Measurements of the Flavor Composition and Inelasticity Distribution of High-Energy Neutrino Interactions in IceCube

By

Gary Binder

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Doctor Spencer R. Klein, Co-chair
Professor Marjorie D. Shapiro, Co-chair
Professor Kam-Biu Luk
Professor Karl A. van Bibber

Fall 2017

Measurements of the Flavor Composition and Inelasticity Distribution of High-Energy Neutrino Interactions in IceCube

Abstract

Measurements of the Flavor Composition and Inelasticity Distribution of High-Energy Neutrino Interactions in IceCube

by

Gary Binder

Doctor of Philosophy in Physics

University of California, Berkeley

Doctor Spencer R. Klein, Co-chair
Professor Marjorie D. Shapiro, Co-chair

At the IceCube Neutrino Observatory, neutrino interactions are observed across a broad energy range extending from $\sim 5\,\mathrm{GeV}$ to $\sim 1\,\mathrm{PeV}$. Above $\sim 100\,\mathrm{TeV}$, neutrinos primarily originate from distant astrophysical objects whose identity is yet unknown. At lower energies, most neutrinos result from cosmic-ray interactions in Earth's atmosphere. In this work, two studies of contained neutrino interactions in IceCube are presented.

In the first study, three years of IceCube data recorded from May 2010 to May 2013 are analyzed to measure the flavor composition of astrophysical neutrinos in the energy range from 35 TeV to 1.9 PeV. The flavor composition is found to be consistent with the ratio of $(f_e : f_\mu : f_\tau)_\oplus \approx \left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ expected from the oscillation of neutrinos produced by complete pion decay in distant astrophysical sources. Limits are placed on non-standard flavor compositions that could arise in exotic physics scenarios. The energy spectrum of astrophysical neutrinos is also measured and is well-described by a falling power law with index $\gamma = 2.6 \pm 0.15$.

In the second study, five years of IceCube data from May 2011 to May 2016 are analyzed with improved background rejection and event classification techniques, and the inelasticity in charged-current ν_μ interactions is reconstructed. Improved limits on the flavor composition of astrophysical neutrinos are obtained, and there is consistency with the expected flavor ratio of $\approx \left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$. Additionally, the energy spectrum of astrophysical neutrinos is found to be consistent with a power-law with index $\gamma = 2.62 \pm 0.07$ in the energy range from 6.6 TeV to 2.2 PeV. Limits are placed on a power-law flux from a second population of astrophysical sources that may explain the harder power-law index of 2.13 ± 0.13 found in an IceCube analysis of up-going muons at higher energies.

The inelasticity distribution of charged-current ν_μ interactions is also obtained and is found to be consistent with the calculation of Cooper-Sarkar et al. across an energy range from 1 TeV to 100 TeV. The inelasticity distribution can also be used measure the neutrino

to antineutrino ratio of the atmospheric neutrino flux. A scaling factor on the ν_μ to $\bar{\nu}_\mu$ flux ratio calculated by Honda et al. is found to be $0.77^{+0.44}_{-0.25}$ in the energy range from 770 GeV to 21 TeV. Lastly, the inelasticity distribution can also be used to perform an indirect search for charged-current charm production by neutrinos. Assuming a leading-order calculation of the inelasticity distribution for charm production events, zero charm production can be excluded at 91% confidence level in the energy range from 1.5 TeV to 340 TeV. In the future, it is expected that the techniques developed for this analysis can be applied to search for neutrino interactions beyond the Standard Model in IceCube and next-generation neutrino telescopes.

Contents

Acknowledgments

First of all, I have been very fortunate to work with my adviser, Spencer Klein. He has always put my interests first and constantly worked to increase the visibility of my work. He also gave me considerable freedom to pursue my own interests and was extremely encouraging and patient as I navigated through the many obstacles of physics research. I will be grateful for the many opportunities he has given me for decades to come.

I am also thankful for my co-adviser Marjorie Shapiro, who took on the burden of supervising a student working outside the physics department and was always helpful and easily accessible. Likewise, thanks to my other committee members Kam-Biu Luk and Karl van Bibber, who were always considerate and encouraging during my qualifying exam and the writing of this dissertation. I am also thankful for assistance from Anne Takizawa and other members of the physics department staff since I began my studies in 2011.

It has been a privilege to work with the other members of the IceCube group at LBNL. Many thanks to former postdocs Lisa Gerhardt and Chang Hyon Ha who helped me get started on IceCube research. Previous grad student Sandy Miarecki set an excellent precedent that made it much easier to complete my studies. Her company always lifted our spirits during much needed coffee breaks. Former postdoc Joulien Tatar has been immensely supportive and a good friend. Current postdoc Tomasz Palczewski was also a helpful source of feedback in numerous discussions.

This work would not have been possible without the dedicated efforts of the IceCube collaboration. Many thanks to the members the Diffuse working group who reviewed the analyses in this work and provided invaluable feedback. It was a pleasure to work with the many hard-working and easy-going people in the collaboration that are too numerous to name here.

I would not have survived graduate school without the support of friends and family. Long-time friends N. and Z. have kept me sane during difficult times. Countless colorful conversations with C. have kept me afloat as well. Lastly, I am thankful for my mother, who has provided loving support in every phase of my life, and my father, who imparted the intellectual drive that pushed me through my academic career.

Chapter 1

Introduction

The concept of high-energy neutrino astrophysics was first proposed in 1960 [1, 2, 3], and following the completion of the IceCube Neutrino Observatory at the South Pole in 2010, the field has finally come to fruition. Today, neutrinos with energies up to $\sim 2 \times 10^{15}$ eV are routinely observed interacting within the cubic kilometer of ice instrumented by IceCube's optical sensors. In this work, two properties of neutrino interactions in IceCube, flavor composition and inelasticity distribution, are measured for the first time. The results have wide-ranging implications in both high-energy astrophysics and fundamental particle physics.

Neutrinos exhibit a rich phenomenology that has been full of surprises since they were first detected in 1956 [4]. As IceCube continues to detect neutrinos with unprecedented energies, the trend for surprises may yet continue. In Ch. 2, the current theoretical understanding of neutrinos in the Standard Model is reviewed with special emphasis on the two phenomena relevant for IceCube, neutrino oscillations and neutrino deep inelastic scattering. Neutrino oscillations–the ability of neutrinos to transform among three flavors, ν_e, ν_μ, and ν_τ, after traveling through space–is the most recent surprising discovery and was awarded the 2015 Nobel Prize physics [5, 6]. The current understanding of neutrino oscillations is tested in a measurement of the flavor composition, the proportion of each flavor within the high-energy neutrino flux detected by IceCube. A precise understanding of high-energy neutrino interactions with nuclei is vital to interpret data from IceCube. At high energies, neutrinos probe the quark and gluon structure of nucleons through deep inelastic scattering. The measurement of inelasticity–the fraction of a neutrino's energy transferred in an interaction–can be used to study deep inelastic scattering at energies unreachable at other accelerator-based neutrino experiments.

In addition to testing fundamental particle physics, high-energy neutrinos attract wide interest from their potential to reveal the origin of cosmic rays. When cosmic rays interact, they inevitably produce neutrinos as a byproduct. Most neutrinos detected by IceCube result from cosmic rays impacting Earth's atmosphere. Neutrinos can also be produced when cosmic rays interact with matter or radiation at the site of their production in astrophysical objects. Unlike cosmic rays, neutrinos travel through matter, radiation, or magnetic fields unaffected, and as a result, they carry unique information about the locations where cosmic

rays are produced. In Ch. 3, present knowledge of cosmic rays and how they produce atmospheric and astrophysical neutrinos is summarized.

The IceCube Neutrino Observatory was constructed with the goal of identifying high-energy astrophysical neutrinos. When neutrinos interact in IceCube, Cherenkov light from charged particles propagates through ice and is detected by an array of 5,160 photomultiplier tubes. Although the ice at the South Pole is exceptionally clear, a unique challenge faced at IceCube is understanding the spatially varying optical properties of ice. Another challenge is the overwhelming background of muons also produced by cosmic rays in the atmosphere. In Ch. 4, the complete process of neutrino detection in IceCube is described from neutrino interaction to Cherenkov light propagation to data acquisition. To understand the data acquired by IceCube, detailed simulations of the entire detection process are needed and are described in Ch. 5.

In 2013, the first evidence for a diffuse flux of astrophysical neutrinos was uncovered at IceCube [7], but to date, no neutrino has been conclusively associated with a known astronomical object. While searches for sources continue, further insight into the origin of the diffuse flux can be gained from measuring its flavor composition. After accounting for neutrino oscillations, the astrophysical neutrino flavor composition at Earth is limited to a narrow range close to a ratio of $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_{\oplus}$. If a flavor composition outside of this range were observed, it would be a sure sign of new physics affecting the propagation neutrinos over astronomical distances. A precise measurement within the range allowed by standard oscillations can reveal unique information about the environment where neutrinos were produced. A measurement of flavor composition is of great interest both as a probe of cosmic rays and fundamental neutrino physics. In Ch. 6, the first measurement of the flavor composition of astrophysical neutrinos is presented. The energy spectrum of astrophysical neutrinos, generally expected to be a power law, is also measured.

Like the flavor composition, the inelasticity distribution of neutrino interactions in IceCube is of interest in both particle physics and cosmic-ray physics. Typically, the inelasticity of ν_μ interactions can only be inferred if the energies of the muon and hadrons produced by the interaction are observed. New techniques are presented that enable inelasticity to be reconstructed for the first time in IceCube. Fundamentally, the inelasticity distribution is a probe of the quark and gluon structure of nucleons, and measurements of nucleon structure at electron-proton and proton-proton colliders can be used to predict the neutrino inelasticity distribution in IceCube. In Ch. 7, the inelasticity distribution is measured in the broad energy range from 1 TeV to 100 TeV and compared with predictions based on data from the HERA electron-proton collider.

Inelasticity distributions can also be used to indirectly identify the contribution of specific neutrino scattering processes, for example the production of a charm hadron by a neutrino. In Ch. 7, this technique is exploited to search for the signature of neutrino charm production. It is expected the same technique can be applied in the future to search for other neutrino scattering processes beyond the Standard Model with the extraordinarily high-energies available in IceCube.

The inelasticity distribution can also be used to determine the composition of neutrino

interactions in IceCube. Since neutrinos and antineutrinos produce distinct inelasticity distributions, it is possible to determine the relative number of neutrino and antineutrino interactions. This fact can be exploited to measure the neutrino to antineutrino ratio of the atmospheric neutrino flux at TeV energies, a quantity that is useful to test models of cosmic-ray interactions in the atmosphere. Although usually only the inelasticity of ν_μ interactions can be reconstructed, muons produced from secondary particle decays in ν_e and ν_τ interactions can mimic ν_μ interactions with a higher inelasticity. Thus, the distinct reconstructed inelasticity distribution for each flavor can also aid in the determination of the astrophysical flavor composition. Both of these topics are also covered in Ch. 7.

Throughout Ch. 7, novel techniques from machine learning are applied to improve background rejection, event classification, and energy reconstruction compared to the techniques of the first flavor composition measurement presented in Ch. 6. The measurements of the flavor composition and energy spectrum of astrophysical neutrinos are updated to reflect these improvements. Chapter 8 summarizes all results and concludes with the prospects for future measurements with IceCube and next-generation neutrino observatories.

Chapter 2

Theory and Phenomenology of TeV Neutrinos

Since they were first proposed in 1930 by Wolfgang Pauli to preserve the conservation of energy in beta decay [8], neutrinos have had a profound impact on the development of particle physics. Even 60 years after the first direct detection of neutrinos from nuclear reactors by Cowan and Reines [4], neutrinos are still the subject of intense experimental investigation, and many of their fundamental properties remain to be measured. This chapter will summarize the current understanding of neutrinos with emphasis on the phenomena occurring at energies beyond 1 TeV and across astronomical distances that can be probed by the IceCube detector.

2.1 Neutrinos in the Standard Model

As first proposed by Pauli, neutrinos are electrically neutral, nearly massless, spin-$\frac{1}{2}$ fermions. They interact with other particles only through the weak interaction and presumably also gravitation. In addition to giving neutrinos their name meaning "little neutral ones," Enrico Fermi developed the first theory of the weak interaction, correctly calculating the energy spectrum of beta decay in 1934 [9]. By 1970, the theory of the weak interaction in its modern form was completed by Glashow [10], Salam [11], and Weinberg [12], who achieved a unified description of both electromagnetic and weak interactions in electroweak theory. Along with quantum chromodynamics (QCD) to describe the strong interaction, electroweak theory is a pillar of the Standard Model of particle physics.

In electroweak theory, neutrinos interact through the exchange of the weak gauge bosons, $W^{\pm}$ and Z^0. With masses of $M_W = 80\,\mathrm{GeV}$ and $M_Z = 91\,\mathrm{GeV}$, the W and Z-bosons are the only massive gauge bosons in the Standard Model, a fact responsible for the feeble strength of the weak interaction in ordinary circumstances. Processes involving the W-boson are called charged-current (CC) processes, while those involving the Z-boson are called neutral-current (NC) processes. Charged-current processes include beta decay, but neutral-current

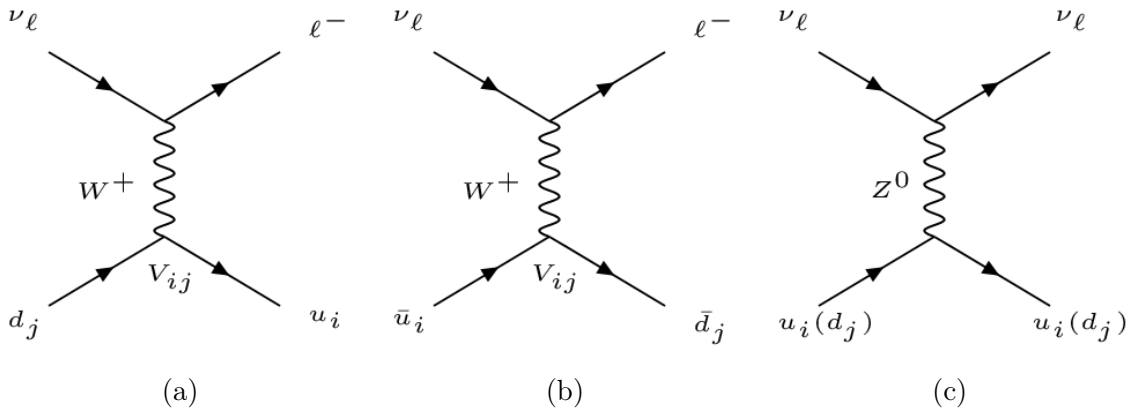

Figure 2.1: (a) The charged-current interaction of a neutrino with a down-type quark. An up-type quark is produced mixed according the CKM matrix V_{ij}. (b) The charged-current interaction of an up-type anti-quark producing a down-type anti-quark mixed according to the CKM matrix V_{ij}. The corresponding processes for antineutrinos may be found through charge conjugation of these diagrams. (c) The neutral-current interaction of a neutrino with either a down-type or up-type quark. The flavor of the quark is unchanged by neutral-current interactions.

processes were not discovered until neutral-current neutrino interactions were observed in bubble chamber experiments at CERN in 1973, verifying a key prediction of electroweak theory [13].

Neutrinos interact with the three generations of charged leptons: electrons, muons, and taus, here denoted by $\ell^\pm \in \{e^\pm, \mu^\pm, \tau^\pm\}$, through both charged-current and neutral-current processes. As shown by the Feynman diagrams in Fig. 2.1, neutrinos also interact with the three generations of quarks, both the three down-type quarks: down, strange, and bottom, $d_j \in \{d, s, b\}$, and the three up-type quarks: up, charm, and top, $u_i \in \{u, c, t\}$. In neutral-current processes, a neutrino scatters off a quark and preserves its type. In charged-current processes, the neutrino creates a charged lepton and a down-type quark is converted into an up-type quark or vice versa. A peculiarity arises from the fact that the quark eigenstates of the charged-current weak interaction differ from the eigenstates of the strong interaction that are typically used to identify quarks. Mixing of weak and strong eigenstates in the first two generations of quarks was first proposed by Cabibbo in 1963 [14], which was later extended to three generations by Kobayashi and Maskawa in 1973 as a means to explain CP violation in neutral kaon decays, even before charm or bottom quarks were discovered [15]. Mixing is quantified through the Cabibbo-Kobayashi-Maskawa (CKM) matrix, V_{ij}. For example, in a charged-current interaction with a down-type quark, d_j, the probability for identifying a given up-type quark, u_i, in the final state is $|V_{ij}|^2$.

Because in charged-current processes there is always an accompanying charged lepton, neutrinos may be described in terms of three corresponding flavor eigenstates: electron neutrinos, muon neutrinos, and tau neutrinos, $\nu_\ell \in \{\nu_e, \nu_\mu, \nu_\tau\}$. Also present are three

antineutrino states, $\bar{\nu}_\ell \in \{\bar{\nu}_e, \bar{\nu}_\mu, \bar{\nu}_\tau\}$. The existence of neutrino flavors distinct from the electron antineutrino produced in beta decay,

$$n \to p^+ + e^- + \bar{\nu}_e, \tag{2.1}$$

was unknown for much of the early history of the weak interaction. For example, a muon neutrino is produced in charged pion decay,

$$\pi^+ \to \mu^+ + \nu_\mu \tag{2.2}$$

and both an electron antineutrino and muon neutrino are produced in muon decay,

$$\mu^- \to e^- + \bar{\nu}_e + \nu_\mu. \tag{2.3}$$

The distinction between electron and muon neutrinos was first demonstrated in 1962, when muons were observed from the interaction of neutrinos produced in pion decay [16]. Tau neutrinos are substantially more difficult to produce and were observed only in 2000 by the DONUT experiment [17] through the decay of the strange D-meson,

$$D_s^+ \to \tau^+ + \nu_\tau. \tag{2.4}$$

To date, only a handful of tau neutrino interactions have ever been observed, most recently by the OPERA experiment [18]. The question of whether neutrinos and antineutrinos are distinct particles or not, i.e. whether they are Dirac or Majorana fermions, is yet unanswered but may be resolved by searching for the rare process of neutrinoless double beta decay [19].

A unique feature of the weak interaction is that it violates parity symmetry, a surprising discovery that was first observed by Wu in 1956 through the beta decay of $^{60}_{27}$Co nuclei whose spins were aligned in a magnetic field [20]. All neutrinos have a left-handed chirality, and likewise all antineutrinos have a right-handed chirality. In the limit that the neutrino is massless, chirality is equivalent to helicity, the projection of spin along momentum. Spin aligned in the same direction as momentum corresponds to right-handed helicity, while opposite spin and momentum corresponds to left-handed helicity. Neutrino helicity was first determined by Goldhaber in 1958 by measuring the polarization of gamma rays emitted following electron capture on $^{152}_{63}$Eu nuclei to infer the helicity of the escaping neutrino [21]. The existence of additional right-handed "sterile" neutrinos that do not participate in the weak interaction is an open possibility of active search, including using the IceCube detector [22].

When it was proposed, the neutrino was thought to be massless. However, if neutrinos have mass, then it is also possible that the mass eigenstates $\nu_i \in \{\nu_1, \nu_2, \nu_3\}$ are not the same as flavor eigenstates, analogous to quark mixing as described by the CKM matrix. As first realized by Pontecorvo in 1957, this gives rise to the phenomenon of neutrino oscillations, the ability of of a neutrino to be detected in a different flavor state with a probability periodic in the distance traveled [23]. Hints of neutrino oscillations were apparent when only a third of the expected amount of electron neutrinos from fusion processes in the

Sun were detected in the Homestake Mine experiment from 1970 to 1998 [24]. Neutrino oscillations as the cause were confirmed when the Sudbury Neutrino Observatory was able to observe neutral-current interactions of all neutrino flavors in accordance with solar models [25]. Independent evidence also came from observing the disappearance of atmospheric muon neutrinos produced by cosmic-ray interactions after they passed through the Earth in the Super-Kamiokande detector [26]. Though oscillations prove neutrino masses must be non-zero, only upper bounds exist on each of their values. Direct bounds on the neutrino mass can be derived from the endpoint of the beta decay energy spectrum, which limits the effective electron neutrino mass to be $\lesssim 2\,\text{eV}$ [27]. Indirect bounds from astrophysics and cosmology also exist, which can constrain the sum of the neutrino masses to be $\lesssim 0.2\,\text{eV}$ [28]. For the purposes of this work, the kinematic effects of neutrino mass in scattering and decays can be safely neglected.

2.2 Neutrino Oscillations

The mixing between neutrino flavor eigenstates, $|\nu_\alpha\rangle$, and mass eigenstates, $|\nu_j\rangle$, is described by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, $U_{\alpha i}$, according to

$$|\nu_\alpha\rangle = \sum_i U^*_{\alpha j}|\nu_j\rangle, \tag{2.5}$$

where Greek letters index flavor eigenstates, $\alpha \in \{e, \mu, \tau\}$, and Latin letters index mass eigenstates, $j \in \{1, 2, 3\}$. The PMNS matrix can be parameterized in terms of three mixing angles, θ_{12}, θ_{23}, θ_{13}, and a CP-violating phase δ_{CP},

$$U_{\alpha i} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.6}$$

where $c_{jk} = \cos\theta_{jk}$ and $s_{jk} = \sin\theta_{jk}$. The mixing angles have been measured in many experiments, with θ_{12} largely determined from solar neutrinos, θ_{23} determined from atmospheric and accelerator-produced neutrinos, and θ_{13} most recently determined from nuclear reactor antineutrinos [29]. Sensitivity to δ_{CP} is poor in all current experiments, and the best constraints on the PMNS matrix are obtained from global fits of all experimental data. Table 2.1 shows the results of the NuFIT global analysis of oscillation parameters that are used in this work [30].

2.2.1 Origin of Oscillations

To see how neutrino oscillations arise, consider a neutrino starting in an initial flavor eigenstate $|\psi(0)\rangle = |\nu_\alpha\rangle$ in a plane-wave state of definite momentum p. The Hamiltonian describing its flavor evolution is easiest to write in the mass eigenstate basis where it is

NuFIT 2.0 (2014)

	Normal Ordering ($\Delta\chi^2 = 0.97$)		Inverted Ordering (best fit)		Any Ordering
	bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range	3σ range
$\sin^2\theta_{12}$	$0.304^{+0.013}_{-0.012}$	$0.270 \to 0.344$	$0.304^{+0.013}_{-0.012}$	$0.270 \to 0.344$	$0.270 \to 0.344$
$\theta_{12}/^\circ$	$33.48^{+0.78}_{-0.75}$	$31.29 \to 35.91$	$33.48^{+0.78}_{-0.75}$	$31.29 \to 35.91$	$31.29 \to 35.91$
$\sin^2\theta_{23}$	$0.452^{+0.052}_{-0.028}$	$0.382 \to 0.643$	$0.579^{+0.025}_{-0.037}$	$0.389 \to 0.644$	$0.385 \to 0.644$
$\theta_{23}/^\circ$	$42.3^{+3.0}_{-1.6}$	$38.2 \to 53.3$	$49.5^{+1.5}_{-2.2}$	$38.6 \to 53.3$	$38.3 \to 53.3$
$\sin^2\theta_{13}$	$0.0218^{+0.0010}_{-0.0010}$	$0.0186 \to 0.0250$	$0.0219^{+0.0011}_{-0.0010}$	$0.0188 \to 0.0251$	$0.0188 \to 0.0251$
$\theta_{13}/^\circ$	$8.50^{+0.20}_{-0.21}$	$7.85 \to 9.10$	$8.51^{+0.20}_{-0.21}$	$7.87 \to 9.11$	$7.87 \to 9.11$
$\delta_{\rm CP}/^\circ$	306^{+39}_{-70}	$0 \to 360$	254^{+63}_{-62}	$0 \to 360$	$0 \to 360$
$\frac{\Delta m^2_{21}}{10^{-5}\ \text{eV}^2}$	$7.50^{+0.19}_{-0.17}$	$7.02 \to 8.09$	$7.50^{+0.19}_{-0.17}$	$7.02 \to 8.09$	$7.02 \to 8.09$
$\frac{\Delta m^2_{3\ell}}{10^{-3}\ \text{eV}^2}$	$+2.457^{+0.047}_{-0.047}$	$+2.317 \to +2.607$	$-2.449^{+0.048}_{-0.047}$	$-2.590 \to -2.307$	$\begin{bmatrix}+2.325 \to +2.599\\ -2.590 \to -2.307\end{bmatrix}$

Table 2.1: The oscillation parameters determined from the NuFIT global analysis [30]. Results are presented for the assumption of a normal mass hierarchy and an inverted mass hierarchy. Here $\Delta m^2_{3\ell}$ represents Δm^2_{31} for the normal hierarchy or $-\Delta m^2_{32}$ for the inverted hierarchy.

diagonal,

$$H = \sum_j E_j |\nu_j\rangle\langle\nu_j|. \tag{2.7}$$

The eigenvalues E_j are related to the momentum and neutrino masses m_j by

$$E_j = \sqrt{p^2 + m_j^2} \simeq p + \frac{m_j^2}{2p}, \tag{2.8}$$

where it is approximated that the neutrino is ultra-relativistic and natural units with $c = \hbar = 1$ are used. At this point, one often further approximates $p \simeq E$, even though it is technically incorrect to assign a definite energy E to a superposition of mass eigenstates with definite momentum. The neutrino's state evolves according to the Schroedinger equation,

$$|\psi(t)\rangle = e^{-iHt}|\nu_\alpha\rangle = \sum_j e^{-iE_j t}|\nu_j\rangle\langle\nu_j|\nu_\alpha\rangle = \sum_j e^{-iE_j t}U^*_{\alpha j}|\nu_j\rangle. \tag{2.9}$$

One can then calculate the transition probability to another flavor eigenstate $|\nu_\beta\rangle$,

$$P(\nu_\alpha \to \nu_\beta) = |\langle\nu_\beta|\psi(t)\rangle|^2 = \left|\sum_j e^{-iE_j t}U^*_{\alpha j}U_{\beta j}\right|^2 = \sum_{jk} U^*_{\alpha j}U_{\beta j}U_{\alpha k}U^*_{\beta k}e^{-i(E_j - E_k)t}. \tag{2.10}$$

Replacing the energy eigenvalues and further approximating that the neutrino travels a distance L in time t at the speed of light, the flavor transition probability is

$$P(\nu_\alpha \to \nu_\beta) = \sum_{jk} U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k} e^{-i\frac{\Delta m^2_{jk} L}{2E}}, \tag{2.11}$$

where $\Delta m^2_{jk} = m^2_j - m^2_k$ is the neutrino mass-squared difference. Note that the absolute values of neutrino masses themselves cannot be inferred from observations of neutrino oscillations alone and experiments only constrain the mass-squared difference as shown in Tab. 2.1. In addition, the ordering of the neutrino masses is unknown, and current measurements leave open only two possibilities, the normal mass hierarchy with $m_1 < m_2 < m_3$ and the inverted mass hierarchy with $m_1 < m_3 < m_2$. Determining the neutrino mass hierarchy remains an outstanding experimental challenge, one that might be tackled by a low-energy upgrade to the IceCube detector, the Precision IceCube Next-Generation Upgrade (PINGU) [31].

2.2.2 Observability of Oscillations

The oscillation phase is a quantity of interest to determine when flavor transformation must be considered. It can be expressed in terms of terrestrial and cosmic length scales,

$$\begin{aligned} \phi_{jk} &= \frac{\Delta m^2_{jk} L}{4E} \\ &= 8.1 \times 10^{-3} \left(\frac{\Delta m^2_{jk}}{10^{-3}\,\mathrm{eV}^2}\right)\left(\frac{L}{R_\oplus}\right)\left(\frac{1\,\mathrm{TeV}}{E}\right) \\ &= 1.89 \left(\frac{\Delta m^2_{jk}}{10^{-3}\,\mathrm{eV}^2}\right)\left(\frac{L}{1\,\mathrm{AU}}\right)\left(\frac{100\,\mathrm{TeV}}{E}\right), \end{aligned} \tag{2.12}$$

where $R_\oplus = 6371\,\mathrm{km}$ is the radius of the Earth and AU is the Astronomical Unit. For neutrinos above 1 TeV, neutrino oscillations are not a large effect even after traversing the entire Earth, but for cosmic neutrinos that originate beyond the Solar System, many oscillations occur. The oscillation pattern for such a 100 TeV cosmic neutrino is shown in Fig. 2.2.

When neutrinos propagate over very long distances, one needs to consider the possibility that neutrino oscillations lose coherence. This can result if neutrinos are emitted out of phase from multiple points at different distances. It also can arise from the fact that a neutrino cannot actually be in a perfect plane-wave momentum eigenstate, but is better described by a wave-packet having a spread in both position and momentum. Since mass eigenstates propagate at slightly different speeds, a wave-packet can disperse by the time it reaches the detector. Either possibility can be described by averaging the transition probability over distance. Separating diagonal and off-diagonal terms in Eq. 2.13, one can obtain

$$P(\nu_\alpha \to \nu_\beta) = \sum_j |U_{\alpha j}|^2 |U_{\beta j}|^2 + 2\mathrm{Re} \sum_{k>j} U^*_{\alpha j} U_{\beta j} U_{\alpha k} U^*_{\beta k} e^{-i\frac{\Delta m^2_{jk} L}{2E}}. \tag{2.13}$$

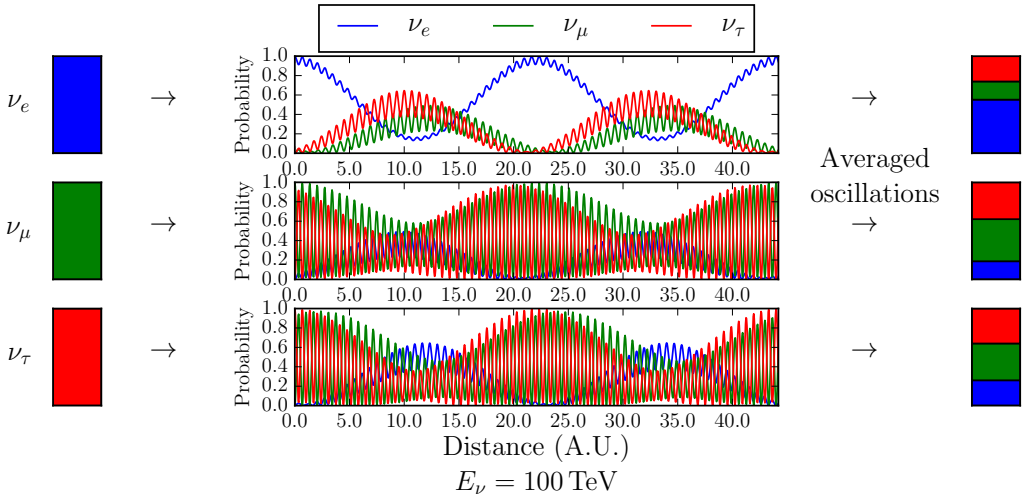

Figure 2.2: The oscillation of a 100 TeV neutrino that starts out in an initial flavor state of ν_e, ν_μ, or ν_τ from top to bottom. The best-fit oscillation parameters from [30] are used. The right-most panels show the oscillation-averaged transition probability to ν_e in blue, ν_μ in green, ν_τ in red. Distances are shown in astronomical units, illustrating a 100 TeV neutrino would undergo many oscillations if it originated outside the Solar System.

The second term averages to zero leaving

$$\langle P(\nu_\alpha \to \nu_\beta) \rangle = \sum_j |U_{\alpha j}|^2 \, |U_{\beta j}|^2 \,. \tag{2.14}$$

As shown in Fig. 2.2, even after averaging there is is still non-zero probability to transition to another flavor, $\alpha \neq \beta$. This fact has important consequences for the flavor composition of cosmic neutrinos that will be described in Sec. 3.2.2.

2.3 Neutrino Interactions

Neutrinos can be detected through rare interactions with both electrons and nuclei. The neutrino interaction cross section for electrons is readily calculated from electroweak theory, but calculating the cross section for composite particles like protons and nuclei is a more demanding task that requires detailed knowledge of their structure. Moving upward in energy, a neutrino becomes an increasingly incisive probe that reveals smaller structure within a nucleus. Eventually at energies above $\sim 10\,\mathrm{GeV}$, neutrinos principally probe the fundamental quarks and gluons in a nucleon, a regime known as deep inelastic scattering (DIS). Before the existence of quarks and gluons was definitively established, the point-like constituents of the nucleon were called partons. Though a neutrino can be thought to scatter off individual partons in the deep inelastic regime, a scattered parton cannot exist freely for long and rapidly produces a jet of hadrons through pair production in a process

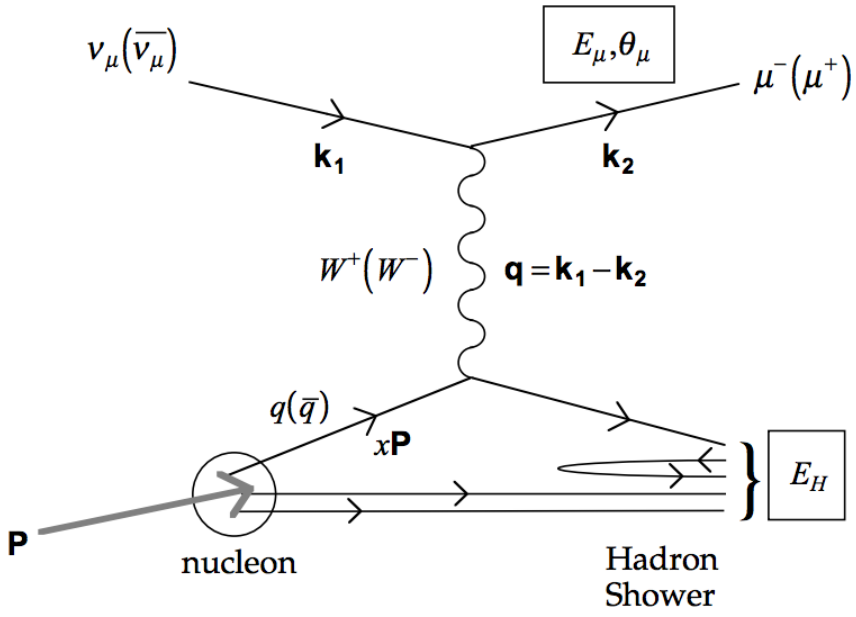

Figure 2.3: A Feynman diagram for charged-current ν_μ deep inelastic scattering on a nucleon, illustrating the creation of a muon and hadronic shower [32].

called hadronization, or fragmentation. Both charged-current (CC) and neutral-current (NC) processes can proceed in this manner to produce an outgoing charged lepton or neutrino,

$$\begin{aligned} \nu_\ell + N &\to \ell^- + X \quad \text{(CC)}, \\ \nu_\ell + N &\to \nu_\ell + X \quad \text{(NC)}, \end{aligned} \tag{2.15}$$

where N represents the target nucleus and X represents the collection of outgoing hadrons.

2.3.1 Deep Inelastic Scattering

Deep inelastic scattering is typically described using a set of Lorentz-invariant variables. These can be written in terms of 4-momenta of the incoming neutrino (k_1), target nucleon (p), outgoing lepton (k_2), and virtual gauge boson ($q = k_1 - k_2$) as illustrated in Fig. 2.3 for a CC ν_μ interaction. In the laboratory frame with the target nucleon at rest, the 4-momenta have components

$$\begin{aligned} p^\alpha &= (M, 0, 0, 0), \\ k_1^\alpha &= (E_\nu, 0, 0, E_\nu), \\ k_2^\alpha &= (E_\ell, E_\ell \sin\theta_\ell, 0, E_\ell \cos\theta_\ell), \end{aligned} \tag{2.16}$$

where M is the mass of a target nucleon, E_ν is the energy of incoming neutrino, E_ℓ is the energy of the outgoing lepton, and θ_ℓ is the angle of the outgoing lepton with respect to the incoming neutrino. The y-axis is chosen to be oriented perpendicular to the scattering

plane of outgoing particles, and both the mass of the neutrino and charged lepton are neglected. The center-of-mass energy squared, s, is important when considering whether there is sufficient energy to produce a massive particle, and can be written as

$$s = (p + k_1)^2 = 2ME_\nu + M^2 \quad \text{(center-of-mass energy squared)}. \tag{2.17}$$

The 4-momentum transfer Q^2 is defined to be

$$Q^2 = -q^2 = 4E_\nu E_\ell \sin^2\left(\frac{\theta_\ell}{2}\right) \quad \text{(4-momentum transfer squared)}. \tag{2.18}$$

Considering the wavelength of the exchanged boson, the 4-momentum transfer is a measure of the resolving power, and structure with size $R \sim 1/Q$ can be probed during scattering. The Bjorken scaling variable, x, is important when considering the picture of point-like partons in a nucleon, and represents the fraction of the target's momentum carried by the struck parton in the infinite momentum frame. If the scattered parton is massless, then its 4-momentum must satisfy $(xp+q)^2 \approx 0$. After neglecting second-order terms in target mass, this yields the definition

$$x = -\frac{q^2}{2p \cdot q} = \frac{2E_\nu E_\ell}{ME_{\text{had}}} \sin^2\left(\frac{\theta_\ell}{2}\right) \quad \text{(Bjorken scaling variable)}. \tag{2.19}$$

The inelasticity, y, has a simple interpretation in the laboratory frame,

$$y = \frac{p \cdot q}{p \cdot k_1} = \frac{E_{\text{had}}}{E_\nu} \quad \text{(inelasticity)}, \tag{2.20}$$

representing the fraction of the neutrino's energy transferred to hadrons. Lastly, the hadronic invariant mass squared,

$$W^2 = (p + q)^2 = M^2 + 2ME_{\text{had}} - Q^2 \quad \text{(hadronic invariant mass squared)}, \tag{2.21}$$

is also of interest since it determines the maximum mass of hadrons that may be produced in the interaction.

With these variables, the most general, Lorentz-invariant way to write the differential neutrino-nucleon cross section on an unpolarized target is

$$\frac{d\sigma^{\nu,\bar{\nu}}}{dxdy} = \frac{G_F^2 ME_\nu}{2\pi} \left(\frac{M_{W,Z}^2}{Q^2 + M_{W,Z}^2}\right)^2 \left[Y_+ F_2(x, Q^2) - y^2 F_L(x, Q^2) \pm Y_- x F_3(x, Q^2)\right], \tag{2.22}$$

where $Y_\pm = 1 \pm (1 - y)^2$, $+$ is used for neutrinos, $-$ is used for antineutrinos, and $M_{W,Z}$ is the W-boson mass for charged-current interactions or the Z-boson mass for neutral-current interactions [33]. All information about nucleon structure is encoded in the three structure functions, $F_2(x, Q^2)$, $xF_3(x, Q^2)$, and $F_L(x, Q^2)$. A key feature of all structure functions is that they become nearly independent of Q^2 in the limit $Q^2 \to \infty$, a property known as Bjorken scaling that was instrumental in building the parton model for deep inelastic scattering [34].

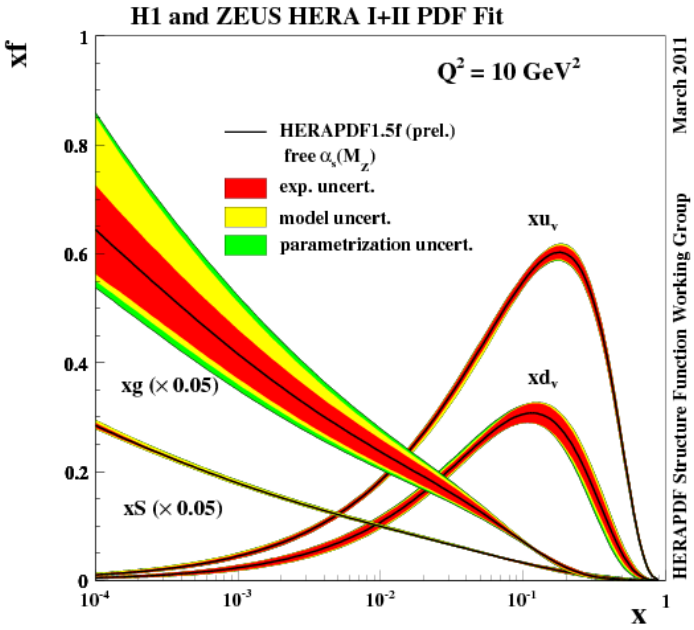

Figure 2.4: The parton distribution functions from HERAPDF1.5 at a fixed value of $Q^2 = 10\,\mathrm{GeV}^2$, showing the contribution of gluons, g, sea quarks, S, and valence quarks, d_v and u_v [35]. Colored bands indicate the estimated uncertainties.

