the particle content and distribution at observation level, we can encapsulate this information in the form of a multivariate probability function. Since different species' shower development varies for a given energy, we consider a single primary particle type. Here we use proton-initiated air shower simulation to build a four-dimensional probability table with bins in zenith angle (θ), primary energy (E), PMT distance from the core (R), and measured PMT signal (Q).

For a simulated proton shower generated with E and θ that triggers the HAWC detector, each PMT contributes a value of 1 to a respective Q, R, θ, E bin in the unnormalized table. This includes PMTs that do not register a hit, as the absence of a signal also provides shower information. The PMT-to-core distance is measured in the shower plane and is called the lateral distance. Since θ, ϕ define an arrival plane, and thus a normal vector $\boldsymbol{n} = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$, the lateral distance is found via a projection from the detector onto the shower plane:

$$R = ||\boldsymbol{x}_d - (\boldsymbol{x}_d \cdot \boldsymbol{n})\boldsymbol{n}||\,, \tag{4.1}$$

where $\boldsymbol{x}_d = \boldsymbol{x}_{\text{PMT}} - \boldsymbol{x}_{\text{core}}$ is the PMT-to-core vector in the detector plane. Exhausting the simulated data set, the table then is normalized along the Q-axis defining the probability density function $\mathscr{f}(Q, R \,|\, \theta, E)$, with the normalization condition $\sum_q \mathscr{f}(q, R \,|\, \theta, E)\,\Delta q = 1 \;\forall\; R, \theta, E$ and $q \in Q$. Hence, $\mathscr{f}(q, r \,|\, \theta, E)$ gives the probability that a primary proton with energy E and zenith angle θ has a PMT register a charge q at a distance r from the shower core. The final form the proton energy table takes is as follows:

- Three zenith bins

  $\theta_0 : 0.957 \leq \cos\theta \leq 1$

  $\theta_1 : 0.817 \leq \cos\theta < 0.957$

  $\theta_2 : 0.5 \leq \cos\theta < 0.817$

- Forty-four energy bins from 70 GeV – 1.4 PeV with bin width 0.1 in $\log E$
- Seventy bins in lateral distance from 0 – 350 m in bins of width 5 m
- Forty bins in charge from 1 – 10^6 photoelectrons (PE) in steps of 0.15 in $\log Q$

The zenith and energy bin widths were defined so as to divide up the simulation into equal statistics bins. The spacing between WCD centers is ∼10 m, so the lateral distance bin widths were assigned half this value. Finally the charge bin spacing was chosen to be of the same order as the estimated PMT charge resolution (∼ 30%). A similar gamma table was created, with the change to forty bins in energy up to 500 TeV spaced 0.1 in $\log E$ as above.

4.1.1 Simulation

The simulated set used to build the tables was generated with identical settings as the nominal simulation, save for the primary spectrum which follows $dN(E)/dE = A\,E^{-1}$. Being harder than the nominal E^{-2} spectrum, this allowed enhanced sampling at the highest energies, ensuring the table was thoroughly filled. Indeed, this spectrum is flat in $\log E$ (figure 4.1), so an equal number of proton showers were generated for each energy bin. This also can be seen by changing variables in the differential spectrum equation:

$$\begin{aligned}\frac{dN(E)}{dE} &= \frac{dN(E)}{d\,\log E}\frac{d\,\log E}{dE} \\ &= \frac{dN(E)}{d\,\log E}\frac{1}{E\ln 10}\,.\end{aligned}$$

Rearranging, we see that the energy terms cancel out:

$$\begin{aligned}\frac{dN(E)}{d\,\log E} &= \frac{dN(E)}{dE}\,E\,\ln 10 \\ &= A\ln 10\,,\end{aligned}$$

i.e. a constant value per unit $\log E$. The angular distribution was thrown on a cosine distribution, and the bin edges were chosen so as to equalize the statistics as well. A total of 5×10^6 proton showers from 5 GeV to 2 PeV were generated in this set to populate the table.

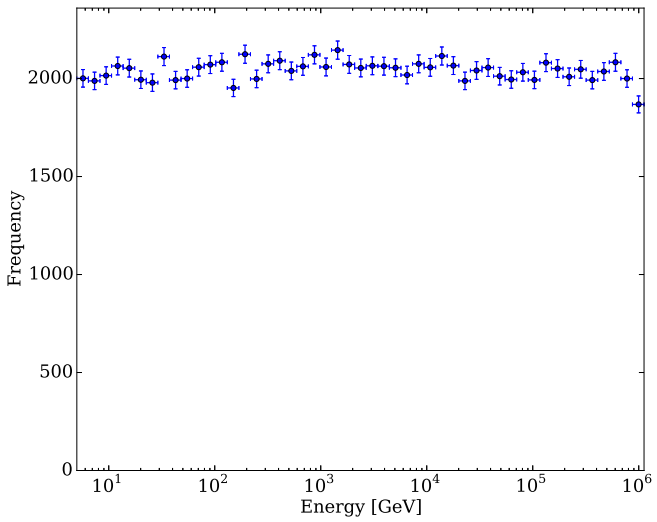

Figure 4.1: Energy distribution from simulation used to build the proton probability table. There are nearly equal numbers of events in each bin from using an E^{-1} differential spectrum and binning in $\log E$.

4.1.2 Table Smoothing

Numerical variations in the tables arise from the statistical nature of producing the simulation as well as from averaging over table bins for normalization. Adjusting the bin sizes merely exacerbates these inherent binning artifacts. Thus the tables are smoothed with a multi-dimensional spline function to ensure that bin-to-bin fluctuations from simulation do not influence the energy estimation. By spline smoothing we also preserve gradient information along several axes that is not available to us for lack of an explicit analytic expression for $\mathscr{E}$. This is paramount to retaining smoothness that has been distorted by the process of discretization. We employ the spline fitting package Photospline [69], which has demonstrated its versatility in smoothing binned probability tables with even higher dimensionality.

The Photospline package uses basis-splines (B-splines) to interpolate between data points from a multi-dimensional surface defined on a grid. A B-spline is a smooth function of order n, where the form of the function approaches a Gaussian as n increases. The progression of the first three basis functions is shown in figure 4.2. In this work, the 3rd order B-spline was implemented, per the recommendation of [69] for both efficiency of the fitting procedure and accuracy of the results. The spline routine solves the best fit to weighted data using a linear combination of B-splines defined on a grid of predefined *knots*, or control points.

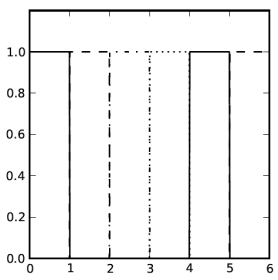

(a) Order-0 B-splines

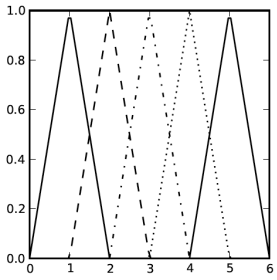

(b) Order-1 B-splines

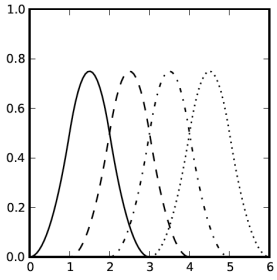

(c) Order-2 B-splines

Figure 4.2: First three basis functions (normalized) used in splines [69]. As n increases in value, the shape of the order-n B-spline approaches a Gaussian. The energy tables were fit using the order-3 B-spline.

The weights applied for the smoothing procedure are the inverse square Monte Carlo uncertainties; hence, filled bins with large uncertainties are assigned small weights and vice versa. Empty bins are assigned zero weight so as not to influence the fit. The fit is performed in the logarithm of the likelihood, as $\log \mathscr{L}$ appears less peaked than $\mathscr{L}$, which also ensures gentle transitions amongst neighboring bins in all dimensions. Overall, it is simply much easier to constrain the fit in $\log \mathscr{L}$. The entire four-dimensional table is not subjected to a single spline fit. Instead, each zenith sub-table is fit independently, making a set of three three-dimensional fits. Figure 4.3 shows two θ, E, R bin-level comparisons of the normalized table before and after spline smoothing, followed by figure 4.4 which depicts a two-dimensional R, Q section at 100 TeV.

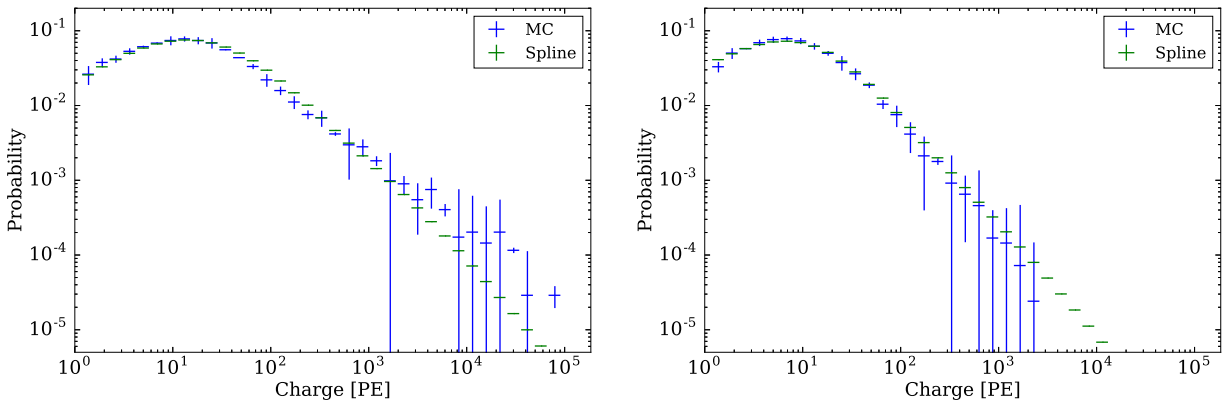

Figure 4.3: Examples of smoothed 1D sections of the MC proton table. The figure on the left corresponds to the θ bin nearest zenith at 71 TeV and 12.5 m (bin centers). The right panel is from the middle zenith bin at 560 TeV and 52.5 m. These slices were taken from the full 3D fit, not fit individually.

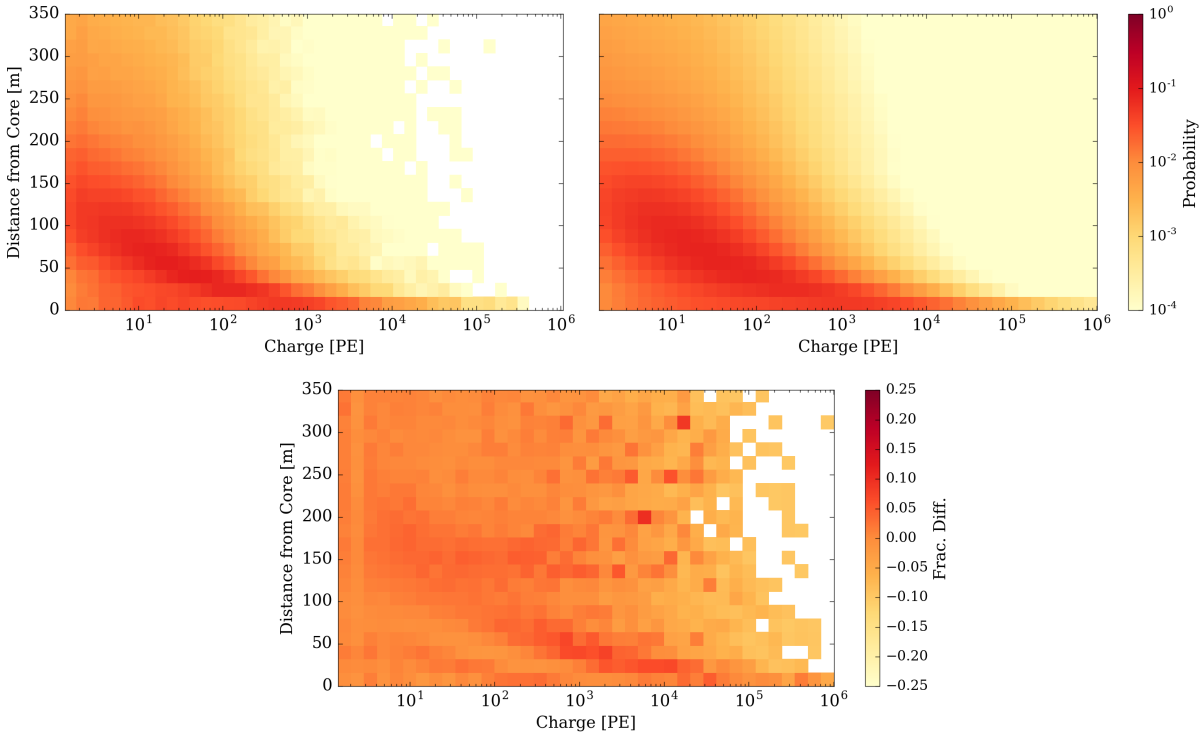

Figure 4.4: Comparison of normalized MC table (left) to spline fit (right) for the energy bin corresponding to 100 TeV. The fractional difference $(\mathscr{f}_{\mathrm{MC}} - \mathscr{f}_{\mathrm{spline}})/\mathscr{f}_{\mathrm{MC}}$ is shown below, and does not exceed more than 10% percent where the MC table is well populated. Note that the spline fit fills in regions with little or no MC simulation. In this particular section, neighboring energy bins that are more populated contribute to smoothing these sparse regions. Large charges are most probable very near the core, and become less probable with increasing distance. Conversely, there is a greater probability for a small charge to be farther from the core in this energy bin.

4.2 Maximum Likelihood Estimation

To use the probability tables for energy estimation, we employ the method of maximum likelihood. Consider a sample of n observations $\{x_i | i \in [1, n]\}$ drawn from a probability density function $f(x|\omega)$, where the set of parameters $\omega \in \Omega$ fully describes the function defined on $x \in X$. For independent and identically distributed x_i, the joint density function for observing the sample is simply the product of f for all observations:

$$f(x_1, x_2, ..., x_n|\omega) = f(x_1|\omega) \times f(x_2|\omega) \times ... \times f(x_n|\omega) = \prod_{i=1}^{n} f(x_i|\omega)\,. \tag{4.2}$$

Here the x_i are considered random variables distributed by the form of the parent distribution $f(x|\omega)$.

We can reverse this perspective and consider the observations to be held while the parameters ω are allowed to vary. Thus the x_i become fixed parameters and the ω are variables, forming the likelihood function $\mathscr{L}$:

$$\mathscr{L}(\omega; x_1, x_2, ..., x_n) = \prod_{i=1}^{n} f(x_i|\omega)\,. \tag{4.3}$$

We see that for a given set of x_i and ω, the values of $\mathscr{L}(\omega; x_1, x_2, ..., x_n)$ and $f(x_1, x_2, ..., x_n|\omega)$ are the same; however, $\mathscr{L}$ is now viewed as a means to identify the ω that best describe the x_i. This is accomplished by finding the value of $\omega_{\rm mle}$ which maximizes $\mathscr{L}$, such that $\mathscr{L}_{\rm max} = \mathscr{L}(\omega_{\rm mle})$, and is defined as the maximum likelihood estimate of ω:

$$\omega_{\rm mle} \subseteq \arg\max_{\omega_{\rm mle}} \mathscr{L}(\omega; x_1, x_2, ..., x_n)\,. \tag{4.4}$$

If $\mathscr{L}$ (and thus f) is analytic and differentiable, then $\omega_{\rm mle}$ is found by

$$\frac{\partial \mathscr{L}}{\partial \omega}\Big|_{\omega_{\rm mle}} = 0\,. \tag{4.5}$$

In practice, the joint likelihood is a rather small numerical value, so it is more convenient to evaluate $\log \mathscr{L}$:

$$\log \mathscr{L}(\omega; x_1, x_2, ..., x_n) = \sum_{i=1}^{n} \log f(x_i|\omega)\,,$$

and equation 4.5 still holds:

$$\frac{\partial \log \mathscr{L}}{\partial \omega}\Big|_{\omega_{\rm mle}} = 0\,. \tag{4.6}$$

The uncertainty on $\omega_{\rm mle}$ is estimated by identifying the values of ω where $\log \mathscr{L}_{\rm max}$ decreases by -0.5.

For a discretized function such as the probability table used in this study, one simply samples the function for each ω under consideration, identifying $\omega_{\rm mle}$ as that which gives the largest likelihood value. Here, the fixed observations x_i are the PMT hit information (q_i, r_i) and the reconstructed zenith angle θ. In the context of energy estimation in HAWC, the energy is the only parameter for which we search a maximum likelihood estimate, and we make the notational change $f(x_i|\omega) \to \mathscr{f}(q_i, r_i, \theta|E)$. In principle, we can allow

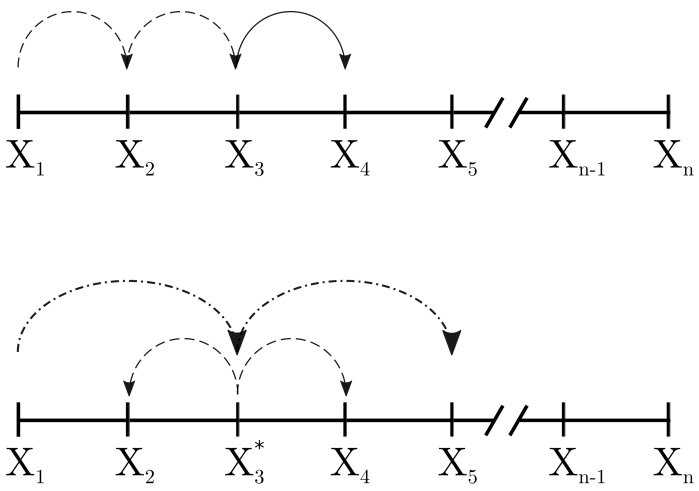

Figure 4.5: Illustration of grid search methods over a general state space spanned by X. *Top*: Simple grid search evaluating $\mathscr{L}$ for all allowable states. *Bottom*: Sparse grid search with stride $s = 1$. Here, X_3 was identified as having the $\mathscr{L}_{\max}$, and subsequently the $s = 1$ neighbors are evaluated.

the core location and zenith angle to vary as well; however, in practice the re-evaluation of the $\sim$ 1200 PMT-core distances in the shower plane becomes a limitation to runtime efficiency.

Since the likelihood function considered here is not analytic, searching the space for $\mathscr{L}_{\max}$ must be performed in an efficient manner. The next two sections present two main classes of maximum likelihood finding methods, the grid search and Markov chain Monte Carlo (MCMC), as well as several of their implementations made available for HAWC energy reconstruction.

4.2.1 Grid Searches

The simple grid search method is the most elementary maximum likelihood finder implemented. For each element in a set of predefined locations in state space, a likelihood value is calculated. Once all bins have been reviewed, the bin corresponding to the maximal likelihood value is identified and chosen as the best estimate. The top panel of figure 4.5 depicts a simple grid search through a general state space $\omega \in \Omega$.

The sparse grid search method introduces a stride, s, to the simple grid search. Instead of checking each bin, every other s-bin's likelihood is evaluated. Once the maximal value and its corresponding bin are identified for this subset, the neighboring s bins are checked as well. This method is illustrated in the bottom panel of figure 4.5. Using this method requires the assumption that the likelihood space is unimodal and strongly peaked. This means that there are no multiple peak structures and the uncertainty in the likelihood estimate is less than the bin width. This is shown to be the case using the energy estimation table $\mathscr{E}$, as demonstrated in the two example simple grid searches in figure 4.10. Since a sparse grid search with stride

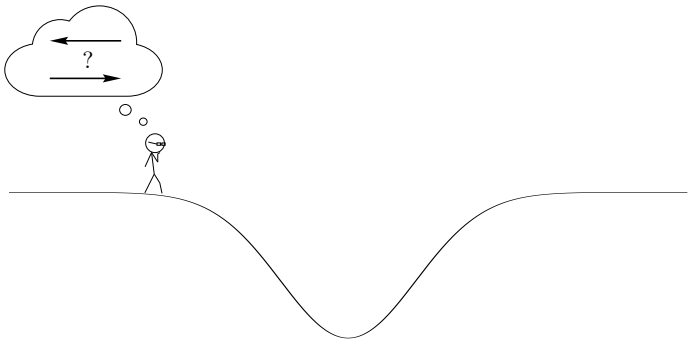

Figure 4.6: Depiction of a random walk on a curved space.

$s = 1$ halves the number of required evaluations, it is the default maximum likelihood energy estimation method used in the reconstruction production.

4.2.2 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) describes a class of stochastic methods to sample probability distributions in a manner satisfying the Markov property, in which sequences exhibit serial dependence. Using MCMC can be an invaluable tool when sampling complicated probability spaces, such as multimodal, highly anisotropic, or non-analytic posterior distributions. In this section, I first present a simple real-world analogy for how MCMC sampling works. In the following subsections, specific MCMC implementations of increasing complexity are discussed, with each method building on concepts from the previous one.

4.2.2.1 MCMC Example

Consider a sidewalk that takes the shape of an inverted normal distribution (figure 4.6), and a walker who is about to embark on a journey on this very sidewalk. So long as the sidewalk is not too steep, i.e. the walker can stay upright, he or she can stumble forward and backward on the path in single steps. If the walker were to begin up at the top, then because of gravity, it would seem easier or more favorable to take a step farther down into this well than to climb out of it. Eventually, the walker would end up at the lowest point. If many attempts are made to climb out in either direction, the walker may become fatigued, coming to a final resting position at the bottom of this strange route.

Here, the distribution $\pi(x|D)$ that is sampled takes the form of the sidewalk, where x denotes the position and D all parameters that describe the shape of the path. Hence, we can later make the association $\pi \to \mathscr{L}$. There is a second distribution, the stepping function $\eta(x|x + \delta x)$, that denotes the probability to go to position $x + \delta x$, given that the algorithm is currently at position x. In this simple example, there is an equal chance of choosing the direction of the step, either forward or backward, so $\eta(x|x \pm \delta x) = 0.5$. Generally

η can take the form of any distribution, however, η typically is chosen to be symmetric with respect to time, i.e. $\eta(x|x+\delta x) = \eta(x+\delta x|x)$. This is a desirable characteristic of stepping functions so as to avoid introducing bias when sampling functions.

Yet despite this symmetry, actually making the step is in fact a weighted decision, as the shape of the path influences whether the proposed step is made. It is also clear that each step in time t is only influenced by the previous step $t-1$, since nothing about the location of step $t-2$ is carried forward. Thus, the only information that is necessary to take another step comes from the current location. The dependence on only the previous step is a property of Markov Chains, and the stochastic nature of the sampling step is known as Monte Carlo sampling, hence the combination MCMC. Eventually, after making many steps using a symmetric η, sampling approaches the region of maximal probability so MCMC can be used as a maximum likelihood finding method. When considering probability distributions, the well from the prior example is simply inverted and now the larger probabilities dictate the more probable stepping directions.

4.2.2.2 Metropolis-Hastings Algorithm

The Metropolis-Hastings (MH) algorithm [70, 71] is one of the most commonly used MCMC methods. It relies on a simple accept-reject evaluation by comparing a random variate to the ratio of the probabilities at the current and proposed step locations. If the ratio is greater than the randomly generated value, then the step is taken, otherwise, a new sampling location is tested. The proposal probability is given by

$$p(X^t \to Y) = \min\left(1, \frac{\pi(Y|\omega)}{\pi(X^t,\omega)}\frac{\eta(Y|X^t)}{\eta(X^t|Y)}\right) \tag{4.7}$$

where X^t is the location in state space at time step t, Y is the proposal location, $\pi(X^t|\omega)$ is the probability function evaluated at X^t with parameters ω, and $p(X^t|Y)$ is the transition probability function to step from $X^t \to Y$. The proposed step to Y is accepted if $r < p(X^t \to Y)$ for random variate $r \in \mathbb{R} \sim [0,1]$, and is otherwise rejected. Thus,

$$X^{t+1} = \begin{cases} Y & \text{if } p(X^t|Y) > r \\ X^t & \text{otherwise.} \end{cases} \tag{4.8}$$

Since r is a random variate, this condition exhibits the random sampling nature of MCMC. If $p = 1$, then clearly the step is favored and it is always taken. However, if $p < 1$, then the step is permitted only in a random manner, dependent on the outcome of generating r. An simple illustration of the MH step is shown in figure 4.7.

Choosing a step-symmetric η, we see that p is dependent only on the shape of the distribution π. A typical symmetric stepping function is the normal distribution, i.e.

$$Y \leftarrow X^t + Z\,,$$

where $Z \sim N_m(0, \Sigma)$ and m is the dimensionality of the space. Since drawing the same random variate Z from position Y to X^t has the same probability, $\eta(Y|X^t) = \eta(X^t|Y)$. Yet, one must define a value for each $\sigma \in \Sigma$, as the manner in which π is sampled is dictated by the form of η, and will determine the direction and distance to step from the current location in state space, X^t.

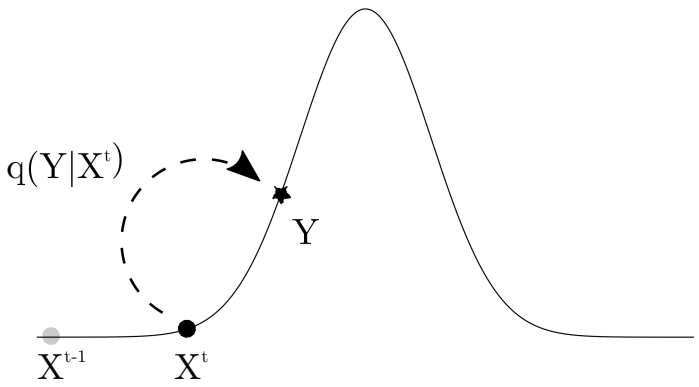

Figure 4.7: Metropolis-Hastings MCMC on Gaussian PDF. The proposed step to Y is made from the current position X^t. The previous location, X^{t-1} does not participate in or influence the current step proposal. Since $\pi(Y) > \pi(X^t)$, this step will be accepted.

4.2.2.3 Affine Invariant Samplers

So far the nature of the probability space defined by π has not been addressed. Indeed, a general probability surface may be harshly peaked or consist of multiple high probability regions. This becomes a major challenge certainly for the grid search methods, and even for the MH algorithm, as peaks can escape sampling via a poor choice of step size. A simple example is a tight two-dimensional Gaussian distribution, where $\sigma_y << \sigma_x$. Further generalizing to an m-dimensional Gaussian with various σ_i, the enhanced complexity exacerbates the difficulty in choosing appropriate step sizes and perhaps even step functions for each dimension. To this end, we describe two MCMC methods from [72, 73] which sample the state space in a manner insensitive to π's structure relative to x, i.e. affine invariant sampling.

Both methods involve ensemble sampling, where k *walkers* evaluate π at their respective locations in state space, and each relies upon a proposal function taking the following form:

$$\eta(z) \propto \begin{cases} \frac{1}{\sqrt{z}} & \text{if } z \in \left[\frac{1}{a}, a\right] \\ 0 & \text{otherwise,} \end{cases} \tag{4.9}$$

where the parameter $a > 1$ is a scaling factor that can be adjusted for performance. Over the range $[1/a, a]$, the proposal function can be shown to be step symmetric, $\eta(1/z) = z\eta(z)$, so again p is independent of η. Furthermore, the nature of the sampling range rescales the proposal step size in a manner that effectively

flattens the probability function, making it affine invariant. Additionally, sampling a random variate from equation 4.9 is simple:

$$z = \frac{a}{1+a}\left(-1 + 2u + au^2\right) \tag{4.10}$$

where $u \in \mathbb{R} \sim [0,1]$ is a uniform variate. Upon inspection, $\eta(z)$ does not enclose the dimensionality m of the space X. Instead of generating m random variates from η for each walker, the ensemble methods use the locations of other walkers to make a step. Clearly, in order to explore the full dimensionality of the state space, $k > m$.

The simplest affine invariant ensemble method is the *stretch* move. For a set of k walkers, their locations at time t are denoted by X_i^t, $i \in k$. To generate a proposal step for walker j, first a walker from the set $[k] = \{i \in k : i \neq j\}$ is chosen at random. Then the unit vector defined by $X_i - X_j$ determines the direction along which the step will occur, with magnitude given by z. This is then repeated for each walker, thus,

$$Y_j \leftarrow X_j + z(X_i - X_j)\ \forall j\,, \tag{4.11}$$

and are accepted/rejected according to equation 4.8. Figure 4.8 illustrates the stretch move.

The second ensemble method incorporates a subset $s \in S \subset k$ and $j \notin S$ of walkers for each step. The mean of the X_i for $i \in S$ is used to define the direction vector with respect to the walkers in S:

$$W = \sum_{i \in S} z_i (X_i - \bar{X}_S) \tag{4.12}$$

$$Y_j \leftarrow X_j + W\ \forall j\,,$$

where $\bar{X}_S = \frac{1}{|S|}\sum\limits_{i \in S} X_i$ and the z_i are random variates drawn from equation 4.10. The case where $|S| = 1$ is equivalent to the stretch move above. Figure 4.9 illustrates the ensemble walk move.

The major advantage of using ensemble MCMC methods is the ability to sample more space by simply employing multiple walkers, and by permitting the walkers to share state space information (via equations 4.11 and 4.12). They effectively provide random nudges from each other in search of the most probable regions of π. Indeed, the walkers gradually approach one another with each iteration, while occasionally sampling π at other random locations. Finally, the complexity of π is scaled away by the affine invariant step function, flattening the state space while preserving both the Markov property and the Monte Carlo sampling. Of course, these methods are well suited for high dimensional problems in which grid or MH samplers are efficiency limited. For energy estimation in HAWC, the grid searches are sufficient, but in future studies, MCMC may be used as an all-purpose angle, core, energy fitting routine.

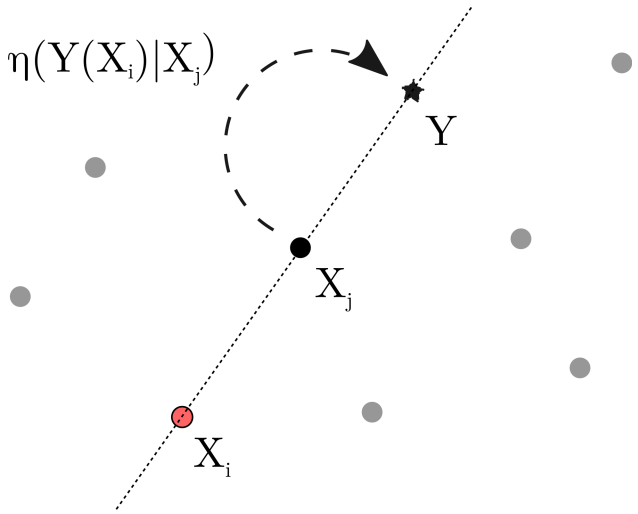

Figure 4.8: The affine invariant stretch move involving an ensemble of walkers. Each walker is denoted by a circle, while the proposal location Y for walker X_j is denoted by a star. This MCMC stepper, like the Metropolis-Hastings sampler, is independent of walkers's histories. However, each step proposal is drawn from the vector connecting the chosen walker X_j and another walker X_i chosen at random. Note the emphasized dependence of Y on X_i.

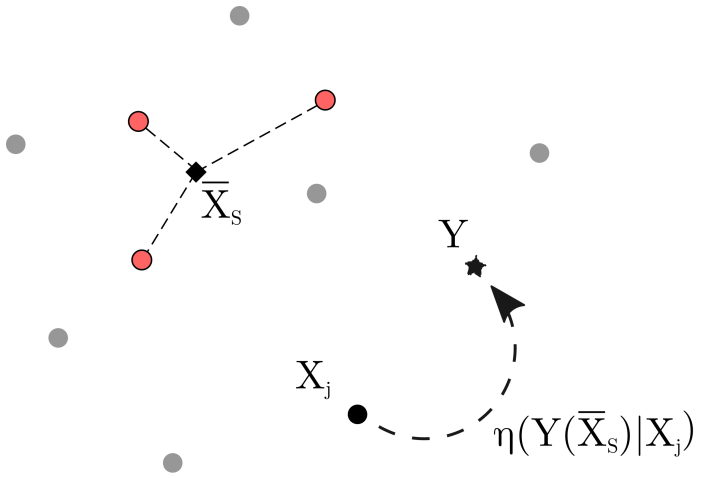

Figure 4.9: Affine invariant walk move involving an ensemble of walkers. Instead of randomly choosing another walker, this sampler proposes a step for walker X_j based on the mean $\bar{X}_S$ from a subset of walkers' current locations, denoted by the red circles. The direction of the walk proposal is chosen based on random multiples of the vectors connecting the red dots to $\bar{X}_S$.

4.3 Performance

We quantify the performance of the likelihood energy estimator via three values: the mean likelihood uncertainty, and the overall estimation bias and energy resolution. The first quantity is based on the shape of the likelihood function for individual events, representing the inherent resolution of the likelihood estimation table. In principle, it can be improved by evaluating a more finely binned table. The remaining two are evaluated using a simulated data set, and quantify the overall effects on the estimator due to modeled shower development and the reconstruction procedure. Thus, these two values can be improved only up to the limits of shower fluctuations and the core and angular resolutions from data selection cuts.

4.3.1 Event-by-Event Uncertainty

On an event-by-event basis, the likelihood values for the energy bins are calculated, and the energy E^* corresponding to $\mathscr{L}_{\max}$ is assigned as the best energy estimate, $E_{\text{reco}} = E^*$. Figure 4.10 shows examples of likelihood evaluations for two shower events. As mentioned at the beginning of section 4.2, the uncertainty σ^* on E_{reco} is estimated by measuring the width of the log-likelihood function where it has dropped by 0.5. For the events in the figure 4.10, the single peak centered on the $\mathscr{L}_{\max}$ is sharp enough such that σ^* is less than half of an energy bin width. This is true for the vast majority (99.99%) of events passing the basic multiplicity and core resolution analysis cuts defined in section 3.4. Since $\sigma^* < 0.05$ in $\log E$, the sparse grid search with stride $s = 1$ is sufficient to identify the maximum likelihood along the energy axis, and reduces the number of required table evaluations by nearly one-half.

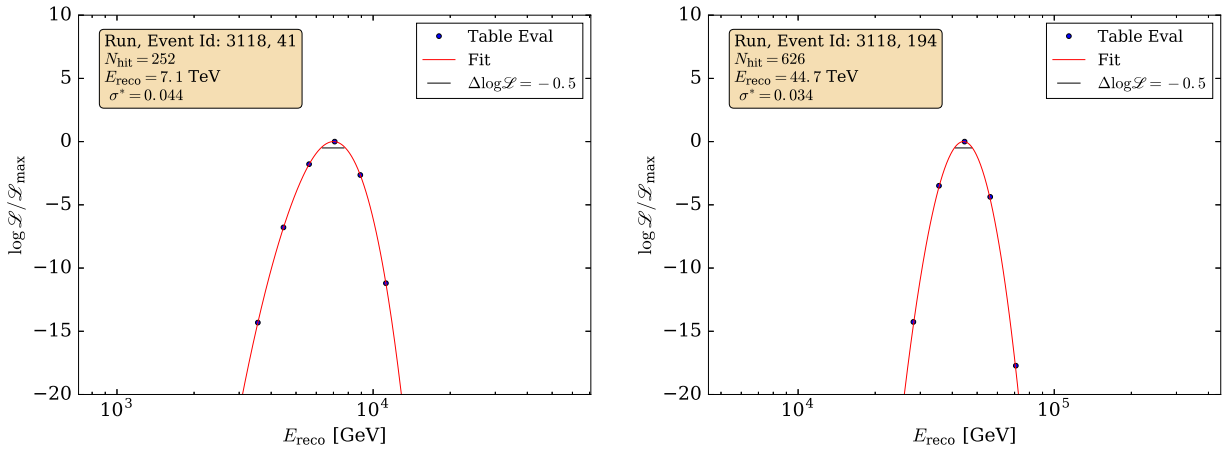

Figure 4.10: Example table evaluation for two events. The curve represents a spline fit to the table evaluation points, used to estimate σ^* shown by the black horizontal line. The likelihood function is clearly peaked less than a bin width, permitting the use of the sparse grid search method.

4.3.2 Bias and Resolution

Using air shower simulation, we define the energy bias and resolution from the mean and width of the difference distribution between the logarithms of the reconstructed and true energy values. Figure 4.11 below shows an example of the bias distribution in a single E_{true} bin, where

$$\text{bias} = \log E_{\text{reco}} - \log E_{\text{true}} \,. \tag{4.13}$$

Simple manipulation of the values converts the width into the more familiar fractional or percent resolution, $\sigma_{\%} = 10^{\sigma_{\log E}} - 1$. In addition to the mean (bias or offset) and width (resolution), we also identify the integral of the distribution as the efficiency $\epsilon(E)$ of that true energy bin. This is the fraction of simulated events that are reconstructed and pass quality cuts, and will be used when measuring energy spectra.

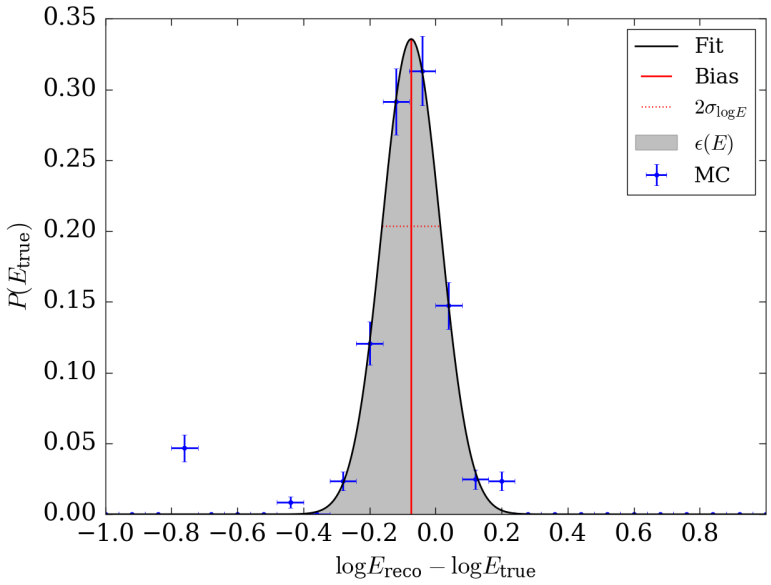

Figure 4.11: Diagram with definitions for the energy estimation bias and resolution. The MC data points come from proton showers in the θ_0 bin centered at $E_{\text{true}} = 6.3$ TeV. A Gaussian fit is shown by the black curve, with the resulting mean and width (red lines). These values are used as the quoted energy bias and resolution, respectively. The distribution has been normalized to the total number of events thrown in this energy bin, so the integral gives the reconstruction efficiency $\epsilon(E)$ of this bin.