2.3.2 Parton Distribution Functions

In QCD, the factorization theorem provides a way to compute structure functions [36]. Nucleon structure can be further described using a set of parton distribution functions (PDFs), $f_i(x, Q^2)$, for every quark type and gluons. According to the factorization theorem, given the differential cross-section between each parton and neutrino, $\hat{\sigma}_i(x, Q^2)$, the differential cross section for a nucleon can be calculated as a convolution over PDFs,

$$\sigma(x, Q^2) = \sum_i \int_0^1 \frac{d\xi}{\xi} f_i(\xi, Q^2) \hat{\sigma}_i \left(\frac{x}{\xi}, Q^2 \right). \tag{2.23}$$

The utility in the factorization theorem is that PDFs are a universal description of nucleon structure for processes beyond neutrino DIS, and PDFs determined in one experiment can be used make cross section predictions for an entirely different process.

PDFs can be calculated using the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution equations [37, 38, 39]. Expressed as a perturbative series in the strong coupling constant, α_s, the DGLAP equations allow PDFs to be computed as a function of Q^2 after providing an initial condition $f_i(x, Q_0^2)$ at a starting scale Q_0^2. At low Q^2, the perturbative approximation used by the DGLAP equations breaks down, and there is no prescription for calculating the initial condition. As a result, $f_i(x, Q_0^2)$ is often parameterized, and the

parameters must be determined from a fit to experimental data. Solutions of the DGLAP equations are typically carried out at a fixed order in α_s, and so PDFs are often labeled leading-order (LO), next-to-leading order (NLO), or next-to-next-to-leading order (NNLO).

Modern PDF fits combine data from a variety of neutrino, electron, and hadron scattering experiments. In this work, all PDFs result from the HERAPDF1.5 fits [35], which used $e^{\pm}p$ deep inelastic scattering data from the HERA collider at center-of-mass energies up to $\sqrt{s} = 318\,\text{GeV}$. These fits were valuable because they included data down to $x = 6 \times 10^{-7}$. Knowledge of PDF behavior at such low values of x is important for making accurate predictions of the neutrino cross section since the propagator term of Eq. 2.22 ensures that typical values $x \sim \frac{M_{W,Z}^2}{2ME_\nu} \sim 3 \times 10^{-3}(1\,\text{PeV}/E_\nu)$ are encountered at ultra-high energies. The resulting PDFs from the HERAPDF1.5 are shown in Fig. 2.4. The PDFs are grouped into gluon, sea, and valence contributions. The gluon, g, and sea, S, contributions can be imagined as the result of virtual gluon splitting into a sea of quarks and anti-quarks in the nucleon. Quark and anti-quark PDFs cannot be equal since the quantum numbers of a proton result from two up quarks and a down quark. The difference between the quark and anti-quark distributions is called the valence contribution, which is $d_v = d - \bar{d}$ for down quarks and $u_v = u - \bar{u}$ for up quarks. The remaining part of quark PDFs is grouped into the sea contribution and all exhibit a power-law behavior at low x, $xS \propto x^{-\lambda}$, which is also exhibited by the gluon PDF.

At this point, only a single nucleon target has been considered and it is worth noting the treatment of targets other than the proton PDFs from HERAPDF1.5. For nuclear targets, isospin symmetry can be assumed, and one can obtain the cross section per nucleon by replacing the PDFs $d \to c_p d + c_n u$ and $u \to c_p u + c_n d$ where c_p and c_n are the fraction of protons and neutrons in the target nucleus, respectively. It is common to calculate cross sections for an isoscalar target with equal numbers of protons and neutrons, $c_p = c_n = 0.5$, which is a good approximation for many materials and will be used throughout this work. Since nuclei are interacting collections of protons and neutrons, violations of the superposition principle for PDFs are expected, an effect called nuclear shadowing. Recently, evidence for nuclear shadowing has been found in neutrino DIS on carbon, iron, and lead targets [40], and nuclear PDF fits indicate nuclear effects may become important at low x and high neutrino energies [41].

2.3.3 Cross Section Calculations and Measurements

Returning to neutrino deep inelastic scattering, one can find expressions for neutrino-parton cross sections appearing in Eq. 2.23. At leading order, only neutrino-quark elastic scattering diagrams contribute as shown in Fig. 2.1. For charged-current scattering, the

cross-sections for these processes are readily calculated to be

$$\begin{aligned}\frac{d\sigma^{CC}}{dy}(\nu_\ell d_j \to \ell^- u_i) &= \frac{d\sigma^{CC}}{dy}(\bar{\nu}_\ell \bar{d}_j \to \ell^+ \bar{u}_i) = \frac{G_F^2 \hat{s}}{\pi}|V_{ij}|^2, \\ \frac{d\sigma^{CC}}{dy}(\nu_\ell \bar{u}_i \to \ell^- \bar{d}_j) &= \frac{d\sigma^{CC}}{dy}(\bar{\nu}_\ell u_i \to \ell^+ d_j) = \frac{G_F^2 \hat{s}}{\pi}(1-y)^2|V_{ij}|^2,\end{aligned} \tag{2.24}$$

where $\hat{s} = 2ME_\nu x$ is the center-of-mass energy squared of the neutrino-quark system within the nucleon and the propagator term $M_W^4/(Q^2+M_W^2)^2$ has been omitted for brevity [33]. The difference in inelasticity distribution between $\nu_\ell d_j$ and $\bar{\nu}_\ell u_i$ scattering can be understood as a consequence of chirality in the weak interaction. This is easiest to see in the center-of-mass frame where the scattering angle between the neutrino and lepton is θ^*, which is related to the inelasticity by $y = (1-\cos\theta^*)/2$ by Eq. 2.20. For $\nu_\ell d_j$ scattering, a spin-0 system is created since both the neutrino and quark are left-handed. As a result, the scattering angle distribution is isotropic and independent of y. On the other hand, $\bar{\nu}_\ell u_i$ scattering must produce a spin-1 system since now the antineutrino is right-handed. The resulting scattering angle distribution is proportional to $(1+\cos\theta^*)^2/4 = (1-y)^2$, which is completely suppressed in the backward direction, as required to conserve angular momentum.

Inserting Eqs. 2.24 into the factorization theorem of Eq. 2.23 and summing over all outgoing quark types that can be produced, one obtains the leading-order, charged-current neutrino cross section,

$$\frac{d\sigma^{CC,\nu}}{dxdy} = \frac{2G_F^2 ME_\nu}{\pi}\left(\frac{M_W^2}{Q^2+M_W^2}\right)^2 \sum_{ij} x\left[d_j(x,Q^2) + (1-y)^2\bar{u}_i(x,Q^2)\right]|V_{ij}|^2, \tag{2.25}$$

where $d_j(x,Q^2)$ is the PDF for every up-type quark and $\bar{u}_i(x,Q^2)$ is the PDF for every down-type anti-quark. The corresponding cross section for antineutrinos may be obtained by replacing $d_j \to \bar{d}_j$ and $\bar{u}_i \to u_i$. So far, it has been assumed that all partons are nearly massless and there is sufficient energy in the neutrino-parton center-of-mass frame to produce the heavy quarks charm ($m_c = 1.3\,\text{GeV}$), bottom ($m_b = 4.2\,\text{GeV}$), or top ($m_t = 173\,\text{GeV}$) [29]. In Sec. 2.3.5 this assumption will be revisited, but it does become a reasonable approximation at ultra-high neutrino energies of interest for IceCube. Taking advantage of the unitarity of the CKM matrix, one can carry out the sum over all quark types to simplify the charged-current cross section to

$$\frac{d\sigma^{CC,\nu}}{dxdy} = \frac{2G_F^2 ME_\nu}{\pi}\left(\frac{M_W^2}{Q^2+M_W^2}\right)^2 x\left[d+s+b+(1-y)^2(\bar{u}+\bar{c}+\bar{t})\right]. \tag{2.26}$$

With this formula, the leading-order structure functions can be found by comparing to Eq. 2.22, which yields

$$\begin{aligned}F_2^{CC,\nu} &= 2x(d+s+b+\bar{u}+\bar{c}+\bar{t}), \quad xF_3^{CC,\nu} = 2x(d+s+b-\bar{u}-\bar{c}-\bar{t}), \\ F_2^{CC,\bar{\nu}} &= 2x(u+c+t+\bar{d}+\bar{s}+\bar{b}), \quad xF_3^{CC,\bar{\nu}} = 2x(u+c+t-\bar{d}-\bar{s}-\bar{b}),\end{aligned} \tag{2.27}$$

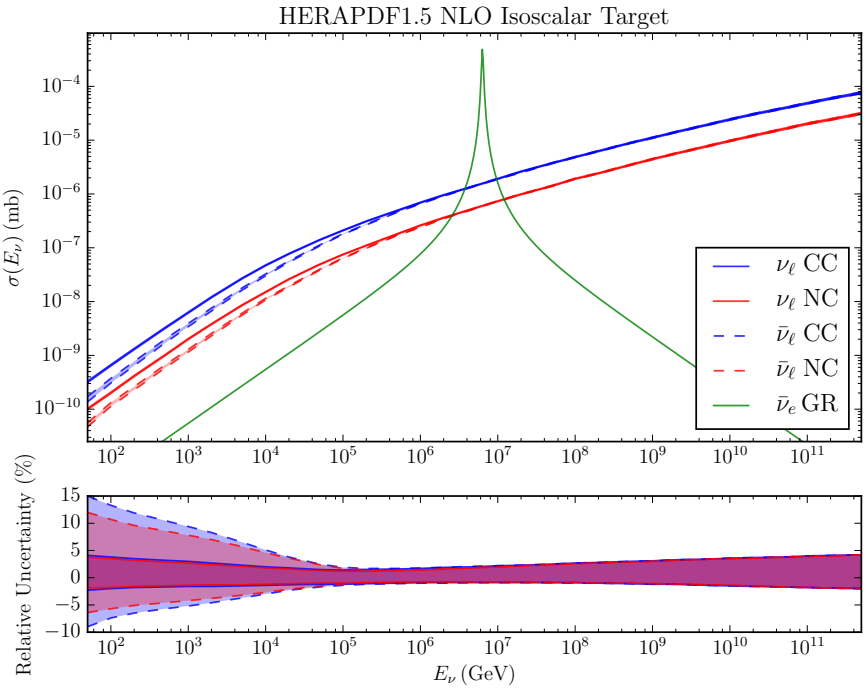

Figure 2.5: The deep inelastic scattering cross sections for charged-current (blue) and neutral-current (red) interactions for neutrinos (solid) and antineutrinos (dashed) on an isoscalar target [42]. The shaded bands indicate estimated PDF-related uncertainties, but do not include uncertainties due to nuclear effects, quark and lepton mass effects, and gluon saturation effects. The propagator term of Eq. 2.22 slows the linear growth of the cross section at energies above $\sim M_{W,Z}^2/2M \sim 10\,\text{TeV}$. The propagator also causes neutrino and antineutrino cross sections to become approximately equal at high energies because lower values of x are probed where PDFs are dominated by nearly equal contributions of sea quarks and anti-quarks. Also shown is the total cross section for the Glashow resonance (green).

and the longitudinal structure function is $F_L^{CC,\nu,\bar{\nu}} = 0$.

The calculation of neutral-current deep inelastic scattering follows a similar approach. The neutrino-quark neutral-current cross sections are

$$\begin{aligned}
\frac{d\sigma^{NC}}{dy}(\nu_\ell q_i \to \nu_\ell q_i) &= \frac{d\sigma^{NC}}{dy}(\bar{\nu}_\ell \bar{q}_i \to \bar{\nu}_\ell \bar{q}_i) = \frac{G_F^2 \hat{s}}{4\pi}\left((v_i + a_i)^2 + (v_i - a_i)^2(1-y)^2\right), \\
\frac{d\sigma^{NC}}{dy}(\nu_\ell \bar{q}_i \to \nu_\ell \bar{q}_i) &= \frac{d\sigma^{NC}}{dy}(\bar{\nu}_\ell q_i \to \bar{\nu}_\ell q_i) = \frac{G_F^2 \hat{s}}{4\pi}\left((v_i + a_i)^2(1-y)^2 + (v_i - a_i)^2\right),
\end{aligned} \tag{2.28}$$

where v_i and a_i are the quark vector and axial couplings, respectively [33]. For down-type quarks, $v_d = -\frac{1}{2} + \frac{2}{3}\sin^2\theta_W$ and $a_d = -\frac{1}{2}$, where θ_W is the Weinberg angle. For up-type quarks, $v_u = \frac{1}{2} - \frac{4}{3}\sin^2\theta_W$ and $a_u = \frac{1}{2}$. This leads to the neutrino neutral-current structure

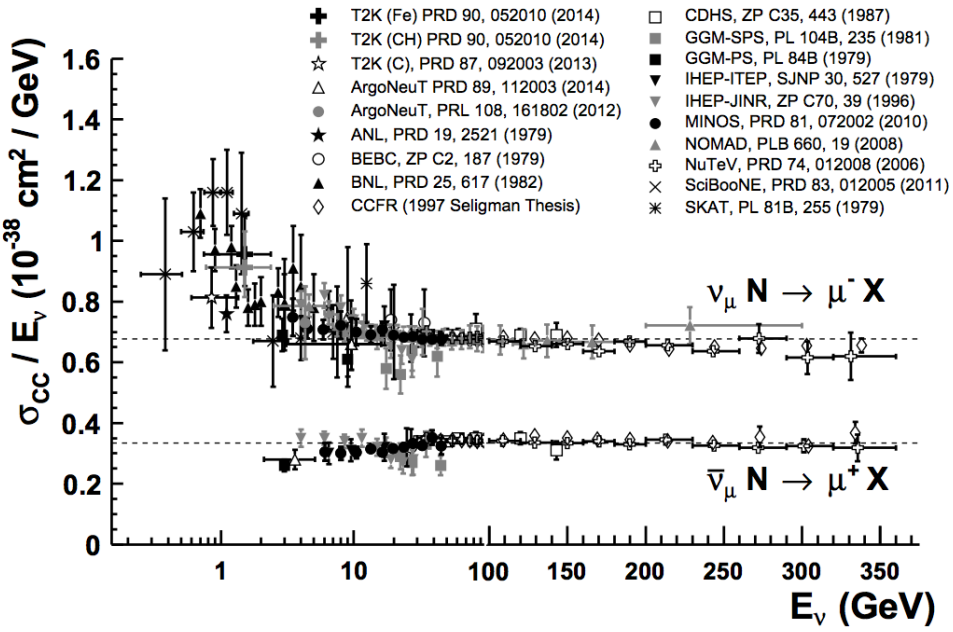

Figure 2.6: A compilation of measurements of the ν_μ and $\bar{\nu}_\mu$ CC DIS cross sections [29].

functions,

$$\begin{aligned} F_2^{NC,\nu,\bar{\nu}} &= \left(a_d^2 + v_d^2\right) x \left(d + \bar{d} + s + \bar{s} + b + \bar{b}\right) + \left(a_u^2 + v_u^2\right) x \left(u + \bar{u} + c + \bar{c} + t + \bar{t}\right), \\ F_2^{NC,\nu,\bar{\nu}} &= 2a_d v_d x \left(d - \bar{d} + s - \bar{s} + b - \bar{b}\right) + 2a_u v_u x \left(u - \bar{u} + c - \bar{c} + t - \bar{t}\right), \end{aligned} \tag{2.29}$$

and $F_L^{NC,\nu,\bar{\nu}} = 0$.

The above is a complete picture for the calculation of neutrino DIS at leading order in QCD, but higher order calculations are desirable for precision predictions. At next-to-leading order, neutrino-parton Feynman diagrams involving gluons enter the calculation of structure functions. Non-trivial convolutions over PDFs must be performed, and the structure function F_L is no longer zero. The most up-to-date calculation of neutrino DIS in the energy range of interest for IceCube has been carried out by Cooper-Sarkar, Sarkar, and Mertsch (CSMS) in 2011 and is the standard that will be used throughout this work [42]. The calculation uses the NLO HERAPDF1.5 fits and carefully quantifies all PDF-related uncertainties entering the cross section. The resulting cross sections and uncertainty bands are shown in Fig. 2.5. The PDF-related uncertainties never exceed a few percent over a large energy range from 10 TeV to 100 PeV, however a number of effects are neglected including nuclear effects, gluon saturation effects beyond the DGLAP formalism [43], and lepton and heavy quark mass effects [44, 45]. In Sec. 2.3.5, heavy quark production will be studied more closely.

As shown in Fig. 2.6, neutrino charged-current DIS cross sections have been measured at many experiments, largely using accelerator-produced neutrinos with energies much below those encountered at IceCube. The highest energy measurements are obtained by the NuTeV experiment and only extend up to 370 GeV [46]. The measurements show the expected linear

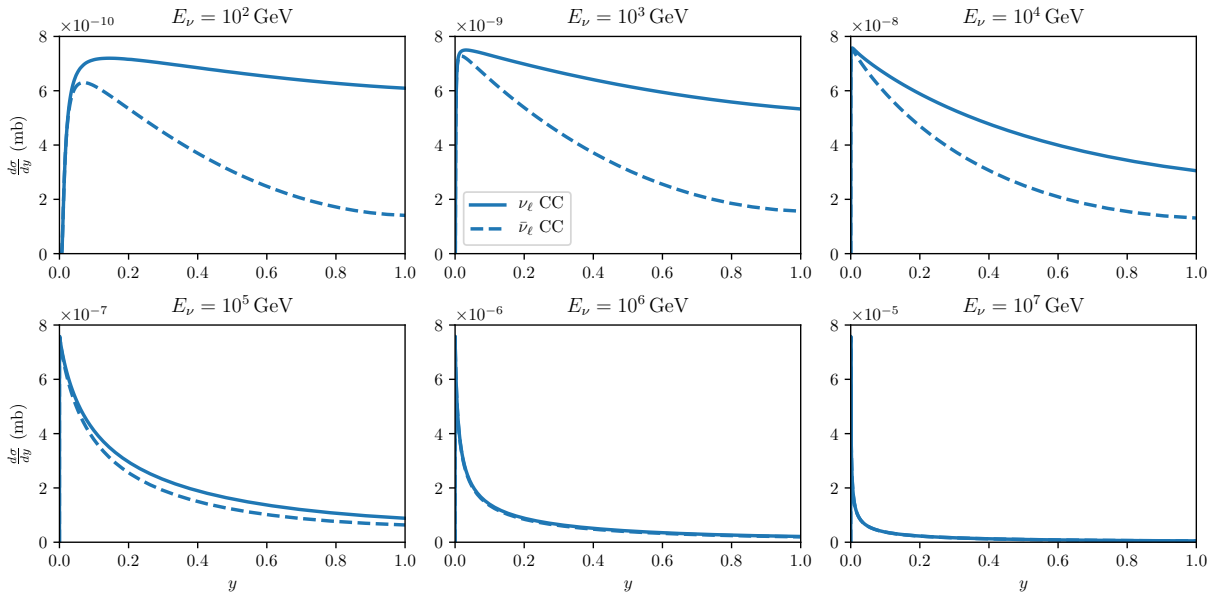

Figure 2.7: NLO calculations of the CC DIS differential cross section $d\sigma/dy$ for several neutrino energies using HERAPDF1.5. Neutrinos are indicated by solid blue lines and antineutrinos are indicated by dashed blue lines. The inelasticity distribution becomes increasingly peaked near $y = 0$ and the difference between neutrinos and antineutrinos vanishes at high energies.

dependence on energy when the propagator term of Eq. 2.22 is negligible. However, many experiments did not have a good calibration of the absolute flux of the neutrino beam and the normalization was chosen to match the world-average total cross section measurement, including the NuTeV results. Recently, the total neutrino DIS cross section was measured at IceCube in a manner independent of an absolute flux calibration by exploiting the absorption of atmospheric neutrinos as they pass through the Earth [47, 48]. The measurement uncertainty of $\sim 40\%$ was large, but the result was compatible with the CSMS prediction for energies from 6 TeV to 980 TeV.

2.3.4 Inelasticity Distribution

Complementing the recent measurement of total cross section in IceCube, the inelasticity distribution of charged-current ν_μ events in IceCube will be measured in Ch. 7. Predictions for the differential cross section, $d\sigma/dy$, at NLO using HERAPDF1.5 are shown in Fig. 2.7 for several neutrino energies. At low energies, the neutrino inelasticity distribution is nearly flat whereas it is proportional to $(1 - y)^2$ for antineutrinos. The difference in inelasticity distributions can be exploited to infer the relative number of neutrinos and antineutrinos in a sample of events. At high-energies, the difference between neutrinos and antineutrinos diminishes as the W-boson propagator starts to have an important role and equal scattering

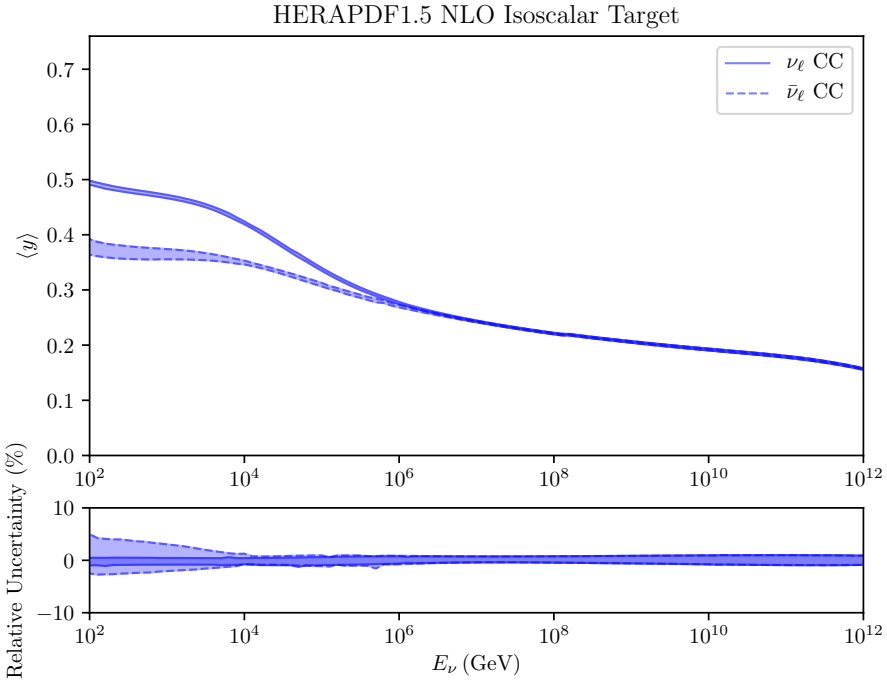

Figure 2.8: NLO calculation of mean inelasticity, $\langle y \rangle$, for CC neutrino DIS using HERAPDF1.5. Neutrinos are indicated by solid blue lines and antineutrinos are indicated by dashed blue lines. PDF-related uncertainties are indicated by the shaded bands.

off sea quarks and anti-quarks occurs. The propagator also acts to push the peak of the inelasticity distribution to $y = 0$. The energy dependence of the inelasticity distribution can be more clearly seen in Fig. 2.8 where the mean of the inelasticity distribution is shown as function of energy. Although the inelasticity distribution becomes increasingly peaked near $y = 0$, the distribution has a long tail that keeps the mean inelasticity at $\langle y \rangle \sim 0.2$ even at ultra-high energies.

2.3.5 Heavy Quark Production

Up to now the mass of heavy quarks has been ignored, but the effect of quark masses in the leading-order calculation of neutrino DIS can be approximated using the slow-rescaling model [49], which has been commonly used in many neutrino charm production experiments, e.g. [50]. The first effect of quark mass, m_i, is to introduce a production threshold on the hadronic center-of-mass energy, $W^2 > m_i^2$. Secondly, the definition of the Bjorken scaling variable must also be adapted since now the fractional momentum carried by the struck quark, x', should satisfy $(x'p + q)^2 \approx m_i^2$. This consideration can be used to define the

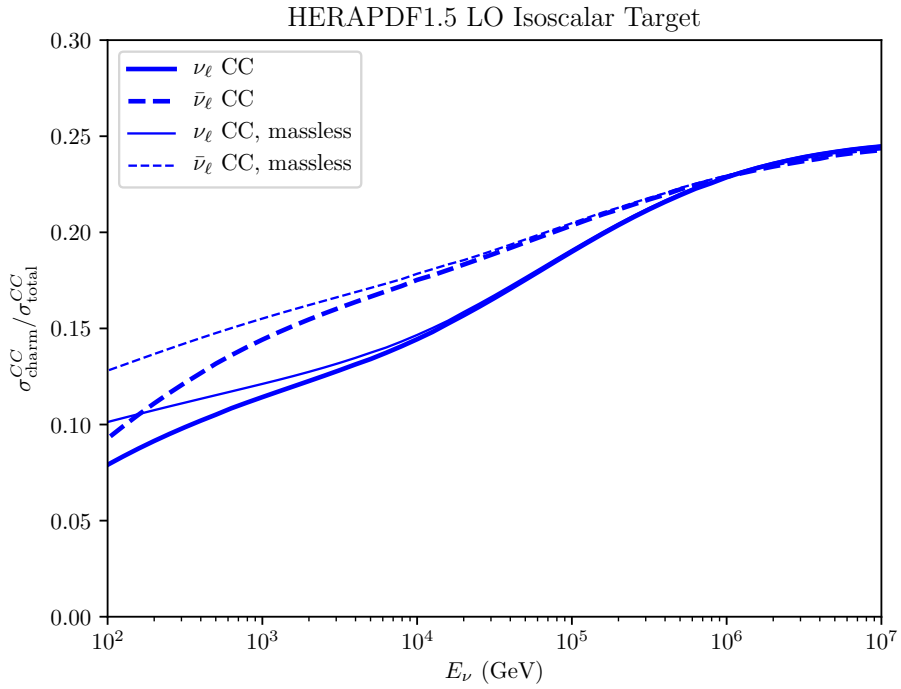

Figure 2.9: LO calculation of the charm production fraction by neutrinos (solid) and antineutrinos (dashed). Calculations using the slow-rescaling model with $m_c = 1.4\,\text{GeV}$ are shown with thick lines and the massless approximation is shown with thin lines. The massless approximation works well for energies above 1 TeV.

slow-rescaling variable,

$$x' = x\left(1 + \frac{m_i^2}{Q^2}\right), \tag{2.30}$$

where x is the usual Bjorken scaling variable as defined in Eq. 2.19. Along with shifting the partonic center-of-mass energy $\hat{s} \to \hat{s} - m_i^2$ in the parton cross section from Eq. 2.24, one can replace $x \to x'$ to obtain the leading-order cross section for the charged-current production of an up-type heavy quark u_i by a neutrino,

$$\frac{d\sigma^{CC,\nu}_{u_i}}{dx'dy} = \frac{2G_F^2 M E_\nu}{\pi}\left(1 - \frac{m_i^2}{2ME_\nu x'}\right)\left(\frac{M_W^2}{Q^2 + M_W^2}\right)^2 \sum_j x' d_j(x', Q^2)|V_{ij}|^2. \tag{2.31}$$

An analogous formula holds for antineutrinos but with $d_j \to \bar{d}_j$. This formula can be readily applied to both charm and top production. The same considerations can also be used for the production of bottom quarks, but due to the small values of $|V_{ub}|$ and $|V_{cb}|$ and the fact that a 5-flavor scheme with no top PDF is used by HERAPDF1.5, it cannot be properly calculated with the present scheme at leading order, but it is expected to be small compared to charm production [44].

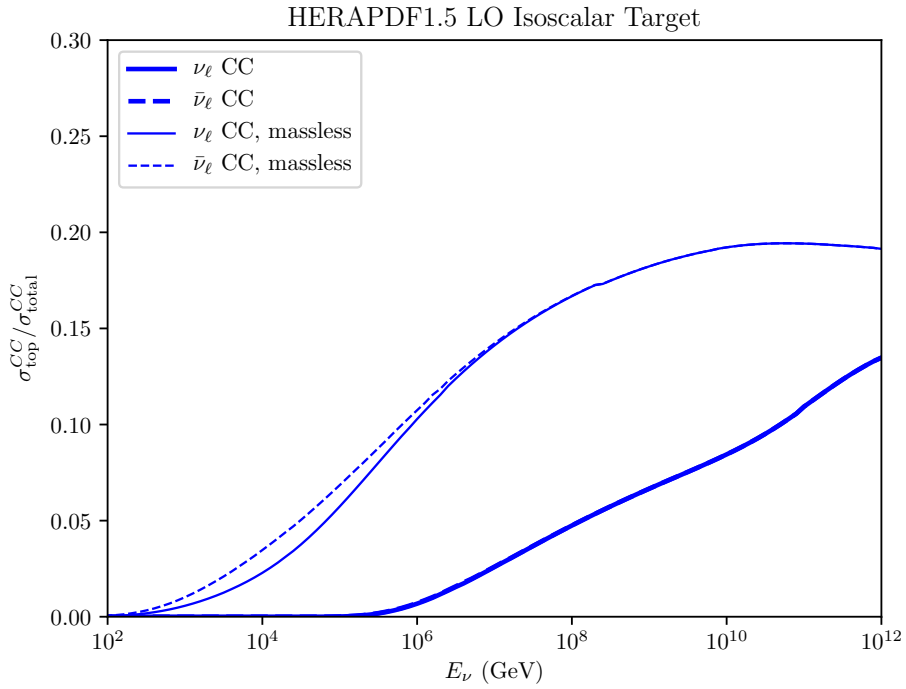

Figure 2.10: LO calculation of the top production fraction by neutrinos (solid) and antineutrinos (dashed). Calculations using the slow-rescaling model with $m_t = 180\,\text{GeV}$ are shown with thick lines and the massless approximation is shown with thin lines. The massless approximation does not work well, even for energies above 1 PeV.

Equation 2.31 can be integrated and divided by the total CC cross section to calculate the fraction of CC interactions producing a charm quark at leading order. A charm mass of $m_c = 1.4\,\text{GeV}$ is used, which is the same value used in the leading-order HERAPDF1.5 fit. The charm production fraction as a function of energy is shown in Fig. 2.9. In the region above 1 TeV where IceCube is sensitive, 10% to 20% of CC events will produce charm hadrons. The suppression caused by charm mass is visible in this energy range, but it can be safely neglected for IceCube as done by CSMS. Equation 2.31 will later be used in Sec. 5.1.2 to simulate charm production in IceCube, and in Sec. 7.4.4 to search for the contribution of charm production to events in IceCube.

For top production, the effects of mass suppression are much more pronounced. Figure 2.10 shows the fraction of CC interactions producing a top quark at leading order. A top mass of $m_t = 180\,\text{GeV}$ is used, which is the same value used in the leading-order HERAPDF1.5 fit. Even for ultra-high energies of $10^{12}\,\text{GeV}$, the slow-rescaling prediction of the top production fraction is suppressed relative to the massless approximation, a consequence of the fact that the top quark is heavier than the W-boson. Even at an energy of 1 PeV, an error of 10% could be introduced in the calculation of the total cross section by assum-

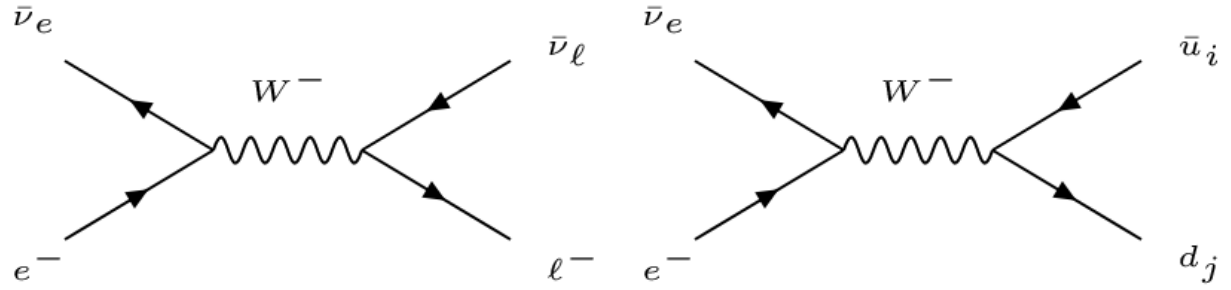

(a) (b)

Figure 2.11: (a) Feynman diagram for leptonic modes of the Glashow resonance, which have nearly equal branching fractions for e^-, μ^- and τ^-. (b) Feynman diagram for hadronic modes of the Glashow resonance. Any up or down-type quark may be produced, with the exception of the top quark, which is too heavy to be produced at the resonant energy of 6.3 PeV. The quarks quickly fragment into observable hadrons.

ing a massless top quark, larger than the estimated PDF uncertainties quoted by CSMS in Fig. 2.5. Similar results were also obtained in an recent independent study of top quark production at LO [51]. This motivates an improved calculation of top quark production at NLO in the future, but presently, top mass suppression effects occur at energies high enough to not strongly affect the results in this work.

2.3.6 Glashow Resonance

The interaction of neutrinos with atomic electrons in the detector medium is typically suppressed relative to neutrino-nucleon interactions due to the electron's much smaller mass [52]. One striking exception is when the $\bar{\nu}_e e^-$ scattering cross section is resonantly enhanced by the formation of an intermediate W-boson. The enhancement is strongest when the center-of-mass energy corresponds to the W-boson mass, which occurs at an antineutrino energy of $E_\nu = M_W^2/2m_e = 6.3\,\text{PeV}$, where m_e is the electron mass. Known as the Glashow resonance, the process was first proposed by Sheldon Glashow in 1960 as a means to detect the W-boson [53]. As shown by the Feynman diagrams in Fig. 2.11, both leptonic and hadronic final states may be formed, which are determined by the branching fractions of W-boson decay [29],

$$\begin{aligned} \mathcal{B}(W^- \to \text{Hadrons}) &= 67.41\%, \\ \mathcal{B}(W^- \to \bar{\nu}_e e^-) &= 10.71\%, \\ \mathcal{B}(W^- \to \bar{\nu}_\mu \mu^-) &= 10.63\%, \\ \mathcal{B}(W^- \to \bar{\nu}_\tau \tau^-) &= 11.38\%. \end{aligned} \tag{2.32}$$

The total cross section for the hadronic mode of the Glashow resonance is given by the Breit-Wigner formula,

$$\sigma(\bar{\nu}_e e^- \to \text{Hadrons}) = \frac{24\pi\Gamma_W^2}{\left(2m_e E_\nu - M_W^2\right)^2 + \Gamma_W^2 M_W^2}\mathcal{B}(W^- \to \bar{\nu}_e e^-)\mathcal{B}(W^- \to \text{Hadrons}) \tag{2.33}$$

where $\Gamma_W = 2.085\,\text{GeV}$ is the decay width of the W-boson [54]. Similarly, the differential cross section for the leptonic modes is given by

$$\frac{d\sigma}{dy}(\bar{\nu}_e e^- \to \bar{\nu}_\ell \ell^-) = \frac{72\pi\Gamma_W^2(1-y)^2}{\left(2m_e E_\nu - M_W^2\right)^2 + \Gamma_W^2 M_W^2}\mathcal{B}(W^- \to \bar{\nu}_e e^-)\mathcal{B}(W^- \to \bar{\nu}_\ell \ell^-) \tag{2.34}$$

where $y = E_\ell/E_\nu$ is the fraction of the antineutrino's energy taken by the outgoing charged lepton, whose mass has been neglected. The $(1-y)^2$ dependence can be seen as yet another consequence of helicity considerations in the decay of the intermediate spin-1 W-boson. Figure 2.5 shows how the total Glashow resonance cross section exceeds the deep inelastic neutrino cross sections near 6.3 PeV, making it a useful tool to search for ultra-high-energy neutrinos.

Chapter 3

Cosmic Rays and the Production of Astrophysical and Atmospheric Neutrinos

High-energy neutrinos typically originate from the weak decays of unstable particles that were produced in earlier hadronic interactions. The highest energy man-made neutrinos are produced when a proton beam from a particle accelerator impacts a fixed target, creating pions and kaons that subsequently decay to neutrinos. Limited by the beam energy and luminosity of particle accelerators, neutrino energies no higher than a few hundred GeV can be studied. To access even higher energies, one must look to the flux of naturally-occurring cosmic rays, which have been observed up to an energy of $3 \times 10^{20}\,\mathrm{eV}$ [55]. Although the origin and acceleration mechanism of cosmic rays remains highly speculative, cosmic rays will interact both in Earth's atmosphere and in any intervening material in the cosmos between Earth and their sources, creating a steady flux of high-energy neutrinos. In addition to being useful in their own right to study neutrino phenomena at high energies, atmospheric and astrophysical neutrinos are also an ideal tracer of the processes that created them. Traversing through matter, radiation, and magnetic fields almost unhindered, they carry information about cosmic-ray interactions both on Earth and in the cosmos.

3.1 Cosmic Rays

In 1912, Victor Hess found the first evidence for a flux of high-energy particles originating from beyond Earth's atmosphere when he showed that the ionization rate in air increases with altitude during a set of balloon flights [56]. In the century since then, many experiments have measured the energy spectrum and composition of cosmic rays. At energies below $\sim 10^{14}\,\mathrm{eV}$, cosmic rays can be directly detected from spacecraft and balloons before they interact in the atmosphere, and they mainly consist of protons and smaller amounts of heavier nuclei up to iron. Over a vast energy range from $\sim 10^{9}\,\mathrm{eV}$ to $\sim 10^{15}\,\mathrm{eV}$, the flux of all nucleons in the

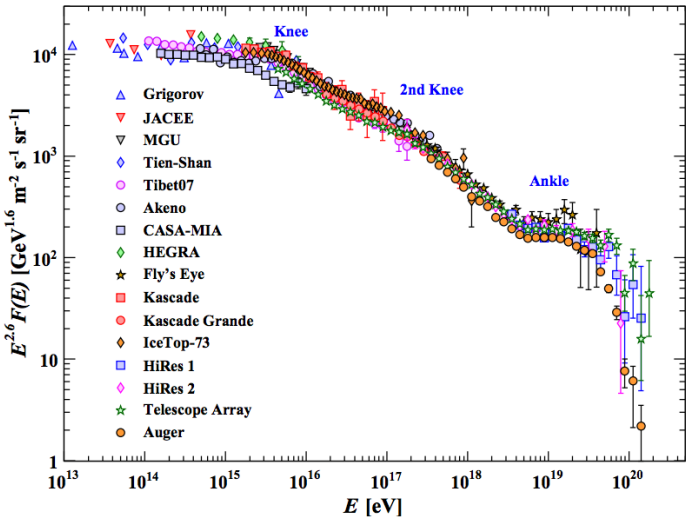

Figure 3.1: A compilation of the latest measurements of the all-particle energy spectrum of cosmic rays from air shower experiments. Each measurement is referenced in [29].

cosmic-ray flux is very well described by a power law,

$$\Phi(E) = 1.8 \times 10^4 \left(\frac{E}{1\,\mathrm{GeV}}\right)^{-2.7} \mathrm{s}^{-1}\,\mathrm{m}^{-2}\,\mathrm{sr}^{-1}. \tag{3.1}$$

At higher energies where the flux is too small to be seen in airborne or space-borne detectors, cosmic rays can only be indirectly detected from the shower of secondary particles they initiate in the atmosphere. As a result, the type and energy of a primary cosmic ray must be inferred from the properties of the cosmic-ray air shower, which are subject to significant systematic uncertainties from modeling high-energy hadronic interactions in the atmosphere. Figure 3.1 shows the energy spectrum of cosmic rays measured by air shower experiments.

The cosmic ray energy spectrum shows several interesting features at high energies. The first is a softening in the power-law index from 2.7 to 3.0 that occurs from 10^{15} eV to 10^{16} eV and is referred to as the "knee." A milder "second knee" also appears at 10^{17} eV. At $10^{18.5}$ eV, the power-law index hardens to 2.6. The most energetic cosmic ray ever observed was found by the HiRes experiment and had an inferred energy of 3×10^{20} eV [55]. At such ultra-high energies, cosmic rays should interact with the cosmic microwave background (CMB) if they originate from more than 50 Mpc outside the galaxy, strongly attenuating the cosmic ray flux above 5×10^{19} eV. This effect was first proposed by Greisen, Zatsepin, and Kuzmin (GZK) in 1966 shortly after the discovery of the CMB [57, 58]. Except for early results from the AGASA experiment [59], later measurements from the Auger [60], HiRes [61], and the Telescope Array [62] experiments appear consistent with the GZK cutoff in the ultra-high-energy cosmic-ray flux.

Despite the large effort to measure the energy spectrum and composition of cosmic rays, very little is known about the objects producing cosmic rays or how they are accelerated

to such prodigious energies. Bent by galactic magnetic fields, cosmic rays cannot be traced back to their sources. There is a small anisotropy in flux of cosmic rays across multiple angular scales at the level of $\sim 0.1\%$ up to 100 TeV [63], which suggests sources exist within the galaxy but offers little clue to their identity. At energies around $\sim 10^{19}$ eV, magnetic deflection becomes small enough that cosmic-ray arrival directions may be clustered. For example, the Telescope Array has recently uncovered a "hot spot" in the arrival directions of cosmic rays that may indicate an astrophysical source, but its radius of 20° is too large to make firm conclusions [62]. With direct identification of cosmic-ray sources so difficult, a simplified model for how cosmic rays might be accelerated is very useful to identify plausible candidate objects while also explaining many features observed in the cosmic-ray flux.

In 1949, Fermi suggested cosmic rays may be accelerated as they repeatedly deflect from moving magnetized gas clouds throughout the galaxy [64]. This idea is called second-order Fermi acceleration since cosmic rays gain energy at a rate proportional to the speed of the clouds squared. A more promising but similar idea is that cosmic rays may be accelerated as they repeatedly cross a shock in a magnetized plasma. Such shocks can originate in a number of objects, for example gamma ray bursts (GRB), supernova remnants (SNR), or the jet of an active galactic nucleus (AGN). This idea is called first-order Fermi acceleration since cosmic rays will gain energy at a rate proportional to speed of the shock-wave. For both Fermi acceleration mechanisms, one can derive that a power-law energy spectrum $E^{-\gamma}$ with $\gamma \approx 2$ should be produced [65, 66]. This is tantalizingly close to the observed power-law index of 2.7, and in fact the difference can be readily attributed to cosmic-ray diffusion in galactic magnetic fields [67]. Though Fermi acceleration is a promising mechanism, many other plausible ideas such as plasma wake-field acceleration [68] or magnetic reconnection [69] have been proposed over the years.

Essential to the idea of Fermi acceleration is that a cosmic ray must be confined in order to make repeated encounters with a shock or other magnetized object. Assuming that the confinement is due to a magnetic field with strength B, this implies that size of the confinement region, r, must be larger than the gyro-radius of the cosmic ray,

$$r > \frac{E}{ZeBc} \tag{3.2}$$

where E and Ze are energy and charge of the cosmic ray, which is assumed to be ultra-relativistic. Alternatively, this can be seen as a constraint on the rigidity, the ratio of the particle's momentum to charge,

$$R = \frac{E}{Ze} < Bcr. \tag{3.3}$$

This restriction is known as the Hillas criterion [71], and using estimates for the size and magnetic field of various astronomical objects, one can identify plausible acceleration sites for a given cosmic-ray rigidity. The magnetic field and size of several classes of objects are shown in Fig. 3.2 in a "Hillas plot."

With these ideas in mind, a simple model, originally from Hillas [73], can describe the features seen in the energy spectrum and composition of cosmic rays. In this model, there

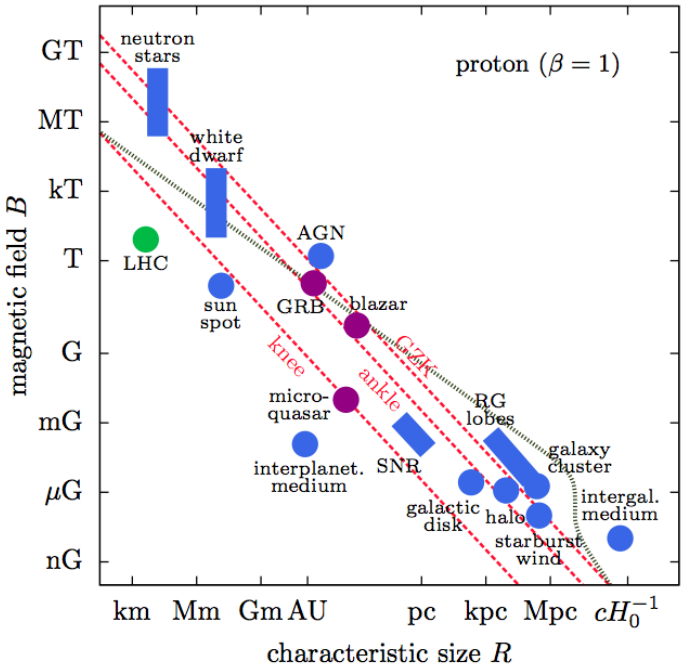

Figure 3.2: A Hillas plot illustrating the magnetic field and characteristic size of various astronomical objects in blue [70]. The dashed red lines indicate the magnetic field and size needed to confine a proton with an energy of 10^{15} eV near the knee, $10^{18.5}$ eV near the ankle, or 5×10^{19} eV where the GZK cutoff occurs. The grey line indicates an additional restriction from synchrotron radiation losses and absorption in the cosmic microwave background.

are multiple populations of cosmic-ray sources that each have a different maximum rigidity. Each population also produces a power-law energy spectrum that is exponentially attenuated above the maximum rigidity. The cosmic-ray flux for each nucleus is then parameterized as a sum over populations by

$$\Phi_i(E) = \sum_j a_{ij} \left(\frac{E}{1\,\mathrm{GeV}} \right)^{-\gamma_{ij}} \exp\left(-\frac{E}{Z_i e R_j} \right), \tag{3.4}$$

where a_{ij} is a normalization constant for each nucleus and population, γ_{ij} is a corresponding spectral index, and R_j is the maximum rigidity for each population. Several attempts to fit this model to available air shower data have been made, most recently in [74], but a three population model called the H3a model will be used throughout this work [72]. The three populations are thought to consist of two galactic populations and an extra-galactic population. The first galactic population explains most cosmic rays up to the knee, and it is thought supernova remnants (SNR) are the most natural candidates for this population. The second galactic population enters above the knee, but it is more speculative as to what sources are responsible. Lastly, the third population produces the highest energy cosmic

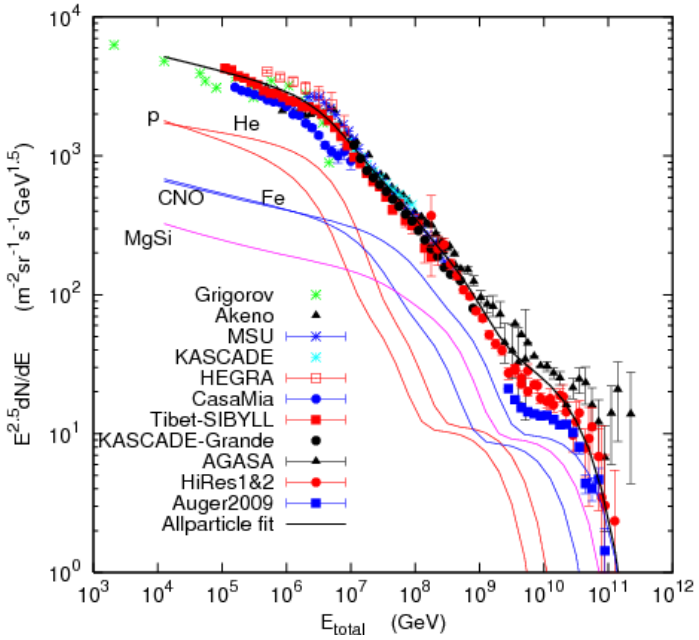

Figure 3.3: Comparison of cosmic-ray air shower measurements to the all-particle energy spectrum of the H3a model (solid black line) [72]. The colored lines indicate the contribution of groups of nuclei to the total flux. It is thought that the discrepancies between different measurements can be partially explained by miscalibration of the absolute energy scale for each experiment.

rays above the ankle, and its maximum rigidity likely requires objects outside the galaxy. A comparison of the H3a model to observed data is shown in Fig. 3.3.

3.2 Astrophysical Neutrinos

Inferring the sources of cosmic rays from their energy spectrum, composition, and anisotropy alone is very limited. To convincingly identify the sources of cosmic rays, an additional messenger particle is needed that is unaffected by galactic magnetic fields. If cosmic rays interact with gas surrounding their source, pions, kaons, neutrons, or other heavier unstable hadrons are produced

$$p^+ + p^+ \rightarrow \pi^{\pm}, \pi^0, K^{\pm}, n, \ldots \tag{3.5}$$

Cosmic rays can also interact with a radiation field surrounding the source, also producing unstable hadrons. Notably, the $p\gamma$ cross section is enhanced at the Δ^+ resonance,

$$\begin{aligned} p^+ + \gamma \rightarrow \Delta^+ &\rightarrow n + \pi^+ \\ &\rightarrow p + \pi^0, \end{aligned} \tag{3.6}$$

which is responsible for the rapid attenuation of ultra-high-energy cosmic rays in the CMB. In both pp or $p\gamma$ scenarios, neutral pions decay to gamma rays,

$$\pi^0 \to \gamma + \gamma, \tag{3.7}$$

which can travel to Earth and provide a signature of cosmic-ray interactions. In fact, this signature has been observed in two supernova remnants by the Fermi Large Area Telescope [75], which to date is the only direct evidence that supernova remnants are a source of cosmic rays. However, gamma rays do have the disadvantages that they are attenuated by intervening matter and radiation in the galaxy and the signature of neutral pion decay can often be obscured by synchrotron radiation from electrons also present in the source. Neutrinos, produced in the decay of charged pions,

$$\pi^- \to \mu^- + \bar{\nu}_\mu \tag{3.8}$$

subsequent decay of muons,

$$\mu^- \to e^- + \bar{\nu}_e + \nu_\mu \tag{3.9}$$

and decay of neutrons,

$$n \to p^+ + e^- + \bar{\nu}_e \tag{3.10}$$

overcome both difficulties. Virtually unaffected by intervening magnetic fields, matter, and radiation, astrophysical neutrinos are a smoking gun for hadronic interactions and the ideal messenger particle to identify the sources of cosmic rays.

Identifying point sources of neutrinos in the sky is the driving goal of the IceCube Neutrino Observatory and the larger field of high-energy neutrino astronomy. The largest obstacle to this goal are the background of neutrinos and muons that are also produced by cosmic rays interactions in Earth's atmosphere. Atmospheric neutrinos will be further described in Sec. 3.3. To date, no point sources of high-energy astrophysical neutrinos have been found [76], and the Sun and Supernova 1987A [77] are the only experimentally confirmed point sources of any neutrinos beyond the Earth, albeit at much lower MeV energies as shown in Fig. 3.4.

While the search for point sources continues, unique information about cosmic ray sources can still be obtained from a diffuse flux of neutrinos from multiple unresolved sources. In 2013, the first evidence for a diffuse astrophysical flux was found at the IceCube Neutrino Observatory in the energy range from $\sim 100\,\text{TeV}$ to $1\,\text{PeV}$ [7, 78]. The origin of the flux is still a wide open question, but so far it is consistent with the isotropy expected from extra-galactic sources and there is no correlation with the galactic plane. As attempts to associate the flux with known astronomical objects continue, valuable insights about its origin and connection to cosmic rays can be obtained through precise measurements of its energy spectrum and flavor composition.

3.2.1 Energy Spectrum

Predictions for a diffuse flux of high-energy astrophysical neutrinos go back to 1970s [80, 81] and have been made for many classes of objects such as gamma ray bursts [82],

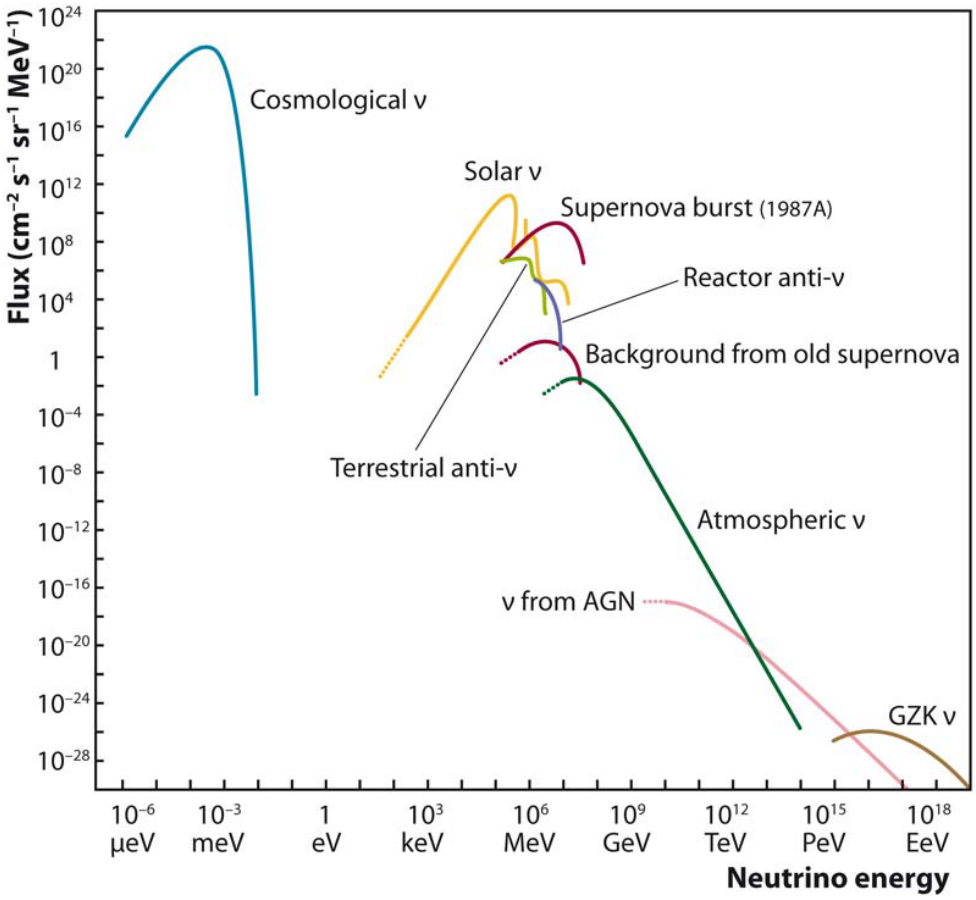

Figure 3.4: The energy spectrum of natural and man-made neutrinos [79]. IceCube is primarily sensitive to neutrinos from 100 GeV to 10 PeV. An estimated diffuse flux of astrophysical neutrinos from active galactic nuclei (AGN) is shown in pink, which begins to dominate over atmospheric neutrinos (green) at around 100 TeV. Cosmogenic neutrinos from the interaction of ultra-high-energy cosmic rays with the CMB in the GZK effect are shown in brown.

active galactic nuclei [83], or supernova remnants [84]. However, they all share several generic features. Since a neutrino from a pp or $p\gamma$ interaction takes on average $\sim 5\%$ of the original cosmic ray's energy roughly independently of energy [85], the neutrino energy spectrum should follow a power law if the cosmic-ray flux from which they originate also follows a power law. If neutrinos are produced by cosmic-ray interactions within their accelerators, this can give a unique view of the energy spectrum of cosmic rays before it is altered by propagation in magnetic fields. Because of the harder power-law index of $\gamma \approx 2$, it is generally expected that a diffuse astrophysical flux will exceed the background of atmospheric neutrinos above ~ 100 TeV, as shown in Fig. 3.4.

Cosmic-ray interactions may also occur in gas and radiation outside the accelerator. For

example, in starburst galaxies that have dense regions with high star formation rates, most neutrinos would originate from interactions in interstellar gas [86]. When cosmic rays interact with the CMB in the GZK effect, a flux of cosmogenic neutrinos is produced [87]. Though the cosmogenic neutrino flux is of great interest as a tracer of the highest energy cosmic rays, it is outside the reach of cubic kilometer detectors like IceCube [88], but it may be detected in larger future experiments utilizing radio technology such as ARA or ARIANNA [89].

A useful benchmark for the diffuse astrophysical flux was calculated by Waxman and Bahcall, who showed there is an upper bound on its magnitude if it originates from the same objects that produce the highest energy cosmic rays [90]. If ultra-high-energy cosmic rays are produced in an optically thin environment where the majority escape before interacting, only a limited neutrino flux of

$$\Phi(E) \lesssim 10^{-8} E^{-2}\,\mathrm{GeV\,s^{-1}\,cm^{-2}\,sr^{-1}} \tag{3.11}$$

per flavor can be produced in order to not exceed the observed flux of ultra-high-energy cosmic rays. Further assumptions include that cosmic rays are produced with an E^{-2} spectrum and that the cosmic ray production rate over time is no more extreme than the star formation rate. It is possible that the Waxman-Bahcall bound can be violated in sources optically thick to cosmic rays, but such sources would not naturally explain the origin of ultra-high-energy cosmic rays. In 2013 when the astrophysical flux was first discovered, its magnitude was found to be very close to the Waxman-Bahcall bound [7, 78].

3.2.2 Flavor Composition

Flavor composition is perhaps an even more powerful probe of the sources of astrophysical neutrinos. Flavor composition, or alternatively called the flavor ratio, is described as a set of three numbers, $(f_e : f_\mu : f_\tau)$, often normalized so that $f_e + f_\mu + f_\tau = 1$, representing the fraction of each neutrino flavor in the astrophysical flux. Because neutrinos oscillate as they travel from their sources to the Earth, it is important to distinguish the flavor composition at sources, $(f_e : f_\mu : f_\tau)_S$, and flavor composition at the Earth, $(f_e, f_\mu, f_\tau)_\oplus$. If neutrinos are detected from multiple sources and/or decohere after traveling a sufficiently long distance, the relationship between the flavor composition at sources and at Earth can be described as a linear transformation,

$$f_{\alpha,\oplus} = \sum_{\beta,i} |U_{\alpha i}|^2 \, |U_{\beta i}|^2 \, f_{\beta,S}, \tag{3.12}$$

as demonstrated in Sec. 2.2.2. Using the best-fit oscillation parameters from [30], this can be written numerically as

$$\begin{pmatrix} f_e \\ f_\mu \\ f_\tau \end{pmatrix}_\oplus = \begin{pmatrix} 0.552^{+0.011}_{-0.008} & 0.187^{+0.050}_{-0.012} & 0.261^{+0.011}_{-0.051} \\ 0.187^{+0.050}_{-0.012} & 0.434^{+0.016}_{-0.052} & 0.379^{+0.008}_{-0.005} \\ 0.261^{+0.011}_{-0.051} & 0.379^{+0.008}_{-0.005} & 0.360^{+0.044}_{-0.013} \end{pmatrix} \begin{pmatrix} f_e \\ f_\mu \\ f_\tau \end{pmatrix}_S , \tag{3.13}$$

where 68% credible intervals are shown that reflect experimental uncertainties in the oscillation parameters. As a consequence of the near tribimaximal form of the PMNS matrix

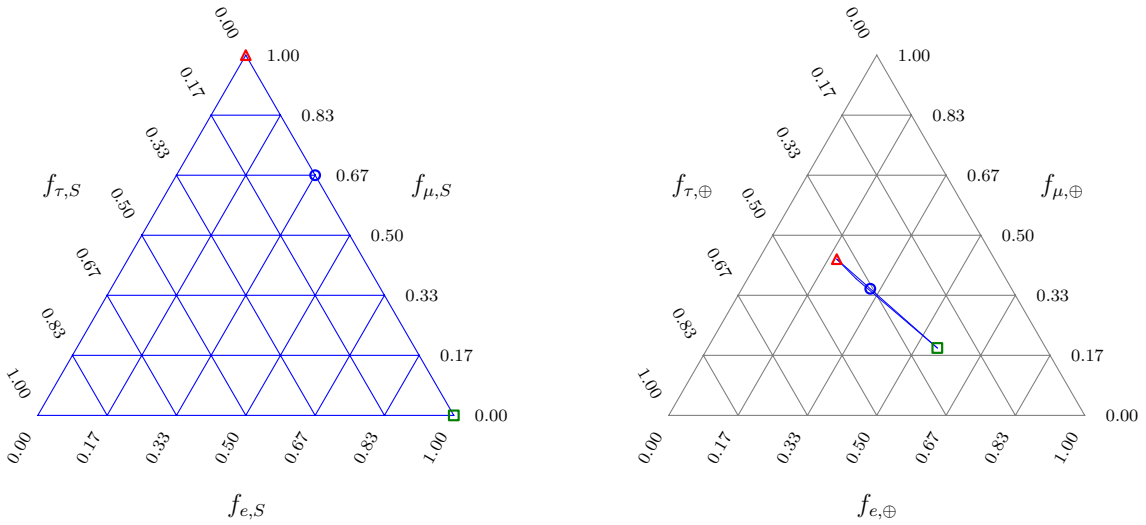

(a) Flavor Composition at Sources

(b) Flavor Composition at Earth

Figure 3.5: The ternary plot on the left illustrates the space of all flavor compositions at astrophysical neutrino sources. The ν_e fraction can be read by following lines up to the right, the ν_μ fraction horizontally to the left, and ν_τ fraction downwards to the right. A pure ν_e composition is at the bottom left corner, pure ν_μ composition at the bottom right corner, and pure ν_τ composition at the top corner. The markers indicate the expected flavor composition from muon-damped pion decay (red triangle), complete pion decay (blue circle), and neutron decay (green square) scenarios for astrophysical neutrino production. On the right, the ternary plot illustrates the space of flavor compositions measured at Earth. Due to neutrino oscillations as described through Eq. 3.12, the space of flavor compositions at astrophysical sources is mapped to an extremely narrow triangle shown in blue. The best-fit neutrino oscillation parameters from [30], which assume the inverted mass hierarchy, are used. Uncertainties on the oscillation parameters are not included here but do moderately broaden the allowed range of flavor compositions [91].

where $\theta_{23} \approx \pi/4$ and $\theta_{13} \approx 0$, it follows that regardless of the initial source composition, the flavor composition at Earth can only occupy a narrow range near $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$. Borrowing an approach from metallurgy, it is convenient to visualize flavor composition symmetrically on a ternary plot. Figure 3.5 shows ternary plots for the flavor composition at sources and at Earth, illustrating the transformation of flavor composition at sources to the narrow range of possibilities at Earth predicted by Eq. 3.12.

Depending on the particles decaying to neutrinos and the environment where the decays take place, there are several possibilities for the flavor compositions at sources and the corresponding flavor composition they produce at Earth:

- **Complete pion decay** $\left(\frac{1}{3} : \frac{2}{3} : 0\right)_S \rightarrow (0.31 : 0.35 : 0.34)_\oplus$
  This is the most commonly considered scenario in which both the muon neutrino from pion decay and the muon and electron neutrino from subsequent muon decay are observed, yielding a source flavor ratio of $(1 : 2 : 0)_S$. It is expected when there is little matter, radiation, or magnetic field to cause pions or muons to lose energy before decaying.

- **Muon-damped pion decay** $(0 : 1 : 0)_S \rightarrow (0.19 : 0.43 : 0.38)_\oplus$
  In this case, the muon neutrino from pion decay is observed, but the electron and muon neutrino from subsequent muon decay are not. This can occur if there is a high matter density in the source causing the muon to lose energy through bremsstrahlung and pair production, or if a high magnetic field causes the muon to lose energy through synchrotron radiation [92].

- **Muon decay** $\left(\frac{1}{2} : \frac{1}{2} : 0\right)_S \rightarrow (0.37 : 0.31 : 0.32)_\oplus$
  If there is a high accelerating gradient within the source, muons may be accelerated to high energies, and all high-energy neutrinos observed result from the muon and electron neutrino produced in muon decay [93]. The same source flavor composition could also be produced by the decay of other particles with similar branching fractions, for example kaons or charm hadrons, if pions lose energy in magnetic fields or matter before decaying [94].

- **Neutron decay** $(1 : 0 : 0)_S \rightarrow (0.55 : 0.19 : 0.26)_\oplus$
  If magnetic fields are so high that charged pions do not decay before losing energy through synchrotron radiation, then electron neutrinos from neutron decay can be the primary source of high-energy neutrinos [95]. Neutron decay can also be a large contributor to neutrino production if cosmic rays are primarily heavy nuclei that undergo photo-disintegration into individual nucleons in a radiation field [96].

These scenarios are shown in Fig. 3.5. In reality, the flavor composition at sources could be a superposition of all scenarios, either in many different sources or within a single source. Thus, it is more appropriate to parameterize the source flavor composition as $(1-x : x : 0)_S$, where x is the fraction of muon neutrinos produced in the source. Tau neutrinos could also be produced in sources, but since they result from the rarer decays of charm or bottom mesons, it is not expected they will be produced in large quantities. Even so, it would be extremely difficult to identify the existence of tau neutrinos produced in sources due to the extreme narrowness of the triangle of allowed flavor compositions produced at Earth shown in Fig. 3.5. Lastly, the flavor composition of astrophysical neutrinos need not be energy independent, and a measurement of flavor composition as a function of energy could provide unique information about the magnitude of magnetic fields or matter density in the object [97, 98].

3.2.3 Neutrino to Antineutrino Ratio

In the description of flavor composition so far, no distinction has been made between neutrinos and antineutrinos. Experimentally, distinguishing the two is a challenging task since as shown in Sec. 2.3 neutrino and antineutrino differential DIS sections become nearly identical at energies above 100 TeV where astrophysical neutrinos begin to dominate over atmospheric neutrinos. Alternatively, the Glashow resonance only occurs for electron antineutrinos, and it may be possible to use the relative rate of Glashow and DIS events to infer the fraction of electron antineutrinos in the astrophysical flux.

Measurement of the antineutrino composition of the astrophysical flux would also provide useful insight into its origin. Specifically, it may provide a means to determine if the flux originates from $p\gamma$ or pp interactions [99, 100, 101]. For example, typically only π^+ are produced in $p\gamma$ interactions near the Δ^+ resonance, whereas pp interactions produce both π^+ and π^-. As a result, different flavor compositions are produced at the Earth for neutrinos and antineutrinos. For example, the complete decay of π^+ will produce one ν_μ, one ν_e, and one $\bar{\nu}_\mu$, and the source flavor compositions are $(1:0:0)^{\nu}_{S}$ for neutrinos and $\left(\frac{1}{2}:\frac{1}{2}:0\right)^{\bar{\nu}}_{S}$ for antineutrinos. This will then produce flavor compositions at Earth of $(0.19:0.43:0.38)^{\nu}_{\oplus}$ for neutrinos and $(0.37:0.31:0.32)^{\bar{\nu}}_{\oplus}$ for antineutrinos. The resulting enhancement of the electron antineutrino rate compared to the equal neutrino and antineutrino rate expected from pp interactions may prove to be detectable via the Glashow resonance. So far, the goal of measuring antineutrino composition has remained out of reach since no Glashow resonance interactions have yet been conclusively identified. For the remainder of this work, an equal neutrino and antineutrino flavor composition for the astrophysical flux will be assumed unless otherwise noted, which has only a minor impact on the results obtained.

3.2.4 Physics Beyond the Standard Model

Although the emphasis so far has been placed on the connection of astrophysical neutrinos to cosmic rays, high-energy astrophysical neutrinos need not originate from cosmic-ray interactions alone and are powerful probes of fundamental physics on their own. Only 5% of all matter and energy in the universe is thought to be in the form of ordinary baryonic matter, and the remaining 25% and 70% are in the mysterious forms of dark matter and dark energy, respectively [102, 103]. The nature of dark energy is highly speculative, but there are many reasons to suspect dark matter may consist of a new weakly interacting massive particle (WIMP) that may self-annihilate or decay, either producing neutrinos directly or through the decays of intermediate Standard Model particles [104]. Astrophysical neutrinos are also an indirect probe of dark matter in the universe, and many searches have been undertaken at IceCube to search for neutrinos from expected regions of high dark matter concentration, such as the Sun [105], Earth's center [106], the galactic center [107], and nearby dwarf galaxies [108]. Evidence for a dark matter origin of astrophysical neutrinos would primarily derive from angular correlation to these objects, but further evidence would include a non-power-law energy spectrum connected to the mass of the dark matter parti-

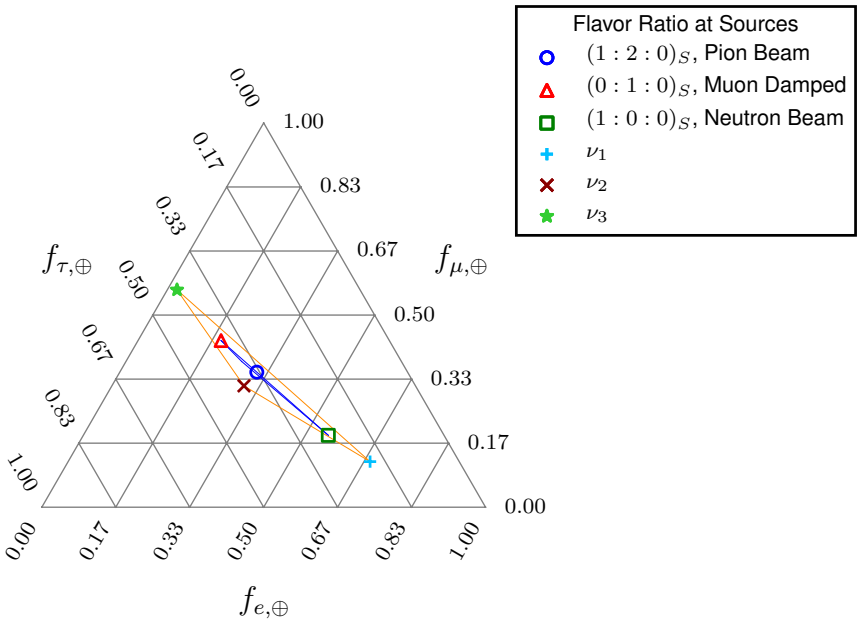

Figure 3.6: The range of flavor compositions allowed if the neutrino flux consists of any incoherent sum of mass eigenstates is shown by the orange triangle. A pure eigenstate, ν_1, might be produced by a total neutrino decay to this lightest mass eigenstate in the normal hierarchy. Likewise a pure eigenstate, ν_3, could be produced in the inverted hierarchy. Many other scenarios fall into this allowed range such as pseudo-Dirac neutrinos and quantum-gravity induced decoherence. Other more extreme phenomena like Lorentz, CPT, and equivalence principle violation could fall outside this range.

cle and a potentially unique flavor composition that could not be produced by cosmic-ray interactions [109].

Even if all astrophysical neutrinos originate from cosmic-ray interactions, their propagation over astronomical distances and enormous energy provide a unique testing ground for fundamental particle physics. Because the range of flavor compositions allowed by standard neutrino oscillations is limited, observing a flavor composition outside this range would be a striking signature of physics beyond the Standard Model. Many examples of non-standard scenarios that can produce such a signature have been proposed over the past decade [110] and are summarized in a recent review [91]. These non-standard scenarios fall into two classes.

In the first class, neutrinos arrive at Earth as an incoherent sum of mass eigenstates, which produce a flavor composition at Earth that can be written in the form $f_{\alpha,\oplus} = \sum_i k_i |U_{\alpha i}|^2$, where k_i represent the fraction of the flux in each mass eigenstate i. Neutrino decay is one mechanism that could achieve this state [111, 112]. For example, if only the lightest neutrino is stable and the neutrino lifetime is sufficiently small that only the lightest neutrino survives

to reach Earth, then $k_i = (1:0:0)$ or $k_i = (0:0:1)$, corresponding to the normal or inverted hierarchy, respectively. Many other scenarios can be represented as an incoherent sum of mass eigenstates, for example pseudo-Dirac neutrinos [113] and quantum-gravity induced decoherence of oscillations [114, 115]. Oscillations to additional sterile neutrino states is another example, but a recent analysis found no large deviations in the flavor composition should be observed within current experimental bounds on sterile neutrino mixing parameters [116]. The range of flavor compositions allowed as any incoherent mixture of mass eigenstates is shown in Fig. 3.6.

In the second class, neutrinos arrive at Earth in a completely general state. This class includes violation of Lorentz symmetry, CPT symmetry, and the equivalence principle [117, 118, 119]. A generalized study of these effects using an effective Hamiltonian approach has shown that it is possible to produce nearly any flavor composition at Earth given current limits from other experiments [120]. With such a rich variety of scenarios producing non-standard flavor compositions, a precision measurement has great potential to find evidence of physics beyond the Standard Model inaccessible to any other experiment.

3.3 Atmospheric Neutrinos

Just as in astrophysical sources, high-energy neutrinos are also produced from the decays of unstable particles produced by cosmic-ray interactions within Earth's atmosphere. Fortunately, the atmospheric neutrino flux differs from the astrophysical flux in several ways. First, cosmic rays are expected to have a softer power-law index after arriving at Earth, and the resulting neutrinos should have a softer power-law index than those produced in astrophysical sources. Secondly, due to time dilation, unstable particles can interact with air before they have a chance to decay to neutrinos if they have a sufficiently high energy, which further acts to soften the power-law index.

To characterize the energy scale where interactions become important, a critical energy, ϵ, for each particle can be defined as the energy where the mean length a particle travels before decaying equals its mean free path length between interactions. In units of column depth ($\mathrm{g/cm^2}$), the decay length for an ultra-relativistic particle is

$$d = \frac{\epsilon \tau \rho}{mc}, \tag{3.14}$$

where m is the mass of the particle, τ is its mean proper lifetime, ϵ is its energy, and ρ is the local density of the atmosphere. The mean free path, or interaction length, is $\lambda = \frac{M}{\sigma}$ where M is the mean molecular mass, and σ is the total inelastic cross section with air for the particle. Because density varies with altitude, the critical energy is usually defined at the altitude where the column depth from the top of the atmosphere is equal to one interaction length. For now, it is useful to approximate the atmosphere as isothermal with an exponential density profile $\rho = \rho_0 \exp(-h/h_0)$. At high altitudes, the scale height is $h_0 \approx 6.4\,\mathrm{km}$. For a particle entering from the top of the atmosphere at zenith angle θ from

Particle	m (GeV)	τ (s)	ϵ (GeV)	Final states	$\mathcal{B}$ (%)
$\pi^\pm$	0.140	2.6×10^{-8}	115	$\mu\nu_\mu$, $e\nu_e$	99.9877, 0.012
$K^\pm$	0.494	1.2×10^{-8}	850	$\mu\nu_\mu$, $e\nu_e$, $\pi^0\mu\nu_\mu$, $\pi^0 e\nu_e$	63.5, 0.0015, 3.4, 5.0
K^0_L	0.498	5.1×10^{-8}	205	$\pi^\pm\mu\nu_\mu$, $\pi^\pm e\nu_e$	27.0, 40.0
K^0_S	0.498	9.0×10^{-11}	1.2×10^5	$\pi^\pm e\nu_e$	0.07
$D^\pm$	1.87	1.0×10^{-12}	4.3×10^7	$K^0\mu\nu_\mu$, $K^0 e\nu_e$	9.2, 8.8
D^0	1.86	4.1×10^{-13}	9.2×10^7	$K^\pm\mu\nu_\mu$, $K^\pm e\nu_e$	3.5, 3.3
$D^\pm_s$	1.97	5×10^{-13}	7.5×10^7	$\tau\nu_\tau$, $\mu\nu_\mu$, $Xe\nu_e$	5.5, 0.0055, 6.5

Table 3.1: Hadron decays of interest for high-energy atmospheric neutrino fluxes. Critical energies are calculated assuming a scale height of $h_0 = 6.4\,\mathrm{km}$.

the vertical, the density at a column depth λ is $\rho = \lambda \cos\theta / h_0$ under the approximation the Earth has no curvature. The critical energy is then

$$\epsilon = \frac{Mmc}{\rho\sigma\tau} = \frac{mch_0}{\tau\cos\theta}. \tag{3.15}$$

The critical energy is usually quoted for vertical particles with $\cos\theta = 1$, but it is worth remembering the critical energy is higher for more inclined particles.

Numerical values of the critical energy for several hadrons are shown in Table 3.1. For energies above 1 TeV of interest here, pions and kaons undergo significant interactions in the atmosphere, and the power-law index of the neutrino flux from these particles is softened from 2.7 to 3.7 above the critical energy. For charm hadrons, the critical energy is beyond 10 PeV, and the resulting neutrino flux follows the same power-law index of 2.7 as the cosmic-ray flux for the energies of interest here. For these reasons, it is typical to split the atmospheric neutrino flux into two components, the "conventional" flux from pion and kaon decay, and the "prompt" flux from charm and other heavy hadron decays.

3.3.1 Conventional Atmospheric Neutrinos

The conventional atmospheric neutrino flux is largely dominated by muon neutrinos from the two-body decays, $\pi^\pm \to \mu\nu_\mu$ and $K^\pm \to \mu\nu_\mu$. The reason almost exclusively ν_μ are produced in these two-body decays is another consequence of chirality in the weak interaction. Both spin-0, π^- and K^- decay to a right-handed antineutrino and a right-handed charged lepton in order to preserve angular momentum. If the charged lepton was massless, the decay would be impossible since only left-handed leptons can couple to the weak interaction. When incorporating lepton masses, the same chirality considerations act to suppress the electron branching fraction, $\mathcal{B}$, by a factor

$$\frac{\mathcal{B}(\pi^- \to e^-\bar{\nu}_e)}{\mathcal{B}(\pi^- \to \mu^-\bar{\nu}_\mu)} \approx \frac{m_e^2}{m_\mu^2}\left(\frac{m_\pi^2 - m_e^2}{m_\pi^2 - m_\mu^2}\right)^2 = 1.2 \times 10^{-4} \tag{3.16}$$

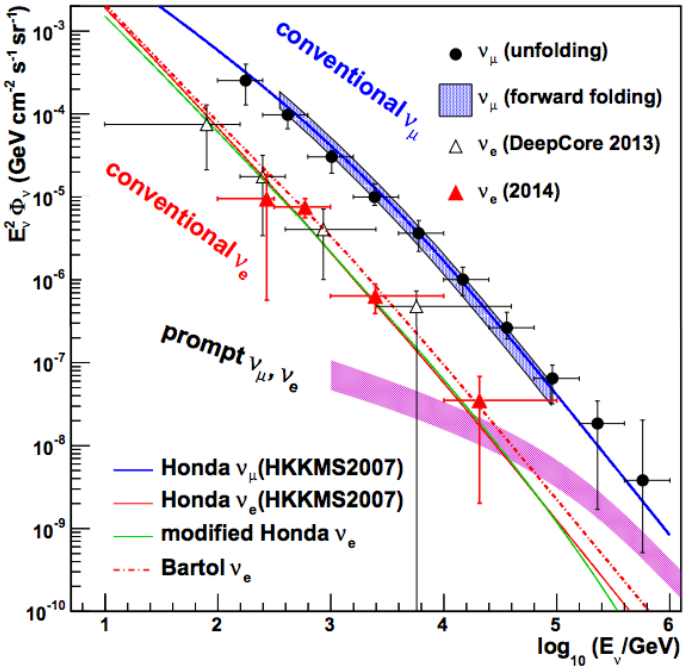

Figure 3.7: Measurements of the conventional atmospheric neutrino flux from IceCube reproduced from [121] compared to the HKKMS [122] and Bartol [123] calculations. Measurements of the atmospheric ν_μ flux are shown in black [124] and two measurements of the ν_e flux are shown with white and red triangles [125, 121]. The ERS calculation of the prompt atmospheric neutrino flux is shown as a pink band, which illustrates the estimated theoretical uncertainties [126].

due to the electron's much smaller mass. However, all kaons can also undergo three-body decays involving pions, which do not have the same suppression. As a result, electron neutrinos from kaons quickly overcome the smaller contribution from pions at high energies. Because of its high critical energy and small branching ratio, the contribution to the electron neutrino flux from K_S^0 has been neglected before, but it was recently pointed out that it may lead to a 30% enhancement of the vertical conventional flux beyond 100 TeV [127]. Muon decays are also a source of atmospheric neutrinos, but with a decay length of $d = 6300(E/1\,\text{TeV})\,\text{km}$, they do not decay before reaching the detector and are not a significant source for the energies considered here.