Since only proton simulation went into building the probability table, a first test of the bias and resolution was done using solely proton showers. As shown in the left panel of figure 4.12, the resulting energy bias for the first zenith bin is strongly dependent on the event quality. Specifically, the most dramatic drop in bias comes from improving the core resolution via the N_{r40} selection cut. This is expected as the table is

built and evaluated using lateral distances from the core location. Including all basic event selection criteria, the bias approaches 0 as expected for estimating proton energies with proton probability tables. A closer look at the proton bias in figure 4.13 reveals a drop to about a bin width, 0.1 in $\log E$, for energies above 3 TeV. This represents an energy reconstruction threshold below which events that just trigger the detector are promoted upwards in E_{reco}. The bias is necessarily negative, as showers that fluctuate downward are less likely to pass selection criteria.

The same behavior is observed considering all simulated species (right panel of figure 4.12), though the bias is slightly lower as E_{reco} tends to overestimate the energies of heavier species. This makes sense as the superposition principle from section 1.2 dictates that the muon content scales as $A^{0.15}N_{\mu}^{p}$ for species of mass number A. Near the core, there is a higher density of charged particles than for a primary proton of similar energy, so the energy estimate using a proton hypothesis will be overestimated.

The width of the bias distribution gives the energy resolution, shown in the right panel of figure 4.13. For both proton and all-particle data sets, the bump around 8 TeV is caused by the detector response approaching full efficiency. Below this energy there are fewer and fewer showers passing the selection cuts, so the offset distribution is thin but sparsely populated. This will be explored more when constructing the effective area necessary for spectral unfolding in chapter 8.

Finally, the all-particle biases and resolutions for events falling into the three zenith bins are shown in figure 4.14, where the basic selection criteria are applied. The first two θ bins are very similar in both quantities, with a slight shift in the resolution peak to 10 TeV for θ_1. This is much more pronounced for θ_2, including an elevated bias at all energies. The main culprit is precisely why the lower zenith bins exhibit a peak just prior to 10 TeV: the detector efficiency is rather low for the highest zenith angles, climbing until leveling out above 100 TeV. Showers must traverse increasing amounts of atomspheric overburden with larger zenith angles, so the energy threshold rapidly shifts to higher energies. The reduced performance is also exacerbated by the large bin size of θ_2, which covers $35^{\circ} \leq \theta \leq 60^{\circ}$, and thus there is a significantly wider energy distribution of showers falling into this bin.

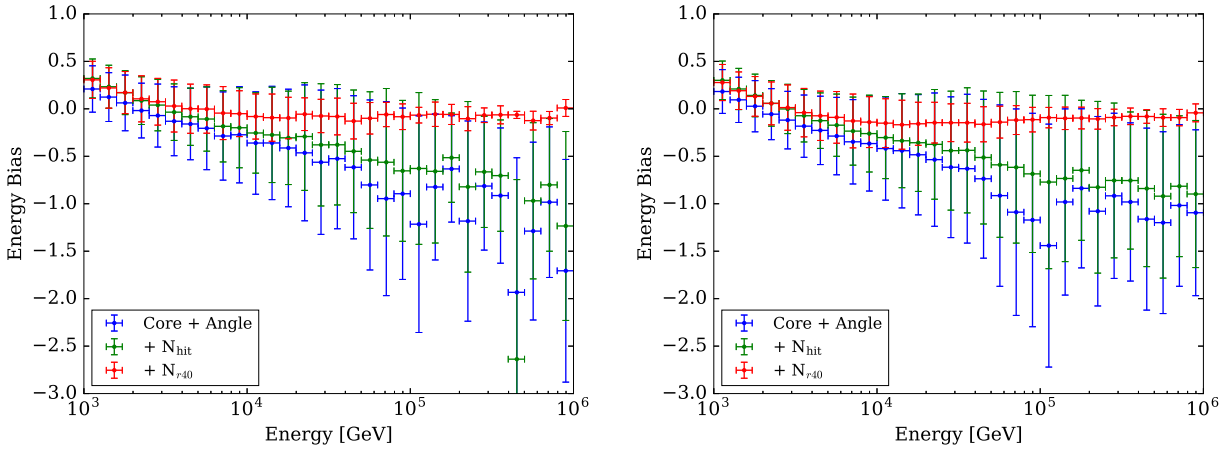

Figure 4.12: Energy bias for proton only (left) and all-particle (using CREAM spectrum, right) energy estimation in the first zenith bin, θ_0. Each figure shows the reduction of bias with improved event selection. The proton only sample exhibits less bias, which is expected since the energy tables are built from proton simulation. The vertical uncertainties are the energy resolution, $\sigma_{\log E}$, shown in figure 4.13

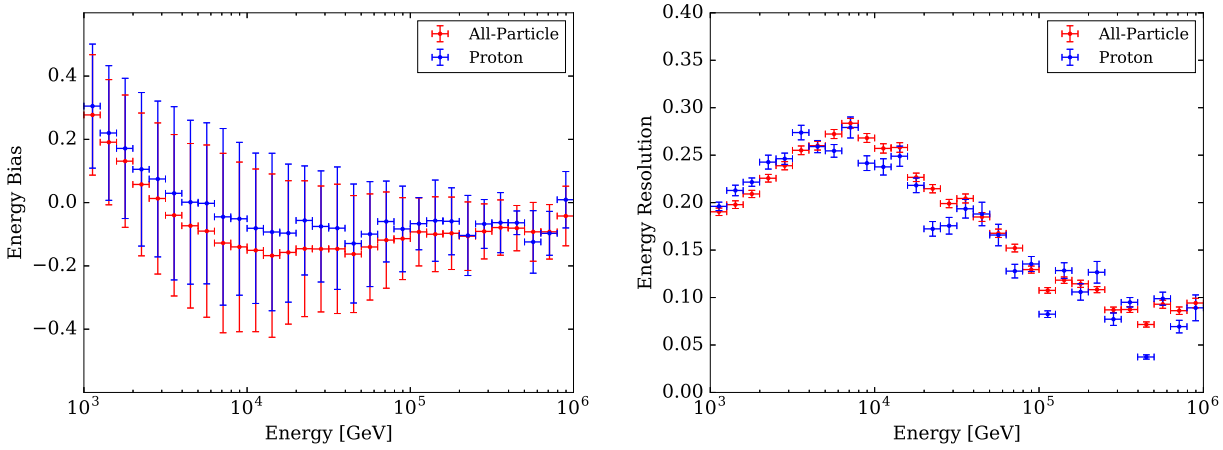

Figure 4.13: The left panel shows the energy bias after event selection for both the all-particle and proton simulations for θ_0. These are also shown by the red curves in figure 4.12. The right panel shows the respective energy resolution, $\sigma_{\log E}$, for each.

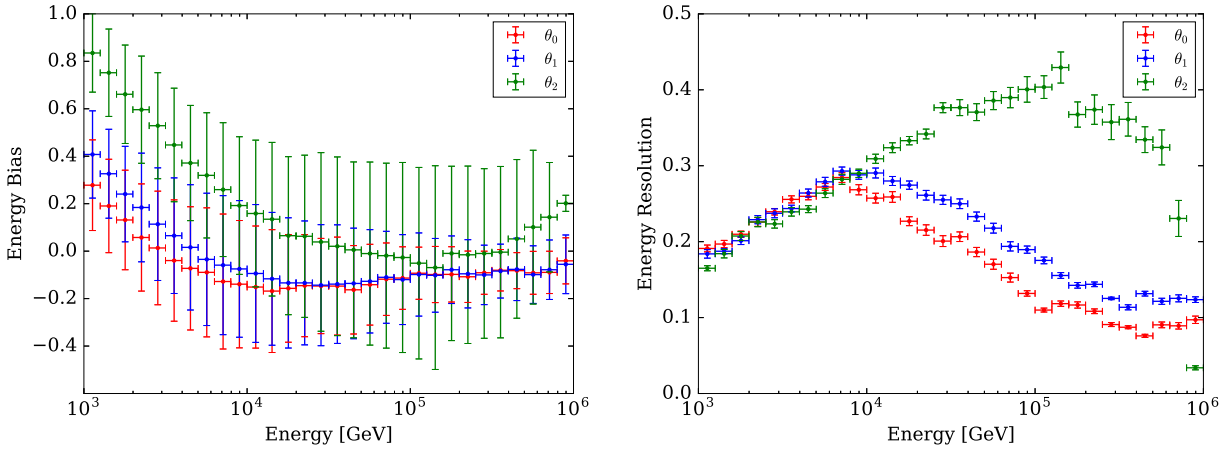

Figure 4.14: The all-particle energy bias (left) and resolution (right) after event selection for the three zenith bins defining the energy tables. The peaks in both plots occur at the same energy for each respective θ bin, and are characteristic of the rise and stabilization of the detector efficiency with respect to energy.

Chapter 5

Sky Maps

The following and final analysis chapters rely on the power of visualizing and quantifying cosmic ray arrival information from the visible sky. Presented here is a review of the methodology used for generating event sky maps. The essentials necessary for this thesis are summarized here, with references to in-depth descriptions and derivations where appropriate.

5.1 Binning the Overhead Sky

As described in chapters 3 and 4, reconstructed cosmic ray air shower events possess a local arrival direction defined by the pair $(\phi,\ \theta)$. Including the Greenwich Mean Sidereal Time (GMST) of an event, we can assign an arrival direction relative to the fixed stars, using the equatorial coordinate system. The quantities defining this system are the right ascension (α) and the declination (δ), as shown in figure 5.1. Using the HEALPix [74] spherical pixelation scheme, the α and δ are binned as equal-area pixels, providing a method to project the numbers of observed events onto the visible celestial sky, i.e. the data map. In a HEALPix map, the unit sphere is divided into twelve equal-area tessellations, each divided into an $N_{\text{side}} \times N_{\text{side}}$ grid, giving a total of $12N_{\text{side}}^2$ pixels. For the maps in this thesis, N_{side} was chosen to be 512 so each pixel occupies an angular size of 0.11 sr. As will be shown, upon normalizing the observed data map to the expected background map, we can make further inferences regarding the underlying cosmic ray arrival distributions.

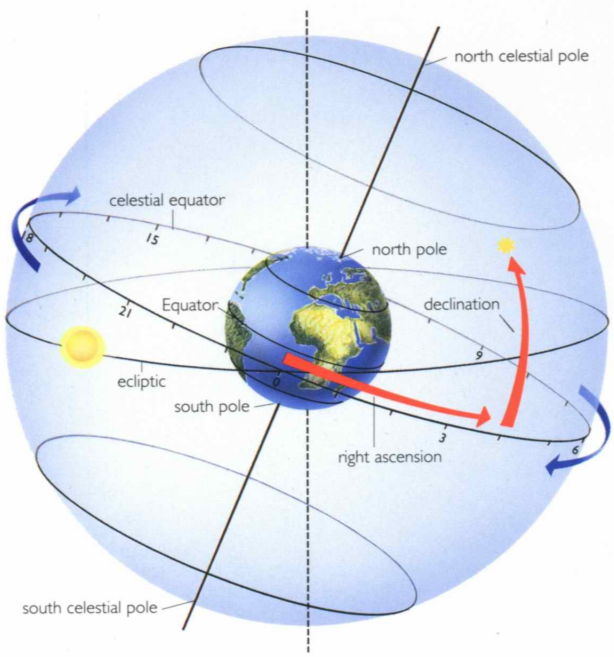

Figure 5.1: Diagram of the equatorial coordinate system. The defining quantities are the spherical coordinates in the fixed-sky system, represented by right ascension and declination. Equatorial coordinates are aligned with Earth's rotational axis, with the Equator serving as the origin for declination. Thus, the northern and southern poles are where $\delta = 90°$ and $\delta = -90°$, respectively. The right ascension increases opposite the rotation of the Earth, and its origin is defined as the vernal equinox, where the Sun crosses the celestial equator. Source: MEMIM Encyclopedia.

5.2 Background Estimation: Direct Integration

In order to identify fluctuations in the nearly isotropic flux of cosmic rays, the data map must be compared to what is expected for a given location in the sky, or a reference map. The HAWC detector's large ∼2 sr coverage of the sky, the variation of data taking conditions, and the observed [75] per-mille level of anisotropy prohibits an exhaustive modeling of the expected background based on simulation. A standard solution to construct a reference map is to estimate it using the data itself. In this study, we accomplish this by the method of Direct Integration (DI) [75].

The DI reference map $\langle N(\alpha,\delta)\rangle$ is computed by the convolution of the local event arrival distribution, $\mathrm{E}_{\Delta t}(h,\delta)$, with the all-sky event rate $R_{\Delta t}(t)$, recorded in a predefined integration duration Δt. Thus, $\langle N\rangle$ will serve to estimate the background and is calculated via

$$\langle N(\alpha,\delta)\rangle = \sum_{\Delta t}\sum_{t=0}^{2\pi} R_{\Delta t}(t)\mathrm{E}_{\Delta t}(h,\delta)\ , \tag{5.1}$$

where the hour angle h, and the Greenwich Mean Sidereal Time t, are related to the right ascension by $h = t - \alpha$. Figure 5.2 shows an example of the DI method for background estimation and signal extraction.

Direct integration effectively averages the number of events within the $\Delta t \times 15^{\circ}\mathrm{hr}^{-1}$ right ascension range, so the reference and data maps are sensitive to structures smaller than this angular scale. Integration durations of two hours and four hours are used in this analysis, corresponding to 30° and 60°, respectively. Lastly, since the data are used to generate the reference map, the inherent arrival distribution and any variations in data taking are present in both maps, and thus normalize when the two are compared in search of signals.

5.3 Region of Interest Masking

The presence of strong, localized signals (strengths $> 1\%$ of the background) can influence the DI reference level. This can be mitigated by excluding the offending region of interest (ROI) during DI, which introduces a correction into calculation of 5.1. Details of the correction factor necessary to account for masked ROI pixels are presented in section 4.5.3 of [77]. Without ROI masking, a signal's observed significance is reduced and an artificial fluctuation will surround it. The extent of this effect on observing the Moon shadow is demonstrated in the next chapter.

5.4 Map Smoothing

To enhance the small signals expected on maps (which are typically more finely pixelated than the angular resolution of HAWC) the data and reference maps are smoothed. In this thesis, a simple 'top-hat' smoothing

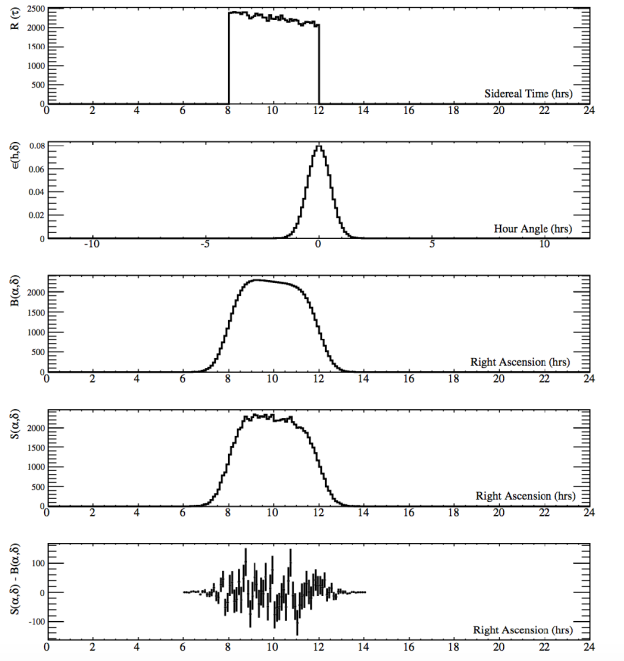

Figure 5.2: Direct integration illustration from [76]. The panels from top to bottom are: the all-sky event rate $R_{\Delta t}(t)$ as a function of sidereal time using $\Delta t = 2$ hr integration period; the *normalized* local arrival distribution $\epsilon_{\Delta t}(h, \delta)$; the estimated background $B = \langle N \rangle$; the raw data map S; the signal map $S - B$.

is applied, wherein the value of pixel i represents the sum of the surrounding pixels within a radius r. Ideally, the smoothing radius is of the same scale as the feature of interest, which correlates the binned event counts in a manner preserving that feature's characteristic angular extent while amplifying its signal. For example, as will be seen with the Moon shadow, a smoothing radius of 1° is sufficient to boost the observed significance 5-fold.

5.5 Relative Intensity and Significance

Armed with data and reference maps, the strength of underlying signals is measured using the relative intensity. This is an unbiased method of measuring the amplitude of deviations from the estimated expectation. The differential relative intensity δI_i of a pixel i at (α, δ) is given by

$$\delta I(\alpha,\delta)_i = \frac{\Delta N_i}{\langle N\rangle_i} = \frac{N(\alpha,\delta)_i - \langle N(\alpha,\delta)\rangle_i}{\langle N(\alpha,\delta)\rangle_i}\,, \tag{5.2}$$

where $\langle N\rangle_i$ is the estimated background counts and N_i the observed counts in bin i with right ascencion and declination α_i, δ_i, respectively.

The significance of the observed relative intensity is calculated by the method of Li & Ma [78]. In the field of gamma ray astronomy, it is a commonly used likelihood test-statistic in the limit of a large number of background events harboring a small number of signal events. For each pixel in the sky map, the Li & Ma significance value, $S_{\text{Li-Ma}}$, is given by

$$S_{\text{Li-Ma}} = \sqrt{2N_{\text{on}} \ln\left(\frac{1+\alpha_{\text{Li-Ma}}}{\alpha_{\text{Li-Ma}}}\frac{N_{\text{on}}}{N_{\text{on}}+N_{\text{off}}}\right) + 2N_{\text{off}}\ln\left(\frac{(1+\alpha_{\text{Li-Ma}})N_{\text{off}}}{\alpha_{\text{Li-Ma}}}\right)}\,. \tag{5.3}$$

The $N_{\text{on}}, N_{\text{off}}$ values refer to the event counts both "on" and "off" the source of interest, which here are related to the observed counts and estimated background counts as follows:

$$N_{\text{on}} = N(\alpha,\delta)$$

$$N_{\text{off}} = \langle N(\alpha,\delta)\rangle/\alpha_{\text{Li-Ma}}\ .$$

The factor $\alpha_{\text{Li-Ma}}$ represents the relative exposure, an overexposure correction required to account for the fact that the local arrival distribution used for the reference map incorporates counts over the integration period, Δt. For a pixel area with $\Delta\Omega$ and declination extent $\Delta\theta$,

$$\alpha_{\text{Li-Ma}} = \frac{\Delta\Omega}{\Delta\theta\ \Delta t\ 15^\circ\ \text{hr}^{-1}\cos\delta}\ .$$

Hence, the estimated N_{off} background counts come from the relative exposure averaged reference map. Provided a large data set and an absence of signals, $S_{\text{Li-Ma}}$ should follow a normal distribution, while regions on a map with significances $> 5\sigma$ consitute discovery of a signal.

Chapter 6

Cosmic Ray Moon Shadow

The Moon blocks the nearly isotropic flux of cosmic rays on their path to Earth, as conceptually illustrated in figure 6.1. If cosmic rays propagated unperturbed, we would expect no events arriving from the location of the Moon, i.e. a perfect Moon shadow 0.52° in diameter. Of course, this also assumes that our ability to reconstruct cosmic ray arrival directions is perfect. In reality, the angular resolution of HAWC above 1 TeV is approximately 0.5°, and we should expect the shadow to be smeared similarly. And because of their charge, cosmic rays interact with the Earth's magnetic field which bends their trajectories. The deflection is dependent on both the particle charge *and* energy, so this has the further effect of shifting the observed Moon shadow in relation to the Moon's true position. However, since we have accurate representations of the Earth's magnetic field where it becomes most relevant to TeV-scale cosmic rays, we can use the evolution of the observed cosmic ray Moon shadow as an independent test of the detector's angular resolution and energy response.

In this chapter, we present the tools and methods used in measuring the energy dependence of the cosmic ray Moon shadow. We first discuss the simulation developed to model particle propagation between the Earth and Moon, including the effect of the Earth's magnetic field, to estimate the shadow's behavior. Next we present results from observing the Moon shadow with 1 year of HAWC data, detailing its measured dependence on energy based on the likelihood energy estimator.

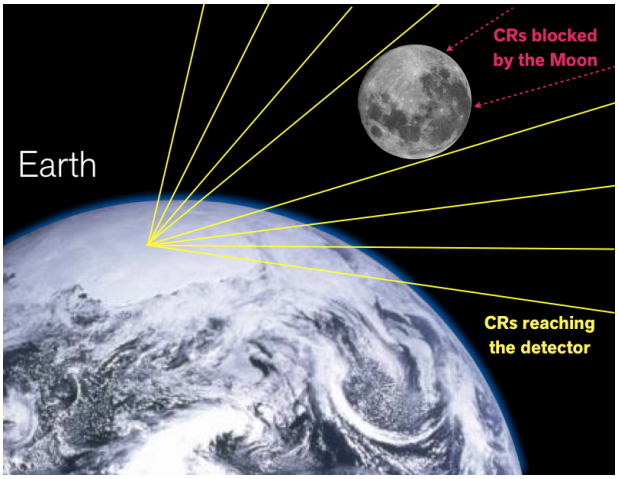

Figure 6.1: An illustration of the Moon blocking cosmic rays forming an apparent shadow. Note that the depiction is not to scale, as the Moon is $\sim 60 R_{\oplus}$ from the Earth. *Source*: HAWC Science website hosted by WIPAC [79].

6.1 Geomagnetic Field

The magnetic field of the Earth is a complex and dynamic structure, yet there are several approaches to estimate its strength. As a first approximation, the Earth's magnetic field resembles a dipole tilted (figure 6.2) from the Earth's rotational axis by $\sim 11.5°$. Analytically, the dipole field takes the following form:

$$B_r = \frac{-2B_0}{R^3}\cos\theta \qquad B_\theta = \frac{-B_0}{R^3}\sin\theta \ ,$$

where θ is the angle from the north magnetic pole, and R is the distance from the Earth's center in units of Earth radii, $R_\oplus$. In the pure dipole form, only the radial (B_r) and azimuthal (B_θ) components are needed to characterize the field, with $B_\phi = 0$. The field strength, $|\mathbf{B}|$, near the surface ranges from about 30 μT near the equator to twice that at the poles, following the relation $|\mathbf{B}| = \frac{B_0}{R^3}\sqrt{1+\cos^2\theta}$.

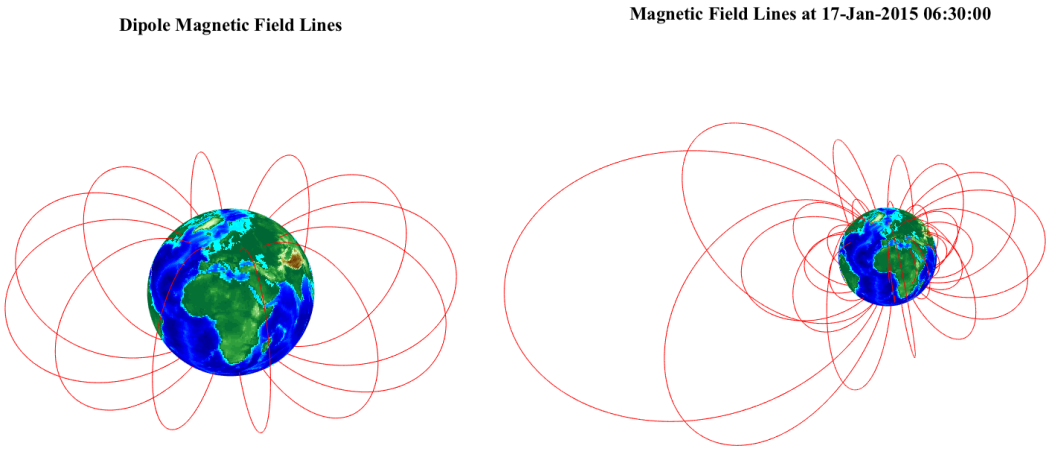

Figure 6.2: Isoclinic diagrams of the geomagnetic field using the dipole (*left*) and the IGRF2015 (*right*) models. Each red path represents a curve of constant magnetic field strength. The deviation of the IGRF from the dipole approximation's symmetry is clear. Figures generated using [80].

A widely accepted model which incorporates the time-dependent asymmetries of the geomagnetic field is the International Geomagnetic Reference Field (IGRF) [81]. It is put forth by the International Association of Geomagnetism and Aeronomy (IAGA) and consists of a set of coefficients updated every 5 years, defining

an epoch. Using observations from satellite borne and ground based sensors, the IGRF coefficients are best fit values to data using a spherical harmonic model of a scalar potential. The set of coefficients for each epoch is defined by "IGRF" and the year. For this study the specification implemented is IGRF2015. An isoclinic (equal magnetic field contours) diagram of the IGRF is shown in figure 6.2, where the deviation from the symmetric dipole approximation is clearly visible.

The analytic model is expressed as a potential $V(r, \theta, \phi)$, represented by a series expansion of order $k = 13$ taking the form:

$$V(r,\theta,\phi) = R_\oplus \sum_{n=1}^{k} \left(\frac{R_\oplus}{r}\right)^{n+1} \sum_{m=0}^{n} (g_n^m \cos m\phi + h_n^m \sin m\phi) P_n^m(\cos\theta)\,, \tag{6.1}$$

where $R_\oplus = 6371.2$ km is the mean Earth radius, (r, θ, ϕ) is a point in geocentric coordinates, g_n^m, h_n^m are gaussian coefficients defined by the epoch, and $P_n^m(\theta)$ are Schmidt quasi-normalized associated Legendre functions of degree n, order m. The geocentric coordinates are defined via $\theta = 90° -$ latitude, $\phi =$ longitude. The magnetic field is the negative gradient of this potential, $\mathbf{B} = -\nabla V$, which in local tangential coordinates is

$$B_r = -\frac{\partial V}{\partial r} = \sum_{n=1}^{k} \left(\frac{R_\oplus}{r}\right)^{n+2} (n+1) \sum_{m=0}^{n} (g_n^m \cos m\phi + h_n^m \sin m\phi) P_n^m(\theta) \tag{6.2}$$

$$B_\theta = -\frac{1}{r}\frac{\partial V}{\partial \theta} = -\sum_{n=1}^{k} \left(\frac{R_\oplus}{r}\right)^{n+2} \sum_{m=0}^{n} (g_n^m \cos m\phi + h_n^m \sin m\phi) \frac{\partial P_n^m(\theta)}{\partial \theta}$$

$$B_\phi = -\frac{1}{r\sin\theta}\frac{\partial V}{\partial \theta} = -\sum_{n=1}^{k} \left(\frac{R_\oplus}{r}\right)^{n+2} \sum_{m=0}^{n} (-g_n^m \sin m\phi + h_n^m \cos m\phi) P_n^m(\theta)\,.$$

According to the IAGA [82], this model is valid up to $\sim 7\ R_\oplus$, beyond which heavily time-dependent solar effects become important for the field's form. However, at this distance the field strength is reduced to a few nT, where the Larmor radius of 1 TeV protons is ~ 20 AU. Even closer to the Earth's surface where the field is of order 50 μT, that radius is 10 $R_\oplus$. Therefore the field structure beyond the model's limitations has a negligible influence on the trajectories of TeV particles.

It should also be noted that the dipole and IGRF2015 models are approximations to the *internal* field of the Earth, which is generated by the planet's core. This component represents $\sim 95\%$ of the total field, with the remaining 5% coming from *external* effects due to solar activity. The semi-empirical Tsyganenko model [83] accounts for this component beyond $\sim 10\ R_\oplus$. However, since its strength is comparable to or below the IGRF, its effect on >TeV particle trajectories has been shown to be negligible [84], so it was not included for this study.

6.2 GPU Simulation

The principal motivation for running cosmic ray propagation simulation on graphics processing units (GPUs) is run time. Indeed, when simulating charged particles traversing electromagnetic fields, the number of particles simulated and the necessary integration duration contribute to the simulation time. The number of simulation steps per particle (integration duration) depends on the required spatial and temporal resolution. For example, if the electromagnetic field is rather weak, perhaps inducing angular deflections $\Delta\theta << 1^{\circ}$, then the number of integration steps will be significantly smaller to probe similar spatial scales than a field that is 10 times as strong. For reference, the typical number of propagation steps in simulating a TeV proton traversing the Earth-Moon system is of order 10^3. While this may not appear to be a significant number for a typical processor, we also must consider the calculation complexity for each step, such as the evaluation of the magnetic field model.

One major advantage of moving such a simulation to the GPU is the fact that this is an "embarrassingly" parallel problem. Since each particle is simulated independently from the others, the processing units that carry out the calculation never communicate. They do not share any information regarding their propagation steps. Hence, there are no bottlenecks related to data transfer as might occur in large matrix inversions or simulations of coupled systems.

6.2.1 Graphics Processing Units

Graphics processing units are specialized compute devices intended for rapid manipulation of vectorized data generally used for image display. The term general purpose GPU (GPGPU) refers to the utilization of such devices for other applications such as scientific computation. The internal structure of GPUs is optimized for tasks requiring high levels of parallelization, such as image manipulation. Typical devices are organized in physical sections that are further divided into blocks of memory and processing units. The language to describe device architecture is dependent on the specific manufacturer; however, most of the terminology is comparable. Though the simulation runs on the UW-Madison cluster which employs GPUs from several vendors, I will use the terminology of the CUDA API [85] for consistency and because of access to exemplary visual aids.

Figure 6.3 shows the architecture of a typical NVIDIA GPU card. Any information defined by the user is done so on the CPU and must be passed to the GPU through one of three types of memory spaces. Any constants needed by the simulation are passed and held in the "Constant Memory" space, while all other input to and output from the GPU is transferred through "Global Memory". The "Texture Memory" space is optimized for graphical video processing, so for the purposes of cosmic ray simulation it is not utilized.

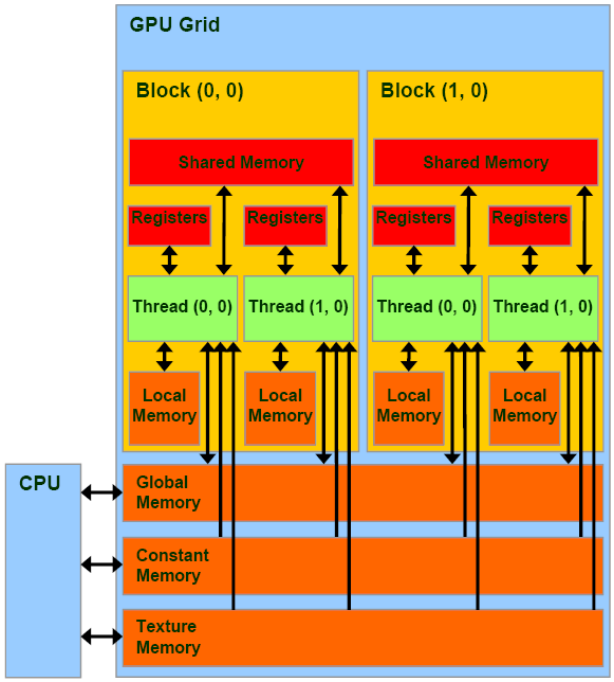

Figure 6.3: NVIDIA GPU architecture per the CUDA API.

The actual computation occurs further withinin the GPU grid, where compute "blocks" are home to a collection of compute "threads". For example, the NVIDIA GTX 980 card, as found on the UW cluster, has 2048 blocks with 1024 threads per block, permitting nearly $2^{21} \approx 10^6$ simultaneous calculations. Strictly speaking, the total parallelization is not as simple as the product of block and threads. Each block runs groups of 16 threads up to the same point in the simulation. Then another group of 16 threads executes catching up to the same point. Since the *blocks* are independent, the maximum number of threads running strictly in parallel is $16 \times 2048 = 32768$.

In the case that the threads run into different evaluations of conditional statements, then the threads serialize into groups having the same conditional value. For other applications where threads share data, the local and shared memory spaces provide the means to do so. In the propagation simulation here, no information is shared among the threads. The few required conditional statements do not influence the particle propagation functions, as they merely test whether the particles are to continue propagating or stop completely. Those that stop do not influence the simulation performance, as they have no more steps to continue.

6.2.2 Available Resources

The UW-Madison cluster is home to 345 graphics cards. Most units are NVIDIA models with 248 GTX 980s, 32 GTX 690s (each containing two GPUs), and one GTX 1080. The cluster also includes 32 AMD Radeon 7970 GPUs. The CUDA [85] proprietary software is meant for NVIDIA GPUs and does not operate on other vendor products. However, we used the open source OpenCL [86] GPGPU framework combined with the host wrapper PyOpenCL [87] API because it is hardware independent. This ensures the code runs regardless of GPU manufacturer and also makes it easier to distribute to others for verification or even for development with similar simulation tasks. One other advantage of the OpenCL-PyOpenCL tag team is the ability to request the calculations be performed on either the host CPU or the GPU, permitting an assessment of performance gain from parallelization.

6.2.3 Magnetic Field Model

The IAGA's implementation of the IGRF2015 model is written in GNU FORTRAN, but the only vendor-independent API available for use on the GPU is OpenCL C [86]. Hence, the algorithm was translated, and is the default magnetic field model used here. In implementing the IGRF2015 for the GPU, we see another advantage of utilizing the GPU for charge particle progagation purposes. A single IGRF2015 calculation involves two nested loops of $k = 13$ iterations for three magnetic field components, amounting to $3k(k + 1)/2 = 273$ evaluations per propagation step. Thus, as the magnetic field model increases in complexity, a serialized computation increases linearly in runtime. The task effectively is rid of this burden by the factor of parallelization.

6.2.4 Particle Propagation

The particle equations of motion for charged particle propagation in an electromagentic field take the form

$$\frac{d\mathbf{p}}{dt} = Z\gamma\Big(\mathbf{E} + \mathbf{v} \times \mathbf{B}\Big) \tag{6.3}$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}\,.$$

Here, Z denotes the particle charge, $\mathbf{v}$ the velocity vector, $\gamma = \frac{1}{\sqrt{1-v^2/c^2}}$ the relativistic boost factor, and $\mathbf{B}$ the geomagnetic field vector. For TeV particles we can consider propagation in the ultrarelativistic limit, where any energy losses during propagation or energy gains from the interplanetary electric field are negligible, and we can set $\mathbf{E} = \vec{0}$.

The following numerical integration methods were considered for the simulation: Euler, Runge-Kutta 4 (RK-4), and Boris with fixed and adaptive time steps. The Euler method is the simplest, requiring a single step. Furthermore, it is an explicit method, i.e. no equations must be solved numerically. However, the global error in position and velocity is proportional to the square of the time step size, $\mathrm{O}(\delta t^2)$, which for long integration periods can induce significant trajectory deviations. The global error for the RK-4 method is $\mathrm{O}(\delta t^5)$, but requires four evaluations of the acceleration function, thus presenting a potential bottleneck to the simulation time if the function itself is complicated as with the 273 evaluations of the IGRF2015 model. Furthermore, it is a non-symplectic integrator, meaning that energy is not necessarily a conserved quantity. Finally, the Boris [88] method is symplectic, explicit, and requires a single evaluation of the acceleration function. Thus, regardless of step size, the particle gyroradius and momentum are perfectly conserved [89]. In addition, its relative simplicity allows for implementation with a fixed or adaptive time step. For this study, the adaptive Boris method was used.

The adaptive time step for charged particle propagation is defined by

$$\delta t = \frac{2\gamma M}{|Z||\mathbf{B}(\mathbf{x})|} \tan \frac{1}{2}\theta_{\mathsf{min}} \,, \tag{6.4}$$

where M, Z, γ are the particle mass, charge and Lorentz factor, respectively, $\mathbf{B}(\mathbf{x})$ is the magnetic field vector at position $\mathbf{x}$, and θ_{min} is the user defined minimum allowed rotation angle. The value used in GPUSim is $\theta_{\mathsf{min}} = 10^{-3}$ rad $\approx 0.05°$. Thus, the integrator adapts the time step based on the value of the gyroradius (function of $\gamma, \mathbf{B}(\mathbf{x})$) evaluated at the current position. To move the particle location, a δt is evaluated per equation 6.4, the velocity vector $\mathbf{v}$ is rotated by angle θ_{min} about the axis collinear with $\mathbf{B}$, and finally the particle's location is updated to $\mathbf{x} \leftarrow \mathbf{x} + \mathbf{v}\ \delta t$.

For studying the Moon shadow, particles are initialized at the geographic location of HAWC and propagated outward to the Moon's orbital radius. This of course necessitates simulating particles of opposite charges, since their velocities are reversed from observations. For the IGRF model requirements, particles were assigned dates between January and February of 2015, for times when the Moon was transiting above the local horizon. Particles were assigned isotropic local direction vectors with zenith angles up to 60°, and energies were sampled from an E^{-1} spectrum. Each integration step is determined via equation 6.4, and a particle continues to propagate if its geocentric position is less than the Earth-Moon distance and its propagation is directed radially outward from the Earth. This last condition assures that the simulation does not waste time with particles below the geomagnetic cutoff which will never escape. Once reaching the Moon's orbit, intersection of the Moon disc is determined, upon which the final trajectory information is read out from the GPU and saved to disk on the CPU. A screen shot of iron nuclei after propagating to three Earth radii is shown in figure 6.4. The visualization tool is not used when generating the simulation set, being for instructional purposes only.

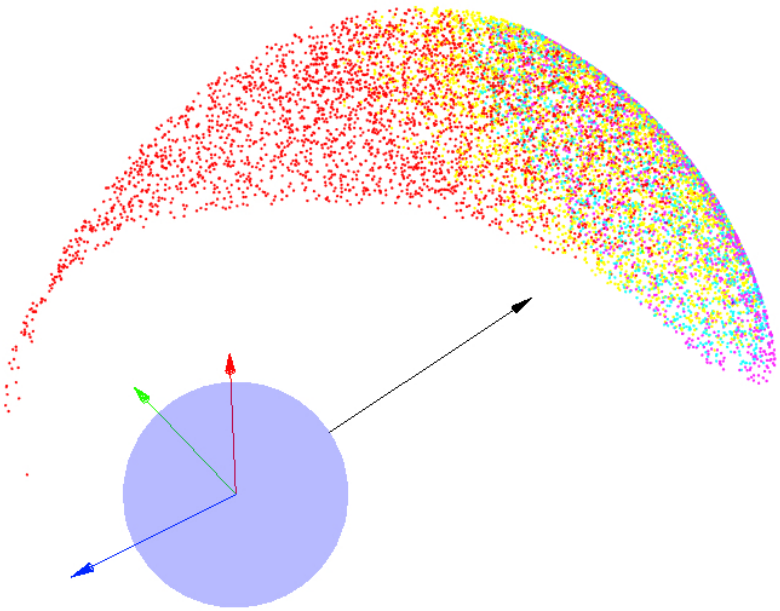

Figure 6.4: Visualization of 10^4 iron nuclei after propagating from the HAWC location outward. The blue sphere represents the Earth, the axes identify geocentric coordinates, and the individual colors indicate the particle energy, where red is the lowest and violet the highest. The longer black vector indicates the normal direction at the HAWC location, from which particles are propagated outward. The lower energy particles experience significant deflection towards the East, translating to an expected negative deflection in right ascension. The simulated energy range of particles is from 500 GeV – 10 TeV to show pronounced deviations. The Moon, being 60 $R_\oplus$ away, is not shown.