To calculate atmospheric neutrino fluxes, a common approach is to solve the so-called cascade equations, a set of first-order differential equations describing the evolution of particle fluxes through the atmosphere [66]. However, the cascade equations assume all particles remain collinear after an interaction and neglect the effect of Earth's magnetic field, which are both important at lower energies. The most accurate calculations of atmospheric neutrino fluxes are carried out by Monte Carlo simulation. Individual cosmic rays are tracked through

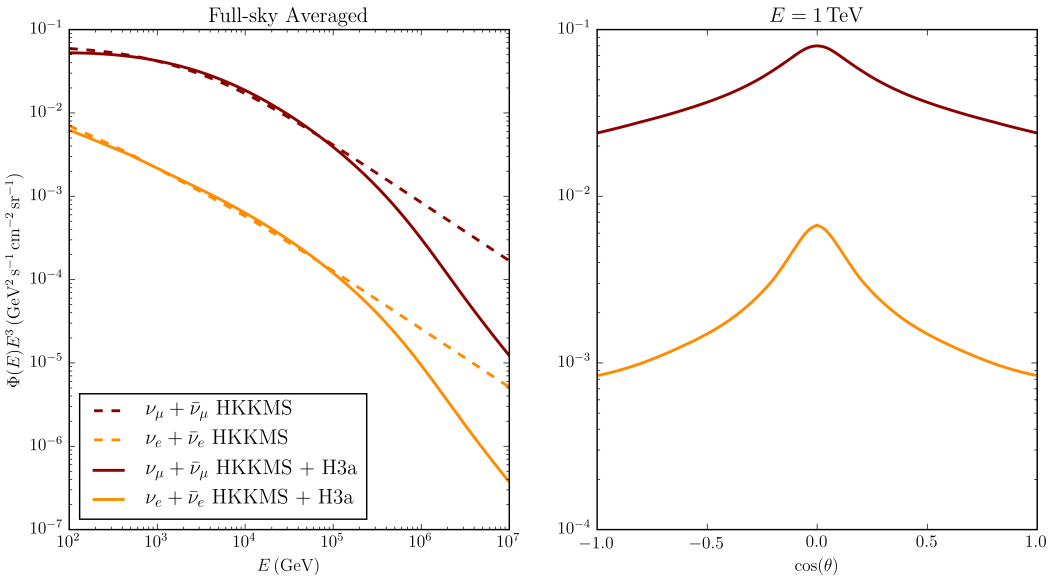

Figure 3.8: Left: The energy spectrum of conventional atmospheric neutrinos averaged over all directions. The dotted line shows the extrapolated HKKMS prediction with no correction for the cosmic-ray knee. The solid line shows the correction for the H3a cosmic-ray model, which significantly steepens the energy spectrum above 100 TeV. Right: The zenith angle dependence of the conventional atmospheric flux at a neutrino energy of 1 TeV.

the atmosphere and their interactions producing secondary particles are simulated using a hadronic model. Secondaries are also tracked through the atmosphere until they decay or undergo further interactions. Throughout this work, the HKKMS Monte Carlo calculation will be the standard [122]. It uses the US Standard Atmosphere 1976 [128] to model the air density profile and a modified version of the DPMJET-III hadronic interaction model [129] tuned to fit measurements of the atmospheric muon flux. As shown in Fig. 3.7, current measurements of both the atmospheric ν_μ flux and ν_e flux from IceCube agree well with the calculation. However, the calculation only extends up to 10 TeV, and it is necessary to extrapolate to higher energies.

In the energy range of interest here, the conventional atmospheric neutrino flux can be parameterized to high fidelity using a simple analytic formula as a sum over the hadron decays contributing to the flux

$$\Phi(E) = \sum_i \frac{A_{i\nu}}{1 + B_{i\nu} E \cos(\theta^*)/\epsilon_i} \left(\frac{E}{1\,\mathrm{TeV}}\right)^{-\gamma}, \tag{3.17}$$

where $\gamma = 2.7$ is the cosmic-ray spectral index, ϵ_i is the (vertical) critical energy for each hadron, and $A_{i\nu}$ and $B_{i\nu}$ are parameters determined from a fit to the HKKMS flux calculation from 700 GeV to 10 TeV separately for ν_e, $\bar{\nu}_e$, ν_μ, and $\bar{\nu}_\mu$. The formula maintains the property

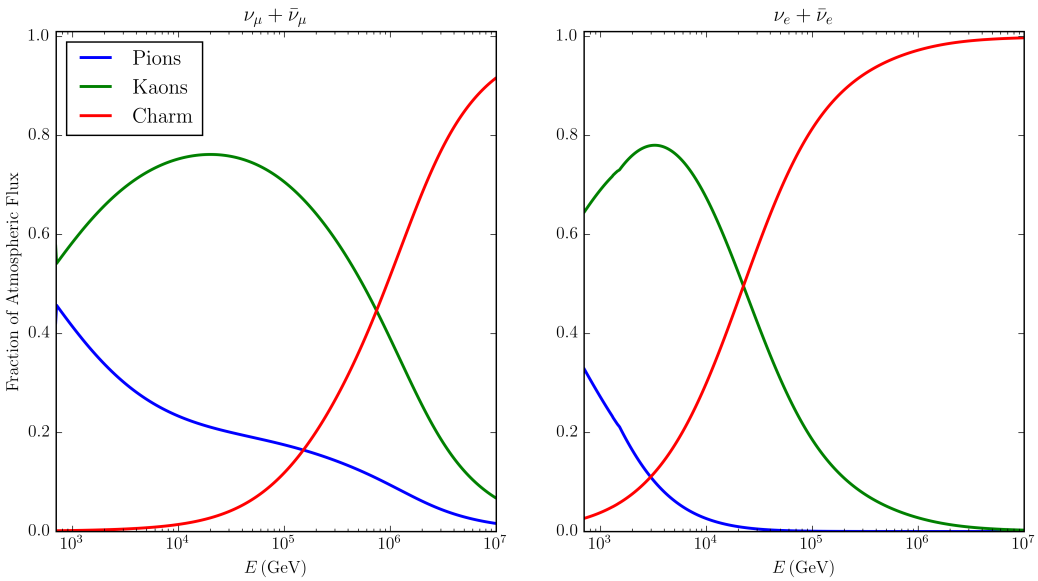

Figure 3.9: The fraction of the atmospheric neutrino flux originating from the decay of pions (blue), kaons (green), and charm hadrons (red) as predicted from the HKKMS and ERS calculations.

that the neutrino flux from each hadron decay steepens from a power-law index γ to $\gamma + 1$ above the critical energy. Likewise, the flux remains isotropic below the critical energy, but acquires a zenith-angle dependence above it. The zenith-angle dependence of the critical energy is described through another parameterization that describes the effects of the Earth's curvature on the column depth through the atmosphere,

$$\cos(\theta^*) = \sqrt{\frac{\cos^2(\theta) + p_0^2 + p_1 \cos^{p_2}(\theta) + p_3 \cos^{p_4}(\theta)}{1 + p_0^2 + p_1 + p_3}}, \tag{3.18}$$

where $p_i = \{0.1026, -0.0683, 0.9586, 0.0407, 0.8173\}$ [130]. Another shortcoming of the HKKMS calculation is that it models the cosmic-ray flux as a perfect power-law that does not account for the changing energy spectrum and composition around the knee. Fortunately, a simple energy-dependent correction factor can be applied to Eq. 3.17 to represent the H3a cosmic-ray model as described in [131]. The atmospheric neutrino fluxes from this parameterization are shown in Fig. 3.8. The parameterization also enables the flux to be broken into contributions from pions and kaons. Figure 3.9 illustrates how kaons tend to dominate the conventional flux at higher energies. The relative contribution of pions and kaons is a considerable systematic uncertainty that will be covered in Sec. 7.3.

Unlike for astrophysical neutrinos, there are better prospects for measuring the neutrino and antineutrino contributions to the atmospheric flux since below the ~ 10 TeV energy scale set by the weak boson propagator, the inelasticity distributions for neutrino and antineutrinos

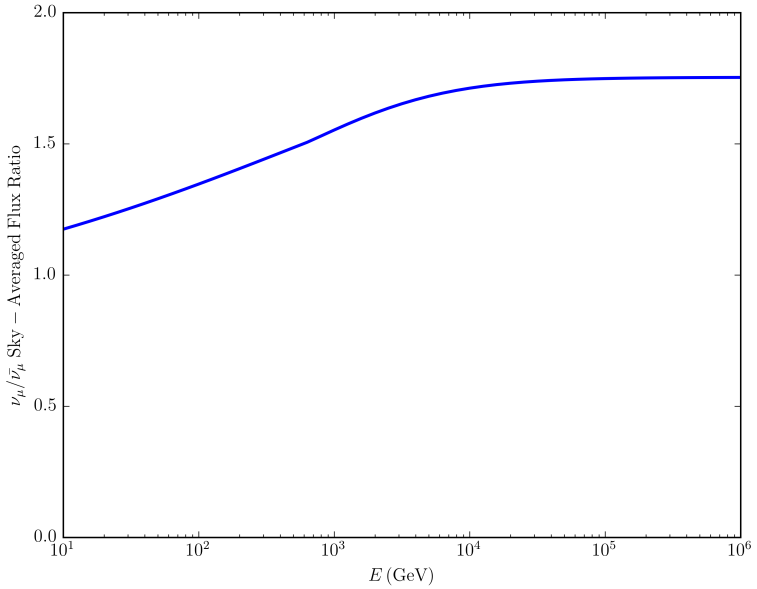

Figure 3.10: The ratio of the full sky-averaged conventional atmospheric ν_μ flux and $\bar{\nu}_\mu$ flux from the HKKMS calculation. The ratio tends to approach a constant at high energies as kaon decays become the dominant source of neutrinos and antineutrinos.

show significant differences. The ratio of atmospheric neutrino to antineutrino fluxes is of interest since it is sensitive to the asymmetry in relative number of positive and negative pions or kaons. This asymmetry is caused by both the excess of protons over neutrons in the cosmic ray flux and the unequal production of $\pi^\pm$ or $K^\pm$ in collisions with air nuclei. Along with the analogous muon charge ratio for the atmospheric muon flux, which has been measured extensively in many experiments up to 10 TeV [132, 133, 134], the neutrino to antineutrino ratio is a sensitive diagnostic for models of cosmic-ray interactions in the atmosphere [72]. The ν_μ to $\bar{\nu}_\mu$ ratio predicted by HKKMS is shown in Fig. 3.10.

3.3.2 Prompt Atmospheric Neutrinos

Despite the lower production cross section for heavy charm hadrons, their high critical energy implies they must produce a nearly isotropic flux with power-law index ~ 2.7 that will eventually overtake the contribution from pions and kaons, as shown in Fig. 3.9. For electron neutrinos, the charm contribution begins to dominate at only ~ 30 TeV. Furthermore, the charm contribution also begins to dominate the atmospheric flux in the same energy range where the astrophysical flux appears. This is especially worrisome since the astrophysical flux could share a similar power-law index of 2.7 and also be isotropic if it is of extra-galactic origin. Unlike the astrophysical flux, the prompt flux contains roughly equal contributions of ν_e and ν_μ from the semileptonic decays of $D^\pm$ and D^0, and only a negligibly small contribution

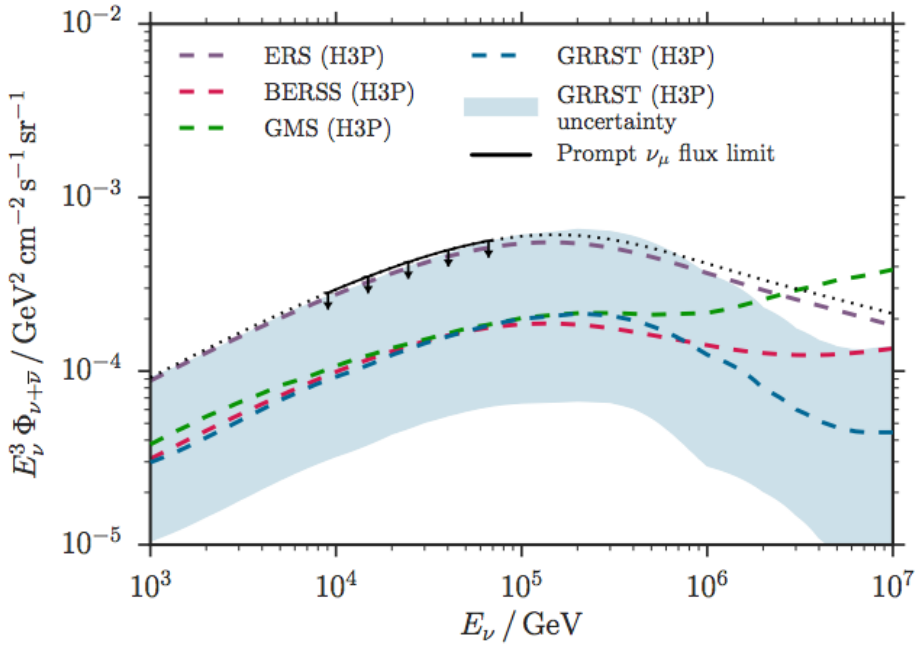

Figure 3.11: Predictions for the prompt atmospheric neutrino flux reproduced from [136]. The purple and red dashed lines show the ERS [126] and BERSS [137] calculations that are used in this work. The dashed blue line shows the GRRST calculation [138], which incorporates the latest data from LHCb. The theoretical uncertainty for the GRRST calculation is marked by the shaded blue region. All calculations here assume the H3a cosmic ray model. The current best limit on the prompt flux at 90% confidence level is represented by the solid black curve with down-going arrows [136].

of ν_τ from the two-body decay of D_s^+ [135]. Although this makes the observation of a ν_τ flux telltale evidence for astrophysical origin, the presence of the prompt flux is a significant confounding background when measuring the energy spectrum and flavor composition of the diffuse astrophysical flux.

In spite of its importance, the prompt neutrino flux is subject to large theoretical uncertainties, and it has never been experimentally observed. Calculations of the prompt flux are mainly limited by imprecise knowledge of the charm production cross section. At high energies, charm production results mainly from gluon fusion, $gg \to c\bar{c}$, and its cross section depends on the gluon parton distribution function, which has increasing uncertainties at the low Bjorken-x. Furthermore, direct measurements of charm production at particle accelerators do not cover the regions of kinematic phase space that contribute most to the prompt flux. As a result, predictions for the prompt flux have varied widely over the years, but have generally trended downwards as accelerator data and calculation techniques have improved. Two calculations will be used here, ERS [126] and BERSS [137], an update by the same authors. The GRRST calculation is the most recent and incorporates the latest data

from LHCb [138]. These calculations for the prompt flux are shown in Fig 3.11. Despite the large uncertainties of about a factor of 3 in the normalization of the prompt flux from these calculations, it is important to note that the shape of the energy spectrum remains very consistent for the energies of interest. Currently, experimental upper limits on the prompt flux are comparable to these calculations. For example, two recent analyses of independent IceCube datasets found 90% confidence level upper limits of 1.52 [139] and 1.06 [136] times the ERS calculation.

3.3.3 Atmospheric Self-Veto Probability

Many features of the energy spectrum, zenith angle dependence, and flavor composition of the atmospheric neutrino flux have been identified that enable it to be discriminated from the astrophysical flux. One additional feature has proven to be very powerful in disentangling atmospheric and astrophysical neutrinos. Atmospheric neutrinos are produced along with muons in cosmic-ray air showers, and if the muons have sufficient energy to penetrate the overburden of an underground neutrino telescope, they may also be detected at the same time as a neutrino. The presence of accompanying muons can then be used to veto down-going neutrinos as being of atmospheric origin.

The probability that detectable muons accompany an atmospheric neutrino is called the atmospheric self-veto probability. It depends on the flavor, energy, and direction of a neutrino. It also depends on the properties of the muons that are detectable, and hence the detailed characteristics of the neutrino telescope. A simplified condition is that a muon with an energy above a given threshold is produced in the same cosmic-ray air shower as the neutrino. The self-veto probability can then be thought of as an effective suppression of the atmospheric neutrino flux that becomes stronger for higher energies and more vertically down-going events. High-energy neutrinos originate from high-energy cosmic-ray air showers that produce high-energy muons likely to reach the detector. Likewise, muons in vertical down-going air showers are more likely to penetrate the overburden to reach the detector.

In general, the self-veto probability is best calculated from Monte Carlo simulations of cosmic-ray air showers, which will be detailed more in Sec. 5.1.3. However, analytic approaches have been developed that capture the necessary features at much lower computational expense. The first calculation of self-veto probability was performed by [140], but only considered muons that arise from the same decay as the detected neutrino and only considered pion and kaon decays. The calculation that is used here updates this result to incorporate muons from other decays in the air shower and included charm hadron decays as well [141]. Figure 3.12 shows the effective suppression of the atmospheric flux as a function of energy and zenith angle. The suppression of down-going atmospheric neutrinos is particularly powerful in discriminating the prompt flux from an astrophysical flux since both would otherwise be isotropic.

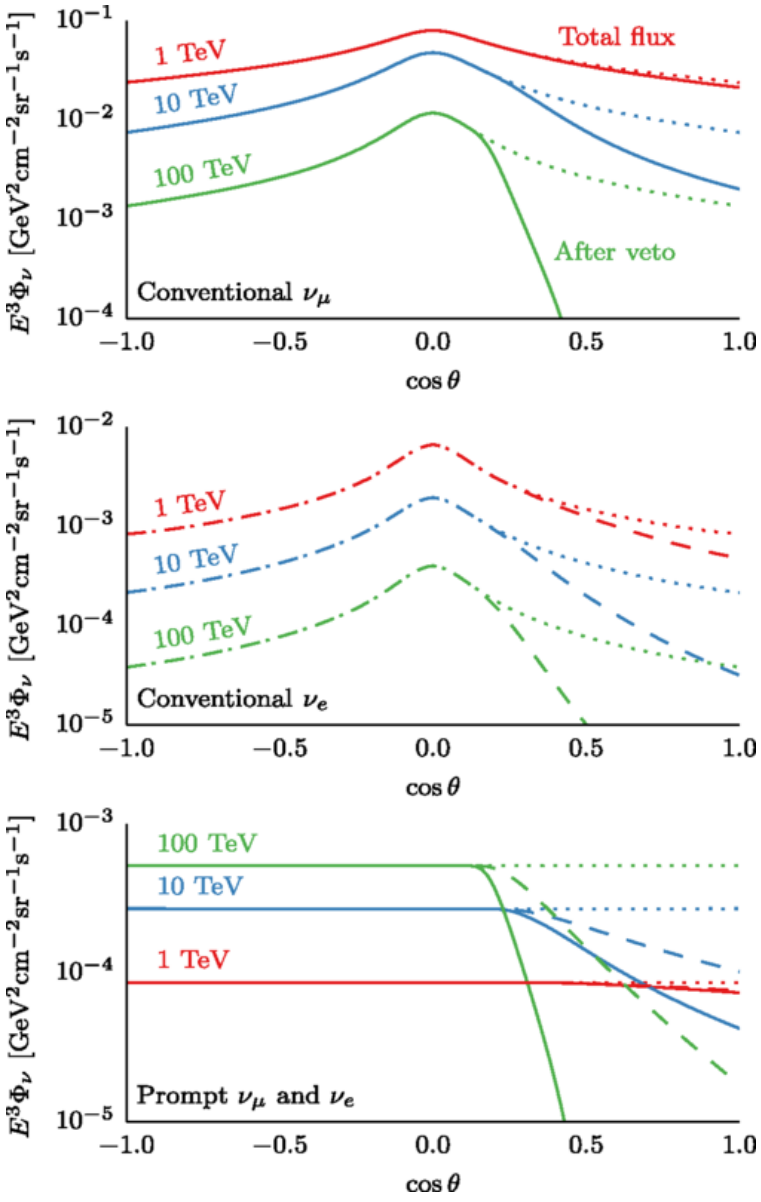

Figure 3.12: The effective conventional and prompt neutrino fluxes after vetoing neutrinos accompanied by a muon with an energy of more than 1 TeV at a depth corresponding to IceCube [141]. The flux with no self-veto effect is shown with dotted lines. An increasingly large fraction of atmospheric neutrinos are vetoed at large energies and small zenith angles.

Chapter 4

Detecting Neutrinos in IceCube

The IceCube Neutrino Observatory consists of an array of optical sensors deployed within a cubic kilometer of clear glacial ice at the South Pole. Why was it necessary to build such a large detector in such an exotic location? To answer this question, consider the problem of detecting a neutrino flux at the Waxman-Bahcall level $\Phi(E_\nu) \sim 10^{-8} E_\nu^{-2}\,\mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}$. Though the flux drops with energy, this is partially counteracted by the fact that the neutrino-nucleon cross section rises with energy. At energies beyond 10 TeV, the total cross section can be parameterized by a power law, $\sigma(E_\nu) \approx 7.84 \times 10^{-36}\,(E_\nu/1\,\mathrm{GeV})^{0.363}\ \mathrm{cm^2}$ [52]. The interaction rate per unit mass for neutrinos with energy larger than E is

$$\dot{n} = \frac{4\pi}{M}\int_E^\infty \sigma(E_\nu)\Phi(E_\nu)dE_\nu \sim 60\left(\frac{E}{100\,\mathrm{TeV}}\right)^{-0.637}\ \mathrm{gigaton^{-1}\,year^{-1}}, \tag{4.1}$$

where M is the mass per nucleon. A gigaton of material is needed to collect dozens of events per year above 100 TeV where atmospheric neutrino fluxes no longer dominate. Amassing such a large amount of material is hopelessly expensive, so one must turn to a naturally-occurring medium. As seen in Ch. 2, a wide variety of secondary particles are created by neutrino interactions. Densely instrumenting such a large volume to identify each particle is not feasible, but in an optically transparent medium, Cherenkov light radiated by charged particles can be detected at large distances with only a sparse array of sensors. The ice sheet at the South Pole has nearly ideal properties for this task. In this chapter, the entire process of neutrino detection will be described from the secondary particles produced in neutrino interactions to the propagation of their Cherenkov light in ice to the recording of this light with IceCube's optical sensors.

4.1 Secondary Particles in Ice

4.1.1 Cherenkov Light

As first observed by Pavel Cherenkov in 1934 [142], charged particles emit light when they exceed the phase velocity of light within the medium they traverse. Analogous to how a

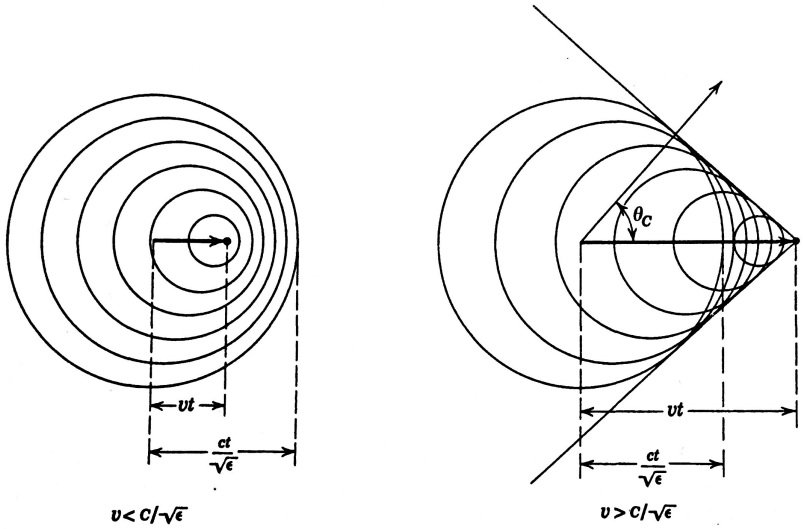

Figure 4.1: Illustration of Cherenkov radiation using expanding wavefronts taken from [145] where $n = \sqrt{\epsilon}$ is the index of refraction.

shock wave forms from a supersonic jet, Cherenkov light is emitted along a cone at an angle $\theta_C = \cos^{-1}(1/\beta n)$ from the direction of travel as shown in Fig. 4.1, where n is the index of refraction of the medium and $\beta = v/c$ is the ratio of the particle's speed to the vacuum speed of light. For ice with $n \approx 1.31$, the Cherenkov angle of an ultra-relativistic particle is $\theta_C \approx 41°$. The number of photons emitted per unit length has a wavelength spectrum given by the Frank-Tamm formula [143],

$$\frac{dN}{dx d\lambda} = \frac{2\pi\alpha z^2}{\lambda^2}\left(1 - \frac{1}{\beta^2 n(\lambda)^2}\right), \tag{4.2}$$

where α is the fine-structure constant, z is the charge of the particle in units of elementary charge, and the index of refraction, $n(\lambda)$, is a function wavelength. The spectrum rises at short wavelengths, explaining the familiar blue glow associated with the Cherenkov light in nuclear reactors. The Frank-Tamm formula is restricted to wavelengths where $\beta n(\lambda) > 1$, which eventually is violated in the near-ultraviolet range for many common materials. Light sensors in the visible wavelength range are well-suited to collect as many Cherenkov photons as possible. In ice, the index of refraction varies smoothly by a few percent in the optical range $\lambda \in [0.3\,\mu\text{m}, 0.6\,\mu\text{m}]$ according to a simple polynomial parameterization [144]. It also depends on the density and temperature of ice, which both vary with depth due to the weight of the ice sheet at the South Pole. However, this variation is at most 0.4% and the refraction of light in the visible range may be safely ignored [144].

4.1.2 Electromagnetic Cascades

As charged particles travel through ice, they lose energy through ionization and the radiative processes of bremsstrahlung, e^+e^- pair production, and photonuclear interactions, creating additional secondary particles. Ionization is the dominant energy loss process at low energies, but radiative losses are proportional to energy and eventually overtake ionization, as shown in Fig. 4.2a. A critical energy where radiative losses begin to dominate can be defined, and for electrons in ice, it is $E_c = 79\,\text{MeV}$ [29]. At multi-TeV energies of interest for IceCube, electrons exponentially lose energy until they drop below the critical energy. The length scale characterizing this is the radiation length, X_0, which is defined as the column depth in units g/cm^2 required for an electron to lose all but a fraction $1/e$ of its energy due to bremsstrahlung. For electrons in ice, $X_0 = 36\,\text{g/cm}^2$, corresponding to a length of 39 cm [29]. As an electron loses energy, the photons produced by bremsstrahlung go on to create additional e^+e^- pairs in the Coulomb field of surrounding nuclei, which can create further photons, electrons, and positrons. The result is a cascade, or shower, of particles that continues to grow until the energies of electrons and positrons fall below the critical energy.

Regardless of the initiating electromagnetic particle (e^-, e^+, or γ), electromagnetic cascades follow nearly the same universal behavior. The longitudinal development of an electromagnetic cascade can be approximately described by a gamma distribution,

$$\frac{dE}{dt} = E_0 b \frac{(bt)^{a-1} e^{-bt}}{\Gamma(a)}, \tag{4.3}$$

where t is the distance along the cascade in units of radiation length and E_0 is the initial energy of the injected photon, electron, or positron. Based on fits to Monte Carlo simulation of electromagnetic cascades in ice using Geant4 [146], a and b are empirical parameters determined to be $a = 2.03 + 0.604 \log_{10}\left(\frac{E_0}{1\,\text{GeV}}\right)$ and $b = 0.633$ [147]. The distance to shower maximum is a useful quantity characterizing the size of electromagnetic cascades, which can be found to be $X_{max} = X_0 \frac{a-1}{b} = 1.25\,\text{m} + 0.37 \log_{10}(E/1\,\text{GeV})\,\text{m}$. Cascades grow logarithmically in size with energy and have a characteristic size of several meters. This justifies a common approximation that light is emitted from a single point, an assumption that will later be used during simulation and event reconstruction.

At high energies, cascade size may be significantly enhanced by the Landau-Pomeranchuk-Migdal (LPM) effect [148, 149]. This refers to suppression of bremsstrahlung and pair production cross sections caused by destructive interference between multiple scattering centers in the medium. This becomes important above a characteristic energy scale $E_{\text{LPM}} = (m_e c^2)^2 \alpha X_0 / 4\pi\hbar c\rho \sim 300\,\text{TeV}$ in ice [150]. With suppressed cross sections, high energy electromagnetic particles can travel longer than the size of any sub-cascades their lower energy secondaries initiate. This results in a highly irregular shower profile that is both elongated and subject to large event-to-event fluctuations.

Though shower size grows slowly with energy, the total Cherenkov light yield from electromagnetic cascades is expected to scale linearly with energy, independent of the initiating

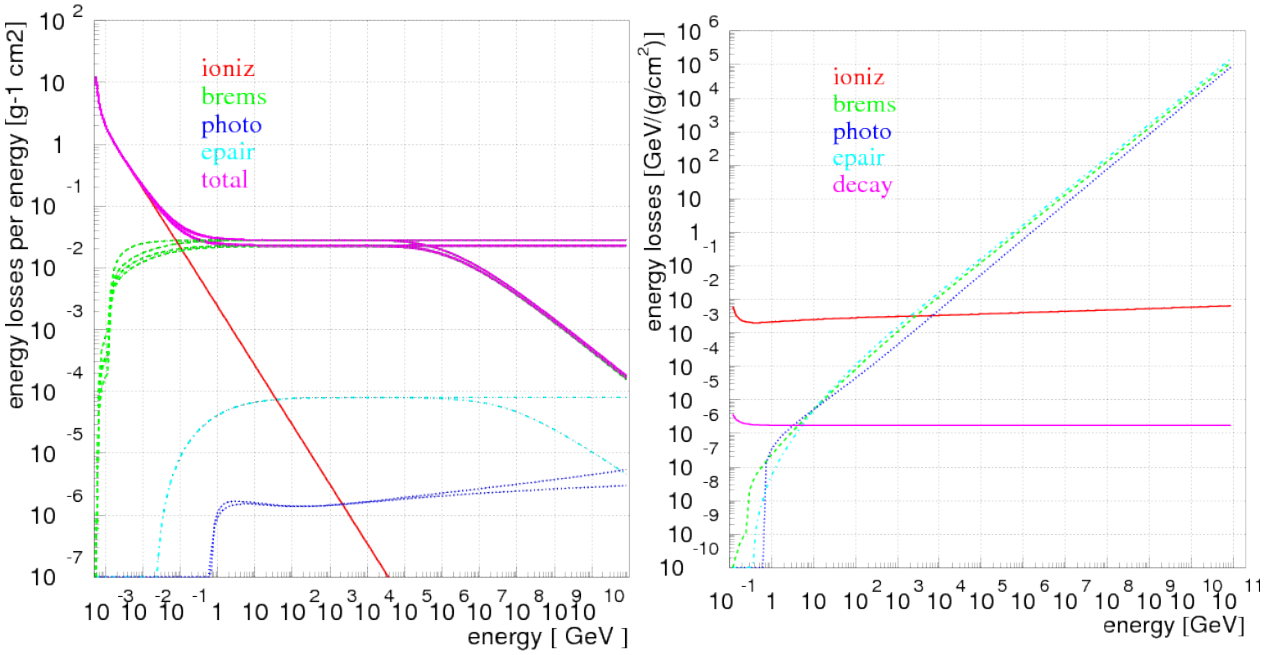

(a) (b)

Figure 4.2: Left: The average energy loss rate divided by energy, $-\frac{1}{E}\frac{dE}{dX}$, for electrons in ice showing the contributions from ionization, bremsstrahlung, photonuclear interactions, and pair production from [152]. Right: The average energy loss rate, $-\frac{dE}{dX}$, for muons in ice showing the contributions from ionization, bremsstrahlung, photonuclear interactions, pair production, and decay from [152].

particle type (e^+, e^-, or γ) with $< 0.1\%$ accuracy [151]. Additionally, the angular light emission profile remains tightly concentrated along the Cherenkov angle independent of energy, a result of the fact that most particles are closely beamed along the axis of the cascade.

4.1.3 Hadronic Cascades

Charged hadrons in ice experience the same electromagnetic energy loss processes as described above, but it is more important to consider their hadronic interactions with nuclei, characterized by a nuclear interaction length of 91 cm in ice [29]. Due to the complexity of final states that arise from hadronic interactions, hadronic cascades are subject to larger fluctuations in the overall development of a particle cascade. Furthermore, some energy goes to neutral particles, such as neutrons and neutrinos, and to overcoming the binding energy of nuclei. As a result, hadronic cascades do not produce as much Cherenkov light as electromagnetic cascades. However, a significant fraction of the cascade's total energy can be among electrons, positrons, and photons since these are produced by electromagnetic sub-cascades initiated by the decay of neutral pions, $\pi^0 \rightarrow \gamma\gamma$. A common approach is to parameterize the light yield of hadronic cascades and its fluctuations relative to electromag-

netic cascades. The parameters are determined from simulated hadronic cascades in ice with Geant4 and using realistic final hadronic states from neutral-current neutrino interactions generated from PYTHIA [153] as initiating particles [154].

The Cherenkov light yield of a hadronic shower relative to an electromagnetic shower initiated by a particle of the same energy can be parameterized as

$$F = F_{\mathrm{EM}} + (1 - F_{\mathrm{EM}}) f_0, \tag{4.4}$$

where F_{EM} is the fraction of the shower's energy in electrons, positrons, and photons and $f_0 = 0.467$ is a factor accounting for the diminished light yield from the remaining hadronic part of the shower. At high energies where a large multiplicity of neutral pions are produced in hadronic collisions and charged hadrons undergo larger bremsstrahlung and pair production losses, the electromagnetic fraction of the shower rises, which is parameterized by

$$F_{\mathrm{EM}} = 1 - \left(\frac{E}{E_0}\right)^{-m}, \tag{4.5}$$

where $E_0 = 0.399\,\mathrm{GeV}$ and $m = 0.13$. To describe fluctuations in F, a Gaussian approximation is used with a standard deviation varying with energy according to

$$\sigma_F = F\sigma_0 \ln E^{-\gamma}, \tag{4.6}$$

where $\sigma_0 = 0.379$ and $\gamma = 1.16$. The relative Cherenkov yield of hadronic cascades showing the size of these fluctuations is illustrated in Fig. 4.3.

Hadronic cascades can also produce long-range muons from the decay and pions and kaons just as in cosmic-ray air showers. Because hadrons are likely to interact before decaying, this is largely suppressed since the length traveled in one mean lifetime ($c\tau \approx 7.8\,\mathrm{m}$ for $\pi^{\pm}$ and $c\tau \approx 3.7\,\mathrm{m}$ for $K^{\pm}$) is larger than their interaction length in ice. Nevertheless, significant numbers of muons can be produced in large hadronic cascades, and the muon energy production spectrum is parameterized by a falling power law proportional to the energy of the hadronic cascade, E_0,

$$\frac{dN}{dE} = N_0 \left(\frac{E_0}{1\,\mathrm{TeV}}\right)\left(\frac{E}{1\,\mathrm{GeV}}\right)^{-\gamma}. \tag{4.7}$$

Parameter values $N_0 = 0.358$ and $\gamma = 2.74$ were found from the simulation of hadronic showers in water using the cosmic-ray Monte Carlo program CORSIKA [155] in [156]. However, newer studies using Geant4 in ice have found $N_0 \sim 0.5$ and $\gamma \sim 2.4$, indicating muon production may be higher than previously thought [157]. The exception to this picture arises from charm hadrons, which may promptly decay to produce muons in ice. Their production, interaction, and decay will be treated in more detail in Ch. 5.

4.1.4 Muons

Muons experience the same ionization and radiative losses as described in Sec. 4.1.2, but due to the muon's larger mass, radiative losses are suppressed, as shown in Fig. 4.2b. As a

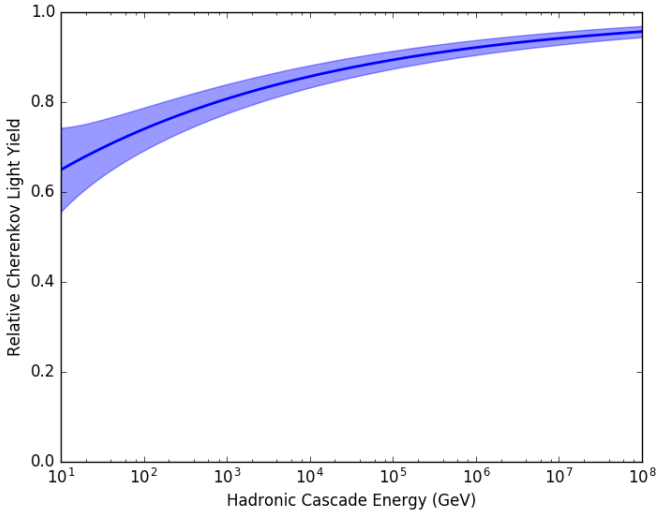

Figure 4.3: The fractional Cherenkov light yield for hadronic cascades relative to electromagnetic cascades as a function of energy. The colored band represents the 1σ fluctuations in light yield.

result, the muon critical energy is 1 TeV in ice [158], and muons can penetrate much more deeply than electrons. Moreover, the stochastic nature of the energy loss process becomes more important to consider since a muon may suffer large energy losses from individual bremsstrahlung and pair production interactions, initiating large electromagnetic cascades spread along its path. Nevertheless, the average energy loss rate may still be approximately parameterized as

$$-\left\langle \frac{dE_\mu}{dX} \right\rangle = a + bE_\mu. \tag{4.8}$$

where $a = 0.00268\,\mathrm{GeV\,cm^2/g}$ and $b = 0.47 \times 10^{-5}\,\mathrm{cm^2/g}$ [152]. This differential equation can be solved for the range of a muon, $R = \frac{1}{b}\ln\left(1 + bE_\mu/a\right)$. For example, the average range of a 1 TeV muon is 2 km in ice.

4.1.5 Taus

With a short lifetime of $\tau = 0.29\,\mathrm{ps}$, ultra-relativistic tau leptons will only travel on average a distance of $50\left(\frac{E_\tau}{1\,\mathrm{PeV}}\right)$ m before decaying, making them difficult to resolve directly except among the highest energy PeV neutrinos IceCube is able to detect. In addition, because of their large mass of 1.77 GeV, taus do not experience significant radiative energy loss within their lifetime except at energies above 100 PeV [152].

A tau has the following decay modes [29],

$$\begin{aligned} \tau^- &\to \nu_\tau + \text{Hadrons}, \quad \text{BR} = 64.79\%, \\ \tau^- &\to \nu_\tau + \bar{\nu}_e + e^-, \quad \text{BR} = 17.82\%, \\ \tau^- &\to \nu_\tau + \bar{\nu}_\mu + \mu^+, \quad \text{BR} = 17.39\%. \end{aligned} \tag{4.9}$$

Hadronic or electromagnetic cascades are produced in 82.61% of tau decays. The neutrinos arising from all decays are also of interest since they can regenerate ν_τ fluxes that would otherwise be attenuated by charged-current interactions within the Earth. The muon decay mode is of particular interest because it provides a mechanism to identify taus when their decay length is unobservable [159]. Taus from ν_τ charged-current interactions are polarized, and in the approximation that the muon is massless, the muon energy spectrum resulting from the decay is

$$\frac{dN}{dx} = \frac{4}{3}\left(1 - x^3\right), \tag{4.10}$$

where $x = E_\mu/E_\tau$ is the fraction of the tau's energy taken by the muon in the laboratory frame [66, 160]. Averaging over this distribution, the muon takes 40% of the tau's energy. Provided muon energy can measured with sufficient resolution to observe this shift, a population of muonic tau decay events might be distinguished among ν_μ events that also produce an outgoing muon. The feasibility of this method will be studied in Sec. 8.2.1.

4.2 The IceCube Neutrino Observatory

After Cherenkov light from secondary particles is emitted, it can travel hundreds of meters before it is detected by one of IceCube's optical sensors. The construction of IceCube began in the austral summer of 2004-2005 when the first 60 optical sensors, called digital optical modules (DOMs), were deployed within a borehole produced by a hot-water drill [161]. The DOMs are separated by a vertical spacing of 17 m at depths from 1450 m and 2450 m. A cable connects the DOMs and runs to the surface. After drilling, the water remaining in the borehole rapidly freezes, and the DOMs can never be removed or serviced. This deployment pattern is called a "string," and 86 total strings were deployed every year from 2005 to 2010 for a total of 5160 DOMs.

The layout of the complete IceCube Neutrino Observatory is shown in Fig. 4.4. Strings 1 through 78 are arranged on an approximate hexagonal grid with a 125 m horizontal spacing. The remaining 8 strings form a denser sensor array called DeepCore with 70 m spacing between strings [163]. The DeepCore strings have 50 DOMs deployed with 7 m spacing from depths of 2100 m to 2450 m where the ice is clearest and 10 additional DOMs deployed from 1750 m to 1850 m to veto vertical cosmic-ray muons.