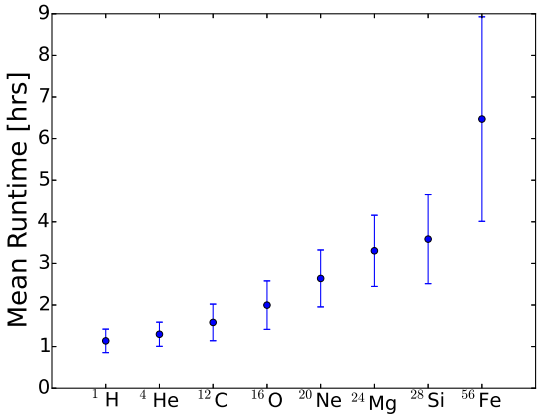

Figure 6.5: GPU mean runtimes in hours for the various simulated species. The increase in both mean value and width is due to particles of increasing Z at similar sub-TeV energies reaching the maximum allowable iterations.

6.2.5 Performance

Each GPU can calculate $\sim 10^6$ trajectories nearly parallel within several minutes. Harnessing the entirety of the UW-Madison cluster allows the simulation of $O(10^9)$ particles on the order of $\sim$hr time scales. Specifically, the simulation production takes $\sim 5-6$ hrs to complete, amounting to a speedup factor of $\times 90$ if 1000 CPU nodes were to generate similar statistics in a serialized manner. Figure 6.5 shows the GPU runtimes for the various species simulated. The times are the mean values with standard deviations for generating one hundred simulation files start to finish, each with 10^7 particles for every species. The overall mean run time is ~ 2.5 hrs, with higher simulation times for species with greater charges.

6.2.6 Simulation Analysis

HAWC is most sensitive to TeV cosmic rays, and we expect the Moon shadow position to deviate commensurate with this energy range. We quantify the total angular deflection $\Delta\theta$ of a particle intersecting the Moon disc by the deviation from its initial direction vector, resulting in apparent right ascension and declination offsets from the true Moon position:

$$\begin{aligned} \Delta\theta^2 &= (\alpha - \alpha_{\text{Moon}})^2 + (\delta - \delta_{\text{Moon}})^2 \\ &= \Delta\alpha^2 + \Delta\delta^2 \,. \end{aligned} \tag{6.5}$$

The simulated energy dependence of $\Delta\theta$ is shown in figure 6.6. The total deflection angle distributions for the various species are parallel comparing the logarithm of both the angle and energy, indeed sharing a common

slope which can be summarized by the relation $\Delta\theta \propto Z/E$. Thus, the total angular deviation is inversely proportional to rigidity. For a particle of charge Z, the proportionality coefficient is found to satisfy

$$\Delta\theta = (1.59^\circ \pm 0.01^\circ)\,\frac{Z}{E[\mathrm{TeV}]}\,, \tag{6.6}$$

consistent with previous studies [77, 90].

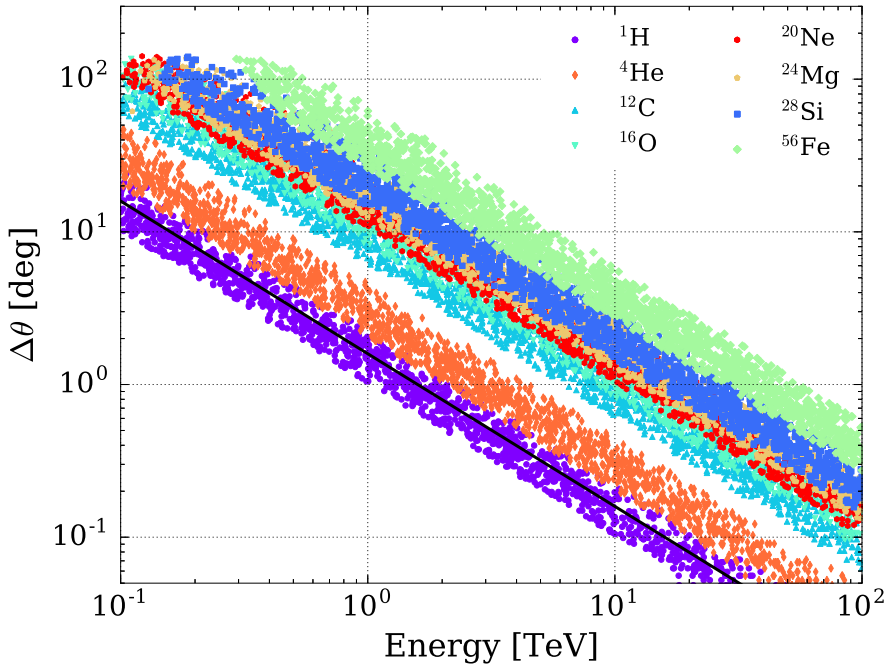

Figure 6.6: Simulated total angular deflection of charged particles intersecting the Moon disc as a function of energy. The slopes of the distributions are equal; thus the deflections are rigidity dependent and follow the form of equation 6.6. The fit for protons with $Z = 1$ is shown via the black line.

Because the Earth's magnetic field lines are nearly aligned in the north-south direction, the angular deflection is primarily along the right ascension direction, i.e. $\Delta\alpha \approx -\Delta\theta = -1.59^\circ Z/E[\mathrm{TeV}]$. While subtle, we expect a slight energy-dependent shift in declination as well. This is due to the tilt of the Earth's rotational axis with respect to the magnetic field polar orientation, which results in a discrepancy between trajectory deviations while traversing the magnetic field when the Moon is rising versus setting. This asymmetric integrated sampling of the geomagnetic field manifests as a similar inverse rigidity relation, and is depicted in figure 6.7, where the right panel highlights this effect via the mean of $\Delta\delta$. The finite asymptotic width of the distribution at higher energies is precisely the angular diameter of the Moon disc, $2r_{\mathrm{Moon}} = 0.52^\circ$. A fit

to the mean value gives the deflection coefficient via the relation

$$\Delta\delta = (0.15^\circ \pm 0.01^\circ)\,\frac{Z}{E[\mathrm{TeV}]}\,. \tag{6.7}$$

It is clear from the similar scale of the $\Delta\delta$ coefficient to r_{Moon} that the more robust means to constrain the energy scale from Moon shadow measurements is via $\Delta\alpha$. The right ascension offset simply has more leverage with which to sample energy dependent deviations. Yet any declination shift below ~ 1 TeV does provide additional, orthogonal rigidity information, and it is assessed in the data analysis.

Of course, these analytic forms describe the behavior for individual species reaching the top of the atmosphere. For the full Moon shadow, the convolution of both composition assumption and detector response for each species produces a mean rigidity scale, which we interpret as $\bar{Z}/E$. Taking into account the nominal composition model from section 3.3.3 as well as simulated efficiencies to the various species, the expected mean charge from simulation is $\bar{Z} = 1.23 \pm 0.02$.

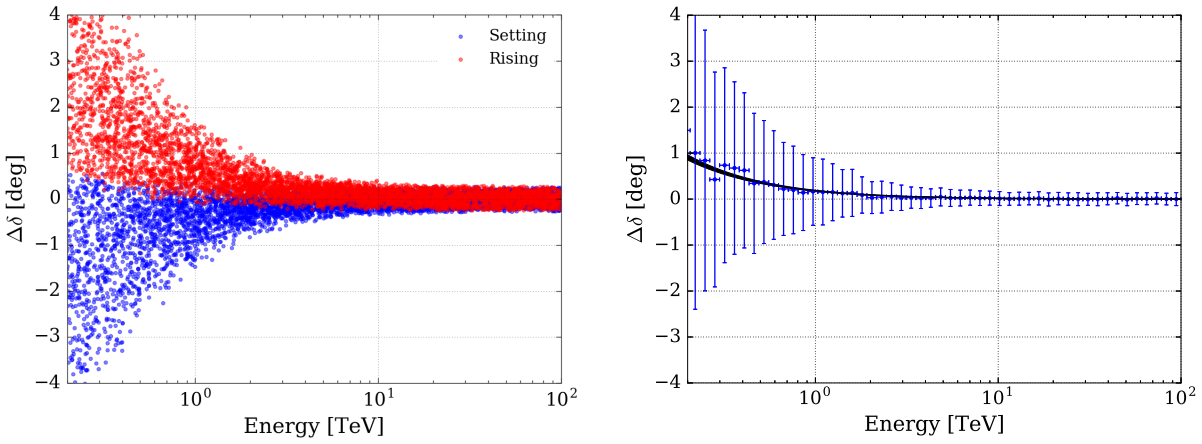

Figure 6.7: Energy dependent declination deflection due to rising-setting asymmetry for protons. The subtle asymmetry is visible when extracting the mean of $\Delta\delta$ as shown in the right panel. The inverse energy relation fit to the mean gives the relation of equation 6.7, and is shown here by the black curve.

6.3 Data

6.3.1 Event Selection

The HAWC detector in its full configuration of 290 WCDs was in stable data taking mode for the runs selected for this analysis, which span 397 days from 16 April, 2015 to 18 June, 2016. The total up-time efficiency was $\sim$ 92% and the mean trigger rate was $\sim$ 25 kHz. For the background estimation used in making maps, a direct integration period of $\Delta t = 2$ hrs was used, so only runs where the detector was stable for times greater than 2 hrs were included. Detector stability was assessed via a χ^2-difference test of the local azimuth and zenith angle distributions as described in [77]. An example of stability for a single day of data taking is shown in figure 6.8.

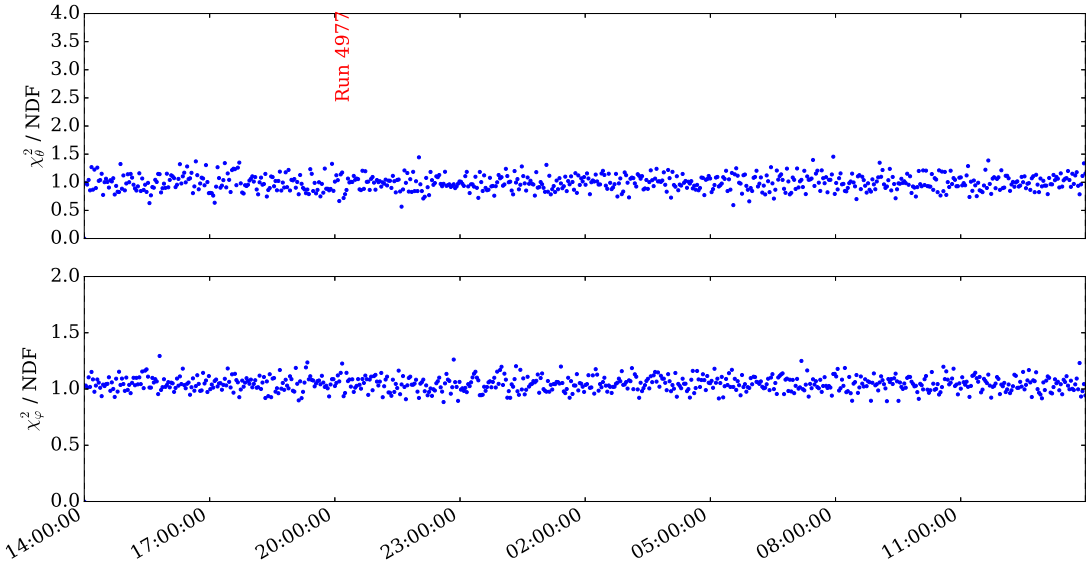

Figure 6.8: χ^2-difference test for zenith (top) and azimuth (bottom) distributions in $\sim$ 2 minute time bins for a single day from the data set. Each point represents the reduced χ^2 difference between the respective angular distribution at the current time on the abscissa to the value of the first time bin in the day.

In accordance with other cosmic ray analyses [1], selected events passed a multiplicity threshold of $\sim$ 6% of the detector. For the nominal detector configuration during this period, the resulting mean multiplicity cut was $N_{\rm hit} \geq 75$. The $N_{r40} \geq 40$ selection cut from section 3.4 was used to ensure the integrated core fit resolution is below 10 m for energies above 10 TeV. Finally, we consider showers with $\theta \leq 45^\circ$ to reduce the influence of increasingly poor energy resolution and bias from sampling the energy table's highest zenith bin θ_2 over its entire zenith range of $35^\circ - 60^\circ$. A total of 4.16×10^{10} events passed the selection cuts, with

1.23×10^8 events falling within 5° of the Moon's position. The mean energy for the entire data set and for events near the Moon is 4.3 TeV.

6.3.2 Energy Binning

Eleven analysis bins were defined using the likelihood energy estimation variable, E_{reco}, with steps of size 0.2 in $\log E_{\text{reco}}$ from 1–160 TeV. We further account for the arrival direction's effect on the estimated energy by sampling from the table represented in figure 6.9, from which an estimate of the true mean energy and resolution was determined for each analysis bin. The estimated true energies are illustrated in figure 6.10 and the values are presented in table 6.1. The normalized counts distribution with respect to the analysis bins is shown in figure 6.11. This includes a comparison of the entire data set to only events falling within 5° from the true Moon position.

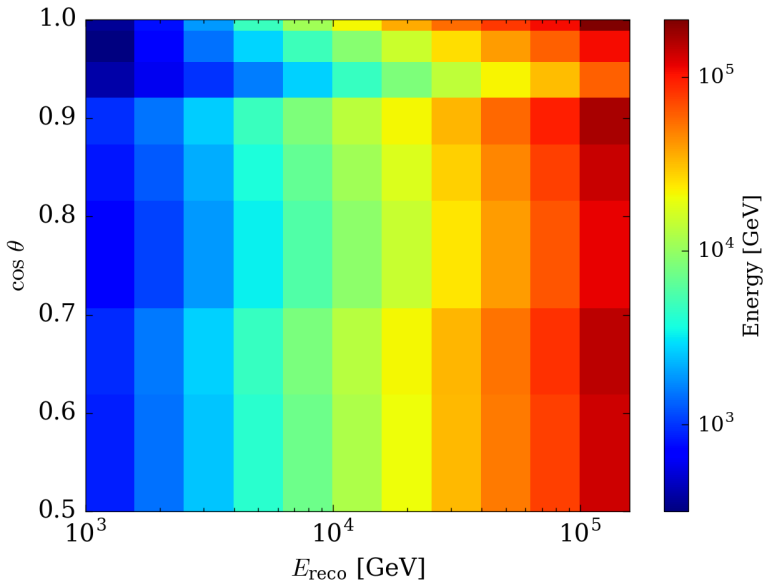

Figure 6.9: Median energy (from simulation) as a function of the reconstructed energy E_{reco} and the cosine of the local zenith angle θ. The mean energy for each analysis bin is determined by sampling this table with the $(E_{\text{reco}}, \theta)$ distributions observed from within 5° of the true Moon location. Three large blocks are also visible in the table, indicating the zenith ranges θ_i of the energy estimation tables.

Analysis Bin	$\log E_{\mathrm{reco}}/\mathrm{TeV}$	$\bar{E}$ [TeV]	σ_E^+	σ_E^-
0	$0.0 - 0.2$	0.793	0.827	0.272
1	$0.2 - 0.4$	1.29	1.33	0.48
2	$0.4 - 0.6$	2.24	2.21	0.85
3	$0.6 - 0.8$	3.94	3.61	1.50
4	$0.8 - 1.0$	6.73	5.63	2.57
5	$1.0 - 1.2$	10.1	8.3	4.2
6	$1.2 - 1.4$	17.9	11.9	6.7
7	$1.4 - 1.6$	28.9	16.4	10.1
8	$1.6 - 1.8$	47.4	22.5	14.8
9	$1.8 - 2.0$	74.9	29.6	20.2
10	$2.0 - 2.2$	133.8	44.9	32.5

Table 6.1: Estimated mean energies for the analysis bins defined by E_{reco}. The bounds indicated by $\sigma_E^{\pm}$ represent the 68% central containment region and are quoted in the same units as E.

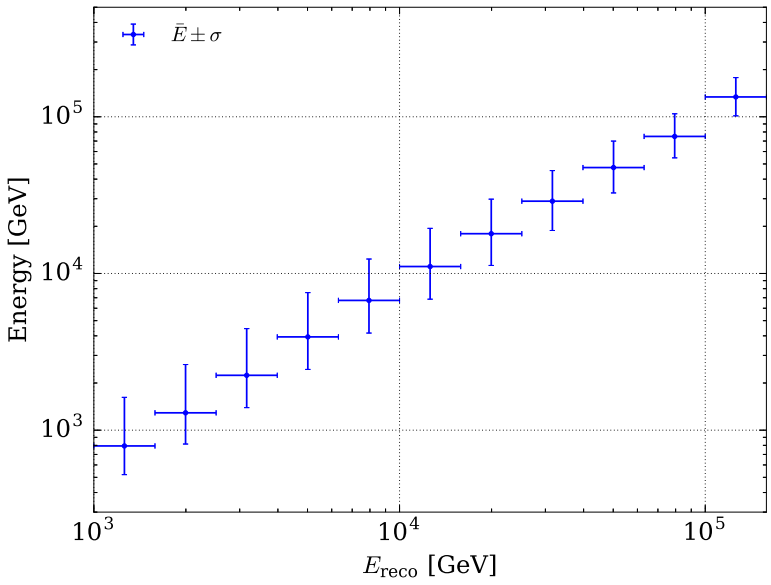

Figure 6.10: Mean energies of analysis bins defined in table 6.1.

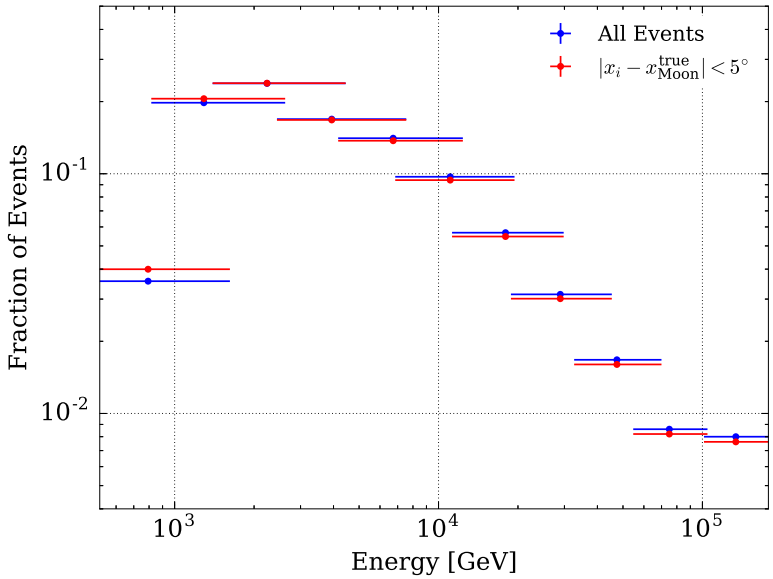

Figure 6.11: Normalized counts distribution as a function of energy for all events in the data set and for events observed within 5° from the Moon's true position.

6.3.3 Moon Maps

An event map was made for each analysis bin, and both relative intensity and significance maps were generated. Unsmoothed example maps centered on the true Moon position are presented in figure 6.12, revealing the presence of the cosmic ray Moon shadow. The same maps after applying 1° top-hat smoothing are shown in figure 6.13. The figure shows the general smoothing out of pixel-to-pixel fluctuations as well as a dramatic increase in signal strength as a result. It is also clear that the cosmic ray Moon shadow is offset by $\sim 0.5^\circ$ from the true Moon position, with primary deflection in right ascension. There is also a noticeable offset in declination of about 0.2° as well. We present the energy dependence of the Moon shadow after first considering several methods to analyze the maps.

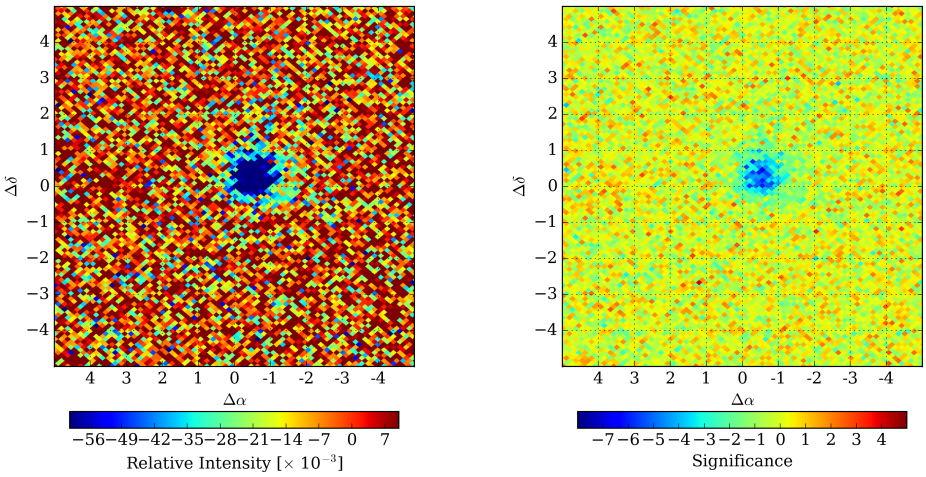

Figure 6.12: Moon shadow maps in relative intensity and significance for bin 4, with a mean energy of 6.7 TeV. The axes indicate the angular distance in α and δ from the true Moon position, i.e. Moon-centered maps. There is a slight offset in right ascension by about $\Delta\alpha \approx 0.5°$. If the sample were purely protons, the expected deflection angle would be $\sim 0.24°$. This discrepancy indicates the presence of other charged particles in this bin. These maps did not include smoothing; yet, the Moon's presense is clear, as its peak deficit in relative intensity is -113×10^{-3} with a significance of -6.8σ, passing the threshold for a significant observation in this energy bin.

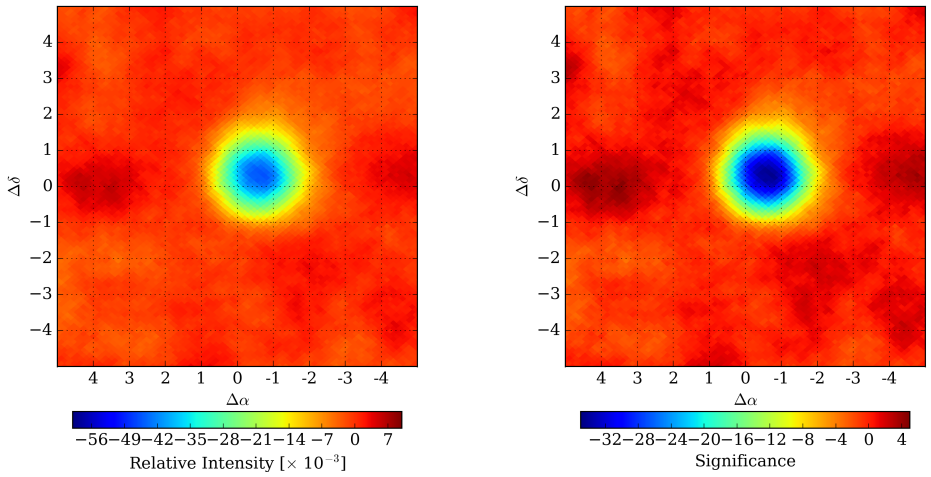

Figure 6.13: Moon shadow maps from figure 6.12 smoothed by 1°. The effect of correlating neighboring bins via smoothing is a reduction of the magnitude of the relative intensity, as here it is -45.7×10^{-3}, but an increase in signal strength, here to -34.5σ. Upward fluctuations on either side of the shadow are present in both maps, which will be addressed via ROI-masking.

6.3.3.1 Fits to the Moon Shadow

As can be seen from the previous figures, the Moon shadow forms a strongly peaked disc and appears to be nearly symmetric. However, as discussed in previous work [77] the shadow's morphology can be affected by the chemical composition, potentially producing asymmetries like that of bin 3 in figure 6.14, which has been smoothed by 1° to enhance the structure. We thus fit the unsmoothed relative intensity maps to a two-dimensional Gaussian function to quantify the evolution of the Moon shadow with energy. The generalized Gaussian takes the form

$$\mathcal{N}(N, A, \delta, \alpha, \sigma_\delta, \sigma_\alpha, \theta) = N + A \exp\left[-a\,(\delta - \delta_c)^2 - 2b\,(\delta - \delta_c)\,(\alpha - \alpha_c) - c\,(\alpha - \alpha_c)^2\right], \tag{6.8}$$

where N is an overall offset, A is the amplitude, the pair $(\delta_c,\ \alpha_c)$ are the centroid positions in declination and right ascension, respectively, and the coefficients a, b, c are defined by the widths $\sigma_{\delta,\alpha}$ and the tilt angle θ from the right ascension axis via

$$a = \frac{\cos^2\theta}{2\sigma_\alpha^2} + \frac{\sin^2\theta}{2\sigma_\delta^2}, \qquad b = \frac{\sin^2\theta}{4}\left(-\frac{1}{\sigma_\alpha^2} + \frac{1}{\sigma_\delta^2}\right), \qquad c = \frac{\sin^2\theta}{2\sigma_\alpha^2} + \frac{\cos^2\theta}{2\sigma_\delta^2}.$$

Thus the centroids give an indication of the rigidity scale of the analysis bin, permitting an estimate of the mean chemical composition of the data set. The widths in δ and α provide further information regarding the point-spread function or the angular resolution with respect to energy, and strong asymmetries can result from composition as well. The resulting best fit values of equation 6.8 to the unsmoothed maps are shown

in figure 6.15. Using the deflection in right ascension, we see that the mean estimated charge of the data set is 1.25 ± 0.06, meaning it lies between proton and helium, and consistent with the simulated $\bar{Z}$ from section 6.2.6.

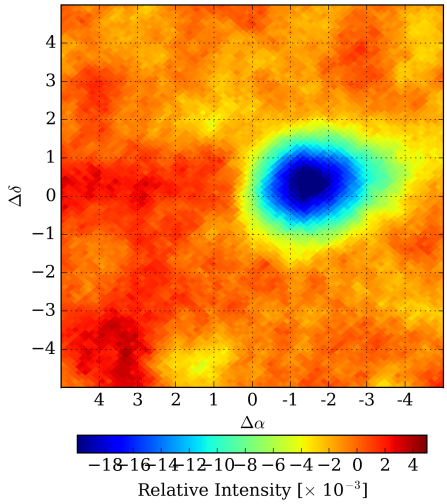

Figure 6.14: Moon shadow relative intensity map from bin 1 ($\bar{E} = 1.3$ TeV) smoothed by 1°. The tail towards the upper right of the figure is potentially caused by the presence of heavier elements, which are deflected more strongly than the dominant proton and helium components. To account for this structure, we fit the unsmoothed maps by the generalized two-dimensional Gaussian of equation 6.8.

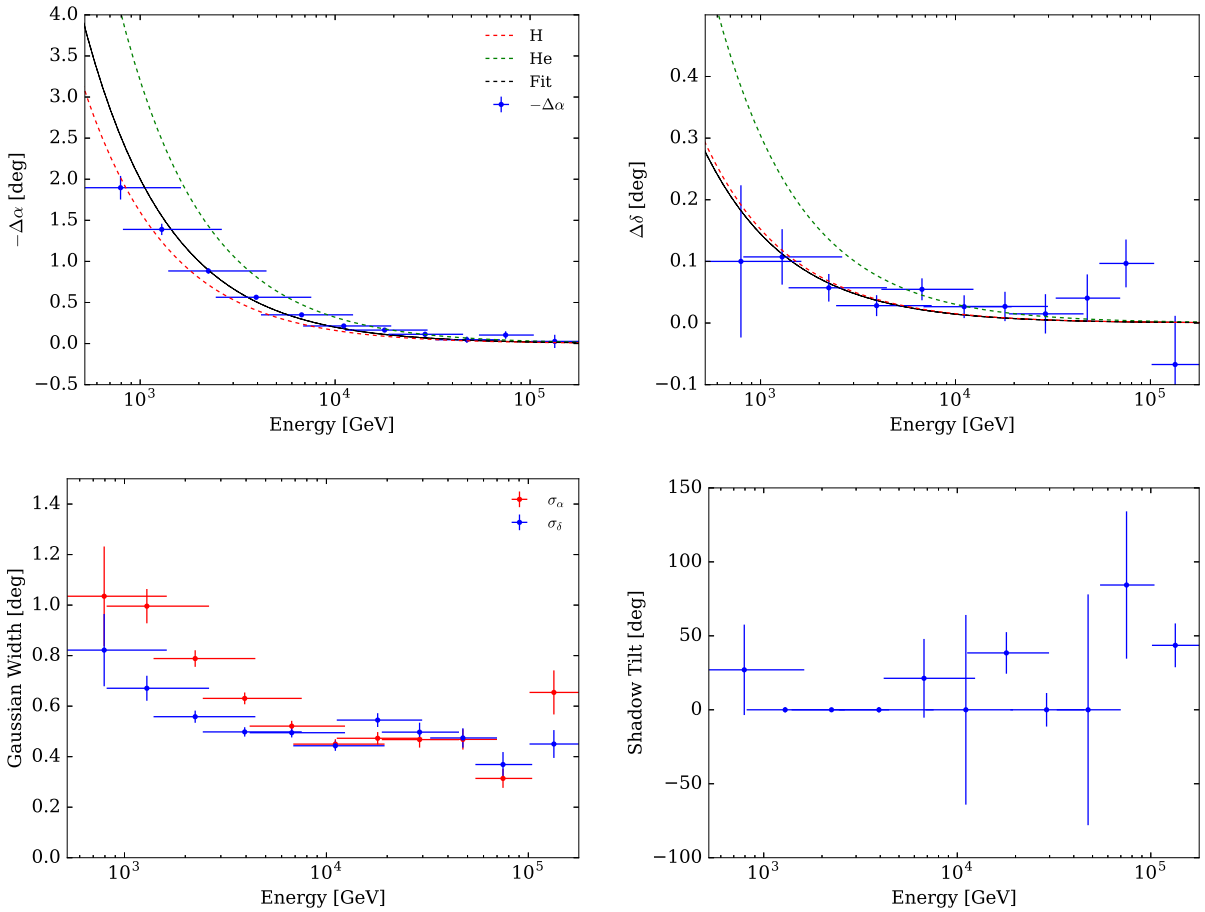

Figure 6.15: Best fit values of equation 6.8 to unsmoothed relative intensity Moon shadow maps for all analysis bins. Top left: centroid offset in right ascension, with Monte Carlo expectations for pure proton and pure helium hypotheses. The fit to the data values results in an estimated mean charge of $\bar{Z} = 1.25 \pm 0.06$. Top right: centroid offset in declination. The fit to the data gives an estimated mean charge value of $\bar{Z} = 0.97 \pm 0.22$. Bottom left: Gaussian widths in δ and α. Bottom right: major axis tilt angle θ.

6.3.3.2 Region of Interest Masking

The effect of masking the Moon from the background estimation as a region of interest was investigated to illustrate its effect on the strength of observed signals. The influence of the Moon shadow as a sink is most prominently seen in figure 6.13. The Moon's presence is so strong that the estimated background changes, causing an over-compensation of relative intensity in neighboring declination pixels, potentially manifesting spurious significant regions. We ameliorate this effect by masking out pixels around the Moon's observed position in the background estimation step.

The resulting fits to the unsmoothed maps are used to define a mask region for each analysis bin. The maps are then regenerated using a masking circle centered on δ_c, α_c with radius $3\,\sigma_{\alpha+\delta}$, where the mean Gaussian width $\sigma_{\alpha+\delta}$ is calculated via the respective uncertainties $e_{\alpha,\delta}$ from the fitting procedure:

$$\sigma_{\alpha+\delta} = \left(\frac{\sigma_\alpha}{e_\alpha} + \frac{\sigma_\delta}{e_\delta}\right) \cdot \left(\frac{1}{e_\alpha} + \frac{1}{e_\delta}\right)^{-1} . \tag{6.9}$$

We thus ensure that each bin's ROI-mask is appropriately centered on its unmasked Moon shadow, and accounts for a region thrice its angular dimensions. Figure 6.17 demonstrates the effect of the ROI-masking procedure, where the fluctuations are absent. Fits to equation 6.8 are performed again using the ROI-masked maps, which are subsequently used for the analysis. The resulting fit values are displayed in figure 6.16 and are consistent with the non-ROI masked map fits, including the mean rigidity scale of the data sample.

In addition to removing artificially adjusted background about the source region, masking a strong signal results in a slight boost to its significance. More correctly, the ROI-mask does not boost the signal, but rather not implementing the mask reduces the signal strength, appropriating a portion to the surrounding pixels. The effect on the peak relative intensity and significance for the analysis bins is presented in figure 6.18. While the peak relative intensity values are not affected by more than a few percent, the underlying significance of the signal is brought out by the use of ROI-masking, and in bins with the greatest statistical power, several units of significance are gained.

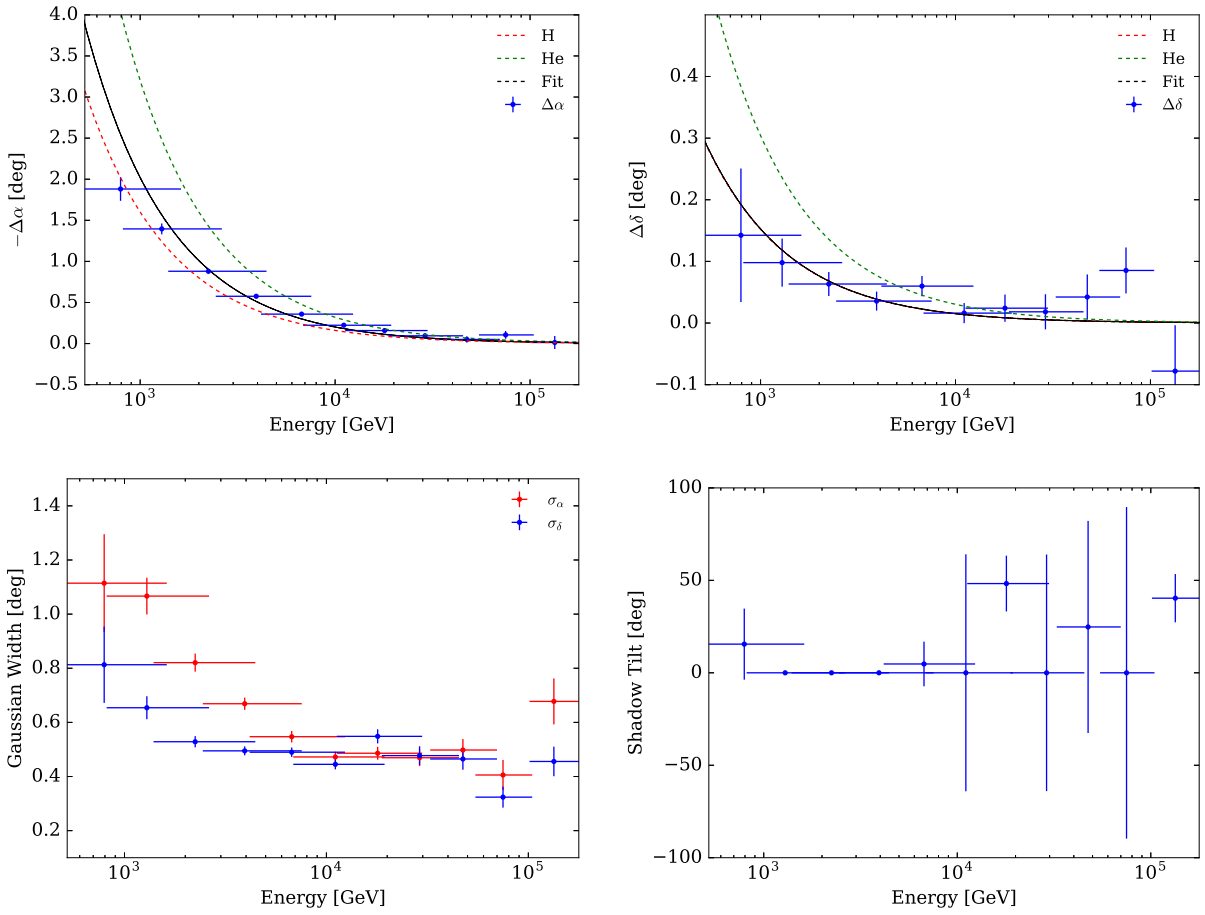

Figure 6.16: Best fit values of equation 6.8 to unsmoothed relative intensity maps that have been masked. Top left: centroid offset in right ascension, again with Monte Carlo expectations for pure proton and pure helium hypotheses. The fit to the data values results in an estimated mean charge of $\bar{Z} = 1.26 \pm 0.06$, consistent with the previous fit. Top right: centroid offset in declination. The fit to the data gives an estimated mean charge value of $\bar{Z} = 1.00 \pm 0.20$. Bottom left: Gaussian widths in δ and α. Bottom right: major axis tilt angle θ.

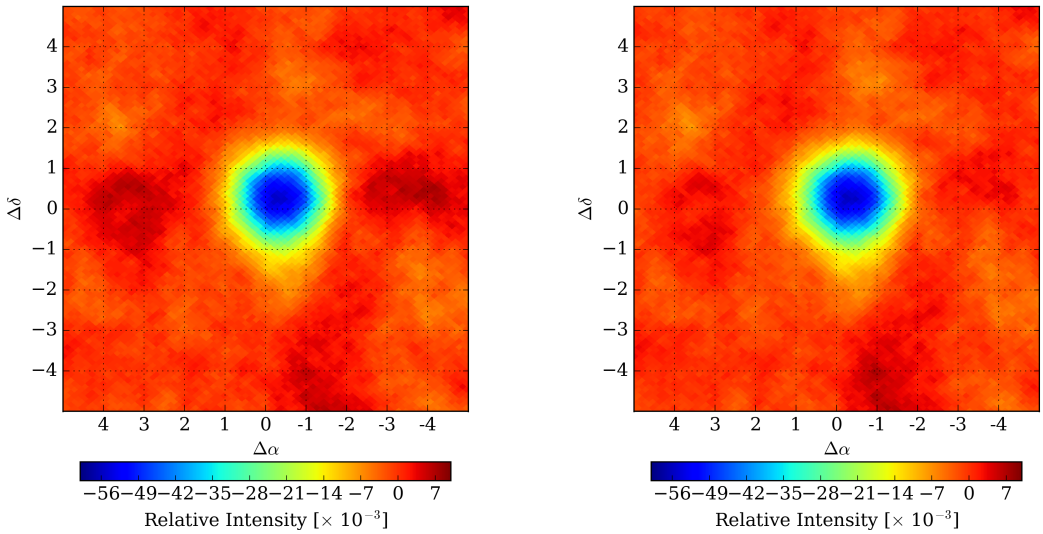

Figure 6.17: Smoothed Moon shadow maps in relative intensity for bin 6 ($\bar{E} = 17.9$ TeV) without (left) and with (right) region of interest masking. The disappearance of the upward fluctuating bands on either side of the shadow indicates the ROI-mask reduces the strong effect of the Moon's presence on the background estimation. The ROI used was a top-hat mask three times the angular size of the Moon shadow, as determined via the Gaussian fit in the next section.

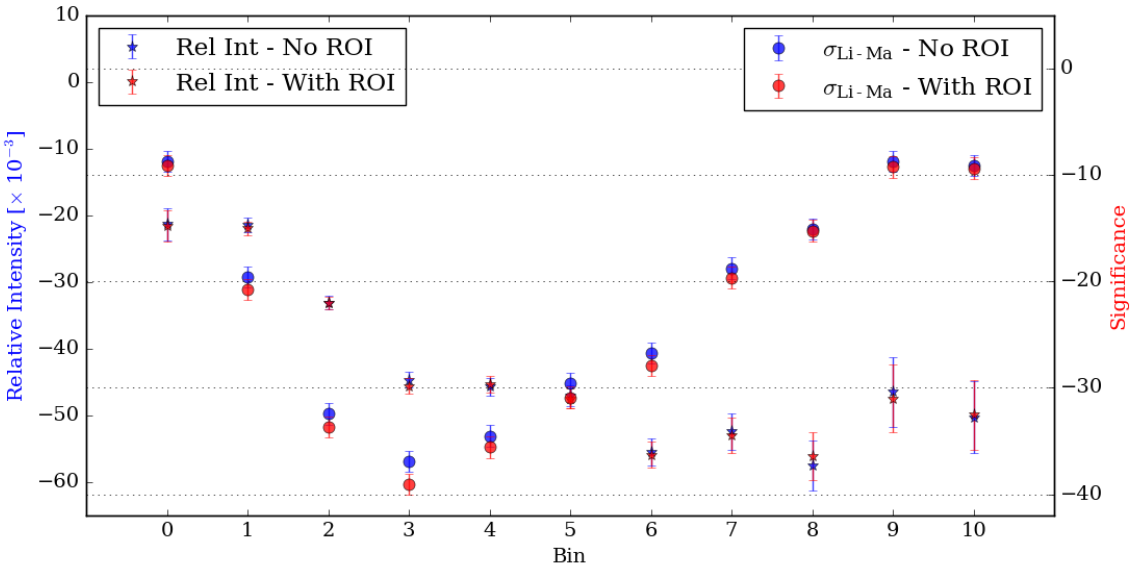

Figure 6.18: Peak relative intensity and significance values with and without ROI-masking for Moon maps smoothed by 1°. The analysis bins are used on the abscissa for clarity of the figure. Thus we see that while the relative intensity magnitude does not change more than a few percent with ROI-masking, a gain of several signifiance units is observed in bins with the largest statistics.

6.3.3.3 Deficit Statistics

The measured width of the Moon shadow serves as an experimental verification of the angular resolution of HAWC. We treat the Moon as a point-like cosmic ray sink that blocks a number of events commensurate with its angular extent, $r_{\text{Moon}} = 0.26°$. We then compare the expected number of blocked events as approximated via the Gaussian fit of the shadow to the observed deficit as calculated from the data map itself. In the equations that follow, x represents a right ascension declination pair, (α, δ).

The expected number of events blocked by the Moon, N_{blocked}, is calculated from the estimated background map:

$$N_{\text{blocked}} = \sum_i \langle N \rangle_i \frac{\pi r_{\text{Moon}}^2}{A_{\text{pixel}} \cdot \sum_i i}\,, \qquad \{i : |x_i - x_{\text{Moon}}^{\text{true}}| < r_{\text{Moon}}\}\,, \tag{6.10}$$

where the summations cover the i pixels within the Moon's true radius and the area per pixel is given by

$$A_{\text{pixel}} = \frac{4\pi}{4N_{\text{side}^2}}\,.$$

The **expected** deficit, $N_{\rm exp}$, at a radius r from the Moon's true position is determined from the fraction of blocked events from the Gaussian shadow observed:

$$N_{\rm exp}(r) = -N_{\rm blocked} \cdot \left[1 - \exp\left(-\frac{r^2}{2\sigma^2_{\alpha+\delta}}\right)\right] . \tag{6.11}$$

The **observed** deficit, $N_{\rm obs}$, is defined as the difference between the data counts N_i and the estimated background counts $\langle N\rangle_i$ within a specified radius r of the *observed* Moon centroid, $x^{\rm obs}_{\rm Moon}$. Hence,

$$N_{\rm obs}(r) = \sum_i \left(N_i - \langle N\rangle_i\right) , \qquad \{i : |x_i - x^{\rm obs}_{\rm Moon}| < r\} . \tag{6.12}$$

Equation 6.12 represents an integral deficit and should converge to the total number of expected blocked events from the Moon for large radii, i.e.

$$N_{\rm obs}(r) \to N_{\rm exp}(r) = N_{\rm blocked} \qquad \text{for} \qquad r >> r_{\rm Moon} .$$

In practice, however, the small angle approximation permitting the use of the Gaussian assumption becomes less viable with increasing r from the source, even just a few degrees from the Moon shadow centroid. Thus, a more appropriate measure is the differential deficit, whereby data counts are integrated in rings centered on the signal, and we normalize this by the background in those rings. For a region contained by two circles a, b where $r_a < r_b$, we define the expected and observed differential fractional deficit as follows:

$$\delta N_{\rm exp}(r) = \frac{\Delta N_{\rm exp}(r)}{\Delta N_{\rm blocked}} = \frac{N_{\rm exp}(r_b) - N_{\rm exp}(r_a)}{N^{i\in b}_{\rm blocked} - N^{i\in a}_{\rm blocked}} \tag{6.13}$$

$$\delta N_{\rm obs}(r) = \frac{\Delta N_{\rm obs}(r)}{\Delta N_{\rm blocked}} = \frac{N_{\rm obs}(r_b) - N_{\rm obs}(r_a)}{N^{i\in b}_{\rm blocked} - N^{i\in a}_{\rm blocked}} , \tag{6.14}$$

where now the indices i are changed to cover pixels within the ring, i.e.

$$\{i\} \to \{i : r_a < |x_i - x^{\rm obs,true}_{\rm Moon}| < r_b\} .$$

Examples comparing the expected Gaussian to the observed deficit for both integral and differential calculations are shown in figure 6.19. The resulting reduced χ^2 values comparing the expected and observed differential excesses for all analysis bins are presented in table 6.2. Finally, the $\sigma_{\alpha+\delta}$ are compared to the angular resolution as calculated via simulation (shown in figure 6.20, first presented in section 3.4 in figure 3.10). The observed consistency between the two implies that the absolute pointing accuracy of the detector behaves as expected from simulation with respect to energy.

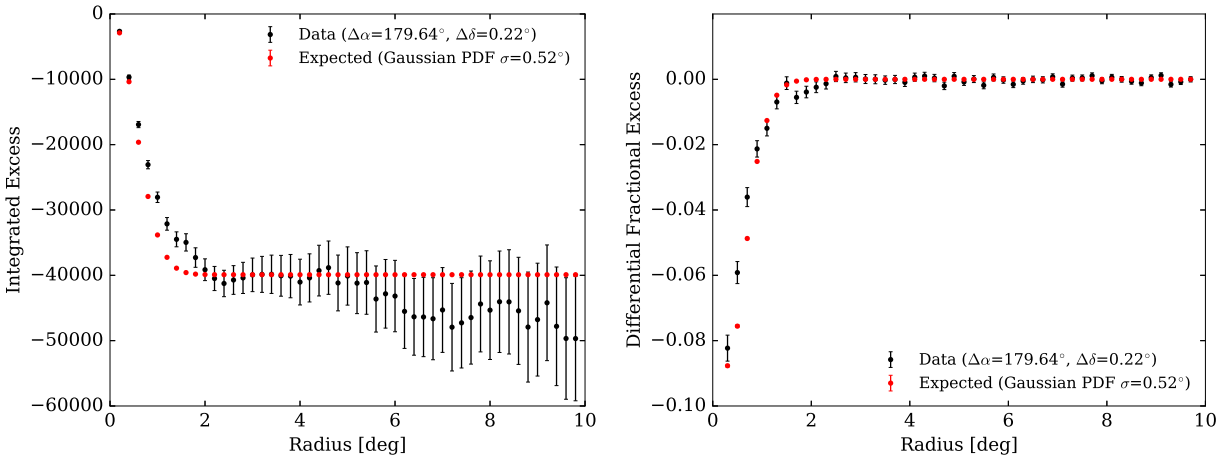

Figure 6.19: Example of expected and observed integral (left) and differential (right) deficits for bin 4, using the unsmoothed Moon shadow map as in figure 6.12 over a range of 10° from the Moon centroid. Beyond a few degrees, the observed integral deficit can fluctuate about the Gaussian expectation, so a more robust measure is the differential fractional deficit.

Analysis Bin	$\sigma_{\alpha+\delta}$ [deg]	χ^2/dof
0	0.94 ± 0.23	34.42/48
1	0.81 ± 0.08	148.99/48
2	0.64 ± 0.04	171.07/48
3	0.57 ± 0.03	52.94/48
4	0.52 ± 0.03	94.61/48
5	0.46 ± 0.03	93.98/48
6	0.52 ± 0.04	49.25/48
7	0.47 ± 0.05	22.99/48
8	0.48 ± 0.06	42.62/48
9	0.36 ± 0.07	54.96/48
10	0.54 ± 0.10	24.10/48

Table 6.2: Mean angular width $\sigma_{\alpha+\delta}$ from centroid fits to ROI-masked maps for each analysis bin, with reduced χ^2 values comparing the observed and expected differential fractional deficits. As shown in figure 6.19, a total of 48 radial bins were used to evaluate $\delta N_{\mathrm{exp,obs}}(r)$ over the range $0° < r < 10°$.

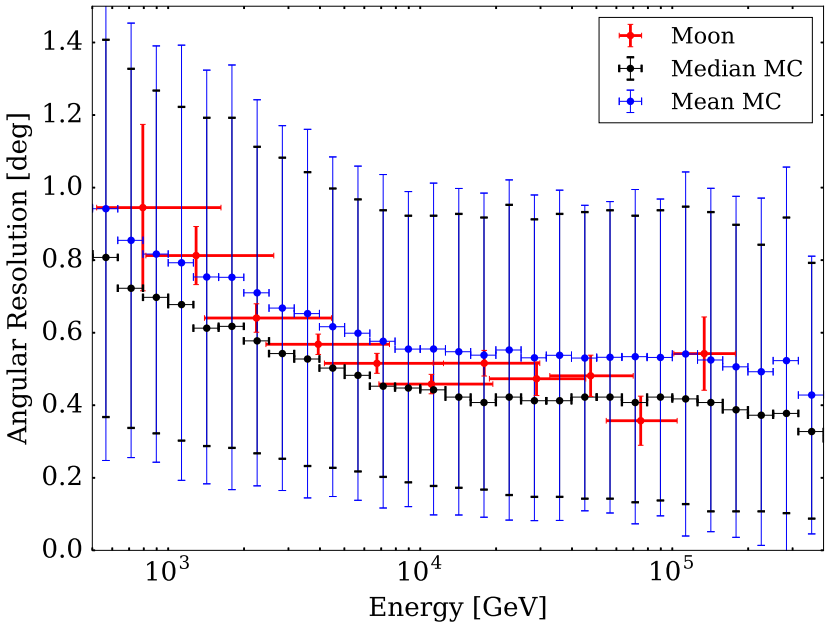

Figure 6.20: Comparison of the measured Gaussian widths $\sigma_{\alpha+\delta}$ of the Moon shadow from equation 6.9 (red) to the expected mean ± standard deviation (blue) and median ± 34% (black) angular resolution from simulation. The vertical bars on the Moon data points are the uncertainties on $\sigma_{\alpha+\delta}$ from the fit.

Chapter 7

Iterative Unfolding

7.1 Inverse Problems

The general class of inverse methods is amongst the physicist's toolbox as a powerful means to connect an experiment's observable variables with true physical quantities. Typically a matrix can be built to encompass the effects of the measurement process on a simulated 'true' distribution and the manifestation of said distribution as an experimenter's desired observable. With this response matrix, a distribution of the observable in an experiment can be **unfolded**, providing an estimate of the true parent distribution.

A variety of unfolding methods exist, each with its respective strengths and weaknesses. For example, the simplest method is the matrix inversion unfolding, which for a well populated, highly linear response matrix can be both efficient and precise. However, even with relatively small off-diagonal elements, this method is unfavorable, as the matrix may be singular or may introduce wildly fluctuating results due to limited statistics. There exist methods to quell such issues, though these require the tuning of various parameters which typically have no physical connection to the experiment at hand.

Here we discuss D'Agostini's unfolding technique presented in [91], a manifestly **inferential** method, and its relation to the cosmic-ray energy reconstruction for the HAWC experiment. Starting from Bayes' theorem, an iterative unfolding procedure is developed, which then can be implemented without too much difficulty for the typical experimenter.

7.2 D'Agostini Unfolding

7.2.1 Method

As discussed in the introduction, the conceptually simplest way to connect true (causes, C_μ) and observable (effects, E_j) variables is via a matrix, R, and its inverse M[1]:

$$\begin{aligned} n(E) &= R\,\phi(C)\,, \\ \phi(C) &= M\,n(E)\,. \end{aligned} \tag{7.1}$$

Due to the aforementioned potential difficulties in matrix inversion, we can take into consideration Bayes' theorem,

$$P(C_\mu|E_j) = \frac{P(E_j|C_\mu)\,P(C_\mu)}{\sum\limits_{\nu}^{n_C} P(E_j|C_\nu)\,P(C_\nu)}\,, \tag{7.2}$$

where n_C is the number of possible causes. Equation 7.2 dictates that having observed the effect E_j, the probability that its origin is due to the cause C_μ is proportional to product of the probability of the cause and the probability of the cause to produce that effect. Hence, the elements $P(E_i|C_\mu)$ represent the probability that a given C_μ results in the effect E_i, and is the response matrix generated via MC simulation. Continuing with $P(C_\mu|E_j)$, we can then connect the measured observed effects to their causes via:

$$\phi(C_\mu) = \sum_{i}^{n_E} P(C_\mu|E_i) n(E_i)\,. \tag{7.3}$$

Stepping back to equation 7.2 for a moment, one identifies $P(C_\mu)$ as the prior cause distribution, representing our current knowledge of the causes. The prior is a normalized distribution such that $\sum\limits_{\mu}^{n_C} P(C_\mu) = 1$. This normalization requirement is not imposed on the response matrix efficiency ϵ_μ: $0 \leq \epsilon_\mu = \sum\limits_{j}^{n_E} P(E_j|C_\mu) \leq 1$, i.e., a cause does not need to produce any effect. Taking this (in)-efficiency into account, we rewrite 7.3 as

$$\phi(C_\mu) = \frac{1}{\epsilon_\mu} \sum_{i}^{n_E} P(C_\mu|E_i) n(E_i)\,. \tag{7.4}$$

Identifying here the explicit form of M, the full matrix (Bayesian) inversion equation is then

$$\phi(C_\mu) = \sum_{j}^{n_E} M_{\mu j} n(E_j)\,, \tag{7.5}$$

[1] Except for C and E, all variables and subscripts related to causes are Greek letters, while Latin letters are used for effects. The only superscript is the iteration number, i.

where

$$M_{\mu j} = \frac{P(E_i|C_\mu)\, P(C_\mu)}{\left[\sum_k^{n_E} P(E_k|C_\mu)\right]\left[\sum_\nu^{n_C} P(E_i|C_\nu)P(C_\nu)\right]}\,. \tag{7.6}$$

The response matrix $P(E_i|C_\mu)$ is generated via simulation, and the $n(E_i)$ provided through measurement, apparently bestowing the freedom to choose the form of $P(C_\mu)$. Again, $P(C_\mu)$ represents the total of our prior knowledge of the parent distribution. Typically an experimenter refrains from introducing bias in the prior. The energy ranges of interest in HAWC span several decades, so the appropriate choice is the Jeffreys Prior [92]:

$$P_{\text{Jeffrey}}(C_\mu) = \frac{1}{\log\left(C_{\max}/C_{\min}\right) C_\mu}\,,$$

keeping in mind that this prior dictates that all cause bins are of equal probability, not that all parent distributions are of equal probability.

We now possess all the necessary machinery to perform an unfolding. Having started with the unbiased Jeffreys Prior, the unfolded result is a best estimate of the true distribution. There is nothing stopping us from using this result as the best knowledge estimate of $P(C_\mu)$ in equation 7.6 for a subsequent unfolding. We can take this any number of steps further, making the process an iterative unfolding. Thus, after calculating $\phi(C_\mu)$ via equation 7.5, we recalculate $M_{\mu j}$ per equation 7.6, returning again to equation 7.5 for an updated $\phi'(C_\mu)$. Since $P(C_\mu) = \frac{\phi_\mu}{\sum_\nu \phi_\nu} = \frac{\phi_\mu}{N_{\text{true}}}$, where N_{true} is the estimated true number of cause events, we can make the change $P(C_\mu) \to \phi_\mu$ in equation 7.6. Adding the iteration superscript and shortening the notation[1], this equates to

$$M_{\mu j} = \frac{P_{\mu j}\, \phi_\mu^i}{\epsilon_\mu \sum_\rho P_{\rho j}\, \phi_\rho^i}$$

$$\phi_\mu^{i+1} = \sum_j M_{\mu j}\, n_j\,.$$

The unfolding proceeds until a desired stopping criterion is satisfied, say by comparing subsequent iterations with a χ^2. The algorithm below outlines the basics to the iterative unfolding scheme:

Algorithm 1 Unfolding Algorithm

$\phi^0 \leftarrow$ Prior

testStatistic$\leftarrow$ Pass

while (testStatistic = Pass) **do**

$M \leftarrow M(P(E|C), \phi^i)$

$\phi^{i+1} \leftarrow M \times n$

testStatistic$\leftarrow$ TS(ϕ^i, ϕ^{i+1})

end while

7.2.2 Unfolding Uncertainties

To begin the excursion into the calculation of uncertainties, we first shorten the notation:

$$P(E_i|C_\mu) = P_{\mu i} \qquad \phi(C_\mu) = \phi_\mu \qquad n(E_j) = n_j.$$

As outlined in [91] (section 4), the covariance matrix $V = V(\phi, \phi')$ from statistical contributions has two components: V^{Data} from the counted measured effects distribution, and V^{MC} due to the limited Monte Carlo statistics in $P_{\mu j}$. This can be seen from considering the uncertainties from n_j and $M_{\mu j}$ in equation 7.5. Since $\phi = M \times n = M(P(E|C)) \times n$, we can identify respectively the aforementioned error contributions as

$$\begin{aligned} V^{\text{Total}} &= V^{\text{Data}} + V^{\text{MC}} \\ &= \frac{\partial \phi}{\partial n} \text{cov}(n, n') \frac{\partial \phi'}{\partial n} \\ &+ \frac{\partial \phi}{\partial P} \text{cov}(P, P') \frac{\partial \phi'}{\partial P}. \end{aligned}$$

7.2.2.1 V^{Data}

D'Agostini argues that since the data sample n_j is a realization of a multinomial distribution, then

$$V^{Data} = M \, \text{cov}(n, n') \, M \tag{7.7}$$

where the $\text{cov}(n, n')$ is the covariance matrix of the measurements with respect to the estimated true number of events $\sum_\mu \phi_\mu = N_{\text{true}}$:

$$\text{cov}(n_k, n_j) = \begin{cases} n_j \left(1 - \frac{n_j}{N_{\text{true}}}\right) & \text{if } k = j \\ -\frac{n_j n_k}{N_{\text{true}}} & \text{if } k \neq j \end{cases}. \tag{7.8}$$

However, Adye ([93] section 5) demonstrates that this error estimation is only valid for the first iteration, as subsequent ϕ^i are **not independent** of n_j. Indeed, we should re-write equation 7.7 appropriately as

$$V^{Data} = \frac{\partial \phi^{i+1}}{\partial n} \, \text{cov}(n, n') \, \frac{\partial \phi^{i+1\prime}}{\partial n}, \tag{7.9}$$

with

$$\frac{\partial \phi_\mu^{i+1}}{\partial n_j} = M_{\mu j} + \frac{\phi_\mu^{i+1}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial n_j} - \sum_{\sigma,k} \epsilon_\sigma \frac{n_k}{\phi_\sigma^i} M_{\mu k} M_{\sigma k} \frac{\partial \phi_\sigma^i}{\partial n_j},$$

where again the superscripts i and $i+1$ refer to the iteration number. The full derivation of $\frac{\partial \phi^{i+1}}{\partial n}$ is found in section B.2. Furthermore, as measuring the cosmic ray spectrum involves counting statistics, it is safe to use the Poisson form of $\mathrm{cov}(n, n')$:

$$\mathrm{cov}(n_k, n_j) = n_k \, \delta_{kj}. \tag{7.10}$$

7.2.2.2 V^{MC}

The contribution from V^{MC} is outlined in [91] and fully derived in section B.3. If one simply implements the equation verbatim into code, the expected time for calculating all elements $\sim$ (number of bins)^7. Thus, here we present the form of V^{MC}, while in section B.3 we show the explicit expansion and further contraction of indices towards a more reasonable, practical calculation.

D'Agostini identifies V^{MC} via $\frac{\partial}{\partial M}$ giving

$$V^{\mathrm{MC}} = n \, \mathrm{cov}(M, M') \, n'. \tag{7.11}$$

Further expansion reveals

$$\mathrm{cov}(M_{\mu k}, M_{\lambda j}) = \sum_{\{\sigma r\},\{\sigma s\}} \frac{\partial M_{\mu k}}{\partial P_{\sigma r}} \frac{\partial M_{\lambda j}}{\partial P_{\sigma s}} \mathrm{cov}(P_{\sigma r}, P_{\sigma s}), \tag{7.12}$$

$$\frac{\partial M_{\mu k}}{\partial P_{\sigma j}} = M_{\mu k} \Big[\frac{\delta_{\mu\sigma} \, \delta_{jk}}{P_{\sigma j}} - \frac{\delta_{\mu\sigma}}{\epsilon_\sigma} - \frac{\delta_{jk} \, M_{\sigma k} \, \epsilon_\sigma}{P_{\sigma k}} \Big], \tag{7.13}$$

$$\mathrm{cov}(P_{\sigma r}, P_{\sigma s}) = \begin{cases} \frac{1}{\tilde{n}_\sigma} P_{\sigma r}(1 - P_{\sigma r}) & \text{if } r = s \\ -\frac{1}{\tilde{n}_\sigma} P_{\sigma r} P_{\sigma s} & \text{if } r \neq s \, . \end{cases} \tag{7.14}$$

In the final expression, $\tilde{n}_\mu$ represents the number of MC events which fell into the true cause bin μ. As each event in the simulation is weighted by a factor w, we identify $\tilde{n}$ with the effective number of events $\tilde{n}_\mu = \frac{(\sum_j w_{\mu j})^2}{\sum_j w_{\mu j}^2}$ for all j events in bin μ.

Once again, Adye [94] shows this is a first order estimate, only valid for the first iteration. Re-writing 7.11 with $\frac{\partial}{\partial P}$,

$$V^{\mathrm{MC}} = \frac{\partial \phi^{i+1}}{\partial P} \mathrm{cov}(P, P') \frac{\partial \phi^{i+1\prime}}{\partial P}, \tag{7.15}$$

we identify $\frac{\partial \phi^{i+1}}{\partial P}$ as

$$\frac{\partial \phi_\mu^{i+1}}{\partial P_{\lambda k}} = \frac{\delta_{\lambda\mu}}{\epsilon_\mu}\left(\frac{n_k\,\phi_\mu^i}{f_k} - \phi_\mu^{i+1}\right) - \frac{n_k\,\phi_\lambda^i}{f_k} M_{\mu k} + \frac{\phi_\mu^{i+1}}{\phi_\mu^i}\frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} - \sum_{\rho,j} n_j\,\frac{\epsilon_\rho}{\phi_\rho^i}\,M_{\rho j}\,M_{\mu j}\,\frac{\partial \phi_\rho^i}{\partial P_{\lambda k}}\,,$$

whose derivation is presented in section B.3 below. Of course, D'Agostini's form of $\mathrm{cov}(P, P')$ remains valid for use with the new construction of the partials. One may also use a Poisson covariance if justified appropriately:

$$\mathrm{cov}(P_{\rho r}, P_{\lambda s}) = \sigma_{\rho r}\sigma_{\lambda s}\delta_{\rho\lambda}\delta_{rs}, \tag{7.16}$$

with $\sigma_{\rho r}$ being the error estimates on $P_{\rho r}$ estimated when filling P with Monte Carlo simulation.

7.2.2.3 Updated Unfolding Algorithm

The afore-outlined unfolding algorithm must be modified to include the propagation of systematic uncertainties. At each iteration we have ϕ^{i+1}, so both $\frac{\partial \phi^{i+1}}{\partial n}$ and $\frac{\partial \phi^{i+1}}{\partial P}$ can be calculated. The results are propagated and saved until the full covariance matrix is required for uncertainty estimates on the final ϕ.

Algorithm 2 Unfolding Algorithm - Including Errors

$\phi^0 \leftarrow$ Prior
testStatistic$\leftarrow$ Pass
while (testStatistic = Pass) **do**
  $M \leftarrow M(P(E|C), \phi^i)$
  $\phi^{i+1} \leftarrow M \times n$
  $\frac{\partial \phi^{i+1}}{\partial n} \leftarrow$ eq. B.4
  $\frac{\partial \phi^{i+1}}{\partial P} \leftarrow$ eq. B.5
  testStatistic$\leftarrow$ TS(ϕ^i, ϕ^{i+1})
end while
$V^{\mathrm{Total}} \leftarrow V^{\mathrm{Data}}(\frac{\partial \phi^{i+1}}{\partial n}) + V^{\mathrm{MC}}(\frac{\partial \phi^{i+1}}{\partial P})$
$\sigma_\phi^2 \approx \mathrm{diag}(V^{\mathrm{Total}})$

7.3 Regularization

After each iteration, the resulting posterior distribution, $P(C_\mu)$, is our new best guess of the (normalized) parent distribution. Using this best estimate as the prior for the next iteration, one can induce large fluctuations in neighboring C_μ bins. It is here the equivalence of matrix inversion techniques and iterative unfolding is seen. After many iterations, wild fluctuations can appear, indicating the granularity in the MC derived $P_{\mu j}$. Furthermore, in using the posterior as the subsequent prior, one is "telling" the unfolding that physical distributions of that nature are allowable priors. Instead, as pointed out in [91] (section 6.3), for an experimenter interested in a particular model's parameters, fitting all but the last posterior is equivalent to performing a maximum likelihood fit to the data.

As physics measurements are expected to be smooth (a safe assumption for the cosmic-ray energy spectrum for example), one can regularize the ϕ^i_μ. In principle one can choose any smoothing function. For the cosmic ray energy spectrum, ϕ^i_μ can be simply fit to a power law, using the fitted function as the input prior for the next iteration. While this could be seen as a loss of information, it is important to remember that this method is based on Bayes' theorem, so **any** improved prior distribution will enhance our estimation method, along with the **prior** expectation that our physics is smooth.

The other possibility is to avoid regularization altogether, and instead ensure that $P_{\mu j}$ is smooth enough. The granularity of the cause and effect bins will dictate the degree of smoothness required to ensure non-fluctuating ϕ^i solutions. The more widely used techniques for smoothing $P_{\mu j}$ include kernel density estimation and penalized spline fitting routines.

Chapter 8

The All-Particle Cosmic-Ray Spectrum

The result of measuring the differential all-particle energy spectrum at TeV energies is presented in this chapter. First we introduce the details of the chosen data set, the event selection criteria, and building the detector response function. Next we show the process of unfolding the all-particle spectrum using the iterative method from chapter 7, and discuss sources of systematic uncertainties. These results then are compared to other experimental measurements spanning the TeV−PeV energy scale.

8.1 Data Set and Event Selection

The HAWC detector in its full configuration of 290 WCDs was in stable data taking mode during the runs selected for this analysis, amounting to a total of 234 days from 8 June, 2016 to 17 February, 2017. The total up-time efficiency was $\sim 92\%$ and the mean trigger rate was ~ 25 kHz. Only runs with angular reconstruction stability with durations greater than 2 hrs (as described in section 6.3.1) were chosen. Furthermore, only runs where the number of active PMTs was above 1100 were considered (figure 8.1), as the HAWC detector simulation employed 1083 PMTs. This amounts to no more than a 2% difference in the effective area estimation and only in the case that the discrepancy in detector configuration is along the perimeter of the array.

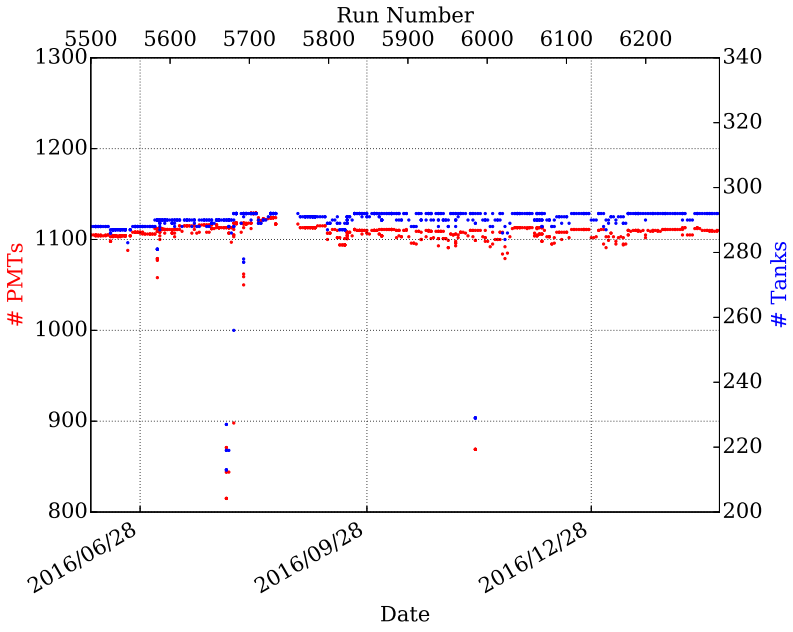

Figure 8.1: Detector configuration stability as a function of time and corresponding run number. Only runs where the number of active PMTs was greater than 1100 were selected, as the nominal detector configuration in simulation used 1083. This corresponded to the last run (5481) when PMT calibration curves were generated prior to the selected data set.

In accordance with previous cosmic ray analyses [1], selected events passed a multiplicity threshold of $\sim 6\%$ of the detector. For the nominal detector configuration during this period, the resulting multiplicity cut was $N_{\text{hit}} \geq 75$. The $N_{r40} \geq 40$ selection cut from section 3.4 was used to ensure the integrated core fit resolution is below 10 m for energies above 10 TeV. Increasingly harsh N_{r40} cuts were also applied to test

the robustness of the cut, with values between 40–100 in steps of 10. The estimator developed in chapter 4 defined the reconstructed energy variable, and we used the fast grid search method for the maximum likelihood finding algorithm. Finally, we only consider the first zenith bin θ_0 where shower development is least sensitive to fluctuations due to the smaller atmospheric overburden, which gives the θ_0 bin the best energy resolution above 10 TeV (figure 4.14). The CREAM model described in section 3.3.1 defined the nominal composition used to build the detector response function. Thus, the following defines the event selection criteria: core and angle fit success, $N_{\rm hit} \geq 75$, $\theta \in \theta_0$ and $N_{r40} \geq 40$. Table 8.1 compares the passing rates from applying these cuts to events in simulation and data.

Prior to constructing the detector response function, the distributions of reconstructed variables involved in the event selection are reviewed to ensure that the applied cuts are geometrical in nature and do not introduce bias in an energy dependent manner. As can be seen from the normalized distributions shown in figures 8.2 – 8.4, neither the Monte Carlo nor the data sets exhibit unexpected structure with respect to the variables of interest, and agree to within the Monte Carlo statistical uncertainties. We can also consider the reconstructed energy distribution itself, which is the quantity to be unfolded; however, we do not necessarily expect that the Monte Carlo and data should match to within at least the different composition models. The systematics from the composition assumptions are discussed in section 8.4.4.

To get a sense of the scale involved with these events compared to HAWC, figure 8.5 presents three sample air shower events from the data set. Each is a typical event from its respective reconstructed energy decade, showing increasing multiplicity and saturation of the detector up to PeV energies. For the lowest energy event, the largest signals defining the core region just span the 40 m containment area, and most of the PMTs within have measured signal. Approximately a third of the array's PMTs do not contain hits. The $E_{\rm reco} = 100$ TeV event spans the entire detector, and the core region extends beyond the 40 m area. For this event, we also see the clumpy depositions indicative of subshowers and energy being transversely propagated away from the core during shower development.

Finally, the most energetic shower completely saturates the detector with large charge deposition across the entire array. Indeed most of the hits are above 100 PEs. Close inspection of the core region reveals PMTs without signal. This is the result of hits which are beyond the calibrated range of the PMT ($q > q_{\rm max}$), and are "dropped" from the reconstruction algorithms, save for the energy calculation. This is taken into account by the knowledge that there was a hit of at least $q_{\rm max}$, so the energy estimation table contributes the sum beyond this value to the likelihood calculation, i.e. $\sum_{q_i > q_{\rm max}} \mathscr{L}(q_i, r_i, \theta | E)$. Despite being used in the energy estimation, these dropped hits are lost to the determination of N_{r40}. Thus, this selection criterion can influence the efficiency for the highest energy events. We find that the energy above which this effect appears in reconstructed data is 600 TeV.

Cut	Passing Percentage		Data Event Rate
	Simulation	Data	[kHz]
No cut (trigger threshold)	100 %	100 %	24.7
Core & angle fit success	99 %	96 %	23.6
$N_{\text{hit}} \geq 75$	31 %	23 %	5.7
$\theta \in \theta_0$	8 %	6 %	1.5
$N_{r40} \geq 40$	2 %	2 %	0.43

Table 8.1: Passing percentages for successive application of event quality cuts in simulation and data, including the observed event rate in data. The percentages represent the fraction of events that passed the previous cut, with the set of triggered events being the reference selection.

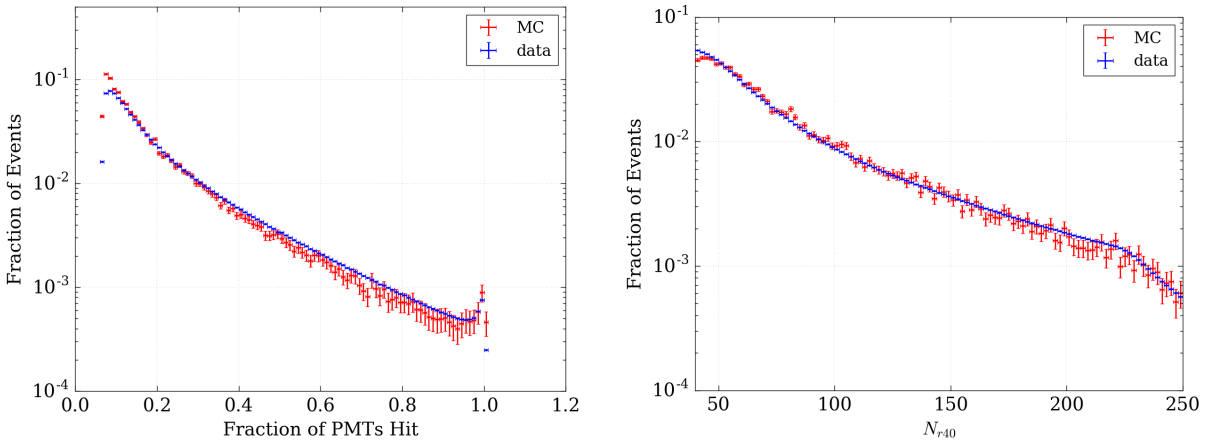

Figure 8.2: Data to Monte Carlo comparisons of reconstructed multiplicity distributions $N_{\text{hit}}/N_{\text{PMTs}}$ and N_{r40}, having applied the full event quality cuts as listed in table 8.1.

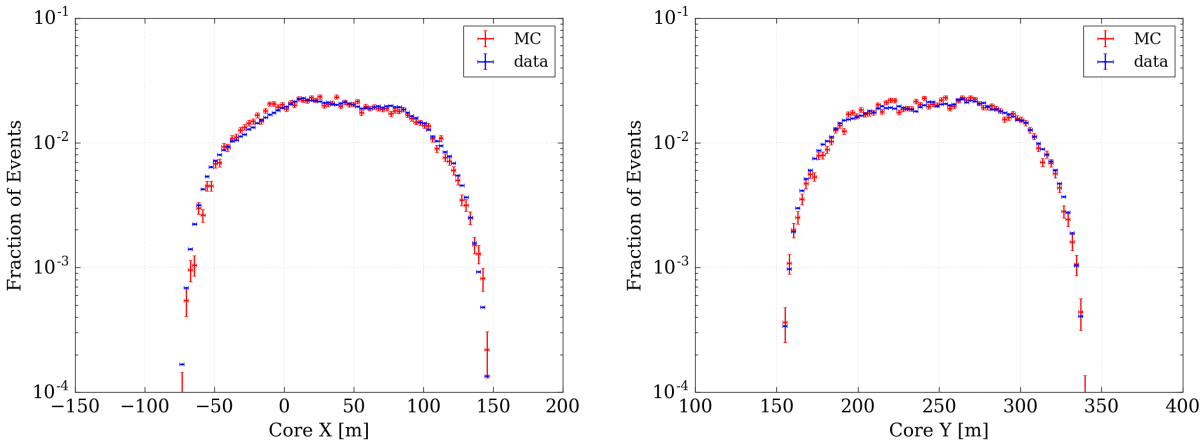

Figure 8.3: Data to Monte Carlo comparisons of reconstructed core position distributions, having applied the full event quality cuts as listed in table 8.1.