During construction, data collection continued between string deployment in successive years. At the start of each year from 2005 to 2010, IceCube operated with 1, 9, 22, 40, 59, and 79 active strings. The completed detector configuration with 86 strings began operation

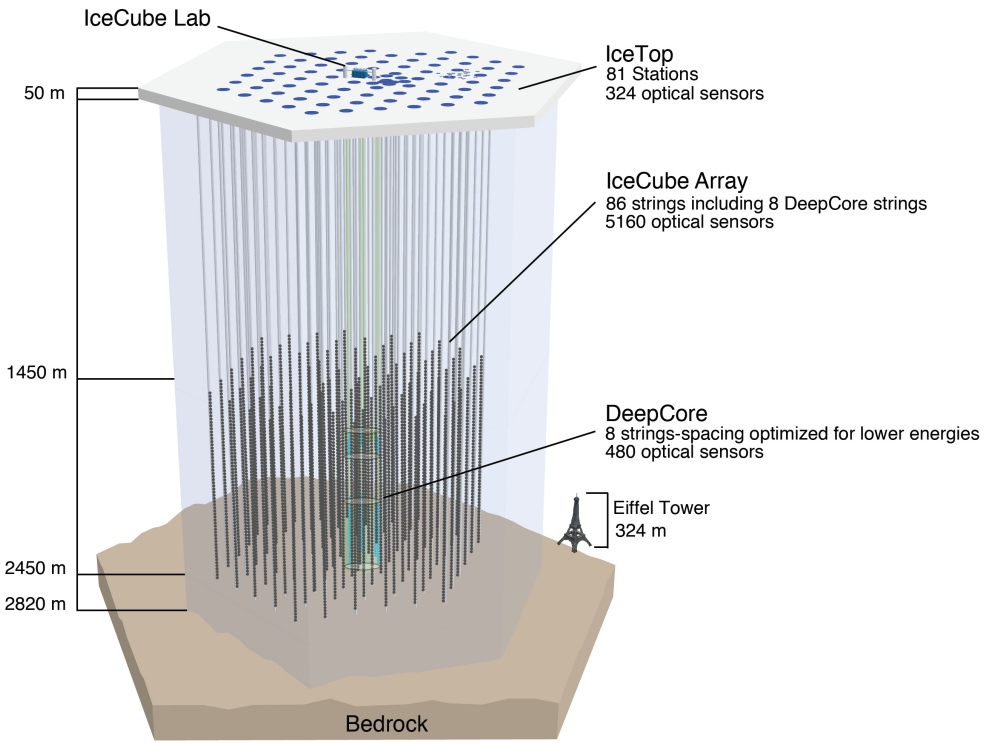

Figure 4.4: Layout of the IceCube Neutrino Observatory [162].

in 2011 and continues to the present. Only data from the 79 and 86 string configurations is used in this work.

On the surface, the IceTop air-shower array [164] consists of 81 stations near the top of IceCube strings. Each station consists of two ice-filled tanks that each contain two DOMs. IceTop is used to detect cosmic-ray air showers and and may be used to veto events detected by IceCube. The IceCube Laboratory is the central location where signals from the DOMs are collected, stored, filtered, and transmitted for further processing.

4.2.1 Digital Optical Module and Photomultiplier Tube

The Digital Optical Module (DOM) is the engineering achievement fundamental to IceCube's successful operation [165]. Illustrated in Fig. 4.5, it is an autonomous data acquisition system that houses a photomultiplier tube (PMT) to detect light and supporting electronics for PMT signal digitization, power, and control, all inside a 35.6 cm-diameter glass pressure sphere. Since it can never be accessed again after deployment, it was constructed to high reliability standards. To date, 98.4% of the DOMs are still operational, and most failures occurred shortly after deployment [166].

To detect light, a photomultiplier tube takes advantage the photoelectric effect to detect the electron created when a photon collides with a metallic cathode. The photoelectron

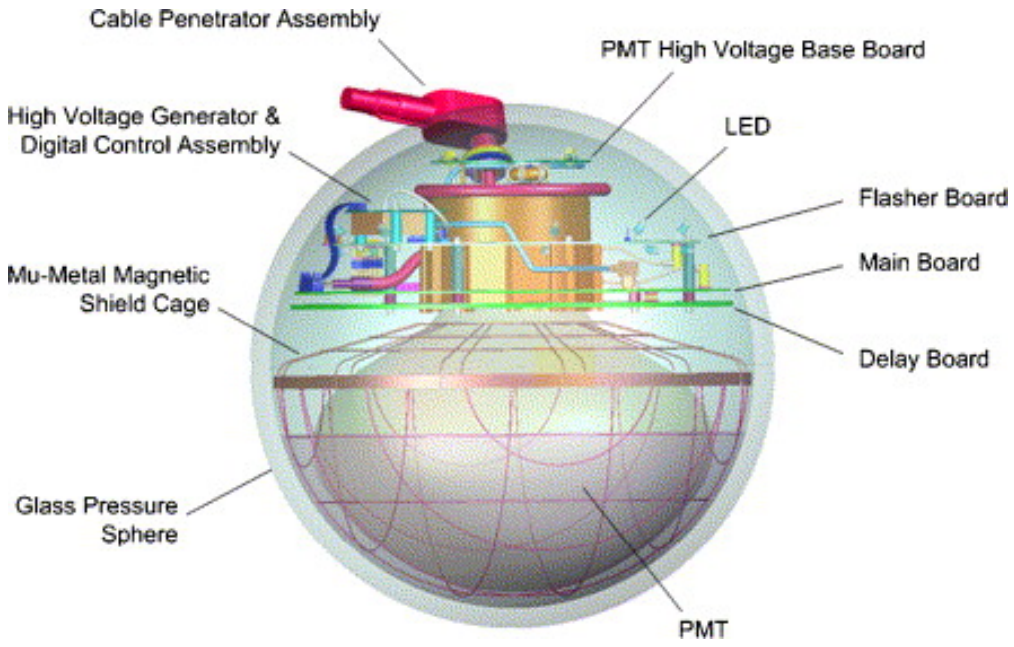

Figure 4.5: A diagram of the Digital Optical Module illustrating the photomultiplier tube in the lower half and supporting electronics in the upper half [161].

(PE) gains energy from a high applied voltage and initiates a cascade of secondary electrons after colliding with successive dynodes in the PMT. This creates a detectable current pulse, multiplying the charge of the original photoelectron by a large factor called the gain.

All IceCube DOMs contain the Hamamatsu R7081-02 PMT whose properties described below have been characterized with laboratory measurements [167]. It has a cathode diameter of 25 cm and a peak quantum efficiency of 25% at a wavelength of 390 nm. DeepCore PMTs have a higher peak quantum efficiency of 40%. Within a DOM, the PMT is oriented downwards to better collect light from up-going neutrino-induced particles. Absorption in the glass limits the sensitivity to wavelengths above 350 nm. Pure glass with low radioactivity limits the dark noise rate to be below 500 Hz. Additionally, a mu-metal cage surrounds the PMT and shields it from the Earth's magnetic field.

Each PMT operates with a gain of 10^7 and a typical applied voltage of 1500 V [167]. Pulses from single photoelectrons (SPEs) are not uniform. Both the integrated charge and arrival time of pulses are subject to fluctuations. The transit time to cross the PMT fluctuates with standard deviation 3 ns, which fundamentally limits the timing resolution of the detector. The charge also undergoes fluctuations of 30% and has a large non-Gaussian tail at high charge. At high photoelectron rates, the current in the PMT becomes saturated. This results in a non-linearity that reaches 2% at 50 mA or 31 PE/ns, an important issue for light from the highest-energy PeV neutrinos IceCube is able to detect.

Three other irregularities affect photoelectron detection: pre-pulsing, late-pulsing, and after-pulsing. A pre-pulse occurs when a photon bypasses the cathode and creates a photo-

electron at a dynode. The resulting pulse may be 30 ns earlier than normal pulses and has a lower charge; however they occur for less than 0.7% of photons. A late-pulse occurs when an electron elastically scatters of the first dynode, but later is accelerated back to the dynode. This initiates a delayed pulse that is otherwise normal and happens for less than 3.5% of photoelectrons. An after-pulse is caused by a photoelectron that ionizes a residual gas atom in the PMT, which accelerates back to the cathode and ejects more electrons upon impact. This causes a large pulse delayed by up to 10 μs that occurs for 5.9% of photoelectrons.

4.2.2 Data Acquisition

To ensure the highest fidelity, pulses from a PMT are digitized on the DOM before being sent to the surface by cable. Continuous digitization would produce a prohibitively high data volume, so a trigger is employed to decide if the PMT output should be digitized. A discriminator fires if the voltage exceeds a threshold calibrated to be 1/4 of the average SPE pulse height. Two types of digitizers are available to record the signal, a high-frequency Analog Transient Waveform Digitizer (ATWD) and a low-frequency fast Digital-to-Analog Converter (fADC). The ATWD sampling rate is 300 Msps, but it cannot operate continuously and is limited to 128 samples that are stored in a buffer. The fADC can operate continuously but at a lower sampling rate of 40 Msps. The ATWD is intended to capture the fine detail of the PMT pulse in the first 427 ns, while the fADC can capture the rest of long signals. Only the first 256 samples from the fADC are kept for a total time window of 6.4 μs. The digitizers have a 10-bit dynamic range, which is not sufficient to capture the full range of potential signals. For this reason, the ATWD has 3 channels receiving at gains of 16, 2, or 0.25. The lower gain channels are retained only if the next lower gain channel exceeded 75% of its range.

To determine if the digitizer readouts should be sent to the surface, a DOM uses information from its neighbors on a string in a system called "local coincidence." If the PMT discriminator fires in a neighboring or next-to-neighboring DOM within a $\pm 1\,\mu$s time window, the full readout from the ATWD and fADC is compressed and sent to the surface. This is called a "hard local coincidence" (HLC) condition. If the HLC condition is not met, it is called "soft local coincidence" (SLC). In this case, the ATWD digitization is aborted and only 3 samples around the peak value from the fADC are sent to the surface.

The IceCube Laboratory (ICL) contains a cluster of computers responsible for collecting and further processing the signals sent to the surface by all DOMs [166]. Though the local coincidence system does lower the data transmission rate, further reduction is needed using information from the whole detector collected at ICL. This is done using several trigger conditions, but the simplest one that captures all events of interest in this work is the Simple Multiplicity Trigger 8 (SMT8). The SMT8 trigger condition requires 8 DOMs to be in hard local coincidence within a sliding 5 μs time window. Data is kept from all DOMs recording in a readout window extending 4 μs backwards and 6 μs forwards from the first time the SMT8 condition was met. The result is an "event," a set of raw digitized PMT pulses recorded by each DOM. The total trigger rate varies seasonally but has a median of

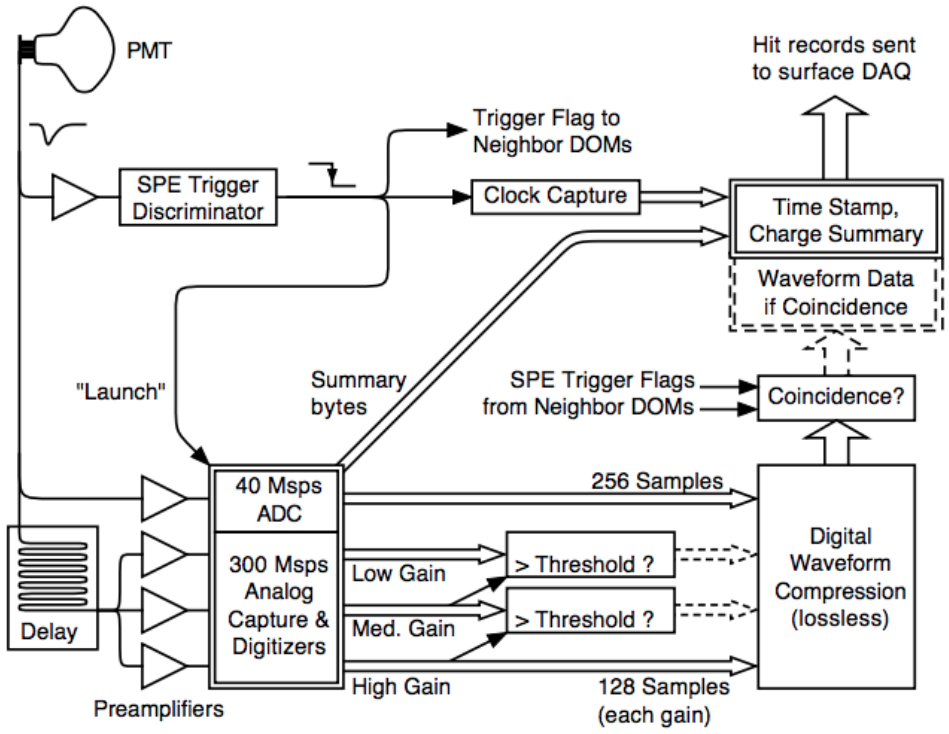

Figure 4.6: A data flow diagram illustrating the digitization of PMT pulses on a DOM [166].

2.7 kHz, and about 1 TB of data is created per day.

4.2.3 Data Processing

After triggering, the raw data from each event is further processed with the goal of reconstructing the timing distribution of photoelectrons emitted from the PMT cathode. This is the fundamental information used by all later reconstruction algorithms to estimate particle energy, direction, and position. To begin, the raw data is corrected to account for a number of distortions including voltage offset, PMT transit time, and transformer droop. The resulting waveform can be thought of as the superposition of pulses from individual photoelectrons. To infer the original distribution of photoelectrons, a deconvolution algorithm is applied using a template average waveform from a single photoelectron [168]. The result is an estimate of the number of photoelectrons observed in discrete time bins. Examples of the deconvolved photoelectron distribution for single and multiple photoelectron waveforms are shown in Fig. 4.7. Because of fluctuations in the shape of a single photoelectron waveform, a fractional number of photoelectrons in a bin can be found by the deconvolution algorithm.

All triggered events are compressed and archived at ICL, but it is also desirable to have more immediate transmission by satellite to the Northern Hemisphere. This is done using the Tracking and Data Relay Satellite System (TDRSS) [169]. The allocation of satellite

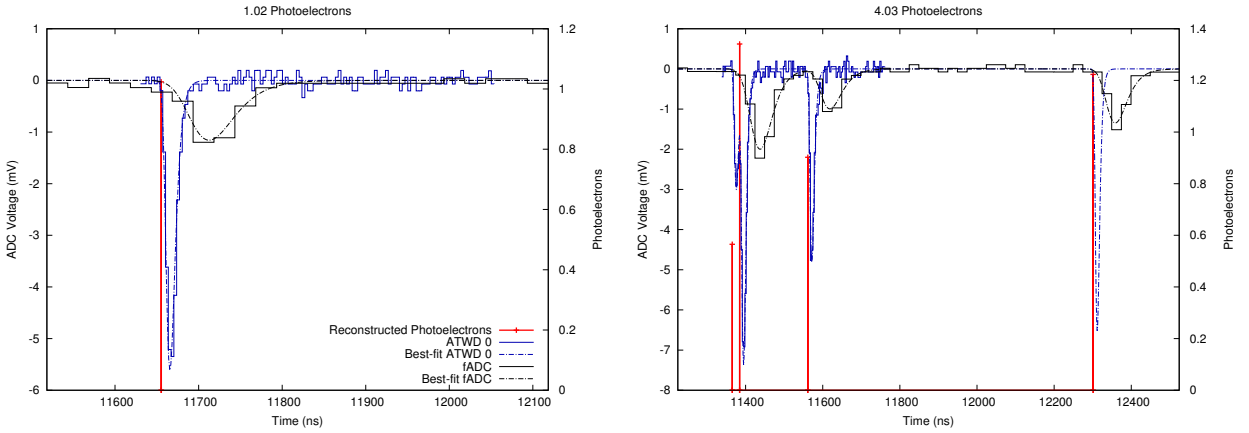

Figure 4.7: Left: An example single photoelectron recorded by the ATWD (blue) and fADC (black). The template functions for the ATWD and fADC are shown with smooth dashed curves and corresponding colors. The red line illustrates the number of photoelectrons found by the deconvolution algorithm, which is nearly one as expected. Right: An example waveform with multiple photoelectrons. The smooth dashed curves represent the best superposition of templates that matches the observed waveform. The red lines illustrate the inferred number of photoelectrons per bin. The variance in their height reflects fluctuations in the PMT response to photoelectrons.

bandwidth limits the amount of data that can be transmitted to $\sim 100\,\mathrm{GB}$ per day. For this reason, additional processing must be done to filter only the most interesting events for immediate transmission. A variety of filters are used that depend on reconstructions of particle energy and direction. Each filter is designed to target different types of events, for example cascades, muons, or extremely high energies. The filters used for the two event selections in this work are described in Sec. 6.1.1 and Sec. 7.1.

4.2.4 Calibration Systems

At IceCube, several calibration systems are available to monitor the detector and control systematic uncertainties. Properties of individual PMTs are monitored on a monthly basis using an on-board low-power light-emitting-diode (LED), oscillator, and voltage reference. Other calibration systems operate over the entire array. Using pulses regularly sent between DOMs and the ICL, a uniform timing standard synchronized to a GPS master clock is achieved with an accuracy of $< 2\,\mathrm{ns}$ [165].

All DOMs are also equipped with 12 LED flashers pointed radially outward; 6 point horizontally and the remaining 6 point upwards at an angle of 48°. They may be flashed individually or jointly with a configurable brightness and timing for a wide variety of applications. Flashers serve as the primary means to measure the spatially varying optical properties of ice at the South Pole described in the next section. Flashers are also used to verify timing resolution, locate the position of all DOMs in the array, and measure the

non-linear response of PMTs. However, their absolute light output is not well known and flashers cannot be used to measure the DOM optical photon detection efficiency, which is very important to correctly estimate a particle's energy.

Cosmic-ray muons are also a useful calibration tool since they have a well-known, nearly constant light output below the critical energy of 1 TeV where ionization dominates their energy loss. This is currently the best method to constrain the DOM optical efficiency to an uncertainty of $< 10\%$ [168]. They have also also been used to study ice properties, the angular dependence of DOM optical efficiency, and further verify the detector timing resolution.

4.3 Optical Properties of Ice at the South Pole

Though ice at the South Pole is clearer than any that can be produced in a laboratory, it still contains dust, air bubbles, and crystal defects that scatter and absorb light. As estimated from radar studies, the ice sheet extends down to bedrock at a depth of ~ 2820 m below the surface [170] and has been slowly accumulating from snowfall over the past 165,000 years [171]. Reflecting climate variations over this time period, the amount of dust deposited by winds has varied considerably. As a result, the degree of light scattering and absorption is a strong function of depth. At depths below ~ 1450 m, air bubbles originally trapped in snow tend to be squeezed into the ice crystal lattice and the ice becomes much clearer. For this reason, the topmost of IceCube's sensors are located at this depth, which also provides a large overburden to shield against down-going cosmic rays.

Ice core samples have been obtained from the South Pole only recently [172], and the first in-situ measurements of optical properties were carried out by the AMANDA experiment, the predecessor to IceCube, using lasers and LEDs that accompanied its optical sensors in the deep ice [173, 174]. IceCube continued this approach by equipping each DOM with LEDs and deploying two calibration lasers. Understanding the inhomogeneous optical properties of ice is one of most difficult challenges in interpreting IceCube data. Currently, a family of ice models called South Pole ICE (SPICE) provides the best available description of calibration data from LED flashers [175].

4.3.1 Absorption and Scattering

Absorption and scattering in ice are characterized by two length scales. The distance a photon travels before it is absorbed follows an exponential distribution, and the mean distance traveled is called the absorption length, λ_a. Similarly, the scattering length, λ_s, is the mean distance a photon travels before it scatters. Since a photon is typically only scattered by a small angle, it is more common to define an effective scattering length that represents the distance a beam of photons will travel before their directions become completely isotropic. This can be written as

$$\lambda_e = \frac{\lambda_s}{1 - \langle \cos\theta \rangle} \tag{4.11}$$

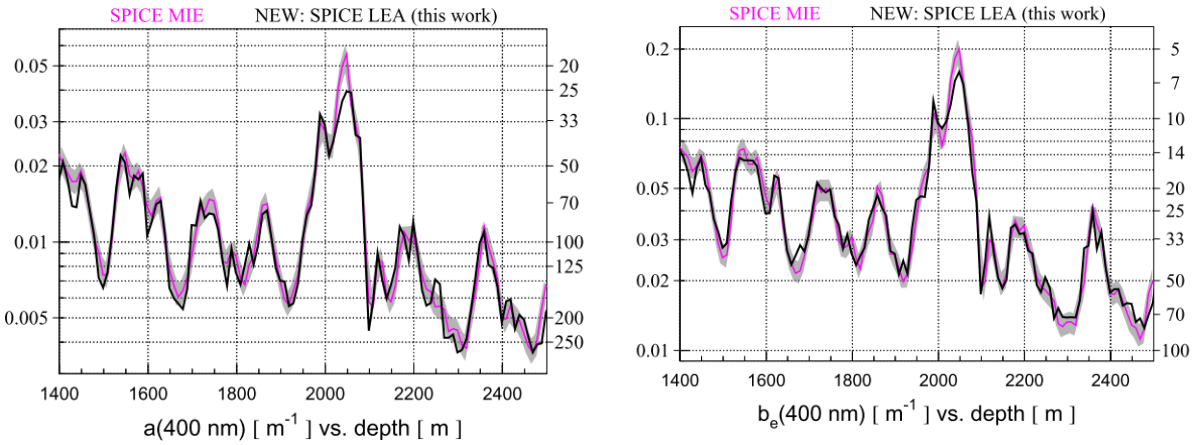

Figure 4.8: Left: The absorption coefficients (left axis) or lengths (right axis) across depth in the SPICE Mie and SPICE Lea ice models [176]. Right: The effective scattering coefficients (left axis) or lengths (right axis) across depth in the SPICE Mie and SPICE Lea ice models [176].

where $\langle \cos \theta \rangle$ is the average cosine of the scattering angle. Absorption and scattering are also often characterized by the coefficients a and b_e instead, which are simply the inverses of the absorption length and effective scattering length, respectively. Furthermore, the absorption and scattering coefficients also depend on wavelength and the scattering angle distribution. Both of these details can be calculated numerically with Mie scattering theory on dust grains [177]; however, in practice a parametric approach is used instead [170].

With the wavelength-dependence and scattering angle distribution parameterized, the absorption and effective scattering coefficients are then specified at a reference wavelength of 400 nm. The absorption and effective scattering coefficients determined from LED flasher data as a function of depth are shown in Fig. 4.8 for the two ice model generations SPICE Mie and SPICE Lea used in this work. Notably, there is a region of high absorption around 2050 m, colloquially called the "dust layer," where the absorption length reaches 20 m. Below this depth, the ice becomes exceptionally pure and light may travel over 200 m without being absorbed.

4.3.2 Tilt and Anisotropy

Tilt and anisotropy are two other important optical properties arising from the fact that the ice sheet is not static but a dynamic object that flows 10 m/yr across the South Pole [176]. Tilt refers to the fact that contours of equal absorption or scattering do not have a constant depth but vary by as much as 56 m over 1 km due the stratigraphy of the bedrock below. Shown in Fig. 4.9, a map of tilt was obtained using data from a dust logger deployed down the boreholes later used for IceCube's sensors [178, 179]. In the ice model generation SPICE Mie, absorption and scattering coefficients were measured across depth using LEDs in

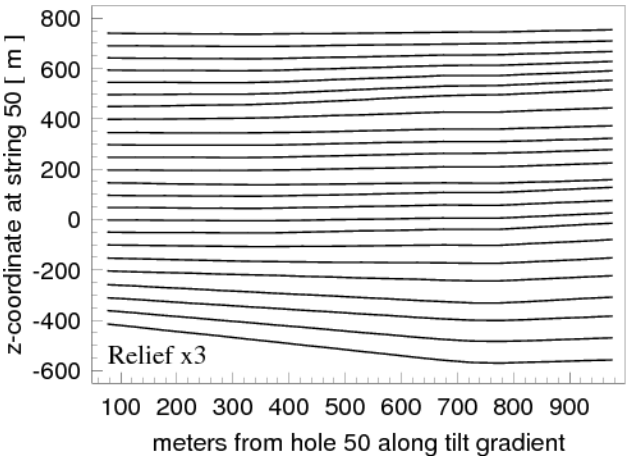

Figure 4.9: The tilt map of ice layers used in the SPICE models [175].

a single borehole and the dust-logger map is used to extrapolate these measurements across the entire array.

The other effect caused by ice flow is scattering anisotropy and was first discovered using IceCube's LED calibration system [176]. It is thought that dust grains become oriented along the ice sheet's direction of flow, and the scattering length varies depending on a photon's direction by up to 8%. Anisotropy can be incorporated by adding two global parameters describing the magnitude and direction of the anisotropy and this is the main advance in the ice model generation SPICE Lea, which is also used in this work.

4.3.3 Hole Ice

So far, properties of the pristine "bulk" ice at the South Pole have been described, but there is also a column of "hole" ice within the borehole for each string. The hole ice contains air bubbles that were trapped as water rapidly froze in the boreholes following string deployment. As a result, hole ice scatters light significantly more than the surrounding bulk ice. The effective scattering length in hole ice has not been well-determined from LED flasher data. An effective scattering length $\lambda_e = 50\,\text{cm}$ is typically used in the SPICE models [175], but $\lambda_e = 30\,\text{cm}$ is also consistent with data. Hole ice acts to scatter a vertically up-going photon away from a PMT and may scatter a down-going photon into a PMT, so it can be thought of as a correction to the angular photon detection efficiency for each PMT. Modeling of hole ice will be further described in Sec. 5.3.1, and it is an important systematic uncertainty that will be encountered again in Sec. 7.2.3.

4.4 Event Signatures

With the full process of data collection outlined, now particle interactions can be revisited to visualize the signature of each neutrino flavor and interaction type in the detector as well

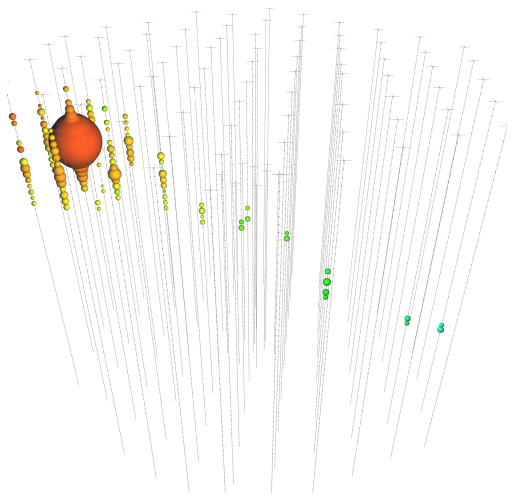

Figure 4.10: An example through-going track where a muon enters from the top of the detector. In this figure and similar event views, the number of photoelectrons detected by each DOM is represented through the size of colored spheres. The time of the first recorded photoelectron in each DOM is represented through color, ranging from early times in red to late times in green. Strings are shown in grey. The muon is estimated to have lost 31.5 TeV in the instrumented volume. Because it is down-going, the event is likely a cosmic-ray muon rather than a ν_μ CC interaction.

as backgrounds from cosmic-ray muons. Although by no means an absolute standard, events detected by IceCube can be typically classified into four topologies, through-going tracks, starting tracks, cascades, and double cascades.

4.4.1 Through-going Tracks

Through-going tracks are created whenever a muon passes through IceCube's instrumented volume. Light is emitted along a muon's path non-uniformly as it experiences stochastic energy losses, creating the pattern of detected photoelectrons shown in Fig. 4.10. Through-going tracks are by far the most common topology encountered since they result from down-going muons produced in cosmic-ray air showers. These events can consist of a bundle of many muons from the same air shower or only a single one that survived to reach the depth of IceCube. It is often difficult to tell these two cases apart since muon bundles typically only have a small lateral spread, but it is possible through analyzing the profile of energy losses along a track [180].

Neutrinos can also produce a through-going track when a muon neutrino undergoes an a charged-current interaction outside the detector, creating a muon that enters the detector. Neutrino-induced events form only 1 part in $\approx 10^6$ of IceCube's total trigger rate of 2.7 kHz [166]. Since neutrinos can pass through the Earth while cosmic-ray muons are blocked, the most straight-forward strategy to overcome the cosmic-ray muon background is to reconstruct

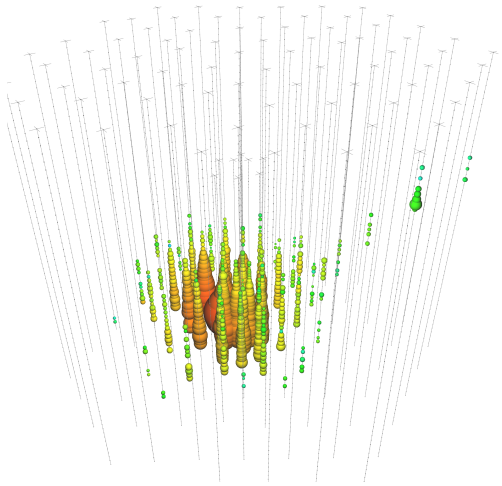

Figure 4.11: An example starting track [7]. A neutrino interaction produces large hadronic cascade and an up-going muon. The energy deposited within the instrumented volume is estimated to be 253 TeV. Most starting tracks are caused by ν_μ CC interactions, but the muonic decay of a τ in a ν_τ CC interaction or a hadronic decay to a muon in any interaction type can also produce this signature.

the direction of through-going tracks and identify up-going events. This strategy has been used to study atmospheric neutrinos and establish the existence of a diffuse astrophysical neutrino flux above 100 TeV [131, 181, 136].

The direction of through-going tracks can be well-reconstructed with angular resolution $< 1^\circ$ [182], a fact verified by observing the cosmic-ray shadow of the Moon [183]. For this reason, they are the focus of most searches for point sources of astrophysical neutrinos [76]. The details of track directional reconstruction will be described further in Sec. 6.1.4 and Sec. 7.2.1. Energy reconstruction typically relies on estimating the average energy loss rate along the track, $\frac{dE}{dX}$, which can be related to the energy of the muon through Eq. 4.8. This approach suffers from the stochasticity of energy losses, which limits resolution on $\log_{10} E_\mu$ to be ~ 0.3. However, since the energy loss distribution has a non-Gaussian tail at high energy, excluding segments of the track where high losses occur can improve the resolution of $\log_{10} E_\mu$ to 0.22 [184]. Even so, if one is interested in the energy of the original neutrino, this is subject to even higher uncertainty since the muon has lost an unknown amount of energy before entering the detector.

4.4.2 Starting Tracks

Starting tracks occur when a muon is created by a neutrino interaction inside IceCube's instrumented volume. Most commonly, starting tracks result from charged-current ν_μ interactions. More rarely, starting tracks can also result for charged-current ν_τ interactions when the tau created decays to a muon, or neutrino interactions of any flavor or interaction

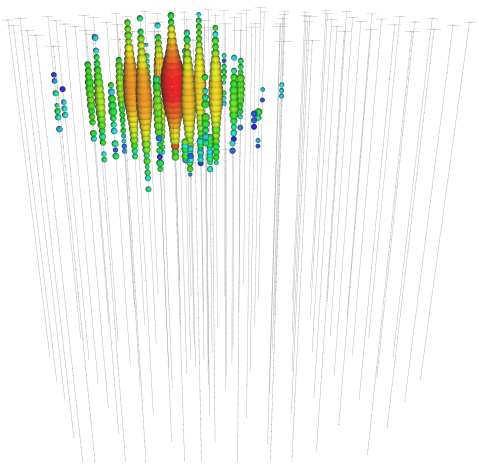

Figure 4.12: An example cascade event [78]. The visible energy is estimated to be 2 PeV, making it the most energetic cascade observed by IceCube. Cascades may arise from ν_e CC interactions, all flavors of NC interactions, or Glashow resonance interactions. The energy of this cascade is too low to be a likely Glashow resonance candidate.

type if a muon is created by decays within the hadronic shower. In any case, light from the hadronic cascade accompanies light from the muon as shown in Fig. 4.11.

Identifying starting tracks is also a powerful method to search for neutrinos that is effective in the down-going direction. One can use light detected by the outermost DOMs in the instrumented volume to veto incoming cosmic-ray muons and isolate starting tracks. Additionally, taking advantage of the self-veto probability in Sec. 3.3.3, atmospheric neutrino backgrounds can be vetoed as well. This strategy was the basis for the first detection of astrophysical neutrinos by IceCube [7] and will be used both in Ch. 6 and Ch. 7.

The angular resolution of starting tracks is comparable to that of through-going tracks, as will be shown in Sec. 7.2.1. Also similarly, the energy of the muon may be estimated through a measurement of $\frac{dE}{dX}$. Provided the energy of the accompanying hadronic cascade can also be measured, a direct estimate of the neutrino's energy can be obtained along with the inelasticity $y = E_{\text{had}}/E_\nu$. The key difficulty is correctly attributing light as originating from the muon or cascade. This will be a challenge addressed in Sec. 7.2.2.

4.4.3 Cascades

Cascades arise from neutrino-induced particle showers. They may arise from the hadronic showers produced in neutral-current interactions of any flavor, the combined electromagnetic and hadronic shower produced by charged-current ν_e or ν_τ interactions, or Glashow resonance events where a hadronic shower or electromagnetic shower is produced. Except at energies above 1 PeV where the LPM effect takes hold, light emission is approximately point-like and directed along the Cherenkov angle. Scattering in ice washes out the Cherenkov, and the detected light pattern is nearly spherical, as shown in Fig. 4.12. As a result, cascade

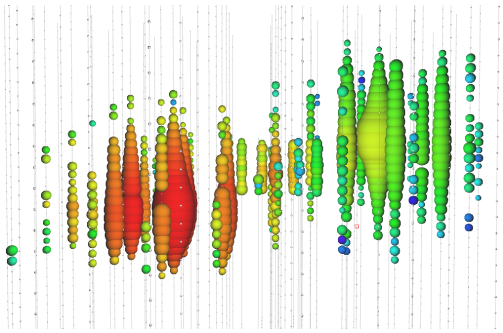

Figure 4.13: A simulated double bang. No double bangs have yet been observed. The neutrino interacts on the left and a tau is produced which travels $\sim 500\,\mathrm{m}$ before decaying to produce a second cascade.

direction has a poorer angular resolution of $\sim 15^\circ$ [168]. Cascade reconstruction will be described in Sec. 6.2.1. No calibration analogous to the Moon shadow is available to verify cascade angular resolution, but in Sec. 7.2.3, a new scheme will be presented to test cascade pointing accuracy using high-inelasticity starting tracks.

Although angular resolution is poor, the energy of an electromagnetic cascade can be reconstructed with $\sim 10\%$ resolution by taking advantage of the linear relationship of energy with the amount of the Cherenkov light emitted [168]. Since it is extremely difficult to differentiate between hadronic and electromagnetic cascades, it is common to define an "electromagnetic-equivalent" energy that represents the energy of an electromagnetic cascade that produces the same amount of Cherenkov light. For charged-current ν_e interactions, this closely approximates the neutrino's energy, differing only because of the lower light output of the hadronic cascade that accompanies the electromagnetic cascade. For neutral-current interactions of any flavor, the electromagnetic-equivalent energy is only a lower bound on the neutrino's energy due to the unknown energy carried away by the outgoing neutrino.

With a topology distinct from the long tracks expected from atmospheric muons and excellent energy resolution, cascades are also highly useful for identifying neutrino interactions. They provide the strongest constraints on the energy spectrum of both atmospheric electron neutrinos [121] and astrophysical neutrinos [139].

4.4.4 Double Cascades

For charged-current ν_τ interactions, the tau produced will usually decay to an electron or hadrons, resulting in a second cascade displaced from the accompanying hadronic cascade. Typically this distance is too small to observe, and most ν_τ interactions will appear as a single cascade. However, at energies above 1 PeV where the decay length becomes comparable to IceCube's string spacing, this displacement may be observed. This gives rise to a rich variety of potential ν_τ event signatures [185], but the most striking is the "double bang" where both cascades are observed inside the instrumented volume [186]. No double bangs have yet been

observed, but a simulated event is shown in Fig. 4.13. Even when the displacement between cascades is smaller than one string spacing, certain orientations may produce a "double pulse" within a single DOM when the wavefront from one cascade arrives shortly after the wavefront from the other cascade [187].

Chapter 5

Simulation

Simulating the IceCube response to neutrinos and other particles is vital to understand the data it collects. By comparing the distribution of an observable quantity between simulated and real data events, conclusions can be drawn about the population of neutrinos detected. Through the use of Monte Carlo methods, every effort is made to accurately model every physical process from secondary particle production to Cherenkov light propagation in ice to the response of individual DOMs.

5.1 Particle Generation

The starting point of any IceCube simulation is a particle generator that simulates a flux of particles impinging on the detector. Three different particle generators are needed for neutrinos, cosmic-ray nuclei, and atmospheric muons.

5.1.1 Neutrino Generator

A program called `neutrino-generator` [188] is used to simulate a flux of high-energy neutrinos as they propagate through the Earth and interact inside the detector. The user may pick from several neutrino cross section models, and the CSMS model [180] described in Sec. 2.3 has become the standard. `neutrino-generator` is based on an earlier program, All Neutrino Interaction Simulation (ANIS) [189]. To begin the simulation, neutrinos of a selected flavor are drawn from a power-law energy spectrum with equal probability of being a neutrino or antineutrino. Neutrino directions are sampled from an isotropic distribution and aimed to intersect with an imaginary surface centered on the detector. The surface is a cylinder with its axis oriented along each neutrino's direction or the axis may be fixed to the vertical. Each neutrino is uniformly distributed within the projected area of the surface perpendicular to the neutrino's direction.

Starting from the surface of the Earth, each neutrino is then propagated in discrete steps through matter encountered on the way to the detector. For each step of length Δx, a neutrino interacts with probability $1 - \exp(-\sigma n \Delta x)$, where σ is the total interaction cross

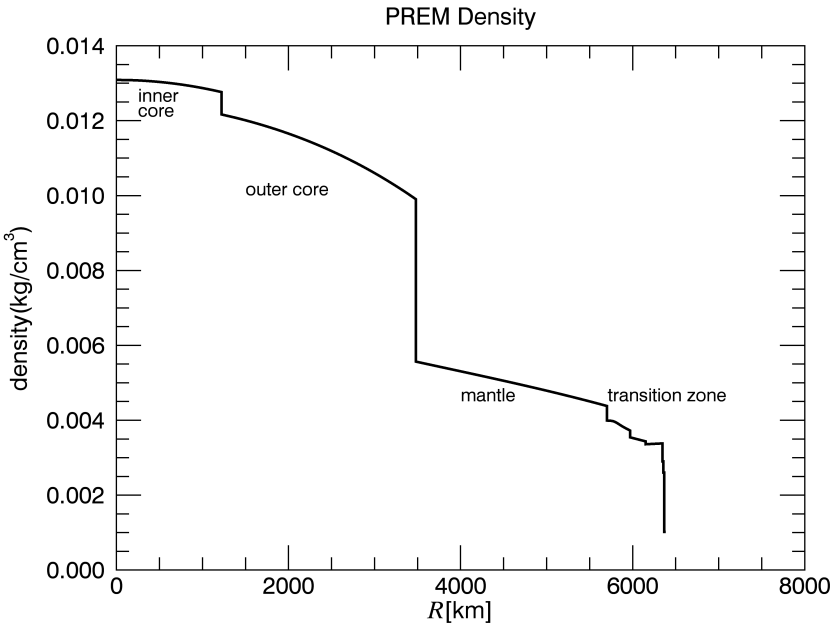

Figure 5.1: The mass density of the Earth as a function of radius from the Preliminary Earth Reference Model. Reproduced from [190].

section and n is the density of nucleons at the location in the Earth. The nucleon density is found from the Preliminary Earth Reference Model (PREM) [191], which provides the mass density of the Earth as function of radius as shown in Fig. 5.1. To convert to a nucleon density, an isoscalar target of equal part neutrons and protons is assumed, which is a good approximation for silicates and iron in the Earth but less so for ice. If an interaction is chosen, an interaction type among charged-current, neutral-current, and Glashow resonance is selected. For electron and muon neutrinos, if a charged-current interaction is chosen, no neutrinos are produced that could yet reach the detector. Terminating the simulation here can drastically reduce event statistics at high energies, so instead the event is forced to undergo a neutral-current (NC) or Glashow resonance (GR) interaction and assigned a weight, p_{prop}, corresponding to the ratio of the sum of the NC and GR interaction cross sections to the total cross section. If an NC or GR interaction is chosen, the energy of the outgoing neutrino is sampled according to the single differential cross section, which will be described below. For tau neutrinos, charged-current interactions are allowed to proceed normally since the tau lepton will decay producing at least one neutrino. One neutrino is randomly selected for further propagation following the same steps outlined so far. Though neutrino absorption in the Earth is simulated by `neutrino-generator`, oscillations are not, limiting its applicability to energies above $\sim 1\,\text{TeV}$.

After enough steps, neutrinos will reach the imaginary cylinder defining the detection volume where they are forced to interact somewhere between the projected entry and exit

points of the cylinder according to a uniform distribution. Since a muon from a CC ν_μ interaction far outside the surface can still reach the detector, the length where interactions are forced to occur is automatically extended backwards according to a parameterization of the muon range in ice. Because the interaction of the neutrino is forced, a weight of $p_{\text{int}} = n\sigma L \exp(-n\sigma L_{\text{before}})$ must be assigned, where L_{before} is distance before the selected interaction point where an interaction may have occurred and L is the total length where an interaction may have occurred. This is often combined with the propagation weight to give a total interaction probability weight $w_{\text{int}} = p_{\text{prop}} p_{\text{int}}$. The interaction type is then randomly selected according to CC, NC, and GR cross sections, and secondary particles are generated for further propagation by another program.