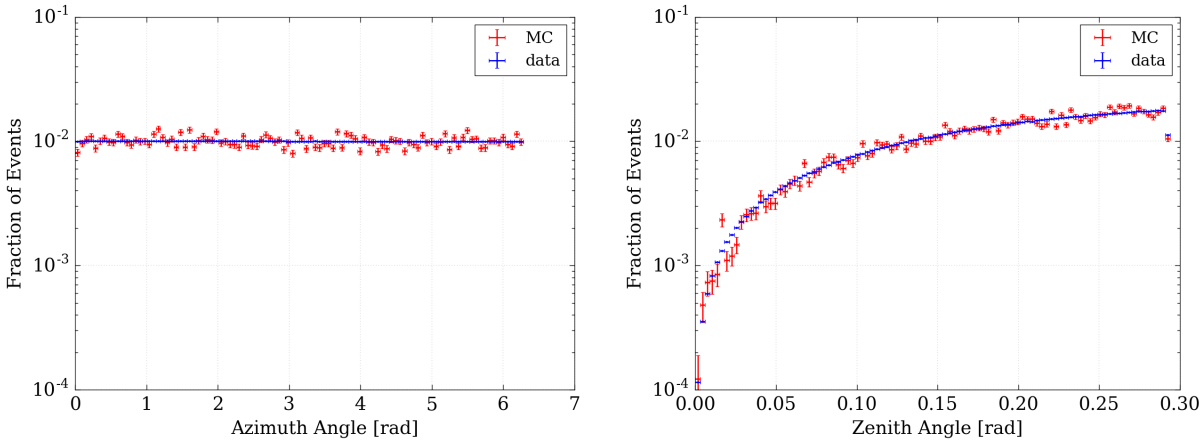

Figure 8.4: Data to Monte Carlo comparisons of reconstructed direction distributions θ, ϕ, having applied the full event quality cuts as listed in table 8.1.

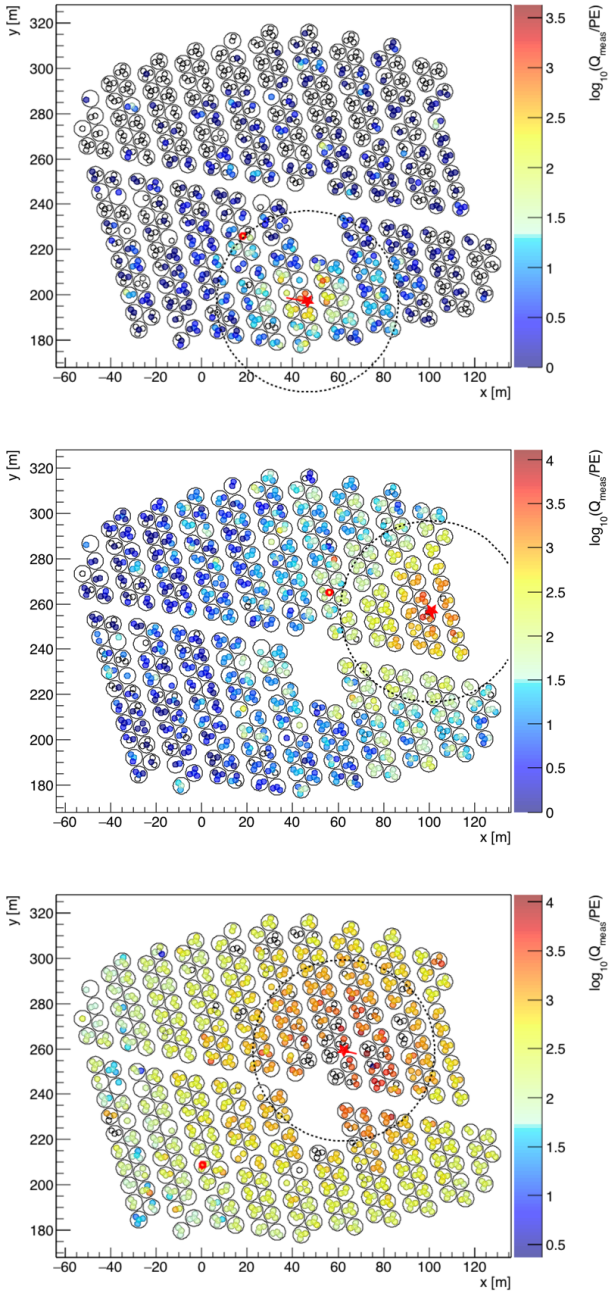

Figure 8.5: Sample air shower event charge distributions from the data sample. The reconstructed energy of the events increases going down the page from top to bottom with E_{reco} values of 10 TeV, 100 TeV, and 1 PeV, where the uncertainty on each is 0.1 in $\log E_{\text{reco}}$. The core location is indicated by the red star, and the 40 meter circle shows the region used for N_{r40}.

8.2 Detector Response

In order to implement the unfolding method of chapter 7, a detector response function must be constructed. This function represents the total knowledge of the detector relating the reconstructed variable to the true quantity of interest. For these spectral measurements, this object is a matrix built from simulation which connects the probabilities of reconstructed energies to the true energies. Indeed, it encapsulates both the efficiency to reconstruct a shower of energy E, and migration to a reconstructed energy $E_{\rm reco}$, with probability given by $P(E_{\rm reco}|E)$. Hence, in accordance with the terminology of the previous chapter, we make the assignment of the causes as $C \rightarrow E$ and the effects as $E \rightarrow E_{\rm reco}$.

Since the detector response function is built from a finite amount of simulation, it is subject to statistical fluctuations. The efficiency $\epsilon(E)$ is smoothed using a spline function so that these fluctuations do not manifest as spurious features in the unfolded spectrum. Figure 8.6 shows the simulated efficiency compared to the resulting fit which is used in the unfolding procedure. The detector efficiency grows rapidly up to 10 TeV, then continues to rise by $\sim 10\%$, reaching a maximum efficiency of about 1.45×10^{-2} above 40 TeV. This value corresponds to an effective area of $\sim 45,000$ m^2, a little over twice the detector area ($A_{\rm det} = 22,000$ m^2). This is expected for the most lenient N_{r40} cut, as showers beyond 40 m from the array edges are being reconstructed closer to the array, boosting $A_{\rm eff}$. This can be compared to the effective area first shown in figure 3.12, where a larger zenith angle range resulted in an $A_{\rm eff}$ closer to $A_{\rm det}$. However, this is due to the decreased efficiencies for events with greater zenith angles, which reduces the mean $A_{\rm eff}$.

The response matrix $P(E_{\rm reco}|E)$ provided the event selection cuts from table 8.1 is shown in figure 8.7. Above 10 TeV, its structure is increasingly linear with true energy, with migration limited to a couple of bins from the diagonal above 200 TeV. Of course, this is simply a two-dimensional representation of the energy bias and resolution from θ_0 in figure 4.14. However, in the response matrix, the normalization within each E bin equals the total efficiency as first suggested in figure 4.11: $\sum_{E_{\rm reco}} P(E_{\rm reco}|E) = \epsilon(E)$. This efficiency is in fact a composition-weighted average of the individual species' efficiencies considered to build the response matrix. Figure 8.8 shows the efficiencies for proton, helium, and iron which upon convolution with the CREAM fits from section 3.3.1 provide the all-particle efficiency. Indeed, it is necessary to assume a composition model in order to build the response matrix for the all-particle spectral measurement. The effects from other models are taken into account as systematic uncertainties, presented in section 8.4.4.

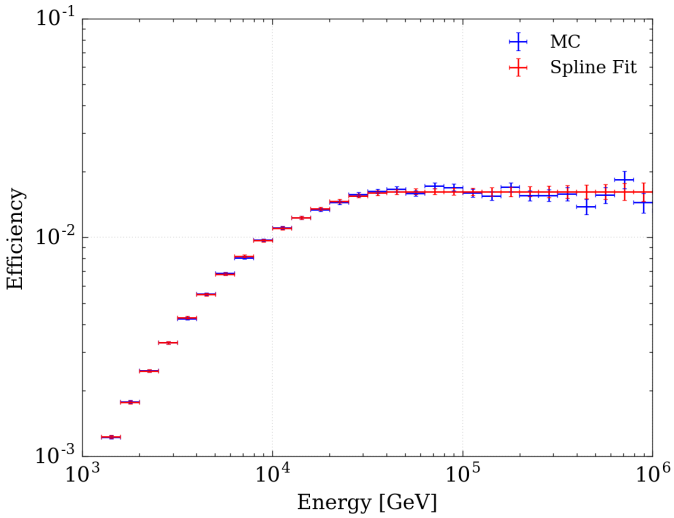

Figure 8.6: Fit to the simulated all-particle efficiency as a function of energy. Within the Monte Carlo (MC) statistical uncertainties, the simulated efficiency is flat above 30 TeV. The fit efficiency is the result of fitting an ensemble of splines to the Monte Carlo points, where for each fit, the "MC" values were varied according to their respective uncertainties. The uncertainties on the final fit efficiency curve represent the 68% containment about the median.

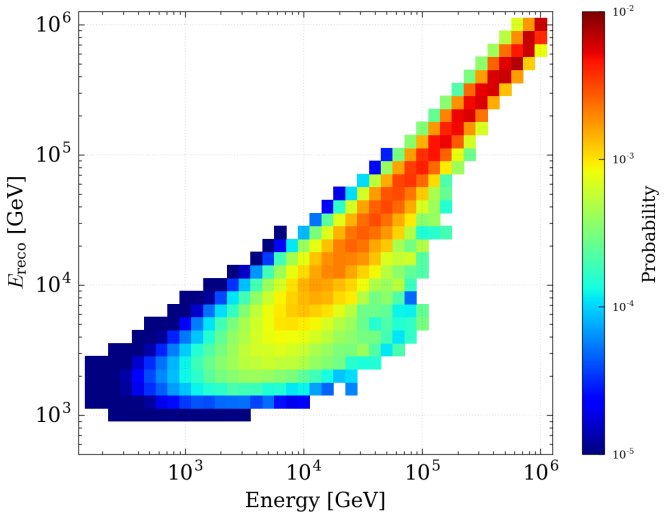

Figure 8.7: The all-particle response matrix, subject to the event quality cuts. The structure is increasingly linear and bin-to-bin migrations decrease with increasing energy, equivalent to improvement in the energy bias and resolution.

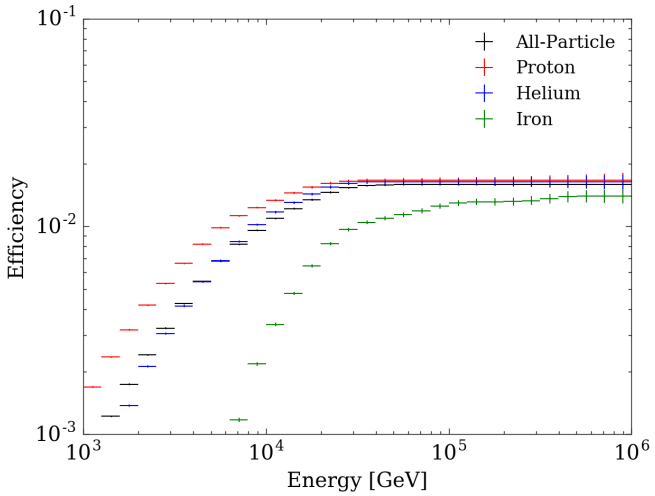

Figure 8.8: Comparison of the proton, helium, and iron efficiencies to the all-particle efficiency. Below $\sim$ 10 TeV, the all-particle efficiency lies between proton and helium, as they are the two most dominant components in this energy range and due to showers from heavier elements not reaching detector level. Above 10 TeV, the influence of the heavier components becomes apparent, as the all-particle efficiency is just below that of both proton and helium, yet above iron. The curves for the other species considered lie between helium and iron, and were omitted for clarity.

There is a tendency to under-reconstruct energies between $10 - 100$ TeV, as evidenced by the presence of more off-diagonal elements below rather than above the diagonal in figure 8.7. There are two potential contributions to this asymmetry: shower cores landing off of the array and the presence of heavier elements. This is found to be primarily the effect of showers with core locations that are reconstructed less well, from cores truly off of the array being reconstructed closer to the edges. This is shown in the response matrices of figure 8.9, where simulated showers with cores landing truly on and off of the array are selected. Since the energy tables are built from proton simulation, the energies of heavier elements are expected to be on average underestimated, as seen in figure 8.10. Yet, comparing the two sets of response matrices, improving the reconstructed core location is expected to have a greater impact on the linearity of the energy estimation than a composition cut.

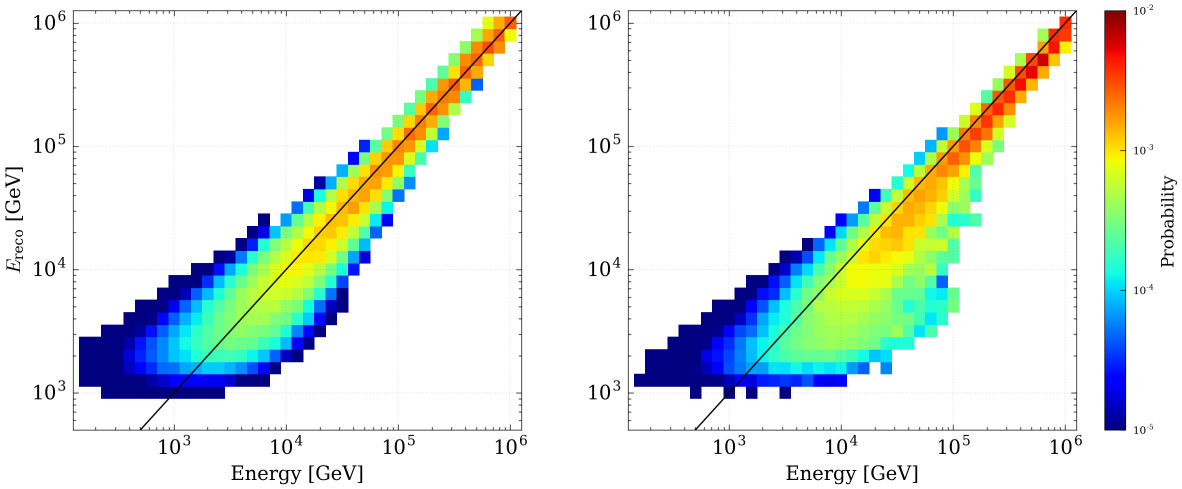

Figure 8.9: Response matrices for all species where the simulated core lands on (left) and off (right) of the array. For cores landing on the array, the energy estimation is symmetric above 5 TeV. The origin of the under-reconstructed energies appears to be from cores landing truly off of the array. The solid diagonal line provides a guide for the eye for 1:1 correspondence.

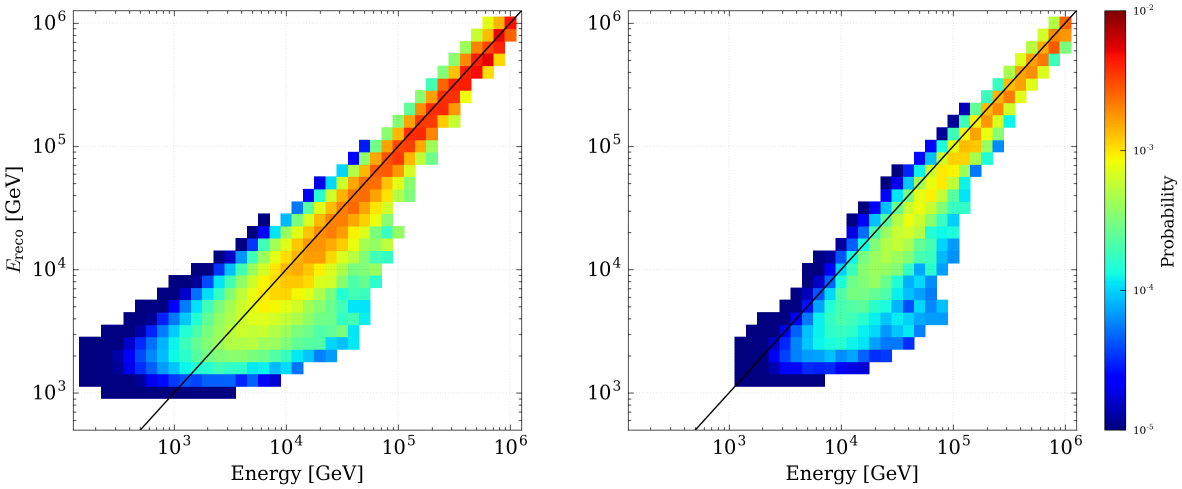

Figure 8.10: Response matrices for light and heavy species with standard event selection criteria. There is a marked turn down from the diagonal below 100 TeV, as reconstructed energies from heavier elements are more frequently underestimated.

8.3 Unfolding

The measured reconstructed energy distribution for the data set given the selection cuts is shown in figure 8.11, and the values are presented in table A.1. The unfolding procedure takes this distribution, and accounts for energy migration and efficiencies providing a distribution of expected counts binned in the true energy, $N(E)$. Of course, cosmic ray fluxes are measured as differential energy spectra. Recalling equation 1.2, we see that to convert the counts distribution to a flux, it is necessary to divide the unfolded $N(E)$ by the differential quantity $dA\,d\Omega\,dt\,dE$, where the differentials $d \to \Delta$ for finite binning, i.e.

$$dA\,d\Omega\,dt\,dE \to A_{\text{eff}}\,\Delta\Omega\,\Delta t\,\Delta E\,.$$

Here the Δt term is the total observational period in seconds. The spanned opening angle of the observation is given by the solid angle $\Delta\Omega = 2\pi(\cos\theta_{\text{min}} - \cos\theta_{\text{max}})$. The width of each energy bin gives ΔE. Finally, the effective area A_{eff} is summarized by the efficiency $\epsilon(E)$ scaled by the area A_{thrown} over which the simulated showers are distributed about the detector, and the inclusion of a geometric term as follows:

$$A_{\text{eff}}(E) = A_{\text{thrown}}\frac{\cos\theta_{\text{max}} + \cos\theta_{\text{min}}}{2}\epsilon(E)\,. \tag{8.1}$$

Since $\theta_{\text{min}} = 0$ and $\theta_{\text{max}} = 16.8°$, the geometric term has a small effect on A_{eff}. The unfolding method, however, already takes $\epsilon(E)$ into account as per equation 7.4, so only the geometric factor of equation 8.1 is required.

The unfolded differential spectrum given the reconstructed energy distribution of figure 8.11 is shown in figure 8.12, and the values are presented in table A.2. The unfolding converged in four iterations, where a Kolmogorov-Smirnov (KS) test [95, 96] defined the convergence criterion. If the KS test statistic comparing the unfolded distributions between iterations i and $i+1$ results in a p-value less than 0.001, the unfolding is said to have converged. For this unfolded energy spectrum, the Jeffreys prior was chosen as the starting distribution, and regularization via a spline fit was applied during the unfolding procedure. Other priors of power-law form were tested with negligible effect on the final spectrum. A starting prior such as a broken power law can induce structure for the first iteration; however, the spline regularization typically smooths this structure with subsequent iterations. Furthermore, we do not start with an *a priori* assumption that the spectrum has a break, so we always unfold starting with a Jeffreys prior. These conditions are used for all unfoldings unless otherwise stated.

At first glance, the unfolded spectrum appears to follow a power law across the depicted energy range. Indeed, it seems rather featureless over the flux versus energy scales provided. However, if we scale the flux by the energy raised to the power 2.6, as in figure 8.13, a slight bend in the spectrum is seen. Since the exponential scaling coefficient is 2.6, the rise of the flux below about 40 TeV implies that the spectrum is

harder than $E^{-2.6}$, while above 50 TeV there it appears slightly softer. This apparent bend in the spectrum will be discussed in section 8.5.2.

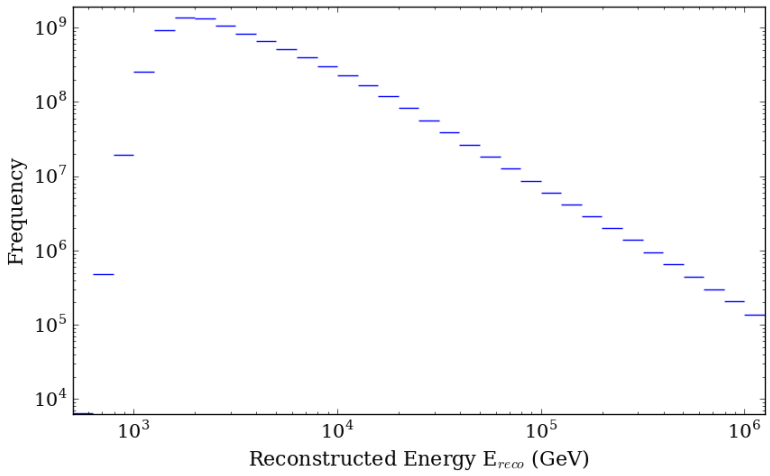

Figure 8.11: Observed distribution of E_{reco} for the data set using $N_{r40} \geq 40$. The distribution is a convolution of the detector efficiency and the cosmic ray spectrum. The peak is located at $E_{\text{reco}} = 2$ TeV, followed by a steep fall to the highest energies. There is also an apparent change in slope at $E_{\text{reco}} = 20$ TeV.

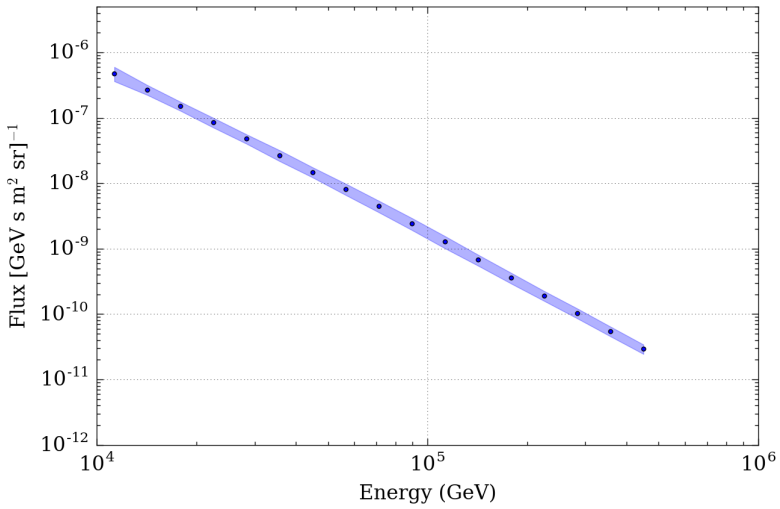

Figure 8.12: Unfolded differential flux of the counts distribution from 8.11 using the detector response function shown in figures 8.6 and 8.7. The spectrum seems to follow a power law form over the given energy range. Only the spectrum up to 500 TeV is shown, as will be addressed in section 8.4.3. The shaded region represents the uncertainty from the Monte Carlo, discussed in section 8.4.2.

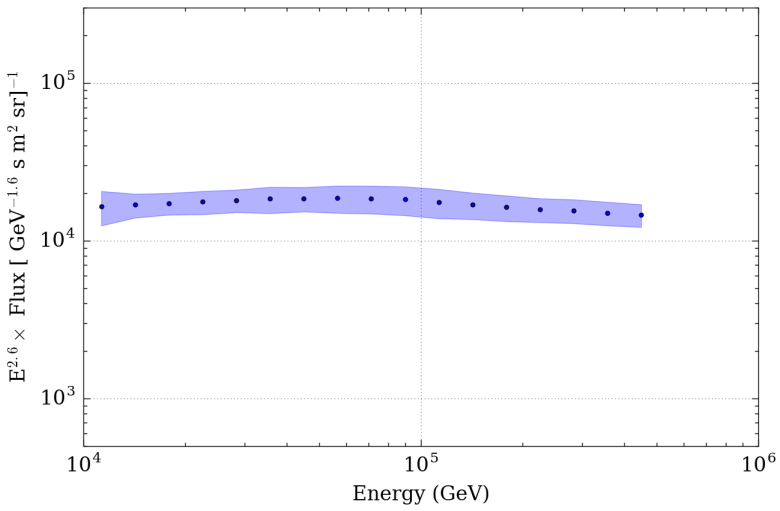

Figure 8.13: Energy scaled differential flux from figure 8.12. A turnover in the spectrum is visible in the 50 TeV range. Only the spectrum up to 500 TeV is shown, as will be addressed in section 8.4.3. The shaded region represents the uncertainty from the Monte Carlo, discussed in section 8.4.2.

8.4 Systematic Uncertainties

A thorough study of the possible systematic effects has been performed for this analysis. The main sources of systematic uncertainty considered in this work are the following:

(i) Effects due to the uncertainty in PMT performance properties. This includes the PMT charge resolution, $Q_{\rm res}$, and the quantum efficiency, QE.
(ii) Effects due to the reliability of the event selection criterion N_{r40}.
(iii) Effects from simulating the detector response, specifically the uncertainty due to the limited Monte Carlo statistics used to build the $P(E_{\rm reco}|E)$.
(iv) Effects related to the assumed composition model used to build the response function.
(v) Effects from the hadronic interaction engines used in the air shower simulations.

Table 8.2 summarizes the various contributions to the overall systematic error in three energy bins. With the exception of the effects from the CREAM model fit uncertainties, each systematic was determined by generating a Monte Carlo data set, building a response function from that set, and unfolding the reconstructed data. The resulting fluxes are then compared to the nominal response for that systematic, providing an estimate of the systematic uncertainty via the fractional difference

$$\sigma_{\rm sys} = -\frac{\mathscr{F} - \mathscr{F}_{\rm nominal}}{\mathscr{F}_{\rm nominal}}\,. \tag{8.2}$$

This ensures that the systematic quoted is in fact a systematic uncertainty on the data themselves. The minus sign is necessary to account for the fact that the data themselves are fixed. If, for example, we were to compare the reconstructed energy distributions from the various Monte Carlo data sets and unfold them with the nominal model for that systematic, the minus sign would be absent. In this case, the shifts in the unfolded fluxes follow the shifts from their respective Monte Carlo $E_{\rm reco}$ distributions. For keeping the data fixed and varying the response functions, the resulting fluxes move in the opposite direction. For the figures in this section, ratios of unfolded fluxes are shown which include statistical uncertainties, though due to the large data sample they are smaller than the marker size.

	10 TeV	100 TeV	1 PeV
PMT QE	±6%	±8%	±9%
PMT Q_{res}	−3%	−5%	−10%
Simulation	±8%	±8%	±8%
Composition Model	−16/+5%	−4/+3%	±3%
Hadronic Int. Model	+5%	+10%	−4/+2%
Total	−20/+12%	−14/+15%	−20/+13%

Table 8.2: Summary of various systematic uncertainties at three energies.

8.4.1 PMT Charge Resolution and Quantum Efficiency

Systematic uncertainties due to the understanding of PMT response arise from two contributions. First is the instrinsic PMT quantum efficiency (QE), which quantifies the conversion of an incident photon to a PE and subsequent collection of that PE in the PMT dynode chain. Typical values of QE range between 20–30%, yet determining the absolute QE is not possible with the calibration system itself, as it requires establishing the efficiency of the calibration system's integrated optical path to the PMTs more precisely than is known. Futhermore, the calibration system's laser emits green light and QE estimation must be extrapolated for application to blue Cherenkov light. The absolute efficiency instead is established by identifying vertical muons passing through WCDs via the PMT signal times. The simulated response to vertical muons is then scaled to match the data. Nevertheless, we conservatively vary the QE in simulation by 10%. The resulting uncertainty in the unfolded flux grows from about 6% to 10 % with increasing energy, as shown in figure 8.14. This is expected as an increase in light collection also increases the detection efficiency for showers of a given energy, and vice versa.

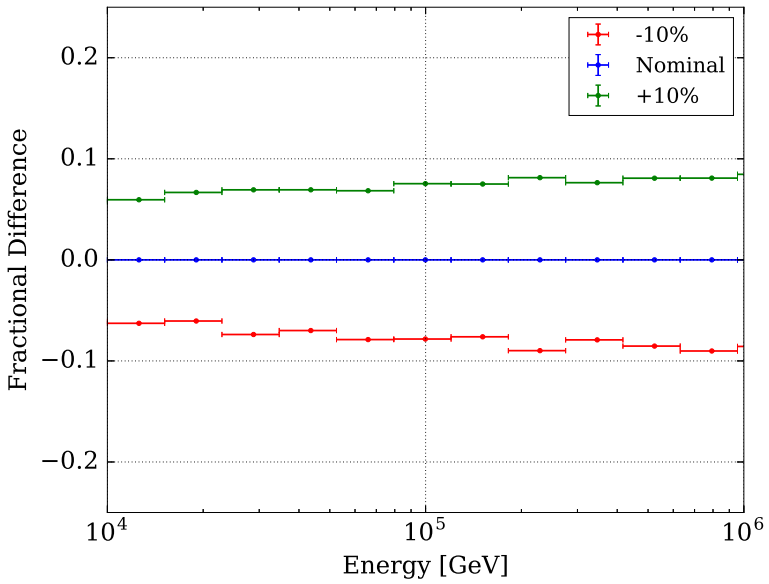

Figure 8.14: Fractional differences of unfolded fluxes from the nominal PMT QE value varied by 10%. The effect of varying the QE slowly changes with energy, as a linear increase or decrease in the PMT efficiency translates directly to a similar change in the detector efficiency or A_{eff}.

PMT charge measurements can also vary for a fixed illumination on the photocathode. We summarize this by the charge resolution, Q_{res}, estimated to be between $10 - 25\%$ from the calibration system. This is evaluated by smearing the measured charges in simulated air showers by 10% and 25% both in charge and in the logarithm of the charge. Despite the wide range of charge resolutions, we find that the resulting unfolded spectrum varies by 5%, except at the highest energy where it reaches 9%. The fractional differences between the smearing functions and the nominal 10% in the logarithm of the charge are shown in figure 8.15. Based on this comparison, the nominal simulation gives a conservative upper limit on the flux.

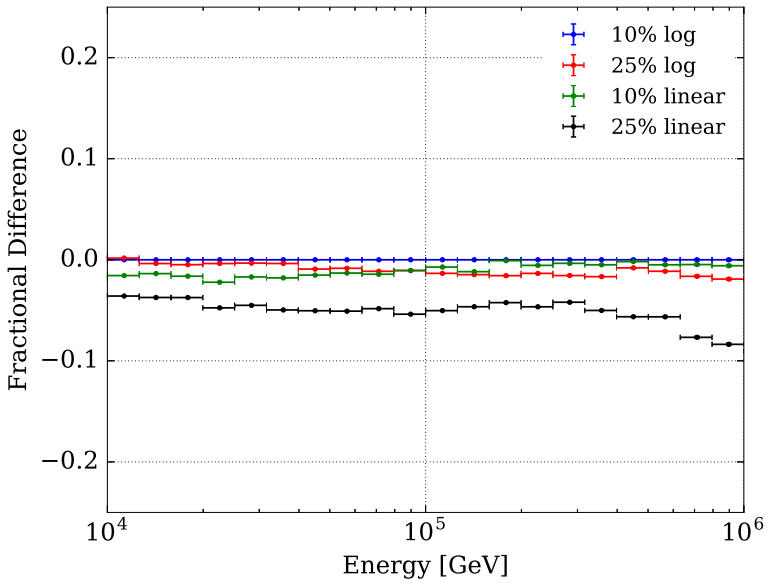

Figure 8.15: Ratios of unfolded fluxes to the nominal 10% $\log Q$ smearing using various charge smearing functions. At the lowest energies, the variation between the smearings is similar because these showers are dominated by $O(1$ PE) hits, as shown in figure 8.5. For such low PE values, smearing by Q or $\log Q$ is similar. However, with increasing charge values, the logarithmic scaling skews smearing towards increased charge relative to the linear scaling.

8.4.2 Detector Simulation

As discussed in section 4.3, the energy resolution above 10 TeV given the base event selection criteria is estimated at $\sim 25\%$. It is not considered to be a source of systematic uncertainty, as it is built into the response matrix for the unfolding procedure. However, the limited Monte Carlo statistics used in determining $P(E_{\text{reco}}|E)$ were taken into account as a source of uncertainty. This is defined as V^{MC} from section 7.2.2.2 and we use the multinomial expression of equation 7.15. We include this as a *pseudo*-systematic because it does not arise from the data themselves; however, it is also clear that an increase in Monte Carlo statistics reduces the effect of V^{MC}. Furthermore, the value itself depends on the selection cuts, as increasingly harsh criteria reduce the size of the Monte Carlo sample for building $P(E_{\text{reco}}|C)$. The estimated Monte Carlo uncertainties for the basic event selection including several N_{r40} cuts are shown in figure 8.16. As will be shown in section 8.5.2, we include the Monte Carlo systematic (in addition to the statistical uncertainty from the data) when fitting the unfolded spectrum to power law forms.

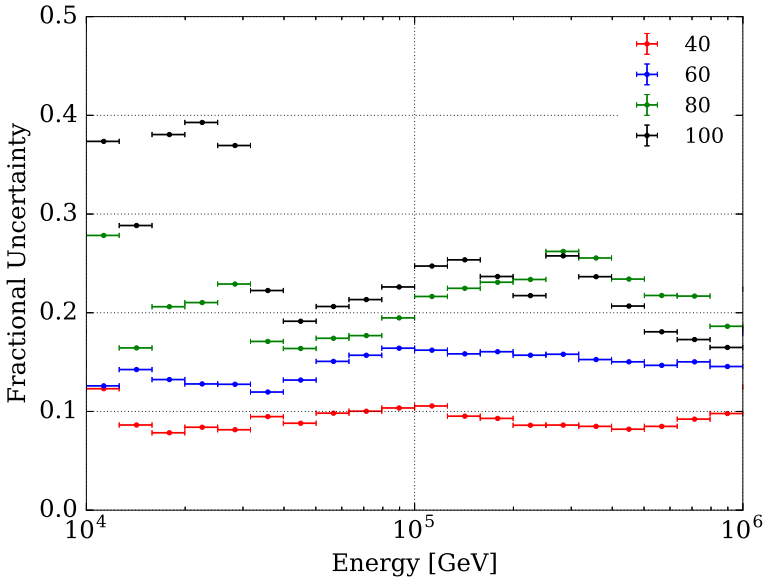

Figure 8.16: Estimated fractional uncertainties from the Monte Carlo data sample for the various N_{r40} cuts. For increasing cut severity, the number of Monte Carlo events used to build the response function drops, so the uncertainty rises. Since the value is calculated via equation 7.15, it is a symmetric uncertainty applied to the unfolded flux.

8.4.3 N_{r40} Selection Criterion

As previously mentioned, the N_{r40} criterion was varied widely to test the robustness of the cut and the stability of the unfolded spectrum. The tested values were from 40–100 in steps of 10, and the ratios of the unfolded fluxes to the baseline cut are shown in figure 8.17. For reference, $N_{r40} = 100$ corresponds to about half of the PMTs within a 40 m radius from the core, as this containment area holds approximately 240 PMTs. With a cut this strong nearly all cores are expected to be reconstructed on the array. This is shown in figure 8.18 where the increased linearity in energy response for $N_{r40} \geq 100$ is clear. Of course, the harsher cuts reduce the total number of events and thus the efficiency of the detector which from figure 8.19 appears to scale linearly with cut severity.

The saturation of PMTs above their calibrated charge thresholds results in hits being omitted in the calculation of N_{r40}. As events with the highest energies tend to saturate the detector, this begins to affect the detector efficiency above 600 TeV. This can be seen as a rapid drop in the fractional difference relative to the $N_{r40} \geq 40$ spectrum in figure 8.17, as well as the drop in efficiency at high energies for high cut thresholds in figure 8.19. The Monte Carlo simulation takes this hit loss into account as seen at the highest energies in figure 8.19, yet, including its effect in the effective area to the precision necessary for spectral measurements is difficult, as the curvature of this drop is abrupt.

Hence, we only consider energies up to 500 TeV, where the effect of this hit dropping remains negligible. For the purposes of this study, we do not consider the variations from the N_{r40} cut as a systematic uncertainty on the reported spectrum. Instead, we use the cuts to demonstrate that the unfolded spectra are stable to within 10% up to 500 TeV, and the resulting fits are consistent to within Monte Carlo uncertainties, as will be discussed in section 8.5.2. For this reason, we use $N_{r40} \geq 40$ for the definitive selection criterion, as its Monte Carlo uncertainties are the least (10% from figure 8.16) among the N_{r40} cuts. Finally, it is clear that with an improved consideration of these hits both in the reconstructed data and the Monte Carlo, the HAWC Observatory has the capability to investigate the knee region.

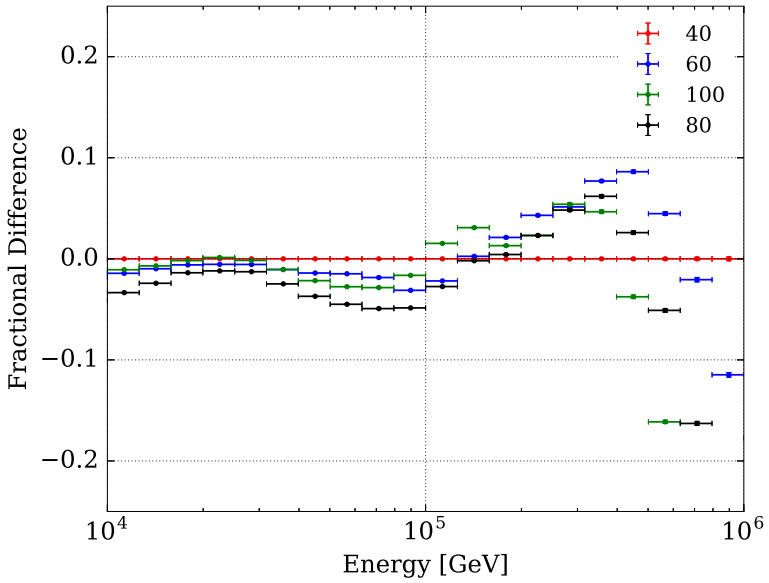

Figure 8.17: Variation of the unfolded energy spectrum from increasingly harsh N_{r40} cuts. Only values in steps of 20 are shown here for ease of viewing, and $N_{r40} \geq 40$ defines the nominal cut value.

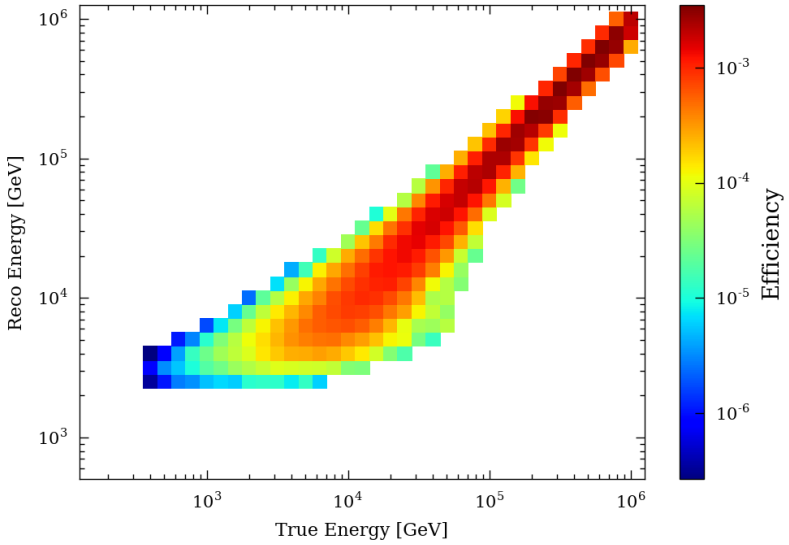

Figure 8.18: The response matrix for $N_{r40} \geq 100$ selection. The improved energy response can be compared to figure 8.9 for the selection of true cores landing on the array.

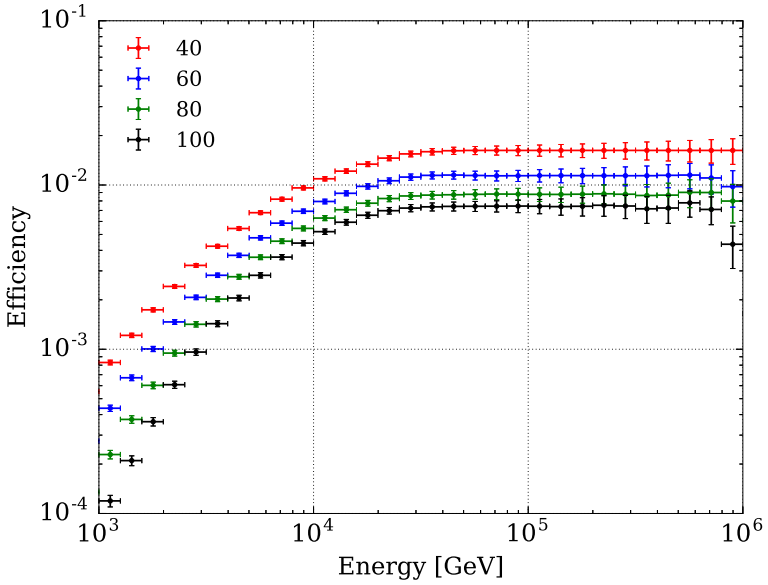

Figure 8.19: Efficiency curves for several N_{r40} cuts. With increasingly harsh cuts, the efficiency scales from loss of events. There is also a drop in efficiency above 600 TeV for the harshest cut as a result of PMT saturation leading to hit loss.

8.4.4 Composition Model

We considered two contributions to the uncertainties from the assumed composition models. The first arises from the fit uncertainties to equation 3.4 in defining the nominal CREAM model. We quantified this by varying the fit parameters to within their estimated uncertainties, obtaining the model's contribution to the uncertainty in the unfolded flux. The distribution of reconstructed energies was weighted for each variation of the elemental spectra, which was then unfolded. The resulting median and 68% containment region defines the resulting flux uncertainty. This amounted to less than 3% uncertainty for all energies, shown as the blue band in figure 8.20. We also considered three other widely used models: H4a [68], Polygonato [67], and the Gaisser-Stanev-Tilav model (GST4-gen) [97]. The H3a model also presented in [68] was considered; however, the unfolded spectrum was within $< 1\%$ of the unfolded spectrum using the H4a model, so we simply quote the H4a result. The compostion models are shown together for comparison in figure 8.21. In all, the systematic uncertainty due to the assumed composition is $\sim 6\%$ between $10 - 100$ TeV, increasing to $\pm 12\%$ above 100 TeV, where the uncertainty in direct detection measurements is maximal.

Figure 8.20 also shows the ratios of the unfolded fluxes using all compositions relative to the baseline CREAM model. We take the full range spanned by the models in each energy band (including the CREAM uncertainty band) as a conservative estimate of the systematic uncertainty, as we have assumed no preference for any one model. Above 100 TeV, the unfolded spectra agree to within 5%, the same order as simply varying the CREAM fit parameters. The greatest deviation from the nominal model comes from H4a, providing an uncertainty of 15% at 10 TeV. This is due to the significantly larger contribution of heavy elements (>He) to the model as compared to the other three as shown in figure 8.21. This has the effect of reducing the effective area since A_{eff} is an abundance-weighted average for all species. The greater presence of heavier elements also induces increased energy migration such that reconstructed events at lower energies are promoted towards higher energies in the unfolded flux.

Figure 8.22 shows the elemental abundances relative to the all-particle spectrum for each model. Indeed, it is the fractional abundances that are most relevant for constructing the response function, as $\sum_{E_{\text{reco}}} P(E_{\text{reco}}|E) = \epsilon(E)$ so the overall flux normalization for each is not important. Yet from figure 8.20 we see that the CREAM, Polygonato, and GST4-gen models are relatively insensitive to variations in the proton to helium ratio, particularly since there is a wide range in their crossing point between models. This suggests that the detector's response to proton and helium is similar enough to probe their combined spectrum.

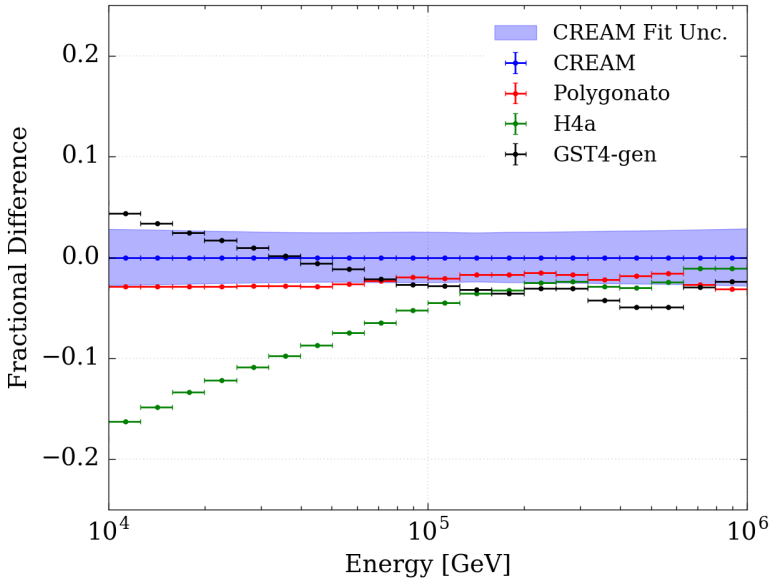

Figure 8.20: Ratios of unfolded fluxes using various compositions to the nominal CREAM model. The blue band indicates the 68% containment range of unfolded fluxes having varied the CREAM model by its estimated uncertainties from table 3.2.

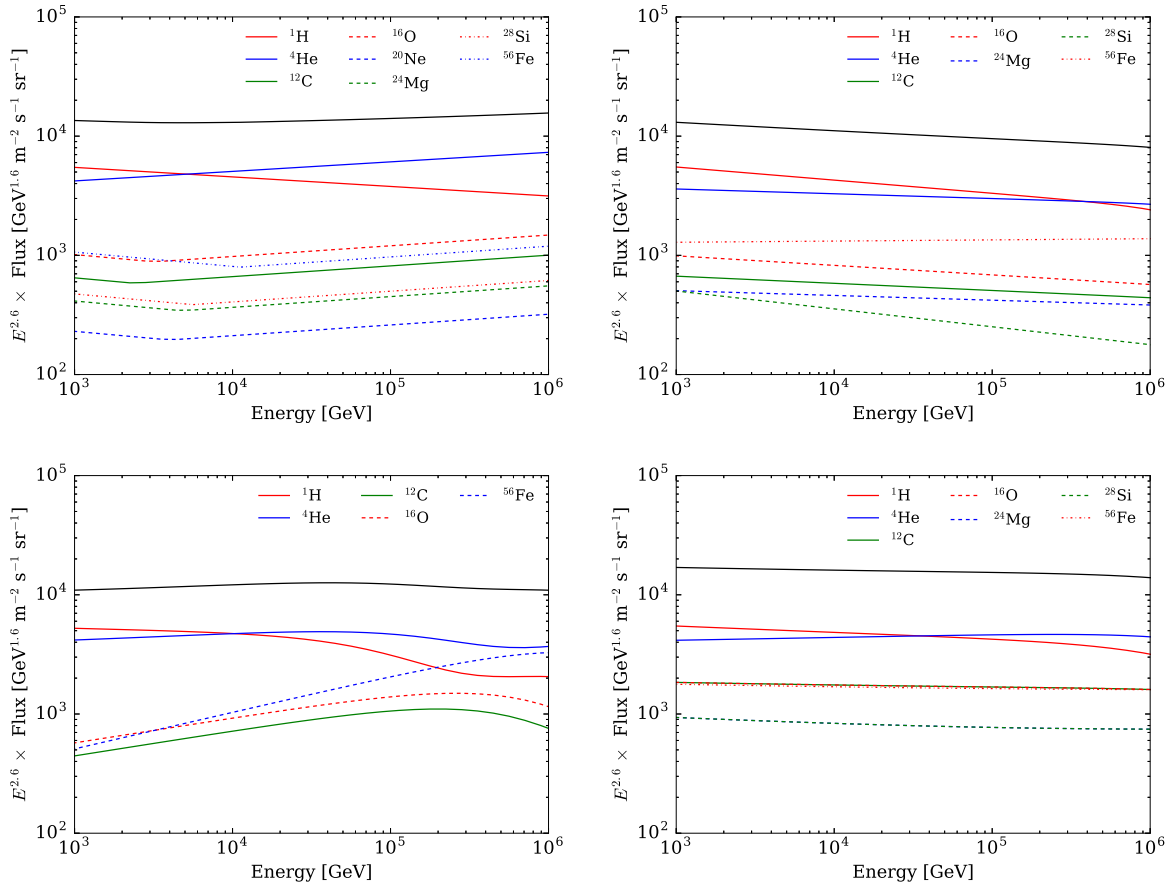

Figure 8.21: Scaled fluxes of the four composition models considered. Clockwise from upper left: CREAM, Polygonato [67], H4a [68], GST4-gen [97]. The black curves in each panel represent each model's all-particle flux.

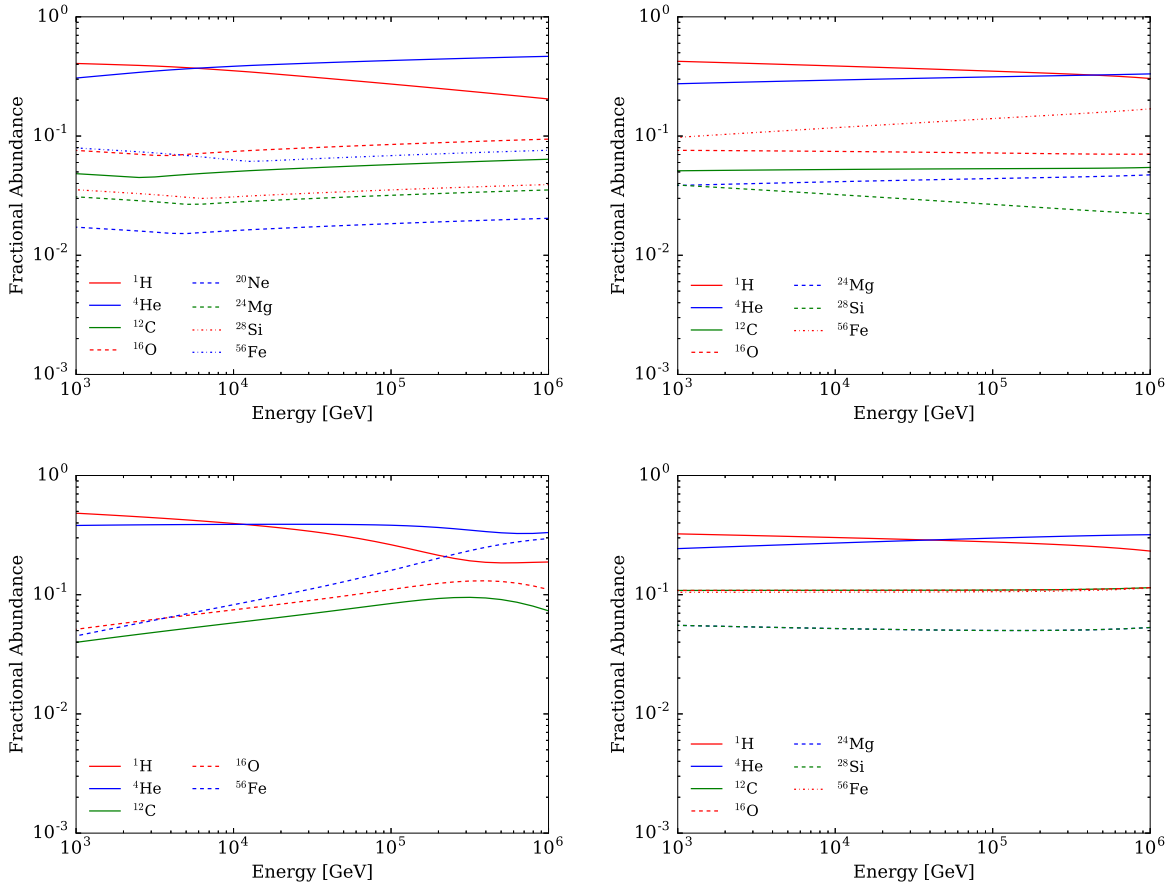

Figure 8.22: Fractional abundances of the species in the various composition models. Clockwise from upper left: CREAM, Polygonato [67], H4a [68], GST4-gen [97].

8.4.5 Hadronic Interaction Model

We also considered the systematic uncertainty from different hadronic interaction models by comparing the nominal simulation using QGSJet-II-03 [61] to the EPOS [62] and SIBYLL [63] high energy models. Each model consists of fits and extrapolations to nucleon cross section data from accelerators. The results of comparing the unfolded spectra are shown in figure 8.23. The EPOS model agrees to within 2% of the nominal simulation, while the SIBYLL model is systematically higher for all energies. Studies from groups such as the GRAPES-3 air shower experiment [30] found that for a fixed composition assumption, the choice of hadronic interaction model influenced the relative abundance of the species arriving at ground level. They found this was primarily due to model differences in determining the point of the first interaction. As the simulated data sets for these models was smaller than the nominal set, a more thorough analysis of the origin of these discrepancies was not possible. Still, we include this as a source of systematic uncertainty using the observed ranges of unfolded spectra.

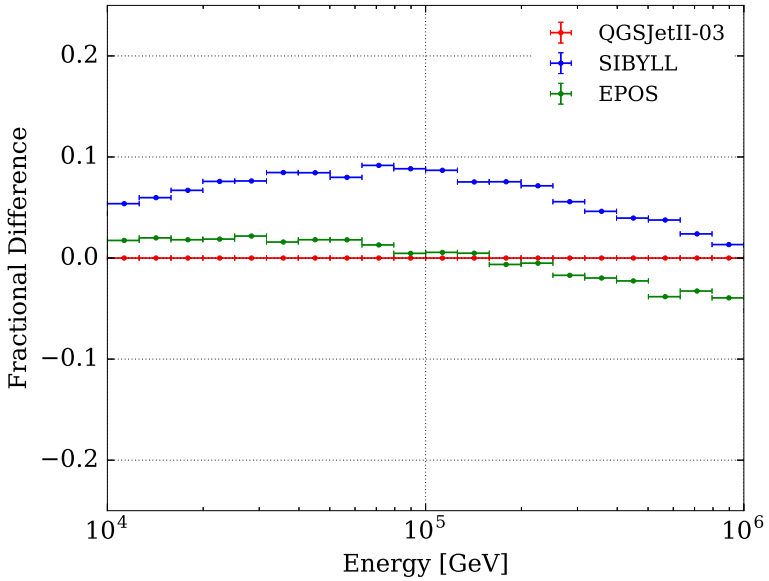

Figure 8.23: Ratios of unfolded fluxes using two hadronic interaction models to the nominal QGSJETII-3 model.

8.5 Discussion

8.5.1 Regularization

The effect of regularization was investigated to ensure that it did not introduce bias in the shape of unfolded spectra. We compared unfoldings with power law regularization, spline smoothing, and no regularization. Figure 8.24 shows the results of three such unfoldings using the same data set and response function. The regularized spectra converged in three iterations, while the unsmoothed unfolding required five. We find that the spectra are spread by no more than 5%, except at the highest energies, where the width of the energy resolution is of the order of the response matrix bin width, i.e. less than 0.1 in $\log E$. We also see that the form across all energies of the unfolded spectrum is unaffected.

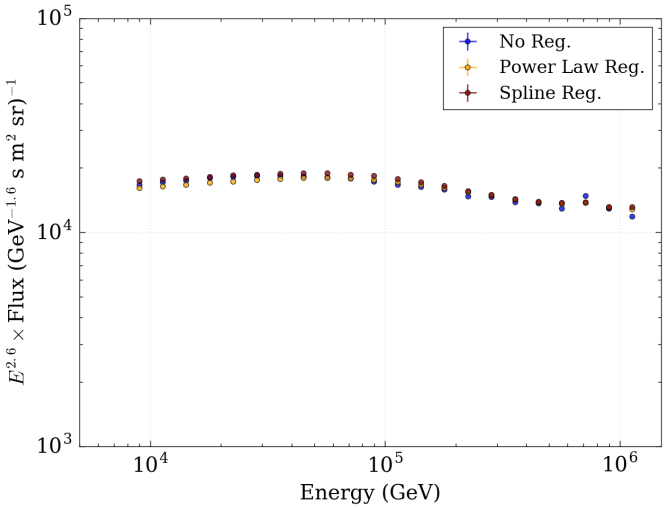

Figure 8.24: The effect of regularization on the unfolded all-particle spectrum. Systematic uncertainties are not shown to highlight only the unfolded data. Regularization was found to not influence the result to within several percent.

If the convergence criterion were to be made extremely harsh, for example requiring the KS statistic to be $< 10^{-10}$, we would see no change in the regularized spectra. However, the un-regularized unfolding attempts to fit precisely to the detector response matrix, which is a discrete function, resulting in sharp peaks. In principle the matrix could be binned more finely. However, separating the Monte Carlo statistics into smaller bins could potentially mask trends in the efficiency curve when smoothing, as well as increase the uncertainty from Monte Carlo statistics for each bin. If instead of being represented by a matrix the response were a smooth, continuous function, this oscillation effect would be absent. Though the three methods agree

to within a few percent, we chose the spline smoothing as a compromise between the potentially shape enforcing power law regularization and the completely unregularized unfolding. The spline function ensures that neighboring points are smoothly connected, so underlying features should manifest in the unfolded spectrum.

8.5.2 Fit to Power-Law Type Functions

The scaled differential flux for all N_{r40} selection cuts reveals a feature which is not well described by a single power law within statistical uncertainties. This structure is shown in figure 8.25 including statistical and Monte Carlo uncertainties. The feature at ~ 40 TeV implies a change in the spectrum, so it was fit to a broken power law of the form of equation 3.4 in addition to a single power law. For the $N_{r40} \geq 40$ set, the difference in goodness of fit is $\Delta\chi^2 = 29.2$, which for a difference in degrees of freedom between the two models ($\Delta\text{dof} = 2$), results in a p-value of 4.6×10^{-7}. Thus the broken power law is the favored model, and the fit suggests a change in spectral index of about -0.2 at a break energy of $E_{\text{br}} = 45.7$ TeV. We performed the fits for the seven cut values of N_{r40}, which are presented in table 8.3.

Since the N_{r40} variable's main effect is to pull reconstructed cores, and by extension true cores, onto the array, the increasing of the cut severity should improve constraints on the spectral shape. If, for example, the under-estimated energies from poor core fits induced a feature in the spectrum, then E_{br} should shift with different cuts. Yet, the fit values from table 8.3 do not deviate more than a few percent, demonstrating that the detector response linearity for even the $N_{r40} \geq 40$ cut is sufficient to unfold the spectrum. Figure 8.7 also suggests this, as the contributions from mis-reconstructed cores are an order of magnitude lower than the peak along the diagonal. Indeed, the normalizations and spectral indices for all selection criteria agree to within 3%, and the location of E_{br} is spread by no more than 15%. This range, $\log 45.7 - \log 40.9 = 0.05$, is of the order of the binning, $\Delta \log E = 0.1$, so the fit E_{br} are within the resolution of the energy spacing. Of course, the estimated uncertainties on the normalization increase with the N_{r40} cut, which simply reflects the fact that the Monte Carlo sample filling the response function is reduced for each cut. Hence, we report the unfolded spectrum using the $N_{r40} \geq 40$ data set.

The efficiencies for all values of N_{r40} are still rising just prior to the break range, as previously shown in figure 8.19. This potentially makes constraining the spectral shape dependent on precisely where the efficiency reaches its plateau value. And since the efficiencies scale nearly linearly with the cut severity, there is little differentiation manifested in the unfolded spectra. Since the efficiencies themselves exhibit curvature in this region, one might expect a similar bending of the unfolded spectrum if the response was not modelled sufficiently well. In this analysis, there is no indication of curvature below E_{br} in the unfolded spectrum. Furthermore, by simply forcing the efficiency to be flat between 10 TeV $\leq E_{\text{reco}} \leq$ 30 TeV, the unfolded flux would maintain the spectral change observed in the measured E_{reco} distribution from figure 8.11. As a final check, we also simulated single power law spectra to determine whether a spectral break is induced by application of the detector response function, or is retained if a broken power law is injected instead. Figure 8.26 shows two such E_{reco} distributions from simulated spectra compared to data. The single power law with $\gamma = -2.7$ does not induce a break in the reconstructed energy distribution; however, the broken

power law with parameters similar to the fit values from data exhibits structure akin to that of the observed distribution.

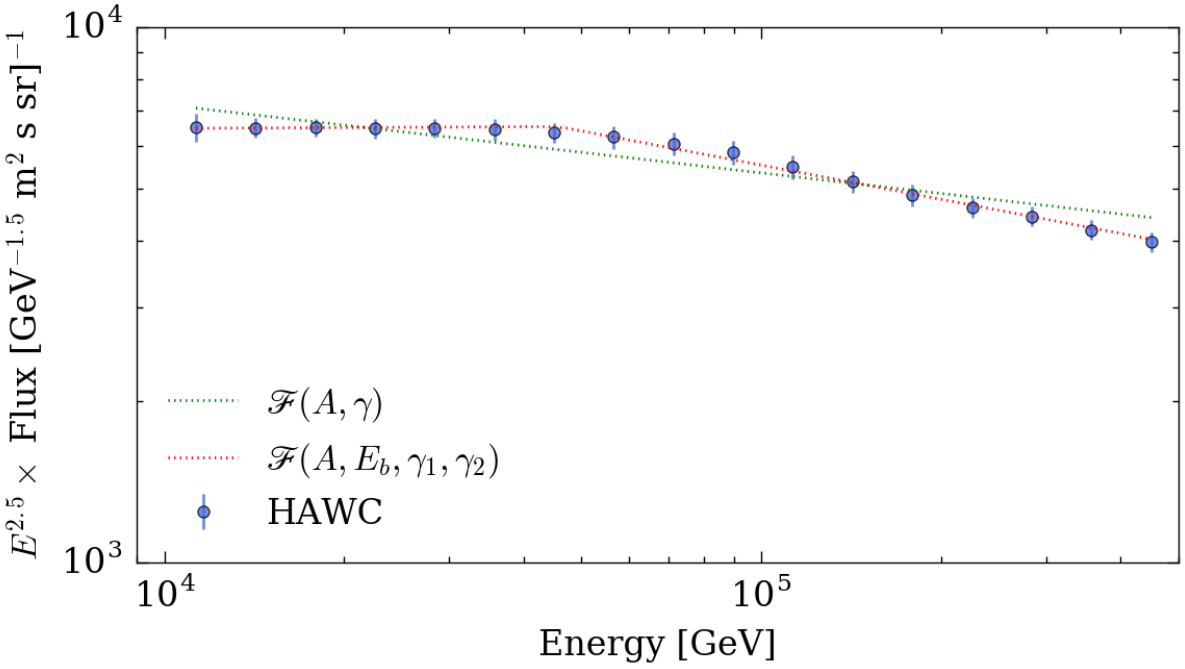

Figure 8.25: Fits to the unfolded flux with $N_{r40} \geq 40$ using single and broken power law forms. The uncertainties on the data are from the Monte Carlo data set. The broken power law is favored, as $\Delta\chi^2 = 29.2$, giving a p-value of 4.6×10^{-7}.

N_{r40} Cut	A[10 TeV] $\times 10^{-7}$ [GeV s m^2 sr]$^{-1}$	γ_1	γ_2	E_{B} [TeV]
40	6.46 ± 0.77	-2.49	-2.71	45.7
50	6.37 ± 1.14	-2.49	-2.69	42.3
60	6.43 ± 1.06	-2.51	-2.67	41.2
70	6.38 ± 1.29	-2.51	-2.66	40.9
80	6.36 ± 1.79	-2.51	-2.67	41.0
90	6.38 ± 2.29	-2.50	-2.68	41.1
100	6.34 ± 2.71	-2.49	-2.70	41.7

Table 8.3: Results to fitting unfolded spectra to a broken power law of form given in equation 3.4. Fits were performed considering only the statistical uncertainties in the flux measurements. The fit uncertainties for both spectral indices were all found to be ± 0.01, and ± 0.1 TeV for the break energies.

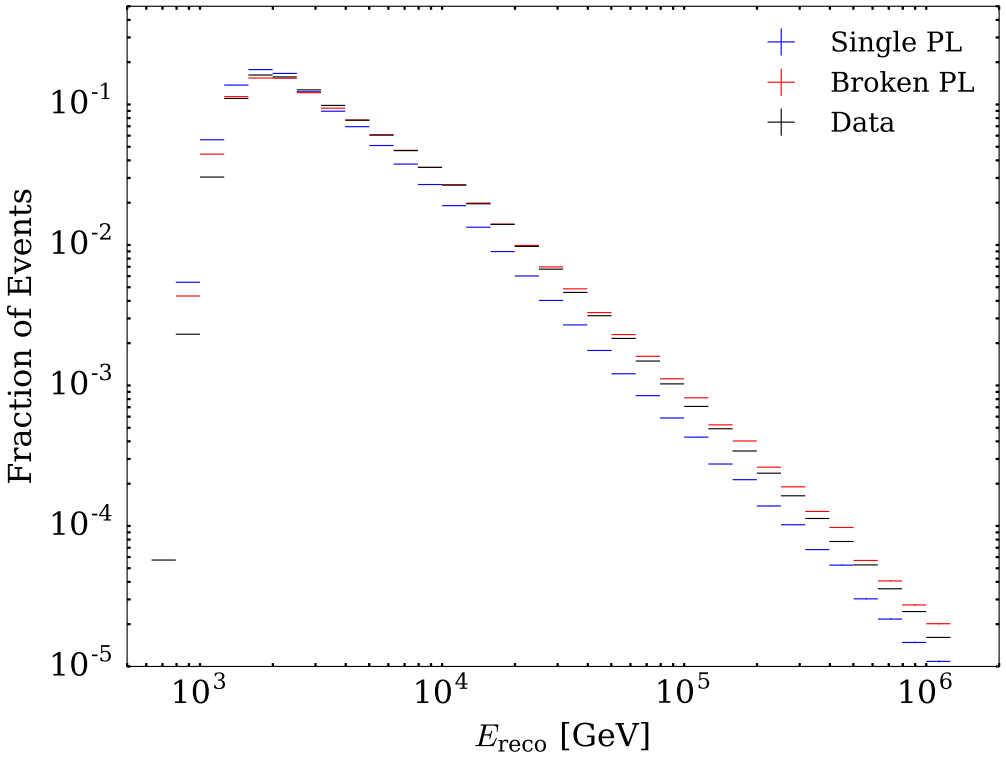

Figure 8.26: Simulated E_{reco} distributions (normalized) from injected power law spectra compared to data. The single power law with an index of $\gamma = -2.7$ does not induce any break-like structure in the reconstructed distribution. The broken power law with $\gamma_1 = -2.5$, $\gamma_2 = -2.7$, and $E_{\mathrm{br}} = 40$ TeV does maintain a break in E_{reco}, similar to that of the observed data distribution. Only the statistical uncertainties are shown for ease of viewing.

8.5.3 Comparison to Other Experiments

A comparison of the unfolded all-particle spectrum to other experimental results is shown in figure 8.27. Including an exhaustive list of other measurements makes comparison on a single frame difficult, so a selection from the most recent results (< 20 yrs) is presented here. Above 100 TeV, there remains agreement with the final ATIC [24] data point though it has statistical uncertainty greater than the HAWC systematics. The HAWC measurement is systematically higher than measurements from other ground based experiments such as GRAPES-3 [30] and the Tibet-III [98] arrays, but consistent within the systematic uncertainties. Indeed, these discrepancies can be understood as arising from differences in the experiments' energy scale calibrations. For example, a 10% systematic shift in energy results in a more dramatic $\sim$ 30% shift in the energy scaled flux. The ARGO-YBJ spectrum [99] is also lower and appears harder with an index of -2.62 ± 0.03 until $\sim$ 700 TeV where it softens, so the energy scaling effect does not address this discrepancy in spectral shape. For comparison, the spectral index as measured by Tibet-III between 150–1000 TeV is -2.68 ± 0.02, where the uncertainties quoted are statistical and whose value is consistent with those of γ_2 presented in table 8.3.

We see that in the 10 TeV range, the HAWC spectrum is consistent with the ATIC all-particle measurement [24], including its rise to 30 TeV. The slightly steeper spectrum below 50 TeV is also mirrored, though not as strongly, in the GST4-gen model, which also depicts a downturn in the 50–60 TeV region. The ARGO-YBJ light-component (proton and helium) measurement [28] does not indicate a spectral hardening at these energies, having a constant slope of -2.61 ± 0.04 from 5–280 TeV. There is evidence reported by CREAM [23] of a hardening of the helium spectrum between 10–30 TeV (shown in figure 1.9), with both proton and helium subsequently becoming softer, though CREAM only reports a single power law index for each species: -2.66 ± 0.02 and -2.58 ± 0.02, respectively. For the nominal composition model used in this work, the broken power law fits from table 3.2 do indicate spectral hardening; however, these fits were merely used to extrapolate composition data into the > 10 TeV regime. The ATIC combined proton and helium measurement also shows evidence of a spectral hardening, though it peaks near 12 TeV [24]. Figure 8.28 shows a closer view of this region, comparing the ATIC, CREAM, and HAWC spectra. The stronger energy scaling ($E^{2.75}$) of the ordinate reveals consistent spectral forms, with a potential kink between 20–40 TeV for ATIC and a clear break at 30 TeV for CREAM. Since the light component comprises $\sim$ 90% of the all-particle flux, these direct detection experiments suggest that proton and helium are responsible for the structure observed in the HAWC measurement.

The HAWC all-particle cosmic ray spectrum exhibits agreement within estimated systematic uncertainties with various experiments from 10–500 TeV. The measurement demonstrates that HAWC can extend the reach of ground-based air-shower arrays into the energy range covered by direct detection experiments.

Furthermore, it is with a single experimental technique that the HAWC spectrum bridges these regions. It is also evident that HAWC has the potential to extend the spectrum up to PeV energies to probe the knee. However, as the current event quality selection limits the range to around 500 TeV, an improved understanding of the detector response to the highest energy events is needed.

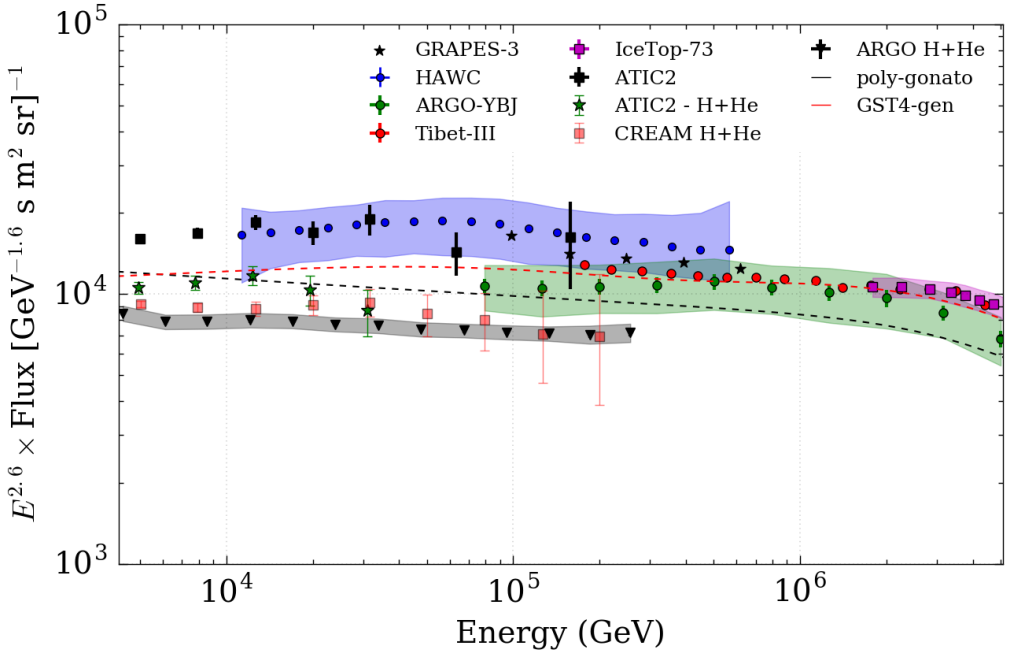

Figure 8.27: Comparison of the measured all-particle spectrum from HAWC to other experiments. The energy flux is scaled by $E^{2.6}$ for ease of viewing over the energy range spanned by all measurements. The HAWC spectrum corresponding to the $N_{r40} \geq 40$ cut has been used. Direct detection measurements are from ATIC [24] and CREAM [23]. Air shower array data are from ARGO-YBJ [99, 28], GRAPES-3 [30], IceTop [100], and Tibet [98]. The two composition models are Polygonato [67], and GST4-gen [97].

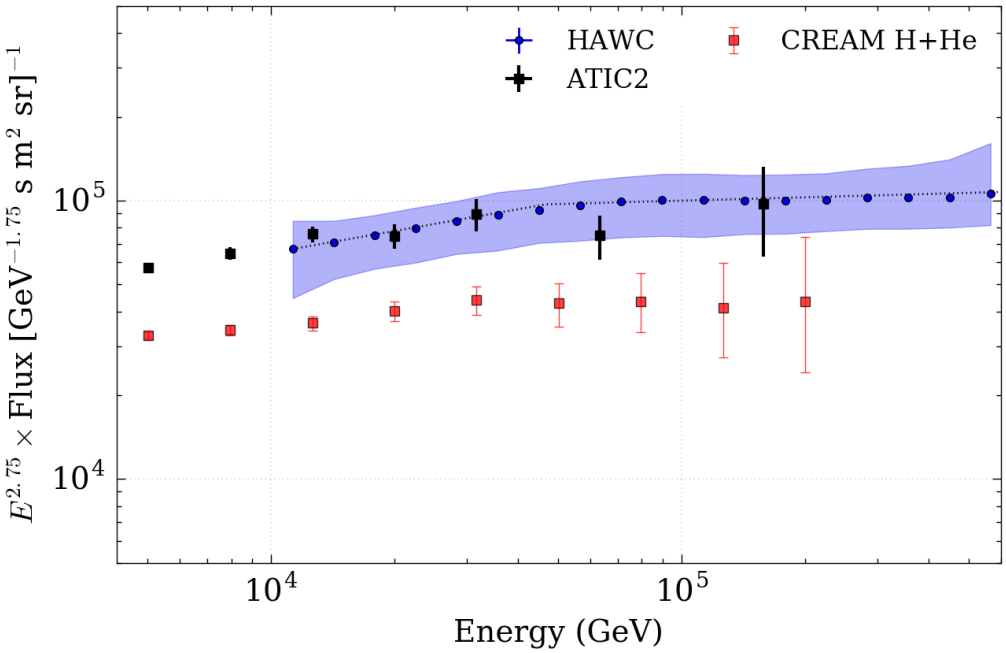

Figure 8.28: Comparison of the HAWC spectrum to the all-particle measurement by ATIC [24] and the light component (proton and helium) by CREAM [23]. The energy flux is scaled by $E^{2.75}$ for ease of viewing, and the dashed line is the broken power law fit to highlight the location of $E_{\rm br}$.

Chapter 9

Cosmic Ray Anisotropy

The sky maps presented in this chapter probe the energy dependence of a previously observed [1, 41, 42] localized excess of cosmic rays, which we refer to as Region A. We again employ the energy estimation method used for the Moon shadow and cosmic ray energy spectrum, permitting an improved measurement of the evolution of Region A with energy. Using an extended data set from that used in the Moon shadow analysis, we first demonstrate an updated view of the anisotropy across the full field of view. Next we show a spectral measurement of Region A relative to the all-particle spectrum across the sky.

9.1 Data

9.1.1 Event Selection

The runs selected for this analysis spanned 622 days from 16 April, 2015 to 17 February, 2017, ensuring that the full 290 WCD detector configuration was in stable data taking mode. The total up-time efficiency during this period was $\sim 92\%$ and the mean trigger rate was ~ 25 kHz. For the background estimation used in making maps, direct integration with periods of 4 and 24 hrs were used. The event sample for the 4 hr integrated maps comprises 6.9×10^{10} events with a total live time of 546 days. The 24 hr maps had a reduced coverage of 200 days, and a total of 2.4×10^{10} events.

The same multiplicity criteria from the Moon shadow analysis were used, namely $N_{\text{hit}} \geq 75$ and $N_{r40} \geq 40$. We also included showers with $\theta \leq 60^\circ$ giving a resulting passing rate of $\sim 2\%$. Increasing the zenith range was important to extend observations beyond Region A, whose lower edge around $\delta = -20^\circ$ lies within a few degrees of the $\theta \leq 45^\circ$ cut used for the Moon shadow.

As cosmic rays are more abundant than diffuse γ rays by a factor of $\sim 10^4$, no gamma/hadron selection cuts were used. As a result, the strongest γ ray source visible to HAWC, the Crab Nebula SNR ($\alpha = 83.5^\circ, \delta = 22^\circ$), is visible in several maps as an extended disc the same dimensions as the smoothing radius, though due to its distance being greater than the applied smoothing radius, it does not influence the spectral measurement of Region A.