For deep inelastic scattering events, an outgoing lepton is generated according to the double differential cross sections in Bjorken-x and inelasticity, y. Pairs of xy values are obtained from a pre-computed table of 3,000 samples of the differential cross section generated at logathmic intervals in neutrino energy, $\log_{10} E = 0.1$. For charged-current interactions, the outgoing lepton angle with respect to the neutrino is calculated from the xy values, whereas the outgoing angle of a neutrino is neglected in neutral-current events. Since the detailed composition of the final state hadrons is not observable, the hadronic system is treated as a single, generic hadron, which is assumed to be collinear with the incoming neutrino. Glashow resonance events are treated in a similar fashion. Hadronic W-boson modes are treated as a single outgoing hadron and leptonic modes are generated with lepton energy sampled according to the differential cross section in Eq. 2.34, collinear to the incoming neutrino.

Since flux models beyond a simple power law are often desired, `neutrino-generator` simulation can be re-weighted to an arbitrary flux, $\Phi(E)$. A weight is assigned to an event that corresponds to the ratio of the desired flux $\Phi(E)$ to the generated fluence of neutrinos. The generated fluence can be written as $\frac{N(E)}{\Omega A_\perp}$ where $A_\perp$ is the projected area of the detection surface for each neutrino and Ω is the solid angle covered by the simulation, typically 4π for neutrinos from the full sky. $N(E)$ is the energy distribution of simulated neutrinos, which for a power law distribution of N_0 events and power-law index γ stretching from energy E_{min} to E_{max} is

$$N(E) = N_0 \frac{E^{-\gamma}}{\int_{E_{\text{min}}}^{E_{\text{max}}} E^{-\gamma} dE}. \tag{5.1}$$

The number of events N_0 should be considered as the number of neutrinos or antineutrinos, not the total, if the flux is also known for neutrinos or antineutrinos separately. This means that N_0 is half of the total number of events generated for each flavor. Lastly accounting for the interaction weight, the total weight (in units of Hz) assigned to each neutrino is then

$$w = w_{\text{int}} \frac{A_\perp \Omega \Phi(E)}{N(E)}. \tag{5.2}$$

By summing these weights over all simulated events surviving an event selection, one can then compute the expected event rate for any flux model desired.

5.1.2 Neutrino Charm Production

One addition made to `neutrino-generator` used in Ch. 7 was the ability to simulate the neutrino production of charm hadrons. This is done using the leading-order cross section for charm production from Sec. 2.3.5. For each xy pair of the sampled CSMS differential cross section, a charm production probability is calculated as

$$p_{\text{charm}}(E, x, y) = K_{\text{NLO}} \frac{\frac{d\sigma^{\text{charm}}}{dxdy}_{\text{LO}}}{\frac{d\sigma}{dxdy}_{\text{NLO}}}. \tag{5.3}$$

where $\frac{d\sigma^{\text{charm}}}{dxdy}_{\text{LO}}$ is the leading-order charm production cross section, $\frac{d\sigma}{dxdy}_{\text{NLO}}$ is the total differential CSMS cross section at next-to-leading order, and $K_{\text{NLO}} = \sigma_{\text{NLO}}/\sigma_{\text{LO}}$ is a K-factor used as a means to approximate next-to-leading order effects using the ratio of total NLO and LO cross sections. This K-factor definition preserves the fraction of charm production events and their xy distribution computed at leading order. The charm quark is assumed to be massless for all calculations, which is a good approximation above 1 TeV as shown in Fig. 2.9.

After an event is selected from the table of pre-computed xy pairs, a charm hadron is generated according to the above probability. To simulate hadronization, the approach from the low-energy neutrino event generator GENIE is adopted [192]. The type of charm hadron is randomly selecting using the fragmentation fractions

$$f_{D^0} = 0.60,\ f_{D^+} = 0.23,\ f_{D_s^+} = 0.15,\ f_{\Lambda_c^+} = 0.02, \tag{5.4}$$

which are obtained from the calculation in [193] at a neutrino energy of 100 GeV. The energy of the charm hadron is chosen according to a fragmentation function

$$D(z) = \frac{4\epsilon}{z}\left(1 - \frac{1}{z} + \frac{\epsilon}{1-z}\right)^{-2}, \tag{5.5}$$

where z is the fraction of the total available hadronic system energy available given to the charm hadron. The parameter $\epsilon = 0.2$ is taken from the fit of neutrino charm production data in [50]. The remaining hadronic system energy is put into another single, generic hadron. The charm hadron is assumed to be collinear with the neutrino just like the remaining hadronic system, so the angular separation distribution between the outgoing lepton and a muon from charm decay is not accurate. The propagation of the resulting charm hadron in ice will be described in 5.2.3.

5.1.3 CORSIKA

To predict the background of atmospheric muons and neutrinos reaching IceCube, it is not sufficient to know the flux of each of these particles individually. These particles are produced in cosmic-ray air showers and often many muons in a tight bundle are detected at

the same time. Thus, treatment of atmospheric backgrounds must involve a full simulation of muon and neutrino production in cosmic-ray air showers. Cosmic-ray air shower simulation has a long history, and the COsmic Ray SImulations for KAscade (CORSIKA) program has become the standard workhorse for most cosmic-ray air shower experiments [155].

To begin a CORSIKA simulation, cosmic-rays are typically drawn from a power-law energy distribution and aimed at an upright cylinder centered around the detector. However, the zenith angle is not selected isotropically, but with an angular distribution proportional to the projected area of the detection cylinder, which improves the simulation efficiency. The type of nucleus is selected from a user-defined list of elements, but typically only 5 are used to represent the full cosmic-ray composition, ${}^{1}_{1}\mathrm{H}$, ${}^{4}_{2}\mathrm{He}$, ${}^{14}_{7}\mathrm{N}$, ${}^{27}_{13}\mathrm{Al}$, or ${}^{56}_{26}\mathrm{Fe}$. The atmospheric model used for IceCube in CORSIKA simulations is derived from the MSIS-90-E global model [194], which provides a density profile at the South Pole that is fit to a 5-layer piecewise exponential function in altitude.

Cosmic-rays are then propagated through the atmosphere, accounting for ionization and radiative losses, until a collision with an air nucleus. A hadronic interaction model is needed to calculate interaction cross sections and simulate the production of particles after a collision. Most often, the SIBYLL 2.1 model is used to describe interactions above 60 GeV and and is calibrated against particle accelerator data [195]. The particles created after each interaction in the atmosphere are further propagated through the atmosphere until they reach surface level at an altitude of 2834 m or fall below a user-supplied energy threshold. This may be set as high as 273 GeV, since particles above this energy cannot produce muons that will penetrate the ice overburden to reach the detector. This considerably lessens the computational burden to simulate cosmic-ray air showers compared to other surface detectors where particles must be tracked down to MeV energies. Hadronic decays are also treated by CORSIKA, and any muons produced during the air shower are propagated through the atmosphere incorporating all ionization and radiative losses in air until they reach surface level.

CORSIKA can also track neutrinos, which are assumed to undergo no interactions after they are produced. neutrino-generator can treat neutrinos produced by CORSIKA after one has been randomly selected from those reaching surface level. The neutrino is selected with a linear energy bias and receives a selection weight, which more efficiently selects the high-energy neutrinos that are likely to interact. The propagation of the neutrino through ice to the detector then follows the same procedure in Sec. 5.1.1. The result is a neutrino-induced event that also may be accompanied by atmospheric muons that reached the detector. This provides the best method to calculate the atmospheric self-veto probability mentioned in Sec. 3.3.3.

Weighting CORSIKA events to an arbitrary cosmic-ray flux model is similar to neutrino-generator, but no interaction weight must be applied for muon bundles. If the flux of each element is give by $\Phi_{A,Z}(E)$, the weight assigned to an event is

$$w = \frac{A_{\perp}\Omega\Phi_{A,Z}(E)}{N_{A,Z}(E)}, \tag{5.6}$$

where $A_{\perp}\Omega$ is the projected area of the upright simulation cylinder integrated over solid angle, and $N_{A,Z}(E)$ is the energy distribution of each element. Often many CORSIKA simulations are run over a variety of energy ranges that may overlap and it is desired to combine the simulations to predict an event rate. In this case the energy distribution of each set is summed before inserting into Eq. 5.6.

5.1.4 MuonGun

Though a cosmic-ray air shower does not need to be simulated to completion for IceCube, the speed of these simulations is still a major bottleneck in producing enough events to adequately model atmospheric backgrounds. To simulate atmospheric muons, one way around this is to parameterize the distribution of muons arriving at the detector and avoid the costly propagation of particles starting from the top of the atmosphere. MuonGun is a tool to compute these distributions from previously run CORSIKA simulations and inject muons directly at a cylinder closely surrounding the detector [196].

In its most general usage, MuonGun can be used to compute the multiplicity, energy, and radial distribution of muon bundles from air showers. However, it is more commonly used to simulate cosmic ray events where a single muon survives to reach the detector, which is often the most pernicious remaining background in many analyses. In this case, the flux of single muons at an upright cylinder surrounding the detector is obtained from a spline fit of CORSIKA simulations from a given primary cosmic-ray flux model.

MuonGun can then be used to inject single muons at the cylinder surface for further propagation. A large number of single muon events can be quickly generated, far larger than the number in the original CORSIKA simulation used to derive the single muon flux. The energy distribution of the injected muons may also be chosen to be a distribution different from the one obtained from CORSIKA, such as a power-law, which is often used to target certain energy ranges of interest. The simulation can then be re-weighted to any other single muon flux obtained from spline fits to CORSIKA distributions from any primary cosmic-ray flux model.

5.2 Secondary Particle Propagation

After a particle generator is used to place secondary particles within or in the vicinity of the detector, they are propagated through ice.

5.2.1 Muon Monte Carlo

Muon propagation in ice was originally simulated using the Muon Monte Carlo (mmc) program written in the Java language [152]. It has since been replaced by PROPOSAL [197], a program written in C++ but producing numerically identical results. PROPOSAL/mmc propagates muons in discrete 1 cm steps and uses the differential cross section of each of the four

interaction processes (ionization, bremsstrahlung, pair production, and photonuclear) to select if an interaction occurs and if so, sample the energy lost by the muon. One technicality is that differential cross sections have a divergence for small energy losses, so it is necessary to introduce a low energy cutoff, E_{cut}, typically chosen to be 500 MeV. Only discrete energy losses above E_{cut} are created, while remaining energy losses below E_{cut} are treated as continuous. For any discrete energy loss experienced by the muon, a corresponding electron, photon, $e^{\pm}$ pair, or generic hadron is created depending if an ionization, bremsstrahlung, pair production, or photonuclear interaction interaction occurred, respectively. Each of secondary particle will then initiate a particle cascade, which is simulated separately.

PROPOSAL/mmc is also used to propagate tau leptons through ice following the same procedure but with interaction cross sections modified for the tau's larger mass. It can also treat tau decay. For leptonic decay modes, the energy of the electron or muon is sampled from the 3-body decay distribution in the center-of-mass frame,

$$\frac{dN}{dx} = \frac{G_F^2 m_\tau^5}{192\pi^3}(3-2x)x^2, \tag{5.7}$$

where x is the fraction of energy taken by outgoing lepton, and the resulting system is boosted to the lab frame. Hadronic decay modes are treated as 2-body decays into one of the following hadrons or resonances: $\pi^{\pm}$, $\rho(770)$, $a1(1260)$, or $\rho(1465)$ [152]. Regardless of the chosen hadronic mode, only a single generic hadron is created for further propagation.

5.2.2 Cascade Monte Carlo

Cascade Monte Carlo (cmc) is a program used to simulate electromagnetic and hadronic cascades in ice. Tracking all the particles produced in a cascade is prohibitively expensive, so the parametric approach described in Sec. 4.1.2 and Sec. 4.1.3 is used.

Electromagnetic cascades initiated by electrons, positrons, or photons are treated identically. If the energy of the particle is below 1 TeV, the extent of the cascade is ignored and no further action is taken. All light emission will later be assumed to be point-like from the origin of the particle. Above 1 TeV, the extent of the cascade is approximated as a sequence of point-like sub-cascades separated by 3 radiation lengths and their energy is distributed according to the longitudinal profile of Eq. 4.3. Above 1 PeV, where the Landau-Pomeranchuk-Migdal (LPM) effect is important, a simplified one-dimensional simulation of pair production and bremsstrahlung interactions is done using the LPM-suppressed cross sections from [150]. Electromagnetic particles and their secondaries are tracked individually until their energy falls below 50 TeV at which point the sum of longitudinal cascade profiles over remaining particles is used to construct a sequence of point-like sub-cascades.

Similarly to electromagnetic cascades, all hadronic cascades are treated identically regardless of the initiating hadron. To account for the diminished light yield of hadronic cascades relative to electromagnetic cascades and its fluctuations, the hadron is replaced by an electromagnetic particle with energy diminished by the light-yield factor described by Eq. 4.5 and randomly fluctuated according to a normal distribution with the standard

deviation found in Eq. 4.6. The electromagnetic particle is then extended into sub-cascades if necessary as described above to account for the cascade's longitudinal profile. The production of muons is also treated by cmc. The total number of muons is selected from a Poisson distribution and their energy is distributed as a power-law according to Eq. 4.7. The muons are assumed to be collinear with the initial hadron, and after they are created they must be further propagated using PROPOSAL.

5.2.3 Charm Hadron Monte Carlo

Charm hadrons require special treatment because they may promptly decay before initiating a hadronic shower. Even if a charm hadron interacts in ice, it may keep a large fraction of its energy and later decay. Charm Hadron Monte Carlo (chmc) was written to simulate both charm hadron decays and their propagation in ice.

Of interest in charm propagation are two length scales, the decay and interaction lengths,

$$d = \frac{E\tau}{mc},\ l = \frac{m_{ice}}{\rho N_A \sigma(E)} \tag{5.8}$$

where E is the energy of the hadron, m is its mass, τ is its proper lifetime, $m_{ice} = 18.02$ g/mol is the molecular mass of ice, and $\rho = 0.917$ g/cm^3 is the density of ice. The total inelastic cross section $\sigma(E)$ is a major unknown for charm hadrons since it cannot be measured directly in particle accelerator experiments. Nevertheless, this problem has been encountered before in propagating charm hadrons within cosmic-ray air showers, and the approach used here will follow a model originally developed for CORSIKA as closely as possible [198].

Parameterizations of the hadron-nucleon cross section for charm mesons and baryons are available in CORSIKA, but are only valid above 1 PeV. To extrapolate below this energy, the kaon-nucleon cross section, $\sigma_{Kp}(E)$ is used and scaled to match this calculation at 1 PeV, so that the charm meson ($D^{\pm}$, D^0, D_s^+) and baryon (Λ_c^+) cross sections are

$$\sigma_{Dp}(E) = \begin{cases} E < 1 \text{ PeV} & 0.87549\sigma_{Kp}(E) \\ E \geq 1 \text{ PeV} & \exp\left[1.891 + 0.2095\log_{10}(E/\text{GeV})\right] - 2.157 + 1.263\log_{10}(E/\text{GeV}), \end{cases} \tag{5.9}$$

$$\sigma_{\Lambda_c p}(E) = \begin{cases} E < 1 \text{ PeV} & 1.22442\sigma_{Kp}(E) \\ E \geq 1 \text{ PeV} & \exp\left[2.269 + 0.207\log_{10}(E/\text{GeV})\right] - 0.9907 + 1.277\log_{10}(E/\text{GeV}). \end{cases} \tag{5.10}$$

The parameterization of $\sigma_{Kp}(E)$ is also taken from CORSIKA [155]. For an ice target medium, the Glauber interpolation scheme from CORSIKA is used for the oxygen nucleus, and the ice cross section is the weighted sum between hydrogen and oxygen [155]. The above cross sections can be used to calculate the critical energy for each charm hadron where the decay length and interaction length are equal in ice,

$$\epsilon_{D^+} = 22 \text{ TeV},\ \epsilon_{D^0} = 53 \text{ TeV},\ \epsilon_{D_s^+} = 47 \text{ TeV},\ \epsilon_{\Lambda_c^+} = 104 \text{ TeV}. \tag{5.11}$$

Above the critical energy it is important to model the interaction of charm hadrons.

The elasticity distribution of the outgoing charm hadron after a collision was also modeled in [198]. Elasticity, z, is the fraction of the incoming charm hadron energy taken by the leading outgoing charm hadron in the lab frame. Elasticity distributions are parameterized with a normal distribution

$$\frac{d\sigma}{dz} \propto e^{-\frac{(z-z_0)^2}{2\sigma^2}}. \tag{5.12}$$

The value z_0 was chosen to match the mean elasticities found in [198]. For charm mesons, $z_0 = 0.56$, and for charm baryons $z_0 = 0.59$, and a standard deviation $\sigma = 0.2$ was used for both. The target in this study was air, but there was little difference when switching from a proton to air target, so these values are kept for ice.

Using these cross section models to calculate an interaction length, a charm hadron can be propagated through ice with a simple one-dimensional method. To start, two distances x_d and x_l are sampled according to exponential distributions, $e^{-x_d/d}$ and $e^{-x_l/l}$. If $x_l > x_d$, the charm hadron interacts and the elasticity is sampled according to Eq. 5.12. A generic hadron is created with the remaining energy that will eventually be propagated with `cmc`. The outgoing charm hadron type is assumed to be the same type as the incoming charm hadron, which is the most likely outcome from the study in [199]. This procedure continues until $x_d > x_l$ or the energy falls below a defined threshold and the charm hadron decays. The threshold, defaulting to 1 GeV, allows users to turn off the effects of interactions below a given energy.

When a decay occurs, the final state is generated from a pre-computed sample of 10,000 simulated charm hadron decays each for D^+, D^0, D_s^+, and Λ_c^+ using Pythia 8 [153]. These decays are stored as a 7-column table. Each of the columns contains the fraction of the charm hadron's energy taken by e, ν_e, μ, ν_μ, τ, ν_τ, or hadrons in the infinite momentum frame. The default branching fractions for charm hadrons from Pythia 8 are used. For each decay, one row of the table is randomly selected and particles are created for each non-zero column entry with a corresponding fraction of the charm hadron's energy. All decay products are assumed to be collinear with the original charm hadron direction.

5.3 Photon Propagation

Simulating the Cherenkov light from every secondary particle is not feasible on a large scale for the high-energy events observed in IceCube, so instead light emission is simplified using two template light sources, one for cascades and another for tracks from muons and taus. Recall that the end result of cascade simulation by `cmc` is a set of point-like sub-cascades that approximates light emission from the complete electromagnetic and hadronic particle showers. The light emitted from each sub-cascade is largely focused along the Cherenkov angle, and the exact angular profile is found from a template based on Geant4 simulations of a 1 GeV electromagnetic cascade [151]. Taking advantage of the high degree of linearity in electromagnetic cascades, the total number of photons is scaled by the energy of the cascade.

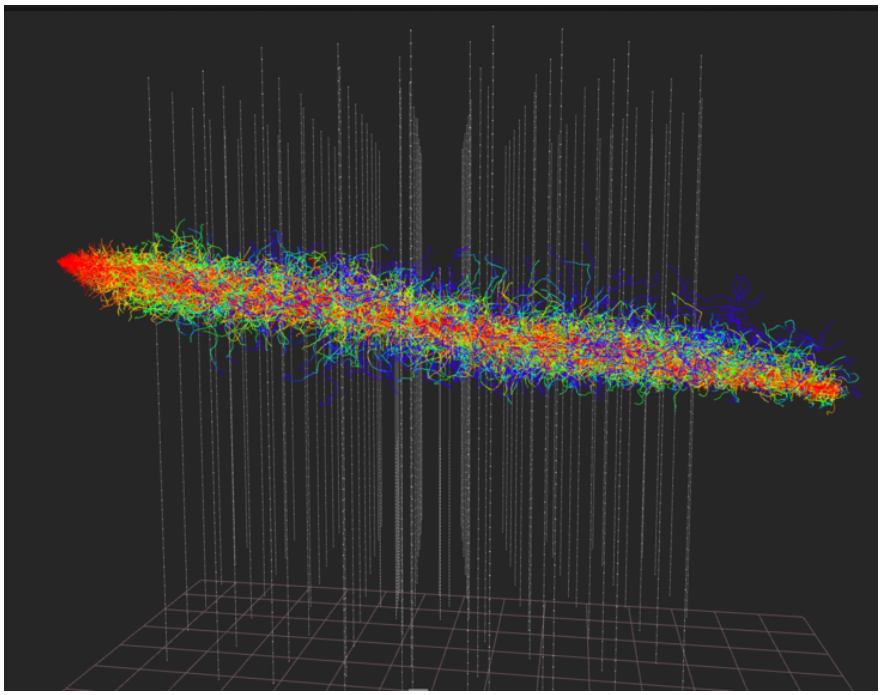

Figure 5.2: Light from a through-going muon as simulated by the direction photon propagation program CLSim [201]. Colored lines show the paths of individual photons as they are tracked through ice. The color shows the delay time with respect to a photon that travels unscattered; red is earlier and blue is later.

For muons and taus propagated by `mmc`/`PROPOSAL`, light is already accounted for from the cascades produced in discrete losses, but additional light is also emitted between each of the discrete energy losses where the energy loss is treated as continuous. In each track segment between discrete losses, the angular light emission profile is again calculated from a template based on Geant4 simulations [200].

With both light emission templates, one can then tackle the larger problem of propagating light through the inhomogeneous glacial ice at the South Pole. Even if the ice were perfectly homogeneous, analytical methods are ill-suited to the problem. This is because the scattering length is comparable to the string spacing of $\sim 125\,\mathrm{m}$. It is too short to consider straight-line paths, but too long to use the approximations of diffusion theory. Monte Carlo methods to directly track individual photons through multiple scatterings are the only way to make precise predictions for the expected number of photoelectrons recorded by a DOM.

5.3.1 Direct Photon Propagation

The Monte Carlo simulation of light propagation in heterogeneous ice was first pioneered in the software package `Photonics` [202]. Since then, it has been supplanted by two newer codes, Photon Propagation Code (PPC) [203] and CLSim [204]. Since optical photons do

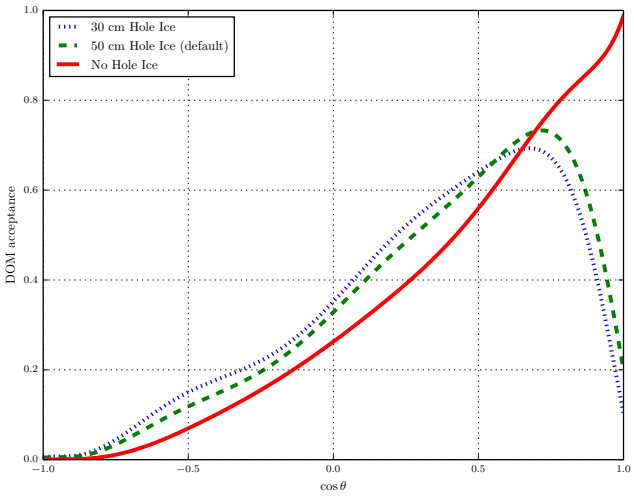

Figure 5.3: The sensitivity to detect photons as a function of the incoming photon angle, measured with respect to the vertically downward direction. The lab-measured sensitivity without any hole ice effects is shown in red. The sensitivity is normalized to a value of 1 at $\cos\theta = 1$. An effective scattering length of 50 cm in hole ice is simulated with direct photon propagation in order to derive the default angular sensitivity model (dashed green). The angular sensitivity for a 30 cm effective scattering length in hole ice is also shown (dotted blue).

not interact with each other, they are easily tracked in parallel, and both of these codes take advantage of the highly parallel architecture of graphics processing units (GPUs) to greatly accelerate the simulation.

To start, a photon is drawn from a Cherenkov wavelength spectrum from either a cascade or muon template light source. Using the SPICE models of Sec. 4.3 to describe scattering and absorption, each photon is followed through multiple scatterings until either it reaches the collection area of a DOM or, much more likely, is absorbed. For example, all photons tracked for a through-going muon are shown in Fig. 5.2. Because a photon reaching a DOM is so rare, a common practice is to oversize the DOM by a factor of 5 transverse to its direction of travel, increasing the speed of the simulation by a factor of 25. After reaching the DOM, the lab-measured wavelength-dependent acceptance and angular sensitivity of the DOM are used to generate a series of photoelectrons at the PMT cathode. The angular sensitivity includes a correction to account for increased scattering in the hole ice surrounding the DOM. The angular correction is derived from simulations of direct photon propagation in a column of ice with a uniform 50 cm effective scattering length throughout. The DOM angular sensitivity with and without the hole ice correction is shown in Fig. 5.3. The scattering length in the hole ice column is not well known, so the angular sensitivity correction for a 30 cm effective scattering length is also shown in Fig. 5.3. Hole ice will be an important systematic

uncertainty covered in Sec. 7.3.2.

Direct photon propagation was also a key tool used to derive the SPICE models [175]. The propagation of light from LED flashers was simulated using PPC, and by comparing the simulated photon arrival distributions to those observed by DOMs in dedicated flasher calibration runs, the absorption and scattering coefficients in 10 m layers were iteratively varied until producing the best agreement between data and simulation.

5.3.2 Tabulated Photon Propagation

For cases when direct photon propagation is too slow, an alternate approach is to store the expected photoelectron timing distribution at a DOM for all possible configurations of the light source and DOM in a look-up table that can be quickly accessed. This may seem to be an intractable problem since in the most general case, the table must have 9 dimensions, 3 for the position of the source, 3 for the position of the DOM, 2 for the orientation of the light source, and 1 for time. However, one can use the approximate horizontal translational and azimuthal symmetry of the ice to reduce the number of dimensions to 6, depth of the source in ice, zenith angle of the source, displacement vector of the DOM from the source, and time. The downside of the approach is that two subtle effects, ice layer tilt and anisotropic scattering, violate these symmetry assumptions and are ignored. Beyond simulation, photon tables are also used for event reconstructions, where the connection between expected photoelectron distribution and light source provided by tables can be inverted to infer the light source configuration that best matches the observed photoelectron distribution in every DOM, a topic that will be further described in Sec. 6.2.

The production of photon tables was first done using `Photonics` to predict and store the expected photon flux in a multi-dimensional histogram structure. However, this was prone to binning issues and results were often inaccurate at great distances where statistics were low and bins were large. Currently, the best approach is to fit the photoelectron distribution from CLSim or PPC to a tensor product B-spline surface [205]. This has the advantage that the photoelectron distribution is a smooth function of all 6 coordinates and unphysical fluctuations from limited statistics can be alleviated through the use of regularization.

5.4 DOM Simulation

The output of photon propagation is a series of photoelectrons produced at the cathode of PMTs in IceCube. At this point, the response of the DOM needs to be simulated. The first part of the simulation, called `PMTResponseSimulator` [206], models the pulses produced by the PMT, accounting for both the transit time and charge fluctuations using the parameters measured for every PMT. It also accounts for saturation, pre-, late-, and after-pulsing as described in Sec. 4.2.1. The second part of the simulation, called `DOMLauncher` [207], models the PMT discriminator, digitization, and local coincidence system as described in Sec. 4.2.2, taking into account all known time delays and distortions of the electronics. The last piece

of DOM simulation is modeling event triggers, as described in Sec. 4.2.3. Although at this point the set of simulated digitized waveforms produced by the DOMs is formatted in the same way as real data, the same software to perform triggering at ICL is not used, which was optimized for a continuous high-volume stream of data. Instead trigger simulation is handled by the `trigger-sim` program [208], which packages the raw digitized data from the DOMs into events. At this point, simulation and data are identical and can be processed further in the same manner in order to compare the distributions of observable quantities in simulation and data.

Chapter 6

Flavor Composition of Astrophysical Neutrinos above 35 TeV

In 2012, two cascades were observed in IceCube with a visible energy of $\sim 1\,\mathrm{PeV}$, providing the first hint of high-energy astrophysical neutrinos [209]. The events were found in an analysis intended to search for cosmogenic neutrinos above 10 PeV, so the high energy threshold of the event selection made it difficult to draw further conclusions about their origin.

In order to determine if the events originated from an astrophysical flux that extended to lower energies, a new event selection called the High Energy Starting Event (HESE) analysis was devised that sought to capture both cascade and track events interacting within IceCube. In two years of data from May 2010 to May 2012, 28 events were found, establishing evidence for an E^{-2} flux of diffuse astrophysical neutrinos above $\sim 100\,\mathrm{TeV}$ with a significance of 4σ [7]. The addition of another year of data further increased the number of events to 36 (28 cascades and 8 tracks) and boosted the significance to 5.7σ [78].

Although the HESE analysis firmly established the existence of a diffuse astrophysical flux above atmospheric backgrounds, many questions remained. No significant clustering or correlations with known astronomical objects were found in the arrival direction or time of the 36 events. The energy spectrum of the astrophysical flux was also highly uncertain. A power-law fit of the astrophysical flux found a spectral index of 2.3 ± 0.3, bridging the gap between the lowest value of 2 expected from Fermi acceleration and the observed index of cosmic rays. The analysis also provided little constraint on the magnitude of the prompt atmospheric flux, leaving open doubts that it could explain a large fraction of the observed events.

Lastly, it was untested whether the observed 28 cascades and 8 tracks were consistent with the $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ flavor ratio expected from complete pion decay and standard neutrino oscillations assumed in the analysis. As elaborated in Sec. 3.2.2, small deviations from a $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ flavor composition are expected for alternate neutrino production mechanisms in astrophysical sources, yielding a clue as to their identity. Large deviations in the flavor composition from $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ must be caused by physics beyond the Standard Model, a

possibility attracting great interest and many independent analyses from outside the IceCube collaboration [210, 211, 212].

To address all of these questions, it is necessary to lower the energy threshold of the event selection even further. Here an extension of the HESE analysis will be documented that enabled the astrophysical flux to be probed down to an energy of 35 TeV, and the first limits to be set on its flavor composition [213].

6.1 Event Selection

Three years of data recorded from May 2010 to May 2013 with a detector live time of 974 days are considered here. While developing the event selection, 90% of data was kept blind. The driving goal of the event selection is to reject as many atmospheric muons as possible while maintaining neutrinos that interact within IceCube. There are three steps in the selection, which follow closely the methods of the original HESE analysis.

1. Select high-energy events with more than 1500 photoelectrons (PE) detected.
2. Reject atmospheric muon background with an outer-layer veto.
3. Classify cascades and tracks and reconstruct their energy and direction.

6.1.1 Filters and NPE Cut

Because there is limited bandwidth for data transfer from the South Pole by satellite, a variety of filters are used to reduce the rate of triggered events while maintaining events of interest for many different physics analyses. Of primary interest to this selection is the extremely high energy (EHE) filter, which selects events where more than 1000 photoelectrons were detected. However, over each of the three years of detector operation considered here, the operation of the EHE filter has changed, most notably the calibration procedure to calculate the number of photoelectrons at the online filter at the South Pole and the procedure to split coincident events where more than one cosmic ray shower is seen in the same trigger window. Normally it would be necessary to simulate the change in the EHE filter for each year of data, but instead it is more convenient to use the calibration and event splitting for the 2012 season and enforce a higher threshold of 1500 PE on both simulation and data. This safely avoids any threshold effects caused by year-to-year calibration differences in the EHE filter while the 2012 detector simulation can be used for all years of data.

6.1.2 Outer-Layer Veto

To reject down-going atmospheric muons, an outer layer of DOMs is designated to detect the light emitted as they enter the detector. Originally designed for the HESE analysis, the geometry of the outer-layer veto is illustrated in Fig. 6.1. A DOM is within the veto layer if it is

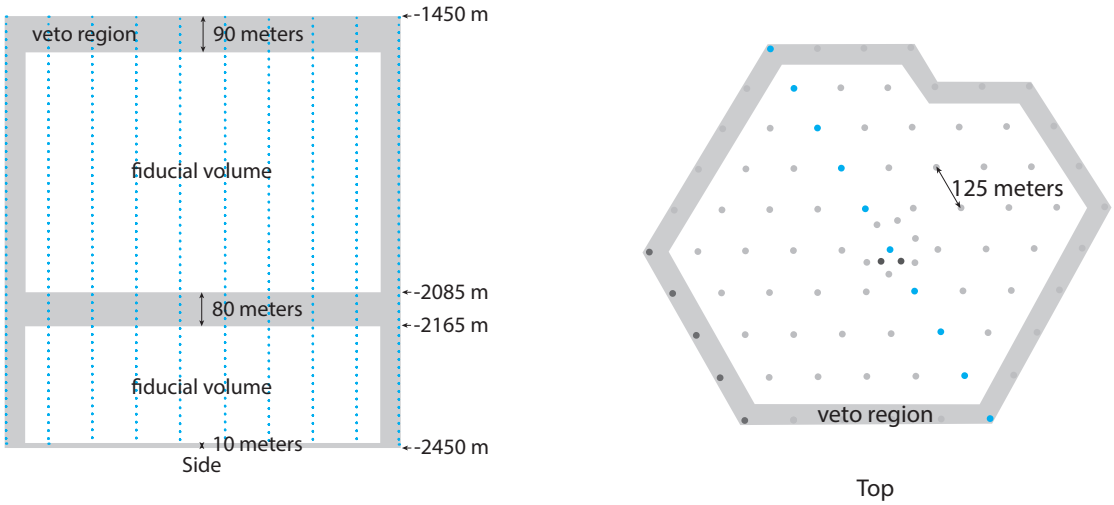

Figure 6.1: The veto layer (grey) as viewed from the side (left) and top (right) [7].

1. On an outermost string
2. Inside a layer 90 m thick at the top of the detector or 10 m thick at the bottom of the detector
3. Between 2085 m and 2165 m in depth

The 90 m thickness of the top layer is needed to increase the veto efficiency for vertically down-going events that slip through "passing lanes" between strings. The region between 2085 m and 2165 m in depth is just below the dust layer and rejects muons where light was not detected in the outermost strings due to high absorption. The fiducial volume contained within the outer-layer has an effective target mass of 0.4 gigatons, corresponding to roughly half of the instrumented volume of IceCube.

With an outer-layer defined, conditions on photoelectrons detected in the layer must be defined that determine if an event is vetoed. First of all, an event starting time is defined. A sliding time window of 3 μs is passed over all time-ordered photoelectrons from every DOM in hard local coincidence, and the first time when the total number of photoelectrons (NPE) within the window first exceeds 250 PE is defined as the starting time of the event. A center of gravity of the resulting first 250 PE is also computed,

$$\mathbf{x}_{\mathrm{COG}} = \frac{\sum_i \mathbf{x_i} q_i}{\sum_i q_i}, \tag{6.1}$$

where $\mathbf{x}_i$ is the position vector of each recording DOM and q_i is the total number of photoelectrons detected within the 3 μs time window for each DOM. Then the number of photoelectrons are counted in the veto layer up to the event start time +50 μs, excluding DOMs that are more than 3 μs travel time from the center of gravity according to the speed of light

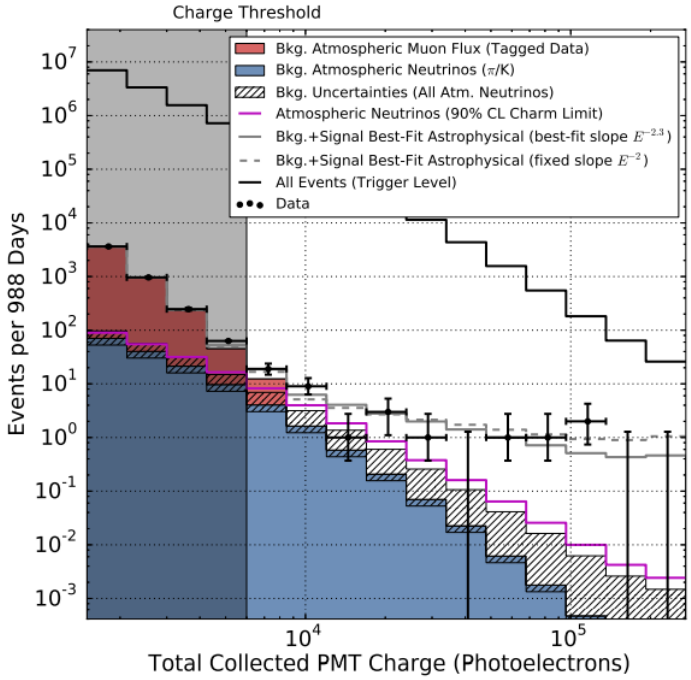

Figure 6.2: The distribution of the number of photoelectrons detected per event after the outer-layer veto [78]. The distribution of all events at trigger level before the outer-layer veto is shown with the solid black curve, illustrating the factor of 10^3 reduction in event rate. Monte Carlo simulation of conventional atmospheric neutrinos according the HKKMS flux calculation is shown in blue. Prompt atmospheric neutrinos at a level of 3.8 times the ERS calculation are shown by the magenta line, corresponding to a previous IceCube upper limit at 90% confidence level [131]. Astrophysical neutrinos with an energy spectrum E^{-2} ($E^{-2.3}$) are shown as solid (dashed) lines corresponding to the best-fit results of [78]. The contribution of atmospheric muons predicted from the tagged inner-layer veto control sample is shown in red, stacked on the atmospheric neutrino contribution. The excess of high NPE events from an astrophysical flux is already visible. In the HESE analysis, only events outside the shaded region above 6000 PE were selected.

in ice. If 3 or more DOMs or more than 3 total photoelectrons in any DOM satisfy this condition, the event is vetoed. This threshold is chosen to be high enough that incoming muons are reliably vetoed, but small enough that there is a negligible chance that a neutrino would be vetoed by random noise.

6.1.3 Muon Background Estimation

The outer-layer veto reduces the rate of atmospheric muons by a factor of 10^3, as shown in Fig. 6.2, but a large contribution still remains. To estimate the size of this remaining

background, CORSIKA or MuonGun simulations would usually be needed, but it can be extremely time consuming to generate enough events at high energies, and the resulting predictions are subject to large systematic uncertainties. As the dominant background of this analysis, it is very desirable to have a data-driven estimate for atmospheric muon background to corroborate expectations from simulation.

To characterize the remaining penetrating muon background, a control sample can be created by designating a second inner-layer veto nested inside the outer-layer veto. The inner-layer veto contains the second outermost strings and cap layers of 90 m and 10 m thickness at the top and bottom, respectively. The inner fiducial volume is geometrically similar to the full fiducial volume but is smaller by a factor of 2.1. Using the outer layer to tag incoming muons with high efficiency, a control sample of muons that pass the inner-layer veto can be obtained. The rate of control sample events can then be scaled up by the volume ratio of 2.1 to obtain the rate of penetrating muons in the full detector.

The geometric similarity of the inner fiducial volume can only be approximate, but the volume scaling does correctly predict the distribution of many variables in regions dominated by muon background. For example, the muon background prediction from the control sample is shown in the distribution of NPE in Fig. 6.2.

6.1.4 Cascade and Track Classification

After the outer-layer veto, events are classified as cascades or tracks. The two topologies are distinguished by a likelihood analysis of the first photoelectron times recorded in every DOM. This is done by considering the first possible arrival time of a photon traveling the shortest geometric distance between the Cherenkov light source and a DOM using the group velocity of light in ice. For a point-like cascade, light expands on a spherical wavefront and the geometric arrival time at DOM i can be written as

$$t^{\text{geo}}_{i,\text{casc}} = t_0 + \frac{n}{c} r_i \tag{6.2}$$

where t_0 is the time of the cascade, $r_i = |\mathbf{x}_i - \mathbf{x}_0|$ is the distance from the DOM to the cascade, and n is the index of refraction in ice. The Cherenkov cone is entirely neglected in this approach, and the direction of the cascade cannot be determined from the geometric first arrival times alone. For a track, a more detailed consideration of the geometry of the Cherenkov light cone is needed taking into account the direction of the track $\hat{\mathbf{n}}_0$. In this case, the geometrical arrival time is

$$t^{\text{geo}}_{i,\text{track}} = t_0 + \frac{\hat{\mathbf{n}}_0 \cdot \mathbf{r}_i + |\mathbf{r}_i - (\mathbf{r}_i \cdot \hat{\mathbf{n}}_0)\hat{\mathbf{n}}_0| \tan\theta_C}{c} \tag{6.3}$$

where $\mathbf{r}_i = \mathbf{x}_i - \mathbf{x}_0$ is the vector from the track's point of origin at time t_0 to a DOM and θ_C is the Cherenkov angle [182].