9.1.2 Energy Binning

The same E_{reco} analysis bins from the Moon shadow analysis were used for studying the cosmic ray anisotropy. Again, each bin was of size 0.2 in $\log E_{\text{reco}}$ from 1–160 TeV. The estimated mean energy of each bin was determined by sampling the same table shown in figure 6.9, and the values corresponding to the 10° area centered about Region A are shown in figure 9.1 and table 9.1. The all-sky maps presented here correspond to the same *reconstructed* analysis bins, which for specific locations on the sky have different local zenith distributions. Thus, there is a slight shift in energy of a few TeV with increasing distance in declination, shown in figure 9.2. This is attributed to an increasing energy threshold with increasing zenith angle, as cosmic ray primaries and their showers must pass through greater atmospheric overburden. For a specific declination such as that for Region A, we estimate the mean energy for each analysis bin, yet when viewing the all-sky maps, one must keep this declination-dependent energy shift in mind.

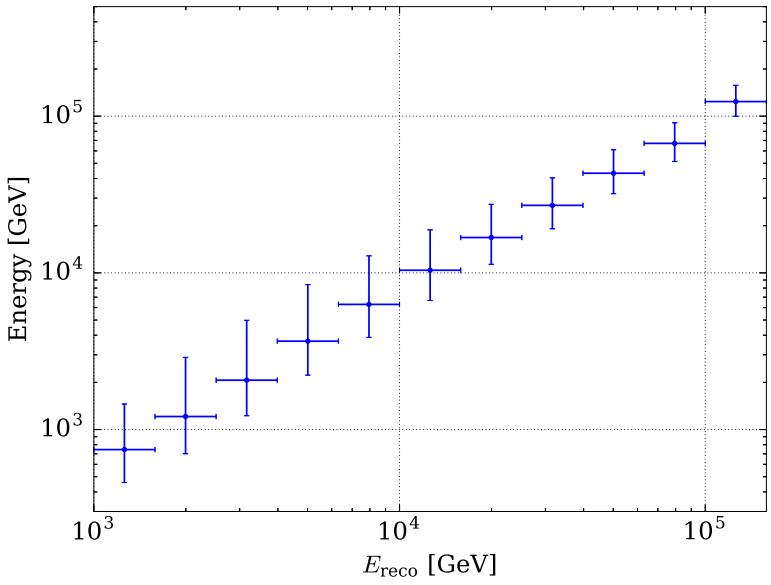

Figure 9.1: Mean energies of analysis bins defined in table 9.1. The bounds on E_{reco} are simply the bin width definitions, while the uncertainties of the ordinate represent the 68% containment region.

Analysis Bin	$\log E_{\mathrm{reco}}/\mathrm{TeV}$	$\bar{E}$ [TeV]	σ_E^+	σ_E^-
0	$0.0 - 0.2$	0.741	0.711	0.286
1	$0.2 - 0.4$	1.21	1.68	0.51
2	$0.4 - 0.6$	2.08	2.91	0.84
3	$0.6 - 0.8$	3.71	4.75	1.44
4	$0.8 - 1.0$	6.32	6.55	2.42
5	$1.0 - 1.2$	10.3	8.42	3.7
6	$1.2 - 1.4$	16.4	10.1	5.5
7	$1.4 - 1.6$	26.3	13.4	7.8
8	$1.6 - 1.8$	43.4	17.8	11.2
9	$1.8 - 2.0$	70.1	23.7	15.6
10	$2.0 - 2.2$	125.5	33.5	24.0

Table 9.1: Estimated true mean energies for the analysis bins defined by E_{reco}. The bounds indicated by $\sigma_E^{\pm}$ represent the 68% central containment region and are quoted in the same units as E.

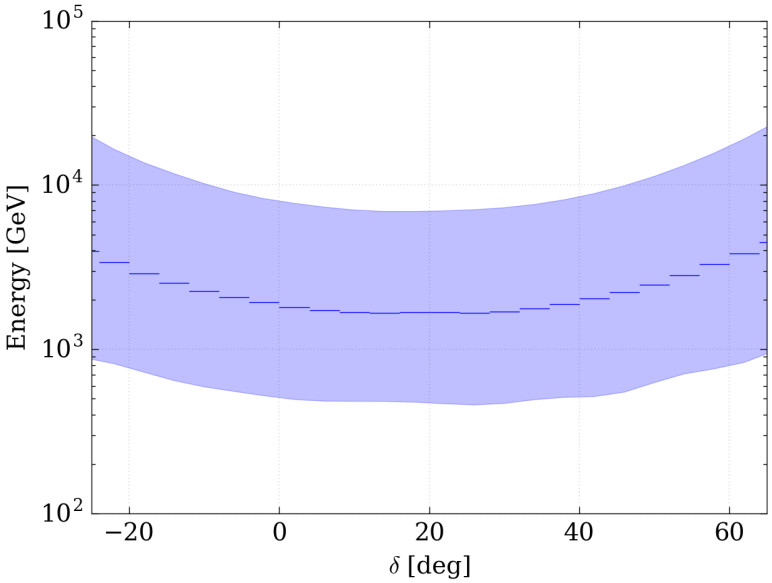

Figure 9.2: Median energy (from simulation) as a function of declination. Farthest away from HAWC ($\delta = 20°$), the median shifts by about a few TeV. The band indicates the 68% central containment region.

9.2 All-Sky Maps

9.2.1 Significance Threshold

The significances shown in the maps are considered pre-trial, as they have been calculated without considering the "look elsewhere" effect. To correctly account for this, a large ensemble of isotropic sky maps with the same exposure as the data would need to be constructed. However, this is computationally prohibitive given the high pre-trial significances observed. Instead we make a conservative correction by estimating that the number of independent pixels in the sky map is of order N_{pix}. The true factor is in fact much smaller because the act of smoothing reduces the number of truly independent pixels by about a factor of the smoothing area. Furthermore, since other experiments including HAWC already have observed Region A this is not a blind search of the data. This correction factor sets a pre-trial significance of 7.12σ as the value necessary to reach a detection threshold of 5σ post-trail significance. Yet despite the large trials factor, Region A maintains a strong post-trials significance value of $\sim 21\sigma$ in both the 4 hr and 24 hr bin-combined maps.

9.2.2 All Sky Anisotropy

An event map was made for each energy bin, from which both relative intensity and significance maps were generated. A value of $N_{\text{side}} = 256$ was used, giving each pixel an angular size of 0.13°, and amounting to a total pixel count in the field of view of $N_{\text{pix}} = 639692$. Unlike the Moon shadow which behaves as a point sink of cosmic rays, the unsmoothed sky maps do not reveal the large regions of cosmic ray anisotropy. Since the scale of these underlying structures is of the order $> 10^\circ$, maps were smoothed by this amount. The smoothed all-sky maps combining all energy bins using 24 hr integration are shown in figure 9.3.

The most prominent feature in both maps is the localized excess centered around $\alpha = 60^\circ$, $\delta = -10^\circ$, which in accordance with figure 1.14 from [1], we call Region A. There is also the significant presence of a large scale feature, the dipole anisotropy, which is most prominently seen as the large area of deficit near the center of the maps. Its relative intensity is of order 10^{-3}, similar to that of Region A. The manner in which to account for the presence of this large scale anisotropy is to fit the monopole, dipole, quadrupole, and octopole components, and subtract them to obtain a residual relative intensity map. The result from removing the large-scale structure from the combined analysis bins is shown in figure 9.4, where during the map generation, a region of interest mask of 20° radius centered on Region A was used. We see that the overall effect on Region A is a reduction in relative intensity and significance, yet it remains above the 7.12σ threshold.

Masking Region A and removing the large-scale components is the appropriate manner to extract the small-scale structure within which Region A resides. Unfortunately, the four lowest energy bins (741 GeV–3.7 TeV) had limited ranges in declination caused by a zenith angle threshold of lower energy events. This resulted in poorly fit multipole harmonics, so the large scale structure subtraction was not feasable. Instead, we used a 4 hr integration time, which effectively removes features of angular size $> 60°$. The $\sim 10°$ scale of Region A prevented the use of the ROI mask, as it grossly exaggerated the estimated significance because the mask itself (40° diameter circle) was of similar order to the integration duration ($4\text{hr} \times 15°\text{hr}^{-1} = 60°$). Yet, from the analysis of the Moon shadow we know that the action of not masking places a lower bound on the strength of Region A. Furthermore, using this shorter integration duration permitted the use of more than twice the number of events, as runs with > 4 hr stability could be included. The all-sky maps using 4 hr integration and 10° smoothing are presented in figure 9.5. Comparison to the 24 hr multipole subtracted maps demonstrates that, indeed, the small-scale structure remains in the 4 hr integrated maps.

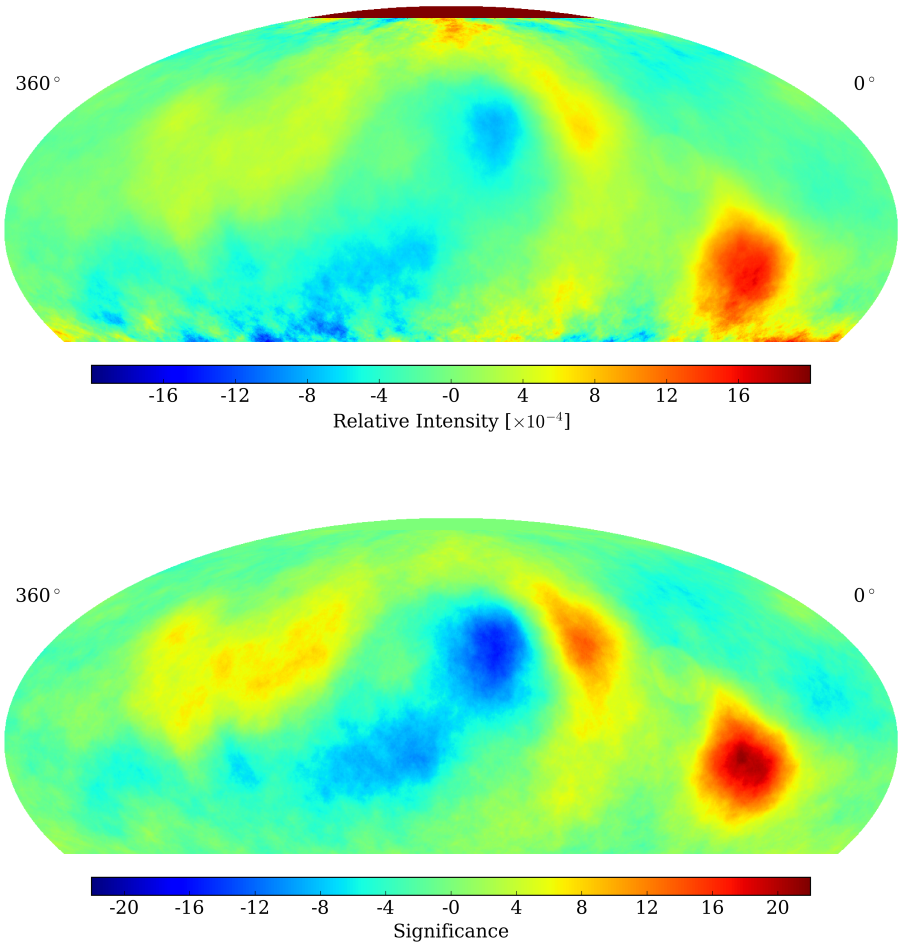

Figure 9.3: All-sky relative intensity (top) and pre-trial significance (bottom) maps for all analysis bins combined. Twenty-four hour integration was used for the reference map estimation, and 10° smoothing was subsequently applied. The strongest excess visible is Region A, which is also the most significant.

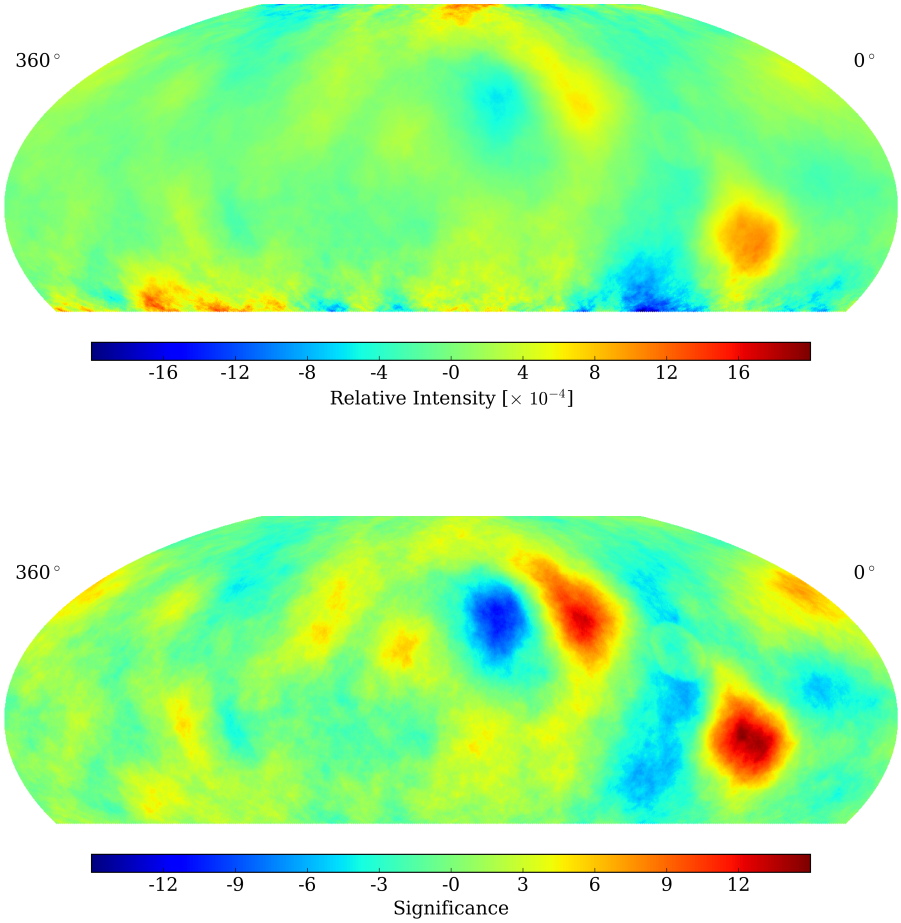

Figure 9.4: All-sky relative intensity (top) and pre-trial significance (bottom) for all analysis bins after using an ROI mask of 20° centered on Region A, and removal of the first four multipole components of the large scale structure. There is a reduction in the relative intensity and significance of Region A, yet it maintains its morphology, and the large scale effects are no longer present.

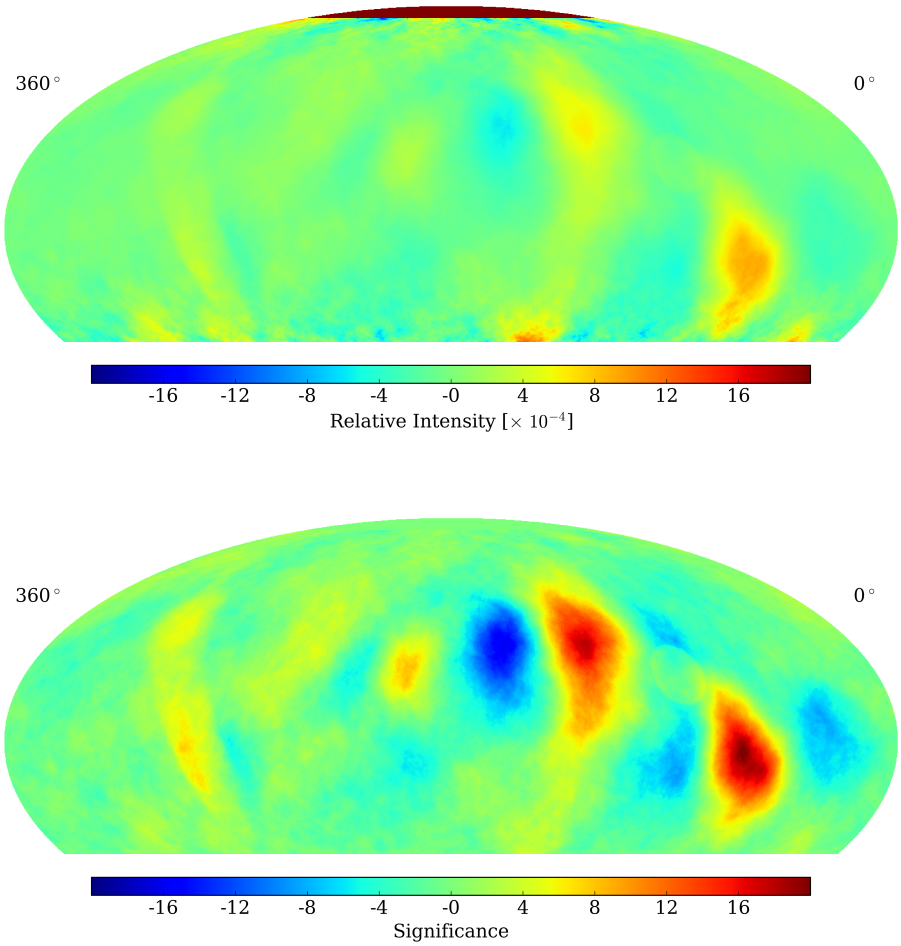

Figure 9.5: All-sky relative intensity (top) and pre-trial significance (bottom) for all analysis bins using 4 hr integration and 10° smoothing. The small-scale structure is qualitatively similar to that for the 24 hr maps with ROI masking and multipole subtraction (figure 9.4). Region A itself becomes slightly elongated in declination in the 4 hr map, yet retains its width in right ascension. We also see the presence of deficits on either side of Region A, a result of its inclusion during the background estimation, as ROI masking was not possible.

9.3 Energy Spectrum of Region A

Previous work using the 111 WCD HAWC configuration [1] indicated that the spectrum of Region A is harder than the all-sky spectrum. This was first suggested by the Milagro Observatory [41], as well as subsequent observations by ARGO-YBJ [42]. These studies also indicate a potential cutoff of the spectrum around 10 TeV; however, the statistics at the highest estimated energies are insufficient to definitively establish the presence of such a feature.

A detailed unfolding of the energy spectrum of Region A is complicated by the need to precisely simulate the detector response over the long observation periods required for significant measurements of the anisotropy. However, the nature of the reference estimation technique and the calculation of the relative intensity inherently normalize these effects. Thus we use the relative intensity as a function of energy as a robust measure of the spectrum of Region A.

Using the map of combined analysis bins as shown in figure 9.6, we chose the point of maximal significance within 10° of the previously reported center of Region A to evaluate the spectrum. This corresponds to the point ($\alpha = 62.2°, \delta = -3.3°$). For reference, the previous maximally significant location of Region A is ($\alpha = 59.2°, \delta = -7.2°$). The lowest energy bin (741.1 GeV) was omitted from the analysis due to its low event count ($< 15\%$ of the next analysis bin) and because its lower declination range was within the 10° smoothing radius from Region A's lowest edge. We also combined the last four energy bins into a single map, as beyond 26.3 TeV Region A was found to be not significant.

The measured relative intensity at the maximally significant point is shown in figure 9.7. The spectrum from HAWC-111 is shown for comparison, where the improved energy resolution is evident. Despite similar statistics (8.6×10^{10} events), the previous measurement did not utilize a core constraining selection criterion, and implemented an energy estimator based on the event multiplicity. This resulted in overlap across multiple energy bins and larger uncertainties in the relative intensity at the highest energies. The likelihood energy estimator bins events in a manner more closely related to the cosmic ray spectrum, constraining the shape both in energy and relative intensity. Indeed, the spectrum from this analysis is less steep considering the highest point from the previous measurement, yet the two are consistent both in shape and normalization.

We fit the relative intensity spectrum using a function of the form $a + b \times \log E$, where the slope b is defined as the spectral shape parameter. The best fit value of b to the spectrum in figure 9.7 was found to be $(2.3 \pm 0.5) \times 10^{-4}$ with a reduced χ^2 of 4.88/5. The fit was also performed for the remaining pixels in the maps to obtain the distribution of slopes across the sky shown in figure 9.8. The distribution of b is approximately Gaussian, with a fit mean of -0.05×10^{-4} and a width of 0.74×10^{-4}. From this we estimate the significance of the slope of Region A as 3.3σ from the all-sky mean. We also fit the spectrum with $b = 0$, effectively a constant normalization condition, resulting in a reduced χ^2 of 27.42/6. The difference $\Delta\chi^2 = 22.54$ for

Δdof $= 1$ degree of freedom, gives a p-value of 2.06×10^{-6}, further supporting the interpretation of a harder spectrum.

While the hardness measured here is not as significant compared to all-sky spectra as reported in [1] (5.2σ), the reduced uncertainties in the relative intensity do represent a stronger constraint on its shape. Indeed, there is an indication in figure 9.7 that there may indeed be a cutoff around 10 TeV; however, a larger data sample is required to futher explore this energy regime. The maps showing the evolution of Region A with energy are presented in figure 9.9.

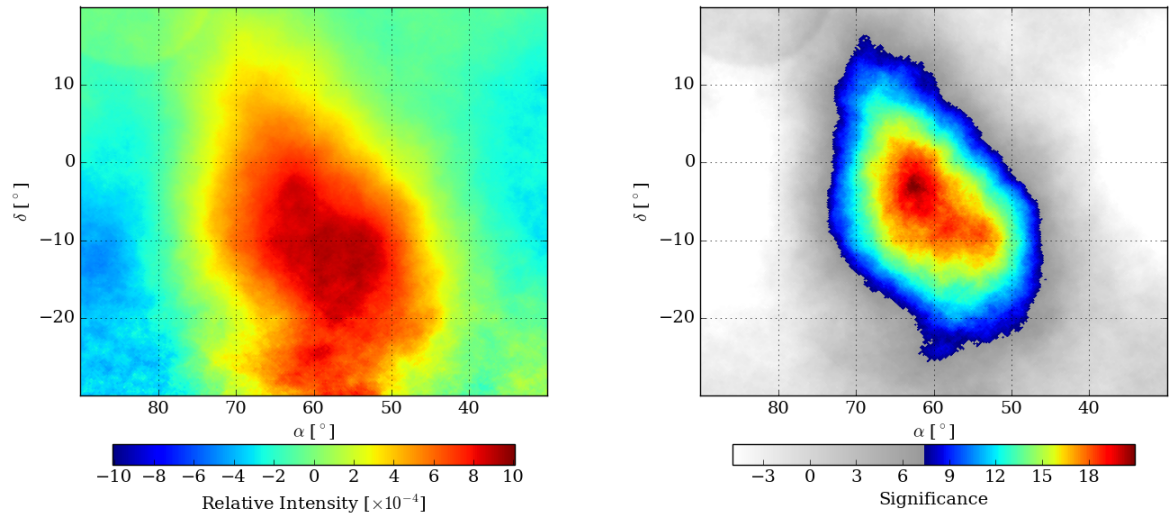

Figure 9.6: The relative intensity (left) and significance (right) of Region A for all analysis bins. The peak significance is 21σ and is located at $(\alpha = 62.2°, \delta = -3.3°)$, which we use for measuring the energy spectrum. For reference, the center of Region A from [1] is $(\alpha = 59.2°, \delta = -7.2°)$. Visible in the upper left corner of the significance figure is part of a circle centered on the Crab Nebula.

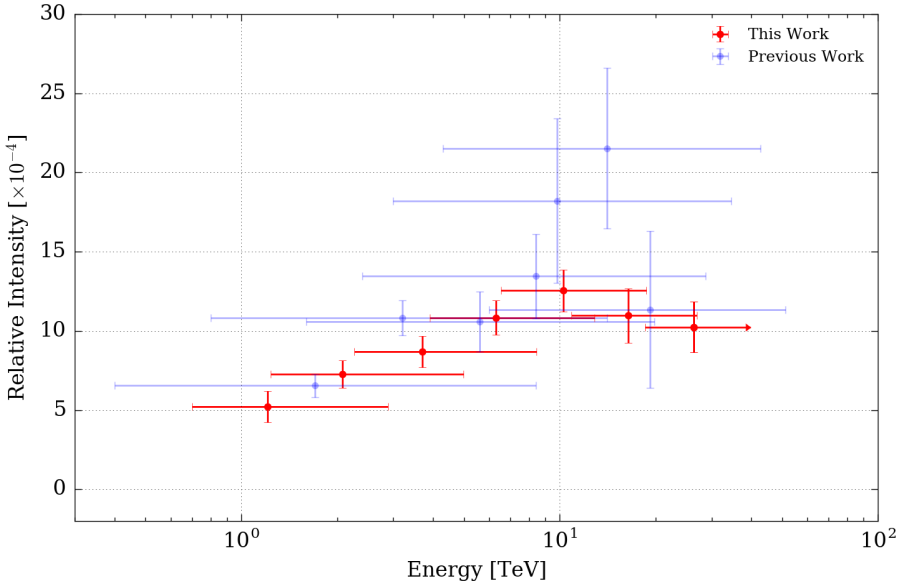

Figure 9.7: Spectrum of Region A in relative intensity as a function of energy. The previous measurement reported in [1] is shown in blue for comparison.

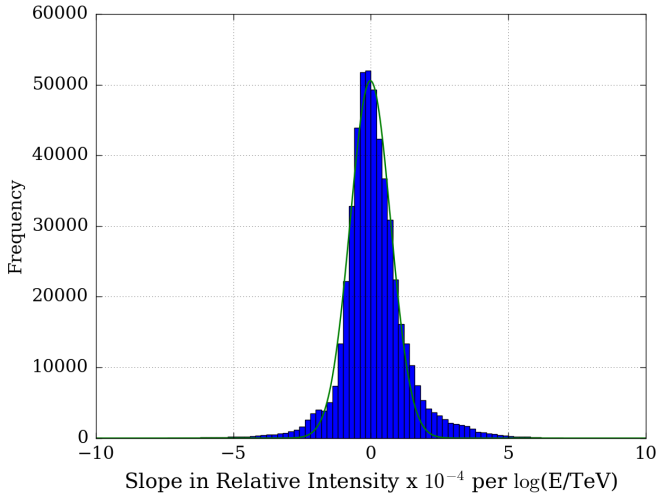

Figure 9.8: Distribution of slopes, b, from fitting $a + b \log E$ to points in the maps. Pixels within 20° of Regions A, B, and C, as well as the Crab Nebula were not included. The green curve represents the Gaussian fit to the distribution, which is used to estimate the significance of the value of b from Region A.

Figure 9.9: Evolution of Region A with energy. The relative intensity (left column) and significance (right column) for the analysis bins in which Region A contained significant pixels. The figures continue for increasing energy over the next two pages.

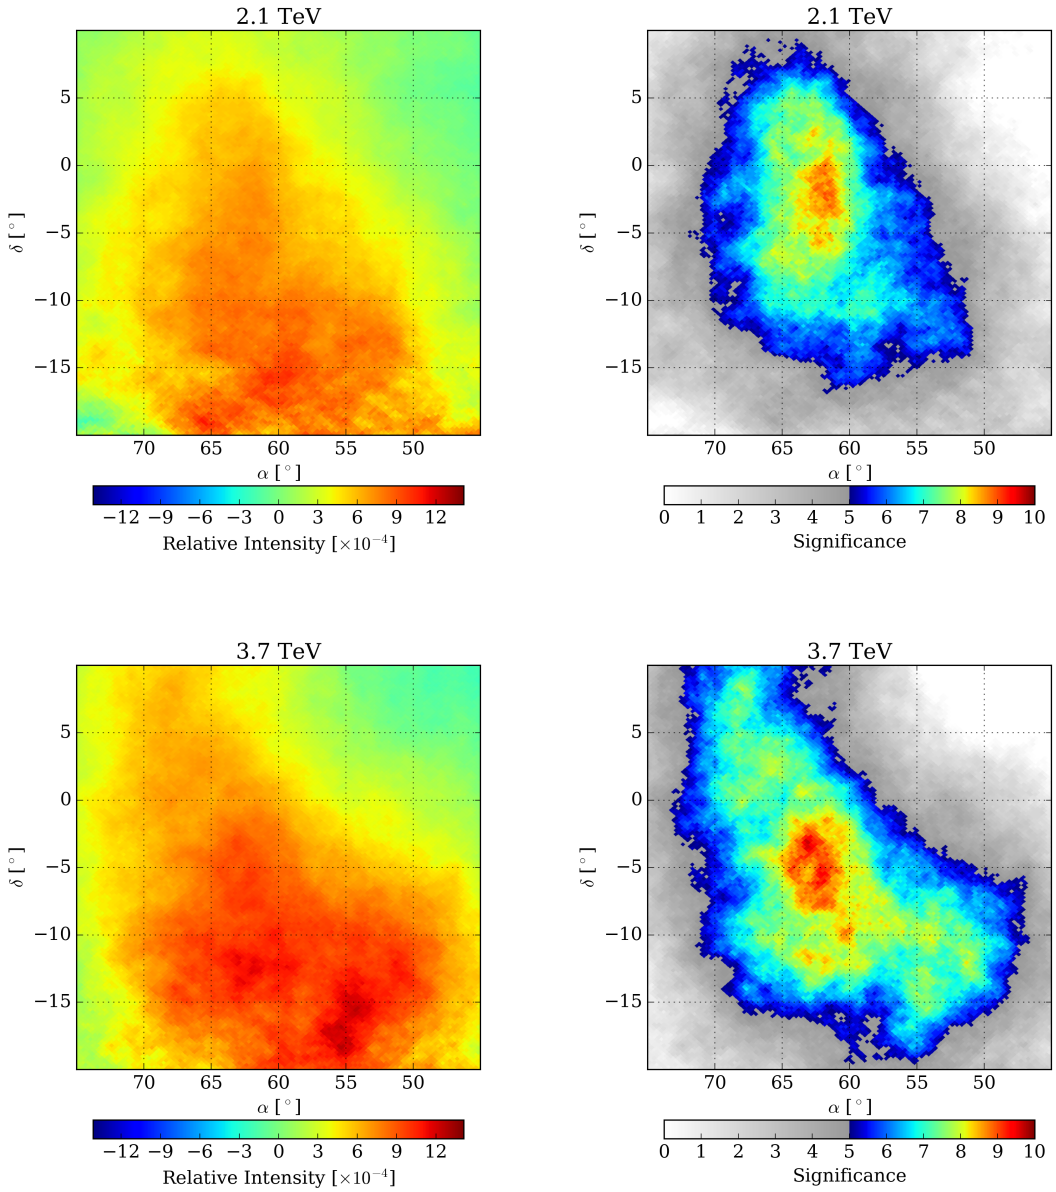

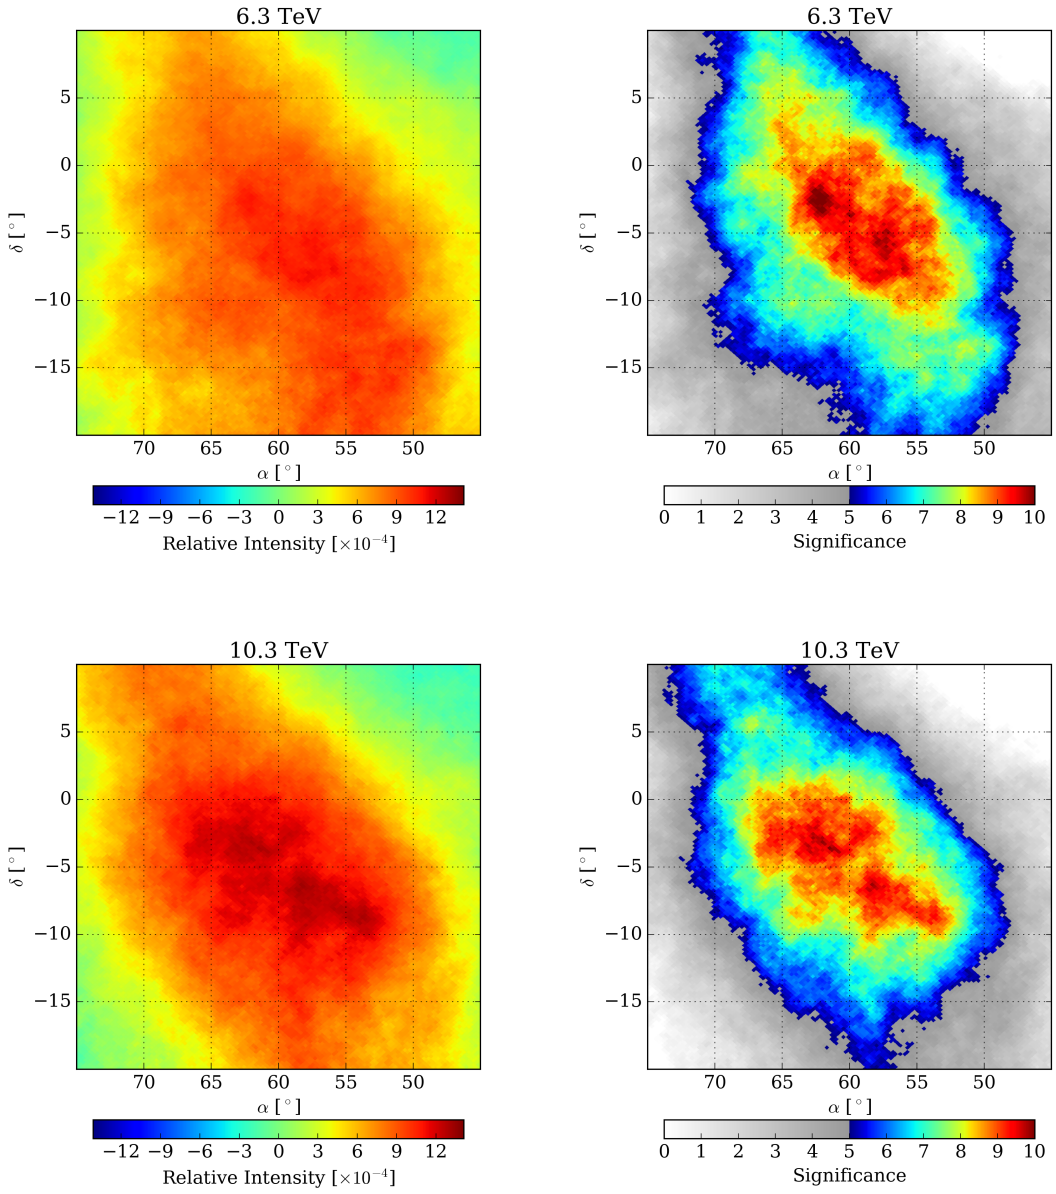
6.3 TeV
δ [°]
α [°]
Relative Intensity [×10⁻⁴]
6.3 TeV
δ [°]
α [°]
Significance
10.3 TeV
δ [°]
α [°]
Relative Intensity [×10⁻⁴]
10.3 TeV
δ [°]
α [°]
Significance

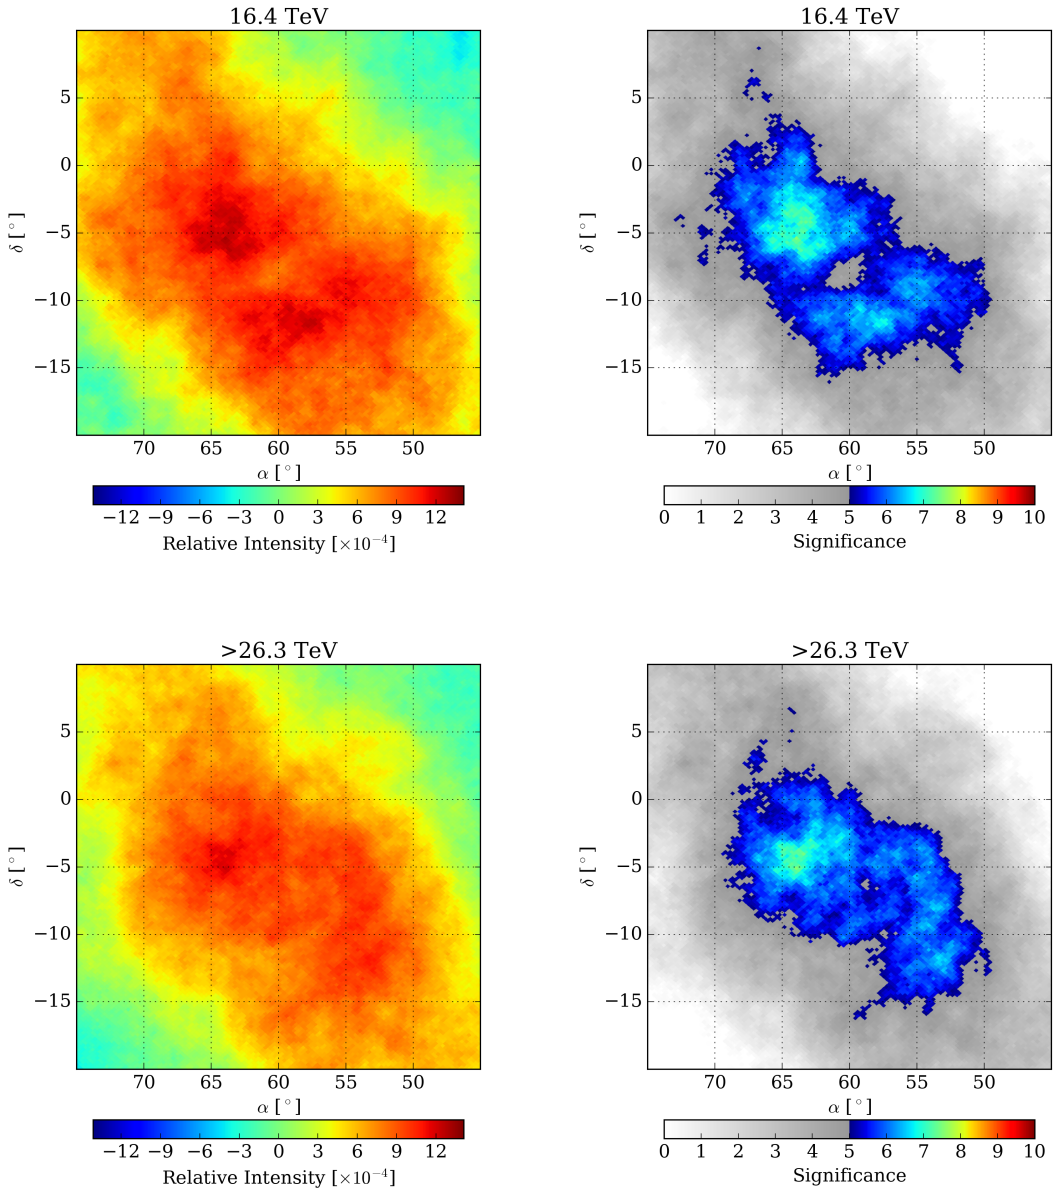
16.4 TeV
δ [°]
α [°]
Relative Intensity [×10⁻⁴]
16.4 TeV
δ [°]
α [°]
Significance
>26.3 TeV
δ [°]
α [°]
Relative Intensity [×10⁻⁴]
>26.3 TeV
δ [°]
α [°]
Significance

Chapter 10

Conclusions

A likelihood based energy estimation technique has been developed, allowing for improved energy reconstruction for cosmic ray studies with HAWC. It has been used in the context of three analyses: evolution of the cosmic ray Moon shadow with energy, the all-particle cosmic ray energy spectrum, and the energy dependence of the most significant excess of the cosmic ray anisotropy.

The cosmic ray Moon shadow was observed in ten bins in reconstructed energy, allowing for a detailed analysis of its deviation in right ascension, including an observed offset in declination. The centroid of the shadow shifted in accordance with charged particle propagation through the Earth's magnetic field, as determined via a novel GPU simulation.

The utility of the GPU simulation extends beyond the study of the Moon shadow. In fact, it is currently being adapted to investigate the much weaker cosmic ray Sun shadow, which is complicated by the addition of the Solar magnetic field. This field is less well-measured than the Earth's, so using the GPU simulation will allow the testing of model uncertainties and their effect on the expected Solar shadow. The GPU toolkit can also be used in the context of testing models of cosmic ray propagation within the Galaxy and investigating mechanisms potentially giving rise to the observed anisotropy.

With an understanding of the energy scale based on the Moon shadow observations, the all-particle cosmic ray energy spectrum was measured between 10–500 TeV. There is an indication of a change in the spectral index around 40 TeV, which is consistent with the evidence of spectral hardening of the helium spectrum, and subsequent softening of both protons and helium which has been observed by the CREAM and ATIC direct detection experiments. HAWC is also capable of investigating the PeV scale in order to probe the spectrum of the knee region; however, an improved core quality criterion is required to reach these energies. With the currently planned and funded HAWC outrigger extension, the simple N_{r40} criterion may well become obsolete, as the larger coverage will permit enhanced event selection.

Finally, we studied the energy dependence of the cosmic ray anisotropy, specifically, the spectrum of the differential relative intensity of Region A. We find that the spectrum of Region A is consistent with being harder than the isotropic cosmic ray background, in agreement with previous observations. We also see

evidence of a cutoff around 10 TeV, and due to the improved energy estimation, a stronger limit on the shape of this feature should be accessible to HAWC in the near future.

LIST OF REFERENCES

[1] A. U. Abeysekara et al. (HAWC Collaboration), *Observation of Small-scale Anisotropy in the Arrival Direction Distribution of TeV Cosmic Rays with HAWC*, Astrophysical Journal **796(2)**, 108 (2014).

[2] V. Hess, *Über Beobachtungen der Durchdringenden Strahlung bei sieben Freiballonfahrten*, Physikalische Zeitschrift **13**, 1084 (1912).

[3] W. Kolhörster, *Messungen der Durchdringenden Strahlung im Freiballon in Grösseren Höhen*, Physikalische Zeitschrift **14**, 1153 (1913).

[4] W. Bothe and W. Kolhörster, *Das Wesen der Höhenstrahlung*, Physikalische Zeitschrift **56**, 751 (1929).

[5] P. M. S. Blackett and G. P. S. Occhialini, *Some Photographs of the Tracks of Penetrating Radiation*, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **139(839)**, 699 (1933).

[6] C. D. Anderson, *The Apparent Existence of Easily Deflectable Positives*, Science **76(1967)**, 238 (1932).

[7] C. D. Anderson and S. H. Neddermeyer, *Cloud Chamber Observations of Cosmic Rays at 4300 Meters Elevation and Near Sea-Level*, Physical Review **50**, 263 (1936).

[8] B. Rossi, *Directional Measurements on the Cosmic Rays Near the Geomagnetic Equator*, Physical Review **45**, 212 (1934).

[9] P. Auger, P. Ehrenfest, R. Maze, J. Daudin, and Robley A. Fréon, *Extensive Cosmic-Ray Showers*, Reviews of Modern Physics **11**, 288 (1939).

[10] M. Schein, William P. Jesse, and E. O. Wollan, *The Nature of the Primary Cosmic Radiation and the Origin of the Mesotron*, Physical Review **59**, 615 (1941).

[11] P. Freier, E. J. Lofgren, E. P. Ney, F. Oppenheimer, H. L. Bradt, and B. Peters, *Evidence for Heavy Nuclei in the Primary Cosmic Radiation*, Physical Review **74**, 213 (1948).

[12] W. Heitler, *Quantum Theory of Radiation*, Oxford University Press (1944).

[13] K. Greisen, J.G. Wilson, et al, *Progress in Cosmic Ray Physics*, Amsterdam, Netherlands (1956).

[14] K. Kamata and J. Nishimura, *The Lateral and the Angular Structure Functions of Electron Showers*, Progress of Theoretical Physics Supplement **6**, 93 (1958).

[15] J. Matthews, *A Heitler Model of Extensive Air Showers*, Astroparticle Physics **22**, 387 (2005).

[16] J. J. Beatty and S. Westerhoff, *The Highest-Energy Cosmic Rays*, Annual Review of Nuclear and Particle Science **59**, 319 (2009).

[17] J. R. Hörandel, *The Origin of Galactic Cosmic Rays*, Nuclear Instruments and Methods in Physics Research A **588**, 181 (2008).

[18] J. Blümer, R. Engel, and J. Hörandel, *Cosmic Rays from the Knee to the Highest Energies*, Progress in Particle and Nuclear Physics **63(2)**, 293 (2009).

[19] C. Amsler et al., *Review of Particle Physics*, Physics Letters B **667(1-5)**, 1–6 (2008).

[20] O. Adriani et al. (PAMELA Collaboration), *PAMELA Measurements of Cosmic-ray Proton and Helium Spectra*, Science **332**, 69 (2011).

[21] H. S. Ahn et al. (CREAM Collaboration), *Energy Spectra of Cosmic-Ray Nuclei at High Energies*, Astrophysical Journal **707**, 593 (2009).

[22] H. S. Ahn et al. (CREAM Collaboration), *Discrepant Hardening Observed in Cosmic-ray Elemental Spectra*, Astrophysical Journal Letters **714**, L89–L93 (2010).

[23] T. S. Yoon et al. (CREAM Collaboration), *Cosmic-ray Proton and Helium Spectra from the First CREAM Flight*, Astrophysical Journal **728(2)**, 122 (2011).

[24] A. D. Panov et al. (ATIC Collaboration), *Energy spectra of Abundant Nuclei of Primary Cosmic Rays from the Data of ATIC-2 Experiment: Final Results*, Bulletin of the Russian Academy of Sciences: Physics **73(5)**, 564 (2009).

[25] V. I. Zatsepin and N. V. Sokolskaya. *Three Component Model of Cosmic Ray Spectra from 10 GeV to 100 PeV*, Astronomy and Astrophysics **458**, 1 (2006).

[26] N. Tomassetti, *Origin of the Proton-to-Helium Ratio Anomaly in Cosmic Rays*, Astrophysical Journal Letters **815**, L1 (2015).

[27] A. Lagutin and A. Tyumentsev, *Proton to Helium Ratio in Cosmic Rays: Analysis of the PAMELA Data*, Journal of Physics: Conference Series **409(1)**, 012043 (2013).

[28] B. Bartoli et al. (ARGO-YBJ Collaboration), *Light-Component Spectrum of the Primary Cosmic Rays in the Multi-TeV Region Measured by the ARGO-YBJ Experiment*, Physical Review D **85(9)**, 092005 (2012).

[29] B. Bartoli et al. (ARGO-YBJ Collaboratio), *Cosmic Ray Proton Plus Helium Energy Spectrum Measured by the ARGO-YBJ Experiment in the Energy Range 3–300 TeV*, Physical Review D **91(11)**, 112017 (2015).

[30] H. Tanaka et al. (GRAPES-3 Collaboration), *Studies of the Energy Spectrum and Composition of the Primary Cosmic Rays at 100-1000 TeV from the GRAPES-3 Experiment*, Journal of Physics G Nuclear Physics **39(2)**, 025201 (2012).

[31] W. Baade and F. Zwicky, *Cosmic rays from Super-Novae*, Proceedings of the National Academy of Sciences **20(5)**, 259 (1934).

[32] E. Fermi, *On the Origin of the Cosmic Radiation*, Physical Review **75(8)**, 1169 (1949).

[33] T. Stanev, *High Energy Cosmic Rays*, Springer-Verlag, Germany (2010).

[34] A. M. Hillas, *The Origin of Ultra-High-Energy Cosmic Rays*, Annual Review of Astronomy and Astrophysics **22**, 425 (1984).

[35] V. L. Ginzburg and S. I. Syrovatskii, *The Origin of Cosmic Rays*, Macmillan, New York (1964).

[36] A. Lukasiak, P. Ferrando, F. B. McDonald, and W. R. Webber, *The Isotopic Composition of Cosmic-Ray Beryllium and its Implication for the Cosmic Ray's Age*, Astrophysical Journal **423**, 426 (1994).

[37] A. Obermeier, P. Boyle, J. Hörandel, and D. Müller, *The Boron-to-Carbon Abundance Ratio and Galactic Propagation of Cosmic Radiation*, Astrophysical Journal **752**, 69 (2012).

[38] G. Di Sciascio and R. Iuppa, *On the Observation of the Cosmic Ray Anisotropy below 10^{15} eV*, ArXiv e-prints: 1407.2144 (2014).

[39] M. Ahlers, *Deciphering the Dipole Anisotropy of Galactic Cosmic Rays*, Physical Review Letters **117(15)**, 151103 (2016).

[40] M. Amenomori et al. (Tibet ASγ Collaboration), *Large-Scale Sidereal Anisotropy of Galactic Cosmic-Ray Intensity Observed by the Tibet Air Shower Array*, Astrophysical Journal Letters **626**, L29–L32 (2005).

[41] A. A. Abdo et al. (Milagro Collaboration), *The Large-Scale Cosmic-Ray Anisotropy as Observed with Milagro*, Astrophysical Journal **698**, 2121 (2009).

[42] B. Bartoli et al. (ARGO-YBJ Collaboration), *Medium Scale Anisotropy in the TeV Cosmic Ray Flux Observed by ARGO-YBJ*, Physical Review D **88(8)**, 082001 (2013).

[43] R. Abbasi, and et al. (IceCube Collaboration), *Observation of Anisotropy in the Galactic Cosmic-Ray Arrival Directions at 400 TeV with IceCube*, Astrophysical Journal **746**, 33 (2012).

[44] M. G. Aartsen, and et al. (IceCube Collaboration), *Observation of Cosmic-Ray Anisotropy with the IceTop Air Shower Array*, Astrophysical Journal **765**, 55 (2013).

[45] G. Giacinti and G. Sigl, *Local Magnetic Turbulence and TeV-PeV Cosmic Ray Anisotropies*, Physical Review Letters **109(7)**, 071101 (2012).

[46] P. L. Biermann, J. Becker Tjus, E.-S. Seo, and M. Mandelartz, *Cosmic-Ray Transport and Anisotropies*, Astrophysical Journal **768(2)**, 124 (2013).

[47] P. Desiati and A. Lazarian, *Anisotropy of TeV Cosmic Rays and Outer Heliospheric Boundaries*, Astrophysical Journal **762**, 44 (2013).

[48] M. G. Aartsen et al. (IceCube Collaboration), *Anisotropy in Cosmic-Ray Arrival Directions in the Southern Hemisphere Based on Six Years of Data from the IceCube Detector*, Astrophysical Journal **826(2)**, 220 (2016).

[49] A. U. Abeysekara et al. (HAWC Collaboration), *On the Sensitivity of the HAWC Observatory to Gamma-Ray Bursts*, Astroparticle Physics **35(10)**, 641 (2012).

[50] Meng Su, Tracy R. Slatyer, and Douglas P. Finkbeiner, *Giant Gamma-Ray Bubbles from FERMI-LAT: Active Galactic Nucleus Activity or Bipolar Galactic Wind?*, Astrophysical Journal **724(2)**, 1044 (2010).

[51] D. Mazin and M. Raue, *New Limits on the Density of the Extragalactic Background Light in the Optical to the Far Infrared from the Spectra of All Known TeV Blazars*, Astronomy and Astrophysics **471(2)**, 439 (2007).

[52] A. Neronov and D. V. Semikoz, *Sensitivity of γ-ray Telescopes for Detection of Magnetic Fields in the Intergalactic Medium*, Physical Review D **80(12)**, 123012 (2009).

[53] A. A. Abdo et al. (Milagro Collaboration), *Milagro Limits and HAWC Sensitivity for the Rate-Density of Evaporating Primordial Black Holes*, Astroparticle Physics **64**, 4 (2015).

[54] J. Goodman for the Milagro Collaboration, *Recent Results from the Milagro Gamma Ray Observatory*, Nuclear Physics B Proceedings Supplements **151**, 101 (2006).

[55] R. Atkins et al. (Milagro Collaboration), *TeV Gamma-Ray Survey of the Northern Hemisphere Sky Using the Milagro Observatory*, Astrophysical Journal **608(2)**, 680 (2004).

[56] A. A. Abdo et al. (Milagro Collabroation), *Discovery of Localized Regions of Excess 10-TeV Cosmic Rays*, Physical Review Letters **101(22)**, 221101 (2008).

[57] D. W. Marquardt, *An Algorithm for Least-Squares Estimation of Nonlinear Parameters*, Journal of the Society for Industrial and Applied Mathematics **11(2)**, 431 (1963).

[58] D. Heck, G. Schatz, T. Thouw, J. Knapp, and J.N. Capdevielle, *CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers*, Report No. FZKA 6019, Forschungszentrum Karlsruhe-Wissenhaltliche Berichte (1998).

[59] G. Battistoni et al., *The FLUKA code: Description and Benchmarking*, In Proceedings of the Hadronic Shower Simulation Workshop 2006, Volume 896 of *AIP Conference Proceedings* page 31 (2007).

[60] A. Ferrari, P.R. Sala, A. Fasso, and J. Ranft, *FLUKA: a Multi-Particle Transport Code*, CERN-2005-10 (2005), INFN/TC_05/11, SLAC-R-773.

[61] S. Ostapchenko, *QGSJET-II: Towards Reliable Description of Very High Energy Hadronic Interactions*, Nuclear Physics B - Proceedings Supplements **151**, 143 (2006).

[62] T. Pierog and K. Werner, *EPOS Model and Ultra High Energy Cosmic Rays*, Nuclear Physics B - Proceedings Supplements **196**, 102 (2009).

[63] S. Fletcher, T.K. Gaisser, P. Lipari, and T. Stanev, *SIBYLL: An Event Generator for Simulation of High-Energy Cosmic Ray Cascades*, Physical Review D **50(9)**, 5710 (1994).

[64] S. Agostinelli et al., *GEANT4: A Simulation Toolkit*, Nuclear Instruments and Methods in Physics A **506(3)**, 250 (2003).

[65] 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*, Physical Review Letters **114(17)**, 131103 (2015).

[66] M. Aguilar et al., *Precision Measurement of the Helium Flux in Primary Cosmic Rays of Rigidities 1.9 GV to 3 TV with the Alpha Magnetic Spectrometer on the International Space Station*, Physical Review Letters **115(21)**, 211101 (2015).

[67] J. R. Hörandel, *On the Knee in the Energy Spectrum of Cosmic Rays*, Astroparticle Physics **19(2)**, 193 (2003).

[68] T. K. Gaisser, *Spectrum of Cosmic-Ray Nucleons, Kaon Production, and the Atmospheric Muon Charge Ratio*, Astroparticle Physics **35(12)**, 801 (2012).

[69] N. Whitehorn et al., *Penalized Splines for Smooth Representation of High-Dimensional Monte Carlo Datasets*, Computer Physics Communications **184(9)**, 2214 (2013).

[70] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *Equation of State Calculations by Fast Computing Machines*, Journal of Chemical Physics **21**, 1087 (1953).

[71] W. K. Hastings, *Monte Carlo Sampling Methods Using Markov Chains and Their Applications*, Biometrika **57(1)**, 97 (1970).

[72] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *emcee: The MCMC Hammer*, Publications of the Astronomical Society of Pacific **125(925)**, 306 (2013).

[73] J. Goodman and J. Weare, *Ensemble Samplers with Affine Invariance*, Communications in Applied Mathematics and Computational Science **5(1)**, 65 (2010).

[74] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, Astrophysical Journal **622(2)**, 759 (2005).

[75] R. W Atkins et al. (Milagro Collaboration), *Observation of TeV Gamma Rays from the Crab Nebula with Milagro Using a New Background Rejection Technique*, Astrophysical Journal **595(2)**, 803 (2003).

[76] A. A. Abdo et al. (Milagro Collaboration), *Observation and Spectral Measurements of the Crab Nebula with Milagro*, Astrophysical Journal **750**, 63 (2012).

[77] D. W. Fiorino, *Observation of TeV-Energy Cosmic-Ray Anisotropy with the HAWC Observatory*, PhD thesis, The University of Wisconsin - Madison (2015). Awesomeness.

[78] T.-P. Li and Y.-Q. Ma, *Analysis Methods for Results in Gamma-Ray Astronomy*, Astrophysical Journal **272**, 317 (1983).

[79] Wisconsin IceCube Particle Astrophysics Center, HAWC Science Webpage (2014).

[80] D. Compston, *International Geomagnetic Reference Field (IGRF) Model*, 2011–2016

[81] E. Thébault et al., *International Geomagnetic Reference Field: the 12th Generation*, Earth, Planets and Space **67**, 79 (2015).

[82] F. J. Lowes, *The International Geomagnetic Reference Field: A "Health" Warning* (2010).

[83] N. A. Tsyganenko, *Modeling the Earth's Magnetosphere Using Spacecraft Magnetometer Data* (2016).

[84] J. M. Santander, *Observation of Cosmic-Ray Anisotropy at TeV and PeV Energies in the Southern Sky*, PhD thesis, The University of Wisconsin - Madison (2013). Awesomeness.

[85] J. Nickolls, I. Buck, M. Garland, and K. Skadron, *Scalable Parallel Programming with CUDA*, Queue **6(2)**, 40 (2008).

[86] J. E. Stone, D. Gohara, and G. Shi, *OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems*, Computing in Science and Engineering **12(3)**, 66 (2010).

[87] A. Klöckner, N. Pinto, Y. Lee, B. Catanzaro, P. Ivanov, and A. Fasih, *PyCUDA and PyOpenCL: A Scripting-Based Approach to GPU Run-Time Code Generation*, Parallel Computing **38(3)**, 157 (2012).

[88] J. P. Boris, *Relativistic Plasma Simulation - Optimization of a Hybrid Code*, Proceedings of the Fourth Conference on Numerical Simulations of Plasmas (Naval Research Laboratory, Washington, D.C.), p. 3, (1970).

[89] H. Qin et al., *Why is Boris algorithm so good?*, Physics of Plasmas **20(8)**, 084503 (2013).

[90] A. Abeysekara et al. (HAWC Collaboration), *The Sensitivity of HAWC to Steady and Transient Sources of Gamma Rays: Contributions to ICRC 2013*, In Proceedings of the 33rd ICRC, Rio de Janeiro, Brazil, July 2013.

[91] G. D'Agostini, *A Multidimensional Unfolding Method Based on Bayes' Theorem*, Nuclear Instruments and Methods in Physics Research A **362(2–3)**, 487 (1995).

[92] H. Jeffreys, *An Invariant Form for the Prior Probability in Estimation Problems*, In Proc. of the Royal Society of London Series A, Mathematical and Physical Sciences **186**, 453, London, England (1946).

[93] T. Adye, *Unfolding Algorithms and Tests using RooUnfold*, In Proceedings of the PHYSTAT 2011 Workshop, pages 313–318, CERN, Geneva, Switzerland (2011).

[94] T. Adye, *Corrected Error Calculation for Iterative Bayesian Unfolding* (2011).

[95] A. N. Kolmogorov, *Sulla Determinazione Empirica di una Legge di Distribuzione*, Giornale dell'Istituto Italiano degli Attuari **4**, 83–91 (1933).

[96] N. Smirnov, *Table for Estimating the Goodness of Fit of Empirical Distributions*, Annals of Mathematical Statistics **19(2)**, 279 (1948).

[97] T. K. Gaisser, T. Stanev, and S. Tilav, *Cosmic Ray Energy Spectrum from Measurements of Air Showers*, Frontiers of Physics **8(6)**, 748 (2013).

[98] M. Amenomori et al. (Tibet ASγ Collaboration), *The Energy Spectrum of All-Particle Cosmic Rays Around the Knee Region Observed with the Tibet-III Air-Shower Array*, Advances in Space Research **42(3)**, 467 (2008).

[99] G. Di Sciasco for the ARGO-YBJ Collaboration, *Measurement of the Cosmic Ray Energy Spectrum with ARGO-YBJ*, In Proceedings, Vulcano Workshop 2014: Frontier Objects in Astrophysics and Particle Physics: Vulcano, Italy, May 18-24 (2014).

[100] M. G. Aartsen et al. (IceCube Collaboration), *Measurement of the Cosmic Ray Energy Spectrum with IceTop-73*, Physical Review D **88(4)**, 042004 (2013).

APPENDIX
All-Particle Spectrum Data

A.1 Reconstructed Energy Distribution

$\log E_{\mathrm{reco}}$/GeV	Number of Events
2.7 – 2.8	6.42×10^3
2.8 – 2.9	4.82×10^5
2.9 – 3.0	1.95×10^7
3.0 – 3.1	2.56×10^8
3.1 – 3.2	9.28×10^8
3.2 – 3.3	1.36×10^9
3.3 – 3.4	1.32×10^9
3.4 – 3.5	1.07×10^9
3.5 – 3.6	8.29×10^8
3.6 – 3.7	6.51×10^8
3.7 – 3.8	5.11×10^8
3.8 – 3.9	3.95×10^8
3.9 – 4.0	3.00×10^8
4.0 – 4.1	2.25×10^8
4.1 – 4.2	1.65×10^8
4.2 – 4.3	1.18×10^8
4.3 – 4.4	8.22×10^7
4.4 – 4.5	5.66×10^7
4.5 – 4.6	3.87×10^7
4.6 – 4.7	2.64×10^7
4.7 – 4.8	1.82×10^7
4.8 – 4.9	1.26×10^7
4.9 – 5.0	8.63×10^6
5.0 – 5.1	5.97×10^6
5.1 – 5.2	4.14×10^6
5.2 – 5.3	2.88×10^6
5.3 – 5.4	2.00×10^6
5.4 – 5.5	1.38×10^6
5.5 – 5.6	9.51×10^5
5.6 – 5.7	6.52×10^5
5.7 – 5.8	4.44×10^5
5.8 – 5.9	3.01×10^5
5.9 – 6.0	2.07×10^5
6.0 – 6.1	1.35×10^5

Table A.1: The number of events in each reconstructed energy bin. This is the effects distribution or $n(E)$ from chapter 7.

A.2 All-Particle Cosmic-Ray Differential Flux

$\log E/\mathrm{GeV}$	N_{events} Unfolded	$\frac{dN}{dE\,d\Omega\,dt\,dA}\pm$ stat $\pm$ $\mathrm{sys_{MC}}+$ sys $-$ sys [GeV s m^2 sr]$^{-1}$
4.0 – 4.1	2.86×10^{10}	$(8.6056 \pm 0.0003 \pm 1.2795 + 0.7913 - 1.6446) \times 10^{-7}$
4.1 – 4.2	2.00×10^{10}	$(4.7968 \pm 0.0002 \pm 0.5901 + 0.4288 - 0.8530) \times 10^{-7}$
4.2 – 4.3	1.42×10^{10}	$(2.6922 \pm 0.0001 \pm 0.2323 + 0.2360 - 0.4467) \times 10^{-7}$
4.3 – 4.4	1.00×10^{10}	$(1.5163 \pm 0.0001 \pm 0.1189 + 0.1315 - 0.2356) \times 10^{-7}$
4.4 – 4.5	7.08×10^{9}	$(8.4947 \pm 0.0007 \pm 0.7137 + 0.7352 - 1.2419) \times 10^{-8}$
4.5 – 4.6	5.02×10^{9}	$(4.7823 \pm 0.0005 \pm 0.3896 + 0.4171 - 0.6614) \times 10^{-8}$
4.6 – 4.7	3.54×10^{9}	$(2.6761 \pm 0.0003 \pm 0.2536 + 0.2377 - 0.3522) \times 10^{-8}$
4.7 – 4.8	2.47×10^{9}	$(1.4823 \pm 0.0002 \pm 0.1305 + 0.1357 - 0.1869) \times 10^{-8}$
4.8 – 4.9	1.71×10^{9}	$(8.1839 \pm 0.0015 \pm 0.8041 + 0.7830 - 0.9947) \times 10^{-9}$
4.9 – 5.0	1.18×10^{9}	$(4.4769 \pm 0.0010 \pm 0.4488 + 0.4547 - 0.5281) \times 10^{-9}$
5.0 – 5.1	8.03×10^{8}	$(2.4193 \pm 0.0007 \pm 0.2504 + 0.2655 - 0.2787) \times 10^{-9}$
5.1 – 5.2	5.34×10^{8}	$(1.2781 \pm 0.0004 \pm 0.1349 + 0.1544 - 0.1447) \times 10^{-9}$
5.2 – 5.3	3.56×10^{8}	$(6.7636 \pm 0.0027 \pm 0.6441 + 0.9164 - 0.7576) \times 10^{-10}$
5.3 – 5.4	2.37×10^{8}	$(3.5835 \pm 0.0017 \pm 0.3331 + 0.5544 - 0.3995) \times 10^{-10}$
5.4 – 5.5	1.59×10^{8}	$(1.9107 \pm 0.0011 \pm 0.1644 + 0.3430 - 0.2134) \times 10^{-10}$
5.5 – 5.6	1.09×10^{8}	$(1.0346 \pm 0.0007 \pm 0.0892 + 0.2184 - 0.1166) \times 10^{-10}$
5.6 – 5.7	7.25×10^{7}	$(5.4882 \pm 0.0047 \pm 0.4659 + 1.3920 - 0.6286) \times 10^{-11}$
5.7 – 5.8	4.87×10^{7}	$(2.9284 \pm 0.0030 \pm 0.2402 + 0.9642 - 0.3441) \times 10^{-11}$

Table A.2: Values of the all-particle cosmic-ray energy spectrum from 10–500 TeV including uncertainties. The second column is the number of events unfolded, or the distribution $\phi(C)$ from chapter 7. The label "stat" represents the statistical uncertainties, "$\mathrm{sys_{MC}}$" is for the uncertainties from the Monte Carlo, and "sys" represents the remaining sources of systematic uncertainty.

APPENDIX
Derivation of Unfolding Uncertainties

B.1 Expansion of Components of V

Recalling the unfolding formulae from before,

$$\phi_\mu^{i+1} = \sum_k M_{\mu k}\, n_k \qquad\qquad M_{\mu j} = \frac{P_{\mu j}\, \phi_\mu^i}{\epsilon_\mu\, f_j},$$

where the efficiency, ϵ, and normalization, f, of M are

$$\epsilon_\mu = \sum_j P_{\mu j} \qquad\qquad f_j = \sum_\mu P_{\mu j}\, \phi_\mu^i.$$

Of note is the presence of ϕ^i, ie, the unfolded cause distribution from the previous iteration, or the prior in the case $i = 0$.

We will be taking derivatives of these objects with respect to n_k and $P_{\lambda k}$, to wit,

$$\frac{\partial P_{\mu j}}{\partial n_k} = 0 \qquad \frac{\partial \epsilon_\mu}{\partial n_k} = 0 \qquad \frac{\partial f_j}{\partial n_k} = \sum_\mu P_{\mu j} \frac{\partial \phi_\mu^i}{\partial n_k} \tag{B.1}$$

$$\frac{\partial P_{\mu j}}{\partial P_{\lambda k}} = \delta_{\mu\lambda}\delta_{jk} \qquad \frac{\partial \epsilon_\mu}{\partial P_{\lambda k}} = \delta_{\lambda\mu} \qquad \frac{\partial f_j}{\partial P_{\lambda k}} = \delta_{jk}\phi_\lambda^i + \sum_\mu P_{\mu j} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}}. \tag{B.2}$$

The explicit forms of $\frac{\partial \phi_\mu^i}{\partial n_k}$ and $\frac{\partial \phi_\mu^i}{\partial P_{\lambda k}}$ will be shown below, but only for $i = 0$ do

$$\frac{\partial \phi_\mu^i}{\partial n_k} = 0 \quad , \quad \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} = 0, \tag{B.3}$$

as no unfolding has been performed. This will clearly not be the case for subsequent iterations when ϕ^i becomes dependent on n_k and $P_{\lambda k}$.

B.2 Expansion of $V^{\mathbf{Data}}$

Making the appropriate substitutions, the index representation of eq. 7.9 is

$$V(\phi_\mu^{i+1}, \phi_\nu^{i+1})^{\mathrm{Data}} = \sum_{jk} \frac{\partial \phi_\mu^{i+1}}{\partial n_j} \operatorname{cov}(n_j, n_k) \frac{\partial \phi_\nu^{i+1}}{\partial n_k},$$

with

$$\begin{aligned}
\frac{\partial \phi_\mu^{i+1}}{\partial n_j} &= \frac{\partial}{\partial n_j} \sum_k M_{\mu k} n_k \\
&= \sum_k (M_{\mu k} \frac{\partial n_k}{\partial n_j} + n_k \frac{\partial M_{\mu k}}{\partial n_j}) \\
&= \sum_k (M_{\mu k} \delta_{jk} + n_k \frac{\partial M_{\mu k}}{\partial n_j}) \\
&= M_{\mu j} + \sum_k n_k \underbrace{\frac{\partial M_{\mu k}}{\partial n_j}}
\end{aligned}$$

$$\begin{aligned}
\frac{\partial M_{\mu k}}{\partial n_j} &= \frac{\partial}{\partial n_j} \frac{P_{\mu k} \phi_\mu^i}{\epsilon_\mu f_k} \\
&= \underbrace{\frac{P_{\mu k}}{\epsilon_\mu f_k}}_{\frac{M_{\mu k}}{\phi_\mu^i}} \frac{\partial \phi_\mu^i}{\partial n_j} - \underbrace{\frac{P_{\mu k} \phi_\mu^i}{\epsilon_\mu f_k}}_{M_{\mu k}} \frac{1}{f_k} \sum_\sigma P_{\sigma k} \frac{\partial \phi_\sigma^i}{\partial n_j} \\
&= \frac{M_{\mu k}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial n_j} - M_{\mu k} \sum_\sigma \epsilon_\sigma \underbrace{\frac{P_{\sigma k}}{\epsilon_\sigma f_k}}_{\frac{M_{\sigma k}}{\phi_\sigma^i}} \frac{\partial \phi_\sigma^i}{\partial n_j} \\
&= \frac{M_{\mu k}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial n_j} - \sum_\sigma \frac{\epsilon_\sigma}{\phi_\rho^i} M_{\mu k} M_{\sigma k} \frac{\partial \phi_\sigma^i}{\partial n_j}
\end{aligned}$$

$$\frac{\partial \phi_\mu^{i+1}}{\partial n_j} = M_{\mu j} + \frac{1}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial n_j} \underbrace{\sum_k M_{\mu k} n_k}_{\phi_\mu^{i+1}} - \sum_{\sigma,k} \epsilon_\sigma \frac{n_k}{\phi_\rho^i} M_{\mu k} M_{\sigma k} \frac{\partial \phi_\sigma^i}{\partial n_j}$$

$$\frac{\partial \phi_\mu^{i+1}}{\partial n_j} = M_{\mu j} + \frac{\phi_\mu^{i+1}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial n_j} - \sum_{\sigma,k} \epsilon_\sigma \frac{n_k}{\phi_\sigma^i} M_{\mu k} M_{\sigma k} \frac{\partial \phi_\sigma^i}{\partial n_j}. \tag{B.4}$$

Recalling eq. B.3, $\frac{\partial \phi_\mu^0}{\partial n_j} = 0$ for the first iteration, eliminating the last two terms of eq. B.4 and recovering $\frac{\partial \phi_\mu^1}{\partial n_j} = M_{\mu j}$ as per [91]. In practice, one need only calculate $\frac{\partial \phi_\mu^{i+1}}{\partial n_j}$ for each iteration, saving the result until the full calculation of V^{Data} is required.

B.3 Expansion of $V^{\mathbf{MC}}$

Similar to $V(\phi_\mu^{i+1}, \phi_\nu^{i+1})^{\text{Data}}$, we identify the contributions to V from the Monte Carlo:

$$V(\phi_\mu^{i+1}, \phi_\nu^{i+1})^{\text{MC}} = \sum_{\lambda j} \sum_{\rho k} \frac{\partial \phi_\mu^{i+1}}{\partial P_{\lambda j}} \operatorname{cov}(P_{\lambda j}, P_{\rho k}) \frac{\partial \phi_\nu^{i+1}}{\partial P_{\rho k}} .$$

Proceeding forward,

$$\frac{\partial \phi_\mu^{i+1}}{\partial P_{\lambda k}} = \frac{\partial}{\partial P_{\lambda k}} \sum_j M_{\mu j} n_j = \sum_j n_j \underbrace{\frac{\partial M_{\mu j}}{\partial P_{\lambda k}}}$$

$$\begin{aligned}
\frac{\partial M_{\mu j}}{\partial P_{\lambda k}} &= \frac{\partial}{\partial P_{\lambda k}} \frac{P_{\mu j} \phi_\mu^i}{\epsilon_\mu f_j} \\
&= \frac{\phi_\mu^i}{\epsilon_\mu f_j} \frac{\partial P_{\mu j}}{\partial P_{\lambda k}} + \underbrace{\frac{P_{\mu j}}{\epsilon_\mu f_j}}_{\frac{M_{\mu j}}{\phi_\mu^i}} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} - \frac{1}{\epsilon_\mu f_j} \underbrace{\frac{P_{\mu j} \phi_\mu^i}{\epsilon_\mu f_j}}_{M_{\mu j}} \Big(f_j \frac{\partial \epsilon_\mu}{\partial P_{\lambda k}} + \epsilon_\mu \frac{\partial f_j}{\partial P_{\lambda k}} \Big) \\
&= \frac{\phi_\mu^i}{\epsilon_\mu f_j} \delta_{\lambda\mu} \delta_{jk} + \frac{M_{\mu j}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} - \frac{1}{\epsilon_\mu f_j} M_{\mu j} \Big(f_j \delta_{\lambda\mu} + \epsilon_\mu \delta_{jk} \phi_\lambda^i + \epsilon_\mu \sum_\rho P_{\rho j} \frac{\partial \phi_\rho^i}{\partial P_{\lambda k}} \Big) \\
&= \frac{\phi_\mu^i}{\epsilon_\mu f_j} \delta_{\lambda\mu} \delta_{jk} + \frac{M_{\mu j}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} - \frac{M_{\mu j}}{\epsilon_\mu} \delta_{\lambda\mu} - \frac{M_{\mu j} \phi_\lambda^i}{f_j} \delta_{jk} - \sum_\rho \epsilon_\rho M_{\mu j} \underbrace{\frac{P_{\rho j}}{\epsilon_\rho f_j}}_{\frac{M_{\rho j}}{\phi_\rho^i}} \frac{\partial \phi_\rho^i}{\partial P_{\lambda k}} \\
&= \frac{\phi_\mu^i}{\epsilon_\mu f_j} \delta_{\lambda\mu} \delta_{jk} + \frac{M_{\mu j}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} - \frac{M_{\mu j}}{\epsilon_\mu} \delta_{\lambda\mu} - \frac{M_{\mu j} \phi_\lambda^i}{f_j} \delta_{jk} - \sum_\rho M_{\rho j} M_{\mu j} \frac{\epsilon_\rho}{\phi_\rho^i} \frac{\partial \phi_\rho^i}{\partial P_{\lambda k}},
\end{aligned}$$

and going back to $\frac{\partial \phi_\mu^{i+1}}{\partial P_{\lambda k}}$ to include the sum over j,

$$\frac{\partial \phi_\mu^{i+1}}{\partial P_{\lambda k}} =$$

$$\begin{aligned}
\sum_j n_j \Big[\frac{\phi_\mu^i}{\epsilon_\mu f_j} \delta_{\lambda\mu} \delta_{jk} + \frac{M_{\mu j}}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} - \frac{M_{\mu j}}{\epsilon_\mu} \delta_{\lambda\mu} - \frac{M_{\mu j} \phi_\lambda^i}{f_j} \delta_{jk} - \sum_\rho M_{\rho j} M_{\mu j} \frac{\epsilon_\rho}{\phi_\rho^i} \frac{\partial \phi_\rho^i}{\partial P_{\lambda k}} \Big] \\
= \frac{n_k \phi_\mu^i}{\epsilon_\mu f_k} \delta_{\lambda\mu} + \frac{1}{\phi_\mu^i} \frac{\partial \phi_\mu^i}{\partial P_{\lambda k}} \underbrace{\sum_j M_{\mu j} n_j}_{\phi_\mu^{i+1}} - \frac{\delta_{\lambda\mu}}{\epsilon_\mu} \underbrace{\sum_j M_{\mu j} n_j}_{\phi_\mu^{i+1}} - \frac{n_k M_{\mu k} \phi_\lambda^i}{f_k} \\
- \sum_j \sum_\rho n_j \frac{\epsilon_\rho}{\phi_\rho^i} M_{\rho j} M_{\mu j} \frac{\partial \phi_\rho^i}{\partial P_{\lambda k}},
\end{aligned}$$

with final form

$$\begin{aligned}\frac{\partial\phi_\mu^{i+1}}{\partial P_{\lambda k}} &= \frac{\delta_{\lambda\mu}}{\epsilon_\mu}\Big(\frac{n_k\phi_\mu^i}{f_k} - \phi_\mu^{i+1}\Big) - \frac{n_k\phi_\lambda^i}{f_k}M_{\mu k} \\ &+ \frac{\phi_\mu^{i+1}}{\phi_\mu^i}\frac{\partial\phi_\mu^i}{\partial P_{\lambda k}} - \sum_{\rho,j} n_j\frac{\epsilon_\rho}{\phi_\rho^i}M_{\rho j}M_{\mu j}\frac{\partial\phi_\rho^i}{\partial P_{\lambda k}}.\end{aligned} \tag{B.5}$$

Again for the first iteration $\frac{\partial\phi_\mu^0}{\partial P_{\lambda k}} = 0$, eliminating the last two terms of eq. B.5, and recovering D'Agostini's version. Again, upon implementation one need only calculate $\frac{\partial\phi_\mu^{i+1}}{\partial P_{\lambda k}}$ at each iteration, saving it until V^{MC} is needed for error estimation.