Due to scattering and absorption in ice, light is randomly delayed from the geometric first arrival time. The first arrival time distribution, $p_1(\Delta t_i, d_i)$, at a DOM i can be parameterized

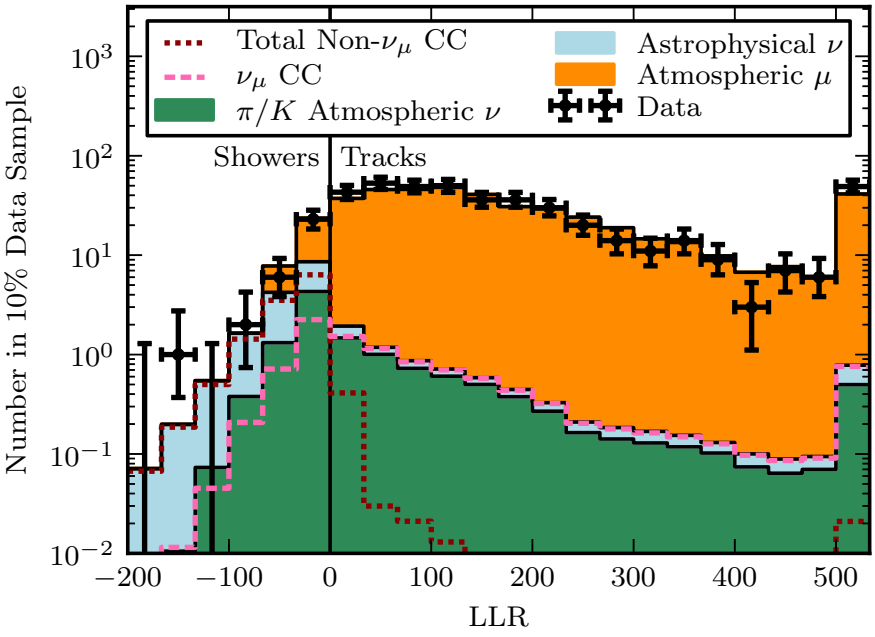

Figure 6.3: The log-likelihood ratio (LLR) between cascade and track reconstructions for veto-passing events with more than 1500 photoelectrons. Events with LLR > 0 are classified as tracks, and events with LLR < 0 are classified as cascades. The solid-colored distributions of muons (gold), atmospheric neutrinos (green), and astrophysical neutrinos (blue) are stacked. The distribution of muons is determined from the muon control sample and is mostly track-like. The dotted lines show the total distributions of ν_μ charged-current events (pink) and all non-ν_μ charged-current events (maroon) and are not cumulative with the solid-colored distributions. The last bin contains all overflow events with LLR > 500. Error bars are 68% Feldman-Cousins intervals [217] and show upper limits for bins with no events.

in terms of the time residual, $\Delta t_i = t_i - t_i^{\text{geo}}$, where t_i is the observed first photoelectron time, and d_i is the distance from the DOM to the point of where light was emitted. For cascades, d_i is simply the distance to the cascade vertex, whereas for tracks it is given by $d_i = |\mathbf{r}_i - (\mathbf{r}_i \cdot \hat{\mathbf{n}}_0)\hat{\mathbf{n}}_0|$. A feature of the arrival time distribution is that both its mean and variance increase with distance due to the increased chance of scattering in ice. A simple analytic parameterization of $p_1(\Delta t_i, d_i)$ was first found by Pandel [214, 215, 216] and corresponds to a modified gamma distribution.

With the probability distribution for time residuals specified, a log likelihood function can then be formed,

$$-\ln L(t_0, \mathbf{x}_0[, \hat{\mathbf{n}}_0]) = \sum_i -\ln p_1(\Delta t_i, d_i), \tag{6.4}$$

which can be minimized to find the best estimate of the time and position for a cascade, or additionally the direction for a track. For veto-passing events, this reconstruction is done

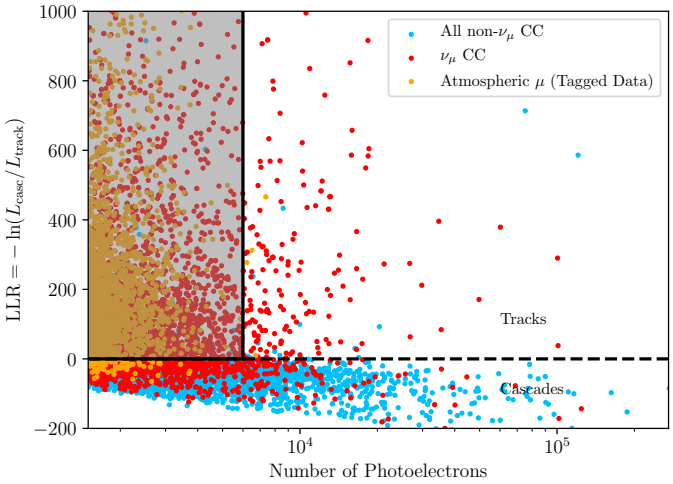

Figure 6.4: The two-dimensional distribution of LLR $= -\ln(L_{\rm Cascade}/L_{\rm Track})$ and NPE for simulated neutrinos (blue) and tagged muons from the control sample (gold). Starting tracks from charged-current ν_μ interactions are shown in red. The dashed line at LLR $= 0$ divides cascades from tracks. The upper left region is excluded by the event selection due to the high atmospheric muon background.

individually for a cascade and track hypothesis, yielding two log likelihood values $-\ln L_{\rm casc}$ and $-\ln L_{\rm track}$. A log-likelihood ratio is then formed in order to discriminate cascades and tracks, LLR $= -\ln L_{\rm casc}/L_{\rm track}$. Larger values indicate more track-like events, as shown in the distribution of LLR in Fig. 6.3. A cut of LLR $= 0$ was chosen to distinguish cascades and tracks. When applied to simulated veto-passing events with more than 1500 NPE weighted to the HKKMS atmospheric flux and an E^{-2} astrophysical neutrino flux with a flavor composition of $\left(\frac{1}{3}:\frac{1}{3}:\frac{1}{3}\right)_\oplus$, this choice correctly identifies ν_μ charged-current events with an efficiency of 70%. The remaining 30% that are missed are largely due to events at the edge of the fiducial volume where the outgoing muon escapes. Furthermore, 95% of all events classified as tracks by this method are ν_μ CC events.

Tracks are subject to larger atmospheric backgrounds than cascades. Atmospheric muons tend to produce tracks, although cascade-like events can still be produced when a muon undergoes a large stochastic energy loss inside the detector. Additionally, most atmospheric neutrinos create track-like ν_μ CC events. As a result, lower energy astrophysical neutrinos are more easily accessed in the cascade detection channel. Thus, all cascade events above 1500 PE are kept while a higher NPE threshold of 6000 PE is enforced for all tracks, corresponding to the same threshold from the HESE analysis used for all events. The result of this selection can be thought of as a 2D cut in the plane of LLR and NPE as shown in Fig. 6.4.

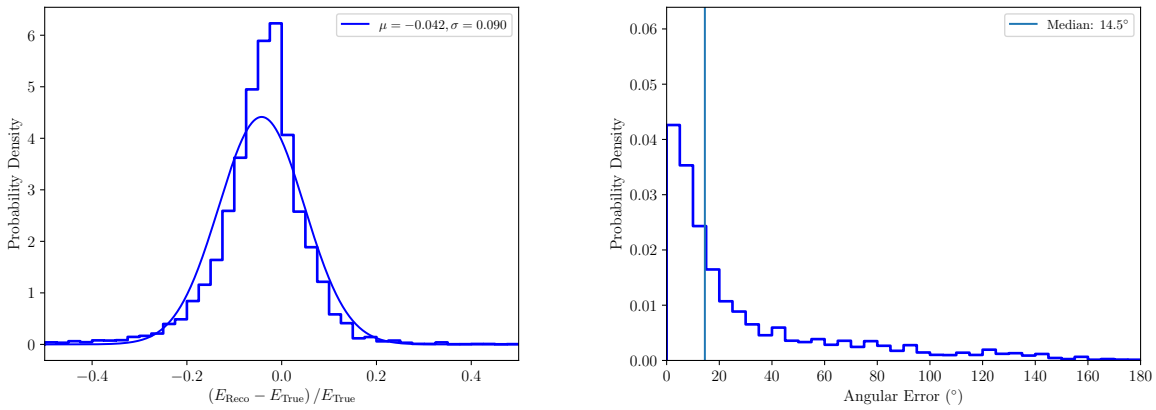

Figure 6.5: Left: The distribution of the relative error between the reconstructed and true visible energy for simulated events. A Gaussian fit to the distribution shows a bias of -4.2% and standard deviation of 9%. The simulation uses a DOM optical efficiency 95% of the nominal value assumed by the reconstruction, causing a non-zero bias. Right: The distribution of the angle between the reconstructed direction and true neutrino direction for simulated events. The median angle is 14.5°.

6.2 Reconstructions

Techniques using the first arrival time per DOM are effective for discriminating cascades and tracks in a simple and robust manner; however, they are unable to provide a cascade direction nor can they provide an estimate of energy for either cascades or tracks. More advanced techniques exploiting all information contained in the full photoelectron timing distribution per DOM are required.

6.2.1 Cascades

As described in Sec. 5.3.2, the photoelectron timing distribution at a DOM can be calculated by Monte Carlo methods and fit to a spline surface so that it is a smooth function of the time, position, and orientation of an electromagnetic cascade. Once the observed photoelectron distribution of an event is split into discrete time bins each containing k_i photoelectrons, the spline can be used to calculate the expected number of photoelectrons per bin, λ_i. For electromagnetic cascades, the expected number of photoelectrons per bin scales linearly with energy, so that $\lambda_i = \Lambda_i E + \rho_i$, where Λ_i is the template for a 1 GeV cascade and ρ_i is an additional term that reflects noise in the PMT. A Poisson log likelihood for each

event can then be constructed,

$$\begin{aligned} -\ln L(t_0, \mathbf{x}_0, \hat{\mathbf{n}}_0, E) &= \sum_i \lambda_i - k_i \ln \lambda_i + \ln k_i! \\ &= \sum_i \Lambda_i E + \rho_i - k_i \ln (\Lambda_i E + \rho_i) + \ln k_i!, \end{aligned} \tag{6.5}$$

and then minimized to find the best-fit seven parameters that describe the cascade – time, t_0, position, $\mathbf{x}_0$, direction, $\mathbf{n}_0$, and electromagnetic-equivalent energy, E. In the absence of noise, Eq. 6.5 can be easily minimized analytically to find the energy of a cascade, $E = \sum_i k_i / \sum_i \Lambda_i$, when the remaining parameters are fixed. In practice, this can greatly improve the convergence and speed of the overall minimization when using a numerical algorithm like MIGRAD from the MINUIT package [218] since it effectively reduces the number of parameters to six. This advantage remains even in the presence of noise since it straightforward to minimize Eq. 6.5 over energy using Newton's method [168].

The performance of the cascade reconstruction can be tested on simulated events with an E^{-2} neutrino spectrum, as shown in Fig. 6.5. For the cascade sample obtained here, the electromagnetic-equivalent energy can be reconstructed with a 9% standard deviation in the fractional error. The direction of the original neutrino can be reconstructed with a median angular error of 14.4°. These performance metrics do not account for systematic uncertainties due to the imperfect modeling of ice properties and DOM optical efficiency. Figure 6.6 shows example photoelectron distributions from the highest energy 2 PeV cascade found in the sample. The observed photoelectron distributions in DOMs do not always line up perfectly with the best-fit predictions, indicating that additional systematic errors are present. Without a calibration source of cascades with known direction, verifying the angular reconstruction of cascades is a challenging task that will be revisited in Sec. 7.2.3.

6.2.2 Tracks

For tracks, the analogous approach of fitting the observed photoelectron distribution to a single template based on the average light yield of a muon is very limited since muons undergo significant event-to-event fluctuations in their energy loss profile. Instead, it is desirable to reconstruct the energy loss as a function of distance through the detector, and this can later be analyzed to infer the energy of the muon if the track is through-going, or the energy of the neutrino if both the hadronic cascade and muon of a starting track are observed.

The energy loss profile of the track can be modeled as a sequence of electromagnetic cascades, uniformly spaced by 10 m and each with an energy E_j. The sequence extends along the track within a $(1200\,\mathrm{m})^3$ box centered on IceCube. The expected number of photoelectrons in a time bin on a DOM can then be found by summing the contributions from each cascade, $\lambda_i = \sum_j \Lambda_{ij} E_j + \rho_i$, where Λ_{ij} is the expected number of photoelectrons for a template 1 GeV cascade for each point along the track. This leads to an expression for

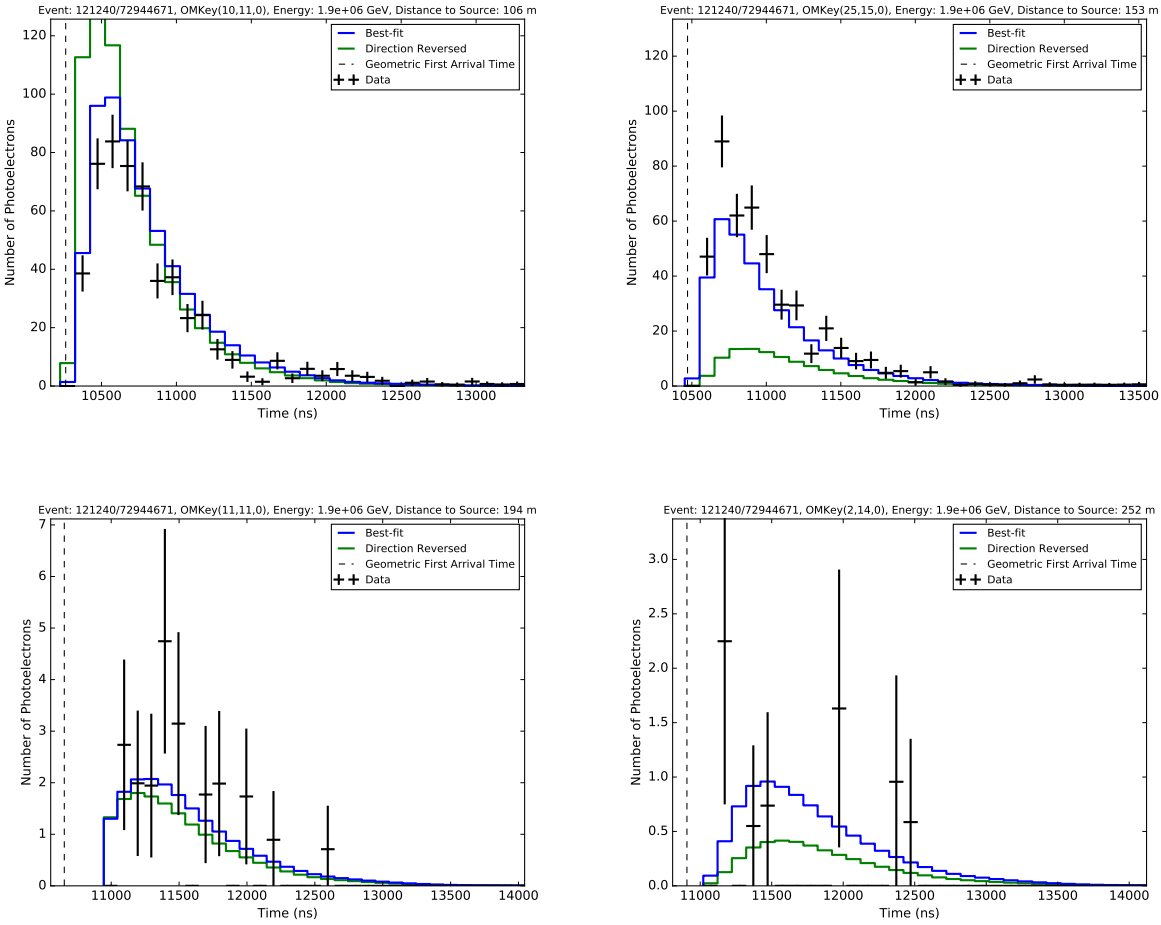

Figure 6.6: Example photoelectron distributions in DOMs from the most energetic 2 PeV cascade observed in the data sample. The observed distributions are shown in black while the best-fit distribution is shown in blue. Reversing the direction of the best-fit cascade produces the photoelectron distribution shown in green, illustrating how cascade direction can be recovered from the photoelectron distributions. The distributions in the upper right panel correspond to a DOM close to the Cherenkov and show the largest difference between the best-fit and reversed directions.

the Poisson log likelihood,

$$
\begin{aligned}
-\ln L(t_0, \mathbf{x}_0, \hat{\mathbf{n}}_0, E_0, \ldots, E_j, \ldots) &= \sum_i \lambda_i - k_i \ln \lambda_i + \ln k_i! \\
&= \sum_i \left(\sum_j \Lambda_{ij} E_j + \rho_i \right) - k_i \ln \left(\sum_j \Lambda_{ij} E_j + \rho_i \right) + \ln k_i!.
\end{aligned}
\tag{6.6}
$$

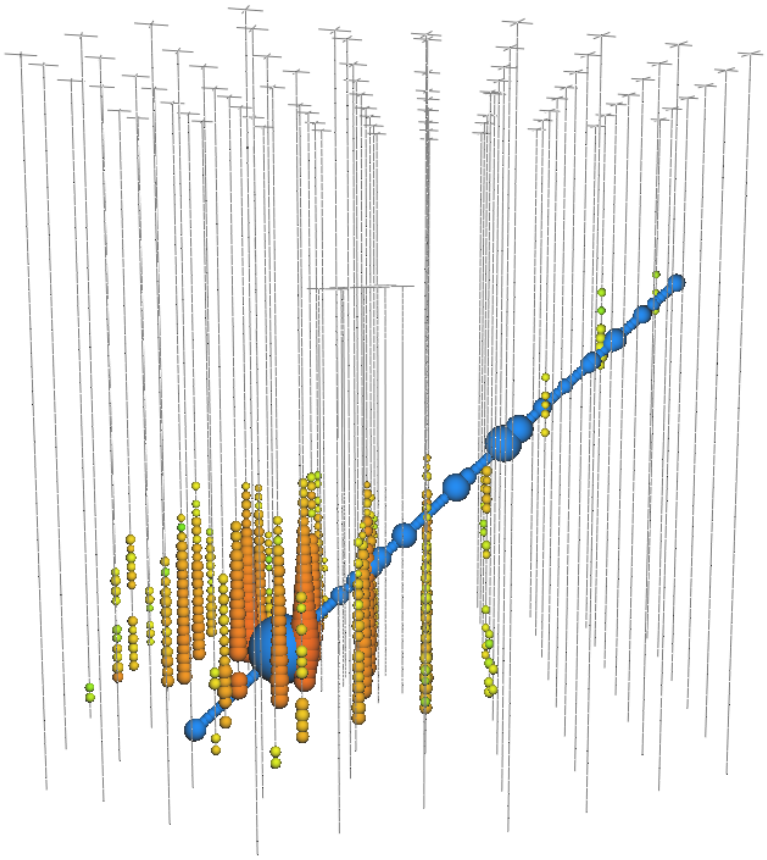

Figure 6.7: An observed up-going starting track event. The observed number of photoelectrons per DOM is illustrated with colored spheres ranging from red to green, corresponding to the first photoelectron time. The blue spheres correspond to the reconstructed sequence of electromagnetic cascades, and their size is proportional to the reconstructed energy of each cascade on a logarithmic scale. The largest sphere likely corresponds to the hadronic cascade at the neutrino interaction vertex. Adding the energies of reconstructed cascades, a total electromagnetic-equivalent energy of 253 TeV was deposited in the detector.

Minimizing this equation to find the energy of each cascade is an example of a common problem encountered in particle physics called an "unfolding." Even when no noise is present, Eq. 6.6 can not be minimized analytically for each cascade's energy, and a pre-conditioned conjugate gradient algorithm is used to perform the unfolding [168]. Because of this expensive step, minimizing Eq. 6.6 over the remaining six parameters of the track is very time consuming and prone to convergence issues. It is not practical to run the full reconstruction over the millions of simulated events needed for the final analysis of data.

However, the time, position, and direction of tracks can be reconstructed well without

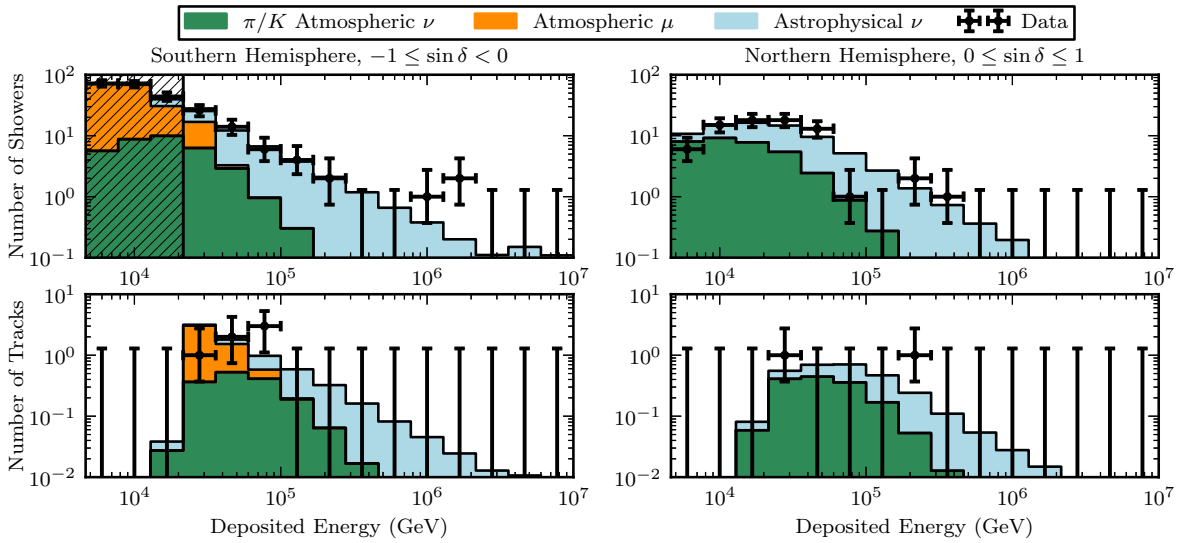

Figure 6.8: The best-fit deposited energy distributions for showers and tracks, divided into the southern (down-going) and northern (up-going) samples, assuming a power-law astrophysical flux with $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_{\oplus}$ flavor ratio at Earth. The contribution of charm atmospheric neutrinos is not shown since the best-fit scaling of the ERS flux is 0. Showers in the southern sky below 20 TeV are dominated by muons and excluded; however, the prediction from the control sample measurement is shown in this region. Though not shown, 4 bins in declination (δ) are used in the fit. The best-fit parameters are found in Tab. 6.1.

using the full photoelectron distribution. For example, the fit of the first photoelectron arrival times from Eq. 6.4 is able to achieve a median angular error of 2.0° for the sample of starting tracks obtained here. The six parameters are fixed to their values from this reconstruction, and only the unfolding step is done to obtain the energy loss profile. An example of the reconstruction is shown in Fig. 6.7.

For starting tracks, further analysis of the reconstructed energy loss profile is needed to reconstruct the energies of the hadronic cascade and muon. This is beyond the scope of the present analysis but will be revisited in Sec. 7.2.2. Instead, only the sum of the energy losses, or deposited energy, in the detector will be used as a proxy for neutrino energy. The deposited energy is only a lower bound on the energy of the neutrino since a muon escapes the detector with an unknown amount of energy.

6.2.3 Final Event Sample

After energy and directional reconstruction, one last selection is applied to the cascade sample. Because there was insufficient simulation of atmospheric muons below 10 TeV, the background dominated region with reconstructed energy less than 20 TeV and reconstructed zenith angle less than 90° was excluded from further analysis. After unblinding all 974 days of data, 129 cascades and 8 tracks were found.

To measure the flavor composition and energy spectrum of the astrophysical flux, a binned maximum likelihood fit is performed jointly on the two-dimensional distributions of reconstructed energy and cosine zenith of the cascade and track samples. The maximum likelihood method is reviewed in Appendix A. Reconstructed energy and cosine zenith are chosen as observables are chosen since they best reflect the differences in the energy and zenith distribution that enable astrophysical and atmospheric neutrinos to be distinguished, as detailed in Ch. 3. The distributions were binned in reconstructed energy from $10^{3.666\ldots}$ GeV to 10^7 GeV in logarithmic bins of width $\Delta \log_{10} E = 0.222\ldots$. In reconstructed cosine zenith, 4 equal-sized bins were used. To aid astronomical interpretation, it is conventional to represent neutrino directions in terms of celestial declination (δ) and right ascension (RA) in equatorial coordinates. Because of IceCube's location at the South Pole, there is a simple relationship between declination and zenith, $\delta = \theta - 90°$, or $\sin\delta = -\cos\theta$. The energy distributions of cascades and tracks are shown in Fig. 6.8, broken in to down-going samples (originating from the Southern Hemisphere, $-1 \leq \sin\delta \leq 0$) and up-going samples (originating from the Northern Hemisphere, $0 < \sin\delta \leq 1$.)

6.3 Model and Likelihood Fit

By comparing the distribution of observed events and the predicted distribution of a model in Fig. 6.8, a maximum likelihood fit can be used to find the best estimates of the flavor composition and energy spectrum of the astrophysical flux. This method relies on an accurate model of not only astrophysical neutrinos, but atmospheric neutrinos and muons as well. A model comes with systematic uncertainties, which may sometimes be described with additional "nuisance" parameters. In this section, the models and parameters used in the maximum likelihood fit will be described in more detail. All parameters are summarized in Tab. 6.1

6.3.1 Neutrinos

The astrophysical neutrino flux is of primary interest, and the flux of each flavor is parameterized with an isotropic angular distribution and power-law energy spectrum,

$$\Phi_\alpha = 3\Phi_0 f_{\alpha,\oplus} \left(\frac{E}{100\,\text{TeV}}\right)^{-\gamma}, \tag{6.7}$$

where $f_{\alpha,\oplus}$ is the fraction of each flavor at Earth, γ is the spectral index, and Φ_0 is the average flux of neutrinos and antineutrinos per flavor at 100 TeV. An equal flux of neutrinos and antineutrinos is assumed, which carries the caveats described in Sec. 3.2.3.

The conventional atmospheric neutrino flux is described by the HKKMS calculation extrapolated above 10 TeV and corrected for the H3a cosmic-ray model as described in Sec. 3.3.1. An overall scaling factor of the HKKMS flux is the only nuisance parameter for conventional atmospheric neutrinos. Additional uncertainties that change the energy and

angular dependence of the flux, such as the relative production of pions and kaons, index of the cosmic-ray spectrum, and differing models of the cosmic-ray knee were found to negligibly impact results. No prior on the overall scaling factor is used since it is well constrained by the data alone.

The prompt neutrino flux is described by the ERS calculation and also incorporates a correction to account for the H3a cosmic-ray model. An overall scaling factor of the ERS flux is sufficient to describe its uncertainties, and no prior was placed on its magnitude so as to be as conservative as possible about the prompt atmospheric contribution to the observed events.

To describe both atmospheric and astrophysical neutrinos, `neutrino-generator` simulation is used. Equal numbers of each neutrino flavor were generated according to a E^{-1} spectrum, and 279,295 simulated events survive the event selection. Event weights are calculated according to Eq. 5.2 using the combined atmospheric and astrophysical flux. This enables the same set of simulated events to describe both astrophysical and atmospheric neutrinos.

Since `neutrino-generator` simulation does not include accompanying muons from cosmic-ray air showers, it is necessary to correct for the atmospheric self-veto probability. The analytic calculation of [141] is used to correct both the conventional and prompt atmospheric fluxes. To approximate the conditions that cause an event to be vetoed, it is required that no muons exist in the air shower that survive to a depth of 1950 m (the center of IceCube) with an energy above 1 TeV.

The muon energy threshold is the main systematic uncertainty associated with this approach since it only approximates the details of the veto conditions. To confirm the choice of 1 TeV, `CORSIKA` simulation of atmospheric neutrinos was performed that included accompanying muons and modeled the full detector response. The true zenith distribution of down-going neutrinos above 10 TeV from the `CORSIKA` simulation was compared to the zenith distribution obtained from `neutrino-generator` simulation weighted to the HKKMS flux suppressed by the self-veto probability of the analytic calculation. As shown in Fig. 6.9, the two distributions show good agreement.

6.3.2 Muons

For atmospheric muons, the control sample of tagged, inner-veto passing muons provides an estimate of the overall rate of atmospheric background muon events, but because of its small sample size, it is insufficient to describe the energy and angular distribution. `CORSIKA` simulation would provide the best available description of muon background, but it is difficult to generate a sufficient sample size. Because most events that pass the outer-layer veto are single muons, `MuonGun` simulation of the single muon flux surrounding IceCube can be used. The muon flux model used is derived from the H3a cosmic-ray model and the SIBYLL2.1 hadronic interaction model.

To describe uncertainties on atmospheric muons, two nuisance parameters are used – two separate scaling factors for events in the track and cascade samples. Two scaling fac-

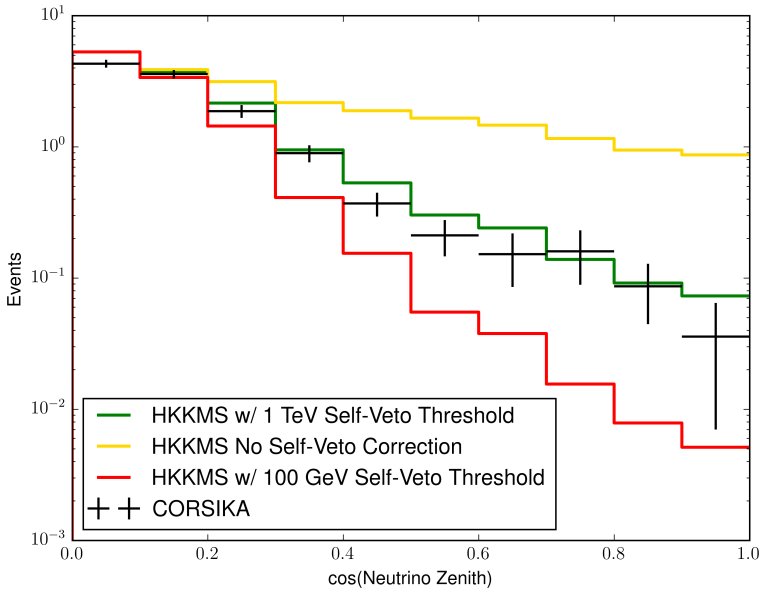

Figure 6.9: The true zenith distribution of down-going neutrinos. The yellow line shows `neutrino-generator` simulation of the HKKMS flux with no atmospheric self-veto probability. The green line shows the same `neutrino-generator` simulation but correcting the HKKMS flux using the analytic self-veto probability calculation with a 1 TeV muon energy veto condition. The red curve shows the same but with a 100 GeV muon energy veto condition. The black error bars show the distribution from a full CORSIKA simulation of neutrino interactions that includes muons from cosmic-ray air showers and the full detector response. The agreement with the green curve illustrates that the analytic self-veto probability with 1 TeV muon energy veto condition can adequately approximate the exact veto condition of the event selection.

tors are needed since uncertainties in the ice model, muon interaction cross sections, and contamination from muon bundles can significantly alter the fraction of muons classified as tracks or showers. Additionally, the muon control sample can be used as prior knowledge on the scaling factors. Since 4 events each remained in the cascade-like and track-like control samples after all selection criteria, the number of muon background events is expected to be 8.4 ± 4.2, respectively, after scaling by the fiducial volume scaling factor of 2.1. This is then implemented as a Gaussian prior on the muon background scaling factors in the likelihood.

6.3.3 Detector Systematic Uncertainties

All simulations described so far also include a model of the detector response that carries additional systematic uncertainties, such as in the DOM optical efficiency, ice properties, and neutrino cross sections. To study each of these uncertainties, several simulations were

generated that represent a variation from the baseline model. For the DOM optical efficiency, simulations with the nominal value varied by ±10% were generated, values consistent with calibration studies of minimum ionizing muons [168]. For ice properties, the nominal ice model is the SPICE Mie model described in Sec. 4.3. Three simulations were generated with a global scaling of $(+10\%, 0\%)$, $(0\%, +10\%)$, and $(-7.1\%, -7.1\%)$ in the depth-dependent bulk ice scattering and absorption coefficients, respectively. These values correspond to uncertainties estimated from LED flasher studies [175]. A simulation with the SPICE Lea model was also generated to study the effect of scattering anisotropy. Lastly, for neutrino cross sections, a simulation was generated with an older calculation of DIS cross sections [219] using the CTEQ5 parton distribution functions [189].

Ideally, each of these systematic uncertainties would be modeled with separate nuisance parameters. However, because of the time-consuming nature of detector simulations, only a small sample size of a few thousand events could be generated for each systematic variation. Instead, it was tested whether a uniform scaling factor of the baseline model distributions used in the likelihood fit could describe the distributions from each systematic variation simulation. This was quantified through a χ^2 fit,

$$\chi^2(\phi) = \sum_i \frac{(\phi\mu_{b,i} - \mu_{s,i})^2}{\sigma_{b,i}^2 + \sigma_{s,i}^2}, \tag{6.8}$$

where $\mu_{b,i}$ is the bin expectation from the baseline simulations, $\mu_{s,i}$ is the bin expectation from the systematic simulation, $\sigma_{b,i}$ and $\sigma_{s,i}$ represent the statistical uncertainty in the bin expectation, and ϕ is the scaling factor of the baseline simulation. For all variations considered here, no scaling factor above 10% was found, and the RMS variation of the scaling factors among the systematic simulations was 4%. Additionally, the reduced χ^2 of the fit was always below 1.3, indicating that the shapes of the distributions were consistent as well. Lastly, no variation more than 5% in the ratio of the expected number of tracks to cascades was observed. Thus, the existing nuisance parameters that are scaling factors on the distributions of atmospheric and astrophysical neutrinos are sufficient to describe the detector systematic uncertainties here. Furthermore, the bias in nuisance parameters introduced by these detector systematic uncertainties will be small compared to their statistical uncertainties derived from the fit, as seen in Tab. 6.1.

6.4 Results on Astrophysical Neutrinos

Before measuring the flavor composition, the energy spectrum of the astrophysical flux was studied under the standard assumption of a $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ flavor composition in order to address three questions left from the original HESE analysis:

1. What is the power-law index of the astrophysical flux?

2. Is the astrophysical flux consistent with a power law?

Parameter	Best-fit value	Prior
Astrophysical γ	2.6 ± 0.15	Flat
Astrophysical Φ_0 $\left(10^{-18}\ \mathrm{GeV^{-1}s^{-1}cm^{-2}sr^{-1}}\right)$	2.3 ± 0.4	Flat
Charm Atmospheric ν ($\times$ERS)	$< 3.4 (90\%\,\mathrm{CL})$	Flat
π/K Atmospheric ν ($\times$HKKMS)	1.0 ± 0.3	Flat
Atmospheric μ Cascades (No. of Events)	$10.9^{+3.7}_{-3.6}$	8.4 ± 4.2
Atmospheric μ Tracks (No. of Events)	$3.9^{+2.7}_{-2.1}$	8.4 ± 4.2

Table 6.1: The best-fit parameters assuming a $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ flavor composition at Earth. Confidence intervals for each parameters are at 68% confidence level.

3. Do prompt atmospheric neutrinos have an influence on these results?

Throughout this section it will be necessary to determine how well a model describes the data, i.e. quantify the goodness-of-fit, or if alternate models are available, select the one that provides the best description of the data, i.e. perform hypothesis testing. With an adequate model, confidence regions for the astrophysical flux parameters can then be calculated. In Appendix A, the techniques for calculating goodness-of-fit, confidence regions, and hypothesis testing are reviewed.

6.4.1 Energy Spectrum

When the flavor composition of the astrophysical flux is fixed at $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$, the parameter estimates and confidence intervals shown in Tab. 6.1 are obtained. The resulting best-fit distributions for tracks and cascades are shown in Fig. 6.8. The best-fit power-law astrophysical flux is

$$\Phi(E) = 2.3 \pm 0.4 \times 10^{-18} \left(\frac{E}{100\,\mathrm{TeV}}\right)^{-2.6 \pm 0.15} \mathrm{GeV^{-1}s^{-1}cm^{-2}sr^{-1}} \tag{6.9}$$

for each neutrino flavor. The two-dimensional confidence region for the normalization of the astrophysical flux, Φ_0, and power-law index, γ, is shown in Fig. 6.10.

The goodness-of-fit $p = 0.15$ shows that the power-law model of the astrophysical flux provides an adequate description of the data, but it is worth checking if a more sophisticated model is also motivated by the data or even favored over a power-law. The power-law index of $\gamma = 2.6 \pm 0.15$ differs from the value of $\gamma = 2$ typically expected from Fermi acceleration with a significance of 3.0σ. An index closer to $\gamma = 2$ could also fit the data if there is cutoff in the energy spectrum that exists because cosmic-ray accelerators are reaching their maximum rigidity. For example, one can parameterize the astrophysical flux with $\gamma = 2$ and an exponential cutoff,

$$\Phi(E) = \Phi_0 \left(\frac{E}{100\,\mathrm{TeV}}\right)^{-2} \exp\left(-\frac{E}{E_c}\right), \tag{6.10}$$

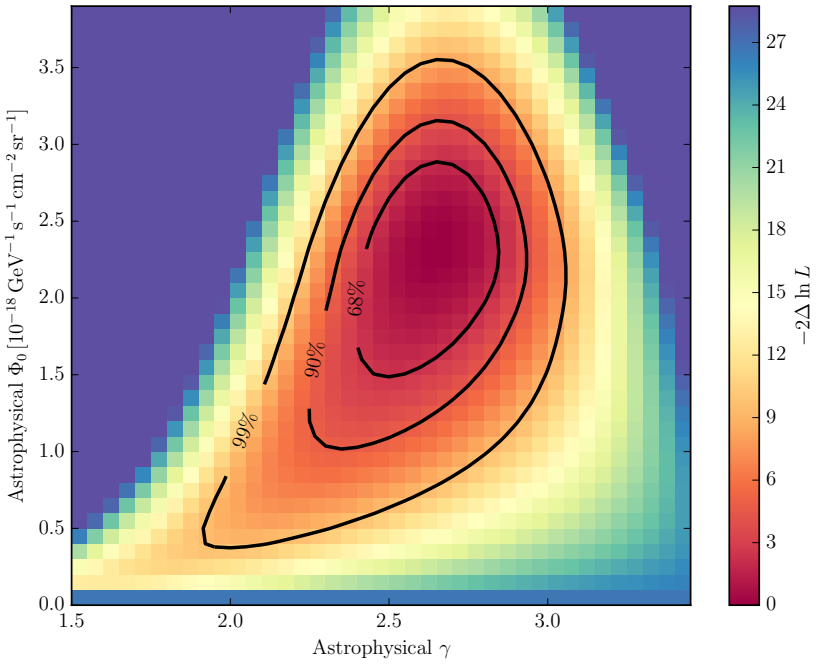

Figure 6.10: Confidence region on the normalization, Φ_0, and power-law index, γ, of the astrophysical flux.

where E_c is the cutoff energy. This hypothesis is rejected with respect to a power-law only hypothesis with a significance of 2.9σ. Thus, the soft value of $\gamma = 2.6$ is driven by low energy events rather than a lack of high-energy events above 1 PeV. The significance for rejecting an E^{-2} spectrum with and without a cutoff was calculated using random pseudo-experiments as described in Appendix A.

A cutoff is only one of many ways the energy spectrum of the astrophysical flux could deviate from a power-law. Instead, one can take a more model-independent approach to search for additional features in the energy spectrum. One simple approach is to measure the astrophysical flux in different energy bins. Here, each bin will be logarithmic and span one third of a decade. Within each bin, the flux will be assumed to follow an E^{-2} spectrum, and the normalization of the flux in each bin will be a free parameter in the likelihood fit. This is another example of an unfolding problem, and the resulting unfolded astrophysical flux is shown in Fig. 6.11. No large deviations from the best-fit power-law flux are observed in any of the energy bins, further indicating that a power-law model of the astrophysical flux is an adequate description of the data.

The observation that the astrophysical power-law index of $\gamma = 2.6 \pm 0.15$ is close to that expected from the prompt atmospheric flux naturally raises the doubt that the astrophysical flux could have been mistaken for an unexpectedly large prompt flux. The clearest evidence

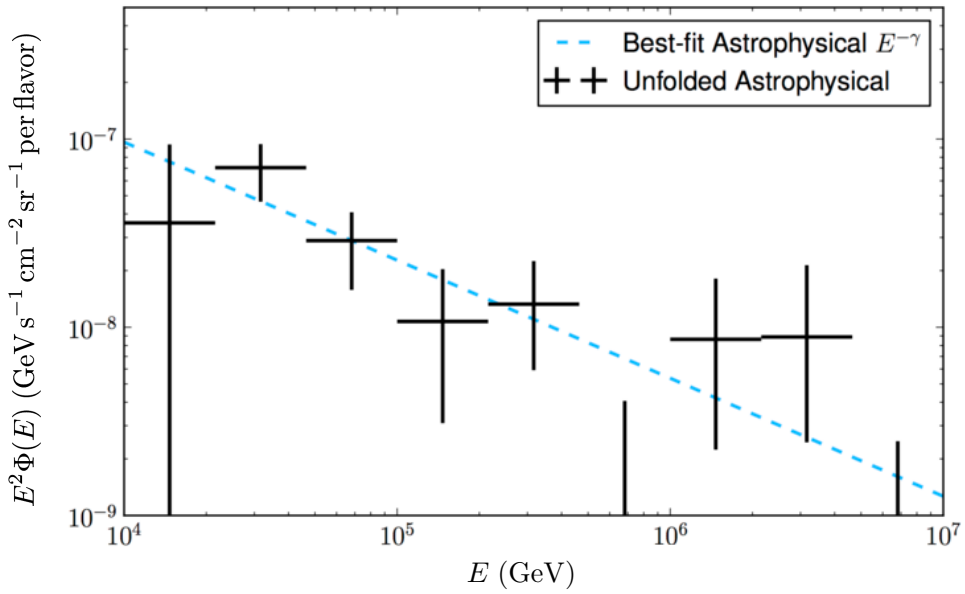

Figure 6.11: The unfolded energy spectrum of astrophysical neutrinos. The blue line shows the best-fit power-law spectrum, and no large deviations from this model are observed.

that this is not the case comes from the zenith distribution of cascade events. As shown in Fig. 6.12, the zenith distribution does not show the suppression of down-going events by accompanying muons expected for prompt atmospheric neutrinos. The best-fit value of the prompt flux scaling is zero, and a 90% confidence level upper limit of 3.4 times the ERS calculation can be derived from this dataset alone. If the prompt flux is fixed at this upper limit, the best-fit astrophysical power-law index only changes to $\gamma = 2.5$. Even when raised to levels far beyond modern calculations such as the BERSS calculation that is a factor of 2 smaller than ERS [137], the prompt flux has only a minor influence on the measurement of the energy spectrum of astrophysical neutrinos.

Lastly, it is important to characterize the energy range where there is sensitivity to an astrophysical flux. At low enough energies, increasing atmospheric backgrounds make it impossible to identify an astrophysical flux. At high enough energies, too few events will be detected to identify a flux. To quantify this idea, the following procedure was used to identify the energy range (E_{low}, E_{high}) where there is statistical evidence for an astrophysical flux.

1. Fix the astrophysical power-law flux parameters to their best-fit values.
2. Set the astrophysical flux to zero below an energy cutoff, E_{low}.
3. Calculate the profile likelihood $L(E_{low})$ allowing all nuisance parameters be vary.

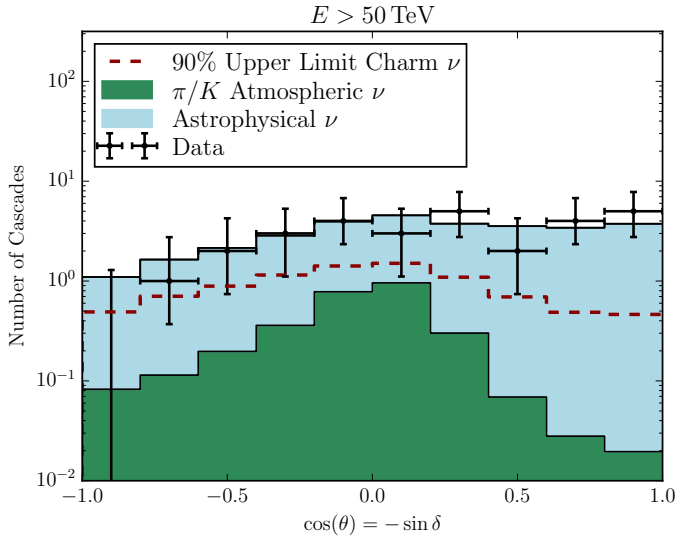

Figure 6.12: The best-fit reconstructed zenith distribution of cascades with reconstructed energy more than 50 TeV. The best-fit contribution from prompt atmospheric neutrinos is zero, but the red dashed line shows the prompt contribution at the 90% confidence level upper limit of 3.4 times the ERS calculation obtained here. The prompt contribution should have have a down-going suppression (right) due to accompanying atmospheric muons, which is not visible in the observed distribution. The suppression of the distribution on the left is due to neutrino absorption in the Earth.

4. Calculate the test statistic $-2\ln(L(E_{low})/L_0)$ where L_0 is the likelihood of the power-law model that extends over all energies.

5. Find the value of E_{low} where the test statistic is $-2\ln(L(E_{low})/L_0) = 1$.

This procedure can then be repeated for the high-energy cutoff E_{high} but setting the astrophysical flux to be zero above this value. As illustrated in Fig. 6.13, the procedure results in an energy range extending from $E_{low} = 35\,\mathrm{TeV}$ to $E_{high} = 1.9\,\mathrm{PeV}$.

6.4.2 Flavor Composition

With the energy spectrum of astrophysical neutrino well-described by a power-law, flavor compositions beyond $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ can be explored. Confidence regions are illustrated by the ternary plot in Fig. 6.14. Due to the small sample size, Wilks' theorem was found not to provide a good approximation, and the test statistic distribution for each point in the flavor composition triangle was found through randomly generated pseudo-experiments. Although a best-fit flavor composition of $(0 : 0.2 : 0.8)_\oplus$ was found, the confidence regions are entirely consistent with all possible flavor compositions allowed by standard neutrino oscillations at 68% confidence level.

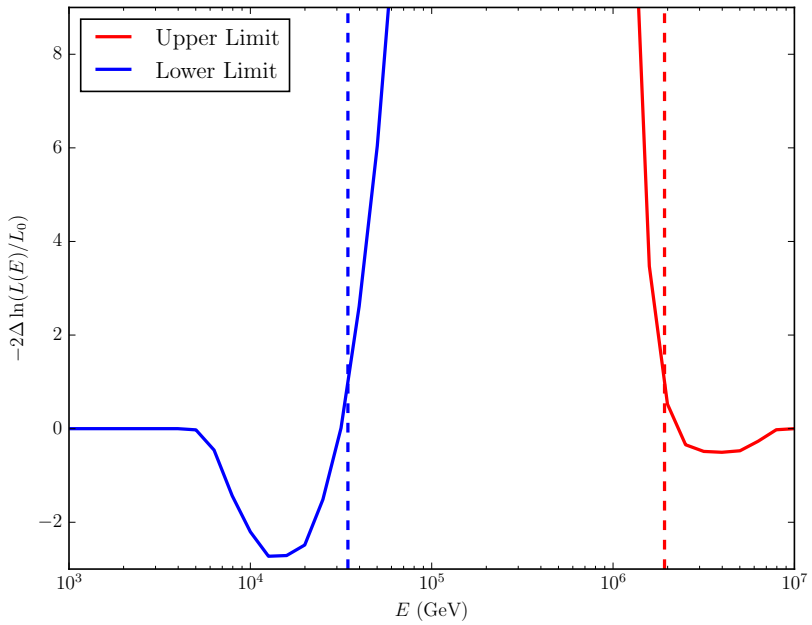

Figure 6.13: The test statistic, $-2\ln(L(E)/L_0)$, between a model with a cutoff in the energy spectrum of the astrophysical flux and a model with no cutoff as a function of cutoff energy. A low-energy cutoff is illustrated in blue and a high-energy cutoff is illustrated in red. The energies where $-2\ln(L(E)/L_0) = 1$ occur at 35 TeV for a low-energy cutoff (dashed blue line) and at 1.9 PeV for a high-energy cutoff (dashed red line) and represent the sensitive energy range for the measurement of the astrophysical flux. Under-fluctuations in data cause the dips where $-2\ln(L(E)/L_0)$ becomes negative.

Since cascades and tracks are the only two identifiers for three neutrino flavors, there is an inherent degeneracy in the determination of the flavor composition. This is reflected in the strong anti-correlation between $f_{e,\oplus}$ and $f_{\tau,\oplus}$, which results from the fact that ν_e and ν_τ both largely produce cascades. The degeneracy is broken mainly by two effects – the shift in the ν_τ deposited energy distribution caused by invisible energy lost to neutrinos in tau decay and the lack of any $\bar{\nu}_e$ Glashow resonance events observed near 6.3 PeV. The preference for a large value of $f_{\tau,\oplus}$ is not statistically significant. Future work to identify ν_τ signatures may resolve this degeneracy.

Since flavor compositions produced by standard neutrino oscillations (narrow blue triangle in Fig. 6.14) are nearly orthogonal to the flavor degeneracy in IceCube, constraints on source flavor composition are possible but not yet significant. After restricting to flavor ratios allowed by standard neutrino oscillations, no source composition can be excluded at $> 68\%$ confidence level, and this remains true even with the additional constraint $f_{\tau,S} = 0$ expected at astrophysical sources.

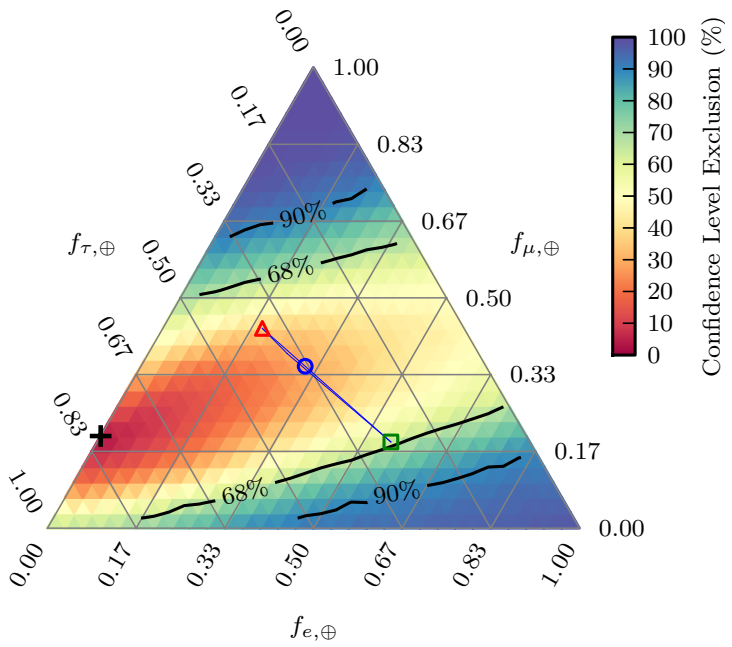

Figure 6.14: The confidence regions for astrophysical flavor ratios $(f_e : f_\mu : f_\tau)_\oplus$ at Earth. Neutrino oscillations map the flavor ratio at sources to points within the extremely narrow blue triangle diagonally across the center. The $\approx \left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ composition at Earth, resulting from a $\left(\frac{1}{3} : \frac{2}{3} : 0\right)_S$ source composition, is marked with a blue circle. The compositions at Earth resulting from source compositions of $(0 : 1 : 0)_S$ and $(1 : 0 : 0)_S$ are marked with a red triangle and a green square, respectively. Though the best-fit composition at Earth (black cross) is $(0{:}0.2{:}0.8)_\oplus$, the limits are consistent with all compositions possible under standard oscillations.

Having found agreement with the predictions of standard neutrino oscillations, constraints are strongest for non-standard flavor compositions producing a large value of $f_{e,\oplus}$ or $f_{\mu,\oplus}$. For example, a maximally track-like, pure ν_μ signature of $(0 : 1 : 0)_\oplus$ is excluded at 3.3σ and a maximally cascade-like ν_e signature of $(1 : 0 : 0)_\oplus$ at 2.3σ.

6.5 Current and Future Flavor Composition Measurements

Although the results reported here are groundbreaking as the first measurement of the flavor composition of astrophysical neutrinos from IceCube, there is little sensitivity to many of the most interesting possibilities for flavor compositions that deviate from $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$. Assuming standard neutrino oscillations, all source flavor compositions are poorly constrained, giving little insight into the origin of astrophysical neutrinos. Furthermore, as illustrated in Fig. 3.6, non-standard flavor compositions involving a general incoherent mixture of mass

eigenstates, such as neutrino decay, pseudo-Dirac neutrinos, or quantum decoherence, are also poorly constrained. To access these possibilities, further experimental improvements are needed.

The biggest limiting factor of the analysis here is the small number of starting tracks. The most straight-forward approach to gather more tracks is to include up-going tracks that interact outside the detector. In 2015, the first evidence for astrophysical neutrinos within a sample of up-going tracks was found [181]. More starting tracks could also be found with improved techniques to reject atmospheric muons, a topic that will be covered next in Ch. 7. Lastly, the lack of any mechanism to identify ν_τ will always leave a degeneracy in the determination of flavor composition. In 2016, a search for a "double pulse" signature from ν_τ was conducted, which could provide one possible identifier to resolve the degeneracy [187]. The possibility of using the inelasticity starting tracks to resolve the degeneracy will be described in Sec. 8.2.1.

6.5.1 Global Fit of IceCube Data

The precise measurement of flavor composition requires the assimilation of data from many different detection channels, namely cascades, starting tracks, through-going tracks, and double cascades. The nature of this problem is well-suited to a joint fit over many distinct event samples. In 2015, a global fit of six different IceCube event samples gathered from 2009 to 2013 was undertaken [220]. The event samples included through-going tracks, cascades, and starting tracks, and the events with more than 6000 PE analyzed here were part of this selection.

In the global fit, the energy spectrum of the astrophysical flux was also found to be statistically consistent with a power-law,

$$\Phi(E) = 2.2 \pm 0.4 \times 10^{-18} \left(\frac{E}{100\,\text{TeV}}\right)^{-2.50\pm0.09} \text{GeV}^{-1}\text{s}^{-1}\text{cm}^{-2}\text{sr}^{-1}, \tag{6.11}$$

when assuming a flavor composition of $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$. Shown in Fig. 6.15, limits were also obtained on the flavor composition. Due to the inclusion of through-going tracks, the constraints on flavor composition become significantly more powerful, but are still consistent with $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ and the results obtained here. Under standard neutrino oscillations, the neutron beam source with source flavor composition $(1 : 0 : 0)_S$ can be excluded with a significance of 3.6σ. Regions described as incoherent mixtures of mass eigenstates are constrained as well. However no event sample with a ν_τ identifier was included, and the degeneracy between $f_{e,\oplus}$ and $f_{\tau,\oplus}$ persists.

6.5.2 Next-Generation Detector

In order to further increase the sensitivity of a flavor composition measurement, a larger detector is needed. A detector upgrade called IceCube-Gen2 has been proposed that would

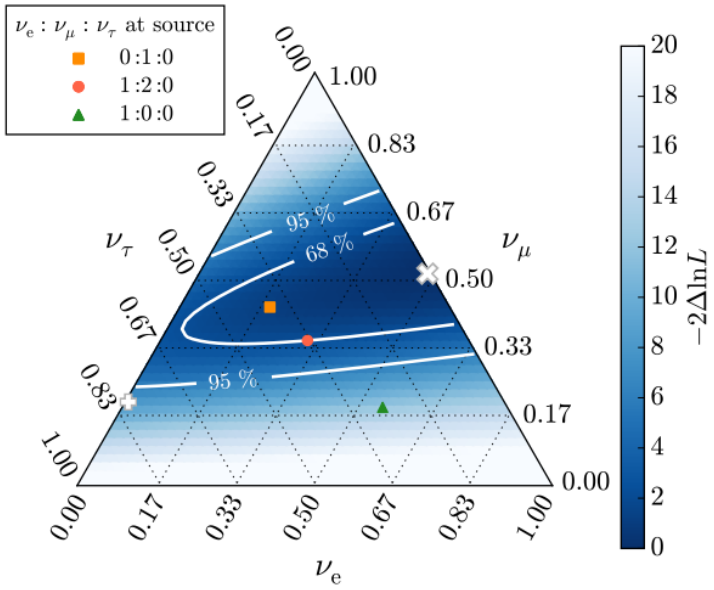

Figure 6.15: Profile likelihood scan of the flavor composition at Earth obtained from a global fit of IceCube data [220]. The best-fit composition is marked with a white "×". The best-fit flavor composition of this work is shown with a white "+" [213].

increase IceCube's instrumented volume to $\sim 10\,\mathrm{km}^3$ [221]. Supposing the effective area can be increased by a factor ~ 6 and data can be collected over 12 years rather than the ~ 3 years covered so far, a factor of ~ 25 increase in the number of events could be expected, which would reduce uncertainties by a factor of ~ 5.

As shown in Fig. 6.16, the contours of the IceCube global fit can be scaled down by a factor of ~ 5 in order to project how well IceCube-Gen2 might measure the flavor composition of astrophysical neutrinos [91]. Assuming standard oscillations, determination of the flavor composition at sources may be possible, but this will depend on reliable measurements of neutrino mixing angles and reducing their uncertainties. Strong limits could also be set on models of new physics that go outside the range allowed by standard oscillations. The construction of IceCube-Gen2, or another similar detector like KM3NET [222], could fully realize the potential of flavor composition measurements, both to gather insight into the sources of astrophysical neutrinos and search for signs of physics beyond the Standard Model.

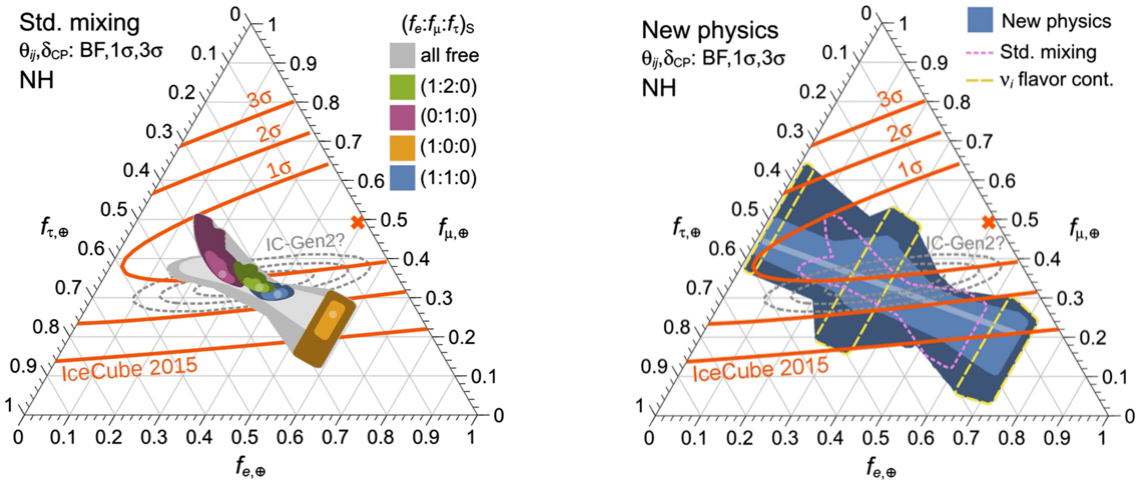

Figure 6.16: Reproduced from [91]. Left: Flavor compositions allowed at Earth for different choices of the flavor composition at sources assuming standard oscillations. The the increasing darkness of each region illustrates the 1σ, 2σ, and 3σ credible region caused by uncertainties in mixing parameters. The orange curves illustrate the confidence regions obtained from the global fit of IceCube data [220]. The dashed grey curves illustrate projected confidence regions for IceCube-Gen2. Right: The blue region shows the allowed flavor composition for models of new physics where neutrinos arrive at Earth as an incoherent mixture of mass eigenstates. The increasing darkness of the region illustrates the 1σ, 2σ, and 3σ credible region caused by uncertainties in mixing parameters. The dashed pink curve illustrates the 3σ region allowed by standard oscillations. The dashed yellow curves illustrate the allowed regions at 3σ for each mass eigenstate.

Chapter 7

Inelasticity Distribution of Neutrino Interactions above 1 TeV

IceCube was built with the main goal of characterizing astrophysical neutrinos, but it is also a laboratory to study the interactions of the highest energy neutrinos found in nature. As elaborated in Sec. 2.3.3, the theory of neutrino-nucleon interactions in the deep inelastic scattering regime relevant for IceCube is thought to be well-understood, but with the center-of-mass energy of $\sim 1\,\mathrm{PeV}$ neutrino interactions now exceeding $1\,\mathrm{TeV}$, observing manifestations of physics beyond the Standard Model is a tantalizing possibility. Until recently, neutrino-nucleon scattering cross sections could only be measured up to neutrino energies of $\sim 370\,\mathrm{GeV}$ at particle accelerator facilities, but an important first step was taken when the total neutrino-nucleon cross section was measured at IceCube in the energy range from $6\,\mathrm{TeV}$ to $980\,\mathrm{TeV}$ using the absorption of neutrinos as they pass through the Earth [47, 48]. However, the total cross section is only one observable characterizing neutrino-nucleon scattering, and by studying neutrino interactions contained in IceCube, one can also measure the distribution of inelasticity–the fraction of a neutrino's energy transferred to a hadronic shower.

In IceCube, it is typically only possible to infer the energy and inelasticity of starting tracks from contained charged-current (CC) ν_μ interactions. Given the energy of the hadronic shower, E_{had}, and muon, E_μ, created in the interaction, the neutrino energy and inelasticity are

$$\begin{aligned} E_\nu &= E_{\mathrm{had}} + E_\mu, \\ y &= \frac{E_{\mathrm{had}}}{E_{\mathrm{had}} + E_\mu}. \end{aligned} \tag{7.1}$$

As described in Sec. 4.1.3, the Cherenkov light yield from hadronic cascades fluctuates, and only the electromagnetic-equivalent energy of the cascade, E_{casc}, is visible, which is a fraction of the hadronic cascade energy. Because of this fact, it is useful to introduce definitions of

the visible energy, E_{vis}, and visible inelasticity, y_{vis},

$$\begin{aligned} E_{vis} &= E_{casc} + E_{\mu}, \\ y_{vis} &= \frac{E_{casc}}{E_{casc} + E_{\mu}}. \end{aligned} \tag{7.2}$$

For ν_μ CC events, the visible energy and visible inelasticity tend to approach the true neutrino energy and inelasticity at high energies since the Cherenkov light emission of hadronic cascades approaches that of electromagnetic cascades. The definition of visible inelasticity has the added advantage that it also has meaning for events that are not ν_μ CC interactions. For example in a ν_τ CC interaction where a tau decays to a muon, the electromagnetic-equivalent energy of the hadronic cascade can be added to the energy of the muon to define a visible energy that ignores the invisible energy lost to neutrinos in tau decay. The definition can also be applied to events where an energetic muon results from a hadronic decay in a cascade. Since hadronic cascades typically produce many low energy muons that cannot be individually resolved, the muon energy in Eq. 7.2 must additionally be defined to be the most energetic muon produced.

In this chapter, the visible inelasticity distribution of neutrinos above 1 TeV will be studied. Techniques from the field of machine learning were applied to both isolate a sample of starting tracks and reconstruct their visible energy and inelasticity. In Sec. 7.1, an event selection using a boosted decision tree will be described that discriminates starting tracks from cascades and reduces atmospheric muon background to a negligible level. In Sec. 7.2.2, an approach based on random forests will be described that enables both the visible energy and inelasticity of starting tracks to be reconstructed.

The distribution of visible inelasticity in IceCube can give many insights into high-energy neutrino scattering. One of the most striking features of the neutrino inelasticity distribution is that there is a difference between neutrinos and antineutrinos, and this provides a method to determine the relative number of neutrinos and antineutrinos in an event sample. In Sec. 7.4.3, this feature will be exploited to measure the neutrino to antineutrino ratio of the atmospheric neutrino flux.

The visible inelasticity distribution can also be used to identify the contribution of different subprocesses to neutrino scattering. For example, CC neutrino interactions where a charm hadron is produced have a flatter visible inelasticity distribution than other events. Muons from charm hadron decays also leave a distinct visible inelasticity distribution. In Sec. 7.4.4, this unique visible inelasticity distribution will be used to search for the neutrino-induced charm production.

The inelasticity distribution also has a strong dependence on neutrino energy, and IceCube collects atmospheric and astrophysical neutrinos over several orders of magnitude in energy. In Sec. 7.4.5, the mean of the inelasticity distribution for ν_μ CC events will be determined across energy from 1 TeV to 100 TeV.

The improved event selection and reconstruction techniques devised for this analysis also enable a much improved measurement of the energy spectrum and flavor composition of astrophysical neutrinos compared to that in Ch. 6. The distinct visible inelasticity distribution

of ν_μ CC events and ν_τ CC events with a muonic tau decay has the potential to break the flavor degeneracy of a measurement using only cascade and track classification. In Sec. 7.4.1, an updated measurement of the energy spectrum and flavor composition of the astrophysical flux will also be presented.

7.1 Event Selection

The goal of the event selection is to obtain a sample of starting tracks and cascades above 1 TeV free of atmospheric muon background. Starting tracks will be of primary interest for studies of their inelasticity distribution, but it is also useful to maintain cascades not only to measure the flavor composition of the astrophysical flux, but to control the systematic error in other measurements caused by uncertainty in the astrophysical flux. Since the astrophysical flux produces largely cascades due to its expected $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ flavor composition and cascades can be reconstructed with excellent energy resolution, cascades provide the most valuable information about the energy spectrum of the astrophysical flux.

The event selection uses five years of data recorded from May 2011 to May 2016 while IceCube was in its complete 86-string configuration. The total detector live time in this period was 1734 days. During the design of the event selection, a 10% burn sample was used for testing, and the remaining 90% of the data was kept blind. The event selection follows four steps:

1. Select events that produce more than 100 photoelectrons (PE) and pass an outer-layer veto.
2. Train a boosted decision tree to classify events as atmospheric muons, neutrino-induced cascades, or starting tracks.
3. Use the boosted decision tree to remove remaining atmospheric muon background to negligible level and split the resulting neutrino sample into a cascade and starting track sample.
4. Reconstruct the energy and inelasticity of starting tracks using a random forest regressor.

7.1.1 Outer-Layer Veto

The event selection begins with all events passing the online cascade or muon filter at the South Pole. Both cascade and muon filters are intended to be very loose cuts that reduce the trigger rate to fit into available satellite bandwidth while keeping as many events with a cascade or muon topology as possible. The design of both filters over as documented at [223, 224, 225] for the cascade filter, and at [226, 227] for the muon filter. During the design of the event selection, it was found that the two filters capture $> 99.5\%$ of all triggered events

that would pass all later event selection criteria, so the changes in the filters are negligible here.

After passing the muon or cascade filters, events were required to produce more than 100 PE, which is a sufficiently low threshold to capture most contained neutrino interactions above 1 TeV. Additionally, it was required that DOMs on more than 3 strings recorded data since otherwise there is generally not enough information for reliable cascade or track reconstructions.

To further reject muon background, an outer-layer veto is used. The design of the veto is very similar to that from Sec. 6.1.2, but because of the lower energy events targeted here, three tweaks are necessary that are documented thoroughly in [196]. The first is a change in the number of photoelectrons observed in a rolling 3 μs time window needed to define the starting time of an event. The previous threshold of $Q_{start} = 250\,\text{PE}$ can occupy most or all of the total number of photoelectrons, Q_{tot}, for the low energy events considered here. Instead, the threshold is chosen to vary as a function of Q_{tot} according to

$$Q_{start} = \begin{cases} 3 & Q_{tot} < 72\,\text{PE} \\ Q_{tot}/24 & 72\,\text{PE} \leq Q_{tot} < 6000\,\text{PE} \\ 250 & Q_{tot} \geq 6000\,\text{PE} \end{cases} . \tag{7.3}$$

The second change is to increase the thickness of the horizontal veto layer just below the dust layer to 120 m. Covering a depth from 2050 m to 2170 m, the veto layer can better reject low energy muons entering through the dust layer. Lastly, the bottom veto layer was redefined to include the bottommost operational DOM on every string. Because strings are deployed at non-uniform depths, the previously defined 10 m thickness left holes in the bottom veto layer. This better excludes neutrino-induced muons entering from below the detector and a rare background where a highly inclined cosmic-ray muon produces a large stochastic energy loss under the detector.

Muon bundles from two coincident air showers in the same trigger window can sometimes fool the outer-layer veto. This can happen if the first low-energy muon bundle does not produce enough photoelectrons in the outer layer to veto the event but does produce enough photoelectrons to determine the start of the event. If an energetic second muon bundle follows soon after, a large total amount of photoelectrons can be recorded. To discriminate against this background, one can use the fact that the photoelectrons produced by one muon bundle are often causally disconnected from the other bundle. The `TopologicalSplitter` algorithm was designed to count clusters of photoelectrons that are causally connected [228]. Here, it is required only one cluster was found, which removes the worst offending cases of coincident muon background.

7.1.2 Boosted Decision Tree

The outer-layer veto is very effective in reducing the rate of atmospheric muons, but at this stage the overall passing rate is 0.39 Hz while the expected neutrino passing rate is

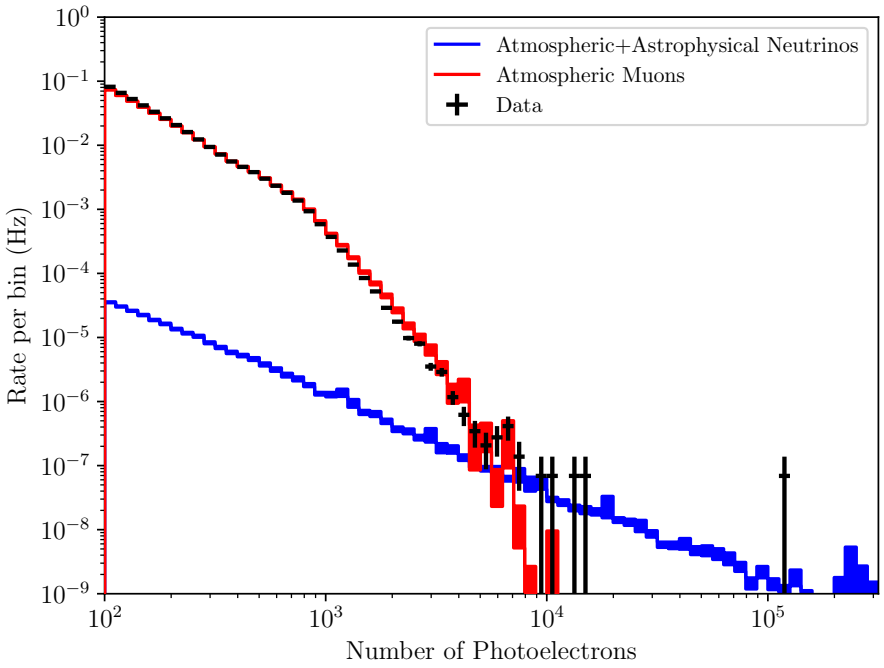

Figure 7.1: The distribution of the detected number of photoelectrons per event for atmospheric muons (red), all atmospheric and astrophysical neutrinos (blue), and experimental data in the 10% burn sample (black) after the outer-layer veto. To access ~ 1 TeV neutrinos that typically produce ~ 100 PE, one must overcome muon background with a rate 10^3 times larger in size.

0.24 mHz as shown in Fig 7.1. Further improvements are still needed to access low energy neutrinos down to 1 TeV.

Down-going atmospheric muons that pass the outer-layer veto can be difficult to reject for several reasons. At lower energies, it becomes more likely for atmospheric muons to slip through the outer-layer veto since they produce less Cherenkov light, and they can easily mimic a starting track. If a muon undergoes a large stochastic energy loss soon after entering the detector, it can also appear to be a cascade. Even the most straight-forward solution of excluding down-going events at lower energies is problematic since occasionally down-going muons can be mis-reconstructed as up-going, especially if there is a second bundle of muons from a coincident air shower in the same trigger window. Due to these multiple types of muon background, it is difficult to devise a series of cuts that will reject all muon background while maintaining a high neutrino signal efficiency.

Even if muon background is rejected, there is the additional problem of classifying cascades and starting tracks. The likelihood ratio approach relying only on first photoelectron times described in Sec. 6.1.4 is effective for high-energy events, but degrades for low-energy

events where fewer photoelectrons are detected. It can also perform poorly on tracks that produce a low energy muon relative to the cascade, events that occur near the edge of the detector, or events where either the cascade or track reconstruction performed poorly. Because the problems of discriminating tracks and cascades and rejecting muon background both rely on a large number of variables describing the location, orientation, topology, and reconstruction quality of an event, it is useful to turn to the field of machine learning for the best solution.

In machine learning, the task of identifying muon background, neutrino-induced cascades, and starting tracks is an example of a multi-class classification problem. In the most general form of a multi-class classification problem, one uses a vector of numbers for each event called "features" and derive a function operating on the features that predicts the class of an event. When a training set of example events is available with known classes, it is known as a "supervised learning" problem. Here, simulated events are used for training, and a boosted decision tree (BDT) is used as the classification algorithm.

The fundamental piece of a BDT is a decision tree. A decision tree is a set of nodes organized into a binary tree structure. Starting at the root node, for each event, a decision, or split, is made based on the value of one of the features. Based on the outcome of the split, an event is sent to one of two child nodes where another split is made. This can continue to an arbitrary depth until a final decision about the class of the event is reached at a leaf node.

Training a decision tree proceeds top-down by searching among all the features at each node to find the one that produces the best split of the training sample into two subsamples according to an appropriate metric. Here, the metric is the Gini impurity, $\sum_j f_j(1 - f_j)$, where j indexes the two subsamples and f_j is the fraction of the majority class in the subsample. Each subsample is then sent to a child node where the above procedure is repeated to produce a new split. In principle, this can can continue to an arbitrary depth until there are no more events left to split. However, this approach does not generalize well to classifying data outside the training sample, and it is usually necessary to stop splitting after meeting certain criteria. The approach here is to limit the depth of the tree and require a minimum number of events in order to split a node. When these criteria are met at a leaf node, the class of an event is determined by a majority vote among the training events. Although simple to understand and implement, decision trees are not very powerful on their own and are prone to over-fitting if the splitting process proceeds too far.

Boosting is a process designed to improve the performance of any classification algorithm. Applied to a decision tree, boosting works by increasing the weights of training events that were misclassified and training a new decision tree with the new weights. This can continue successively to produce an ensemble of typically hundreds of decision trees. By performing a weighted average of the predictions from each decision tree, an output score, s_j, for each class can be assigned to every event that lies between 0 and 1 and satisfies $\sum_j s_j = 1$. The AdaBoost algorithm [229] provides a prescription for how to adjust event weights after training each decision tree and how the weighted average of decision tree outputs is performed. The software implementation of a boosted decision tree from `scikit-learn` [230] is used

here with an ensemble of 400 trees.

For a BDT applied to a the classification problem of identifying atmospheric muons, cascades, and starting tracks, a set of 15 features was selected ad hoc. Each feature was chosen for its ability to identify mis-reconstructed tracks, discriminate cascades from tracks, or pick out signs of a down-going atmospheric muon.

1. **$\log_{10}$ NPE** The base 10 logarithm of the total number of detected photoelectrons. Only DOMs in hard local coincidence and not in DeepCore are considered. This feature acts as a proxy for the energy of each event. Atmospheric muons tend to have a steeper NPE spectrum than neutrinos due to their steeply falling flux and because they become more likely to pass the outer-layer veto at lower NPE.

2. **$\cos\theta_{\text{Track}}$** The cosine of the zenith angle from the track reconstruction. The spline-based track reconstruction method will be described in Sec. 7.2.1. This feature enables down-going muons to be distinguished from up-going neutrinos.

3. **Cascade z-coordinate** The z-coordinate of the spline-based cascade reconstruction measured from the detector center. Down-going atmospheric muons may leave an isolated, stochastic energy loss near the top of the detector, whereas neutrino-induced cascades should be more uniformly distributed through the detector.

4. **Horizontal distance of cascade from detector edge** The smallest horizontal distance of the spline-based cascade reconstruction from the edge of a polygon enclosing the outermost strings. As above, penetrating atmospheric muons are likely to be near the outermost strings.

5. **Distance of first energy loss from detector edge** Using the method described in Sec. 6.2.2, the energy loss profile of an event can be unfolded as a sequence of electromagnetic cascades along the reconstructed track. This feature is the distance of the first unfolded energy loss $> 1\,\text{GeV}$ from the detector entry point as measured along the reconstructed track and estimates the distance of the neutrino interaction vertex within the detector. The detector boundaries are defined by planes $\pm 500\,\text{m}$ above and below the detector center and a polygon enclosing the outermost strings. Like the two features above, this feature discriminates against atmospheric muons that tend to deposit energy near the edge of the detector, but it is more effective for track-like events.

6. **Contained length of deposited energy** The distance between the first and last reconstructed energy losses $> 1\text{GeV}$ contained within the detector. This feature can discriminate between tracks and cascades.

7. **Number of directly hit DOMs** The number of DOMs where a photoelectron was observed in a time window $-15\,\text{ns}$ to $250\,\text{ns}$ around the geometric first arrival from the reconstructed track, as defined in Eq. 6.3. If many DOMs observed photoelectrons

close to the geometric first arrival time, the more likely it is that an accurate track reconstruction was obtained. This feature can be used to discriminate against down-going atmospheric muons that are mis-reconstructed as up-going. It also discriminates between tracks and cascades.

8. **Track length between directly hit DOMs** The length between the first and last directly hit DOMs projected along the reconstructed track. A longer length indicates a long track through the detector with a good reconstruction. This feature also discriminates cascades and mis-reconstructed down-going muons from well-reconstructed up-going tracks.

9. **Reduced cascade log likelihood** The reduced log likelihood, $-\ln L/n$, from the cascade reconstruction using first photoelectron times only, as described in Sec. 6.1.4, where n is the number of degrees of freedom. This feature discriminates cascades and tracks.

10. **Log likelihood ratio of cascade and track fits** This is the same log likelihood ratio described in Sec. 6.1.4. Larger values indicate tracks and smaller values indicate cascades.

11. **Number of photoelectrons consistent with a down-going track** For each direction of 104 different directions from a hemi-spherical HEALPix grid [231], a down-going track ending at the reconstructed cascade vertex is created and the number of photoelectrons in a time window of -15 ns and $+1000$ ns around the geometric first arrival time from the track is counted over all DOMs. The maximum number of photoelectrons over all track directions is then chosen. The intent of this feature is to search for any photoelectrons that are consistent with a down-going atmospheric muon prior to the reconstructed cascade vertex.

12. **Number photoelectrons in the outer veto layer** The amount of photoelectrons recorded by the outer-layer veto, which is already required to be < 3. This feature discriminates against atmospheric muons entering the detector.

13. **Number of photoelectrons consistent with an out-going track** For each direction of 192 different directions from a fully spherical HEALPix grid, a track ending at the reconstructed cascade vertex is created and the number of photoelectrons in a time window of -30 ns and $+500$ ns around the geometric first arrival time from the track is counted over all DOMs in hard local coincidence. The maximum number of photoelectrons over all track directions is then chosen. The intent of this feature is to search for any photoelectrons that are consistent with a muon emanating from the reconstructed cascade vertex.

14. **Angle between reconstructed track and plane-wave velocity** A simplified estimate for the direction of track can be found by fitting the first photoelectron times

in each DOM to a plane wave traveling with velocity $\mathbf{v}$ [232]. The angle between $\mathbf{v}$ and the spline-based track reconstruction is a check on the quality of the reconstruction. This feature acts to discriminate against down-going muons mis-reconstructed as up-going.

15. **Angular uncertainty of track reconstruction** An estimate for the uncertainty in the direction of the spline-based track reconstruction can be found by fitting the profile log likelihood as a function of zenith and azimuth to a two-dimensional paraboloid. The geometric mean of the semi-major and semi-minor axes of the 68% uncertainty ellipse can be used to quantify the angular uncertainty. This feature also serves to identify well-reconstructed tracks.

Training data for the BDT consisted of 4,469,803 simulated events from `CORSIKA` to represent atmospheric muons and 734,797 simulated events from `neutrino-generator` for all flavors of neutrinos. Cosmic-rays were weighted according to the H3a flux model and neutrinos were weighted according to the HKKMS conventional atmospheric flux, BERSS prompt atmospheric flux, and an astrophysical power-law flux with index $\gamma = 2.5$ corresponding to the best-fit results obtained from a global analysis of IceCube data [220].

Because the goal is to reconstruct the energy and inelasticity of all events classified as starting tracks, the track reconstruction should meet several quality requirements to ensure that the unfolding of the energy loss profile along the track is reliable. Specifically, the track should be reconstructed with a small angular error, and its length inside the detector should be long enough that enough information to reconstruct the outgoing muon energy. A minimum track length of about three string spacings, or about 375 m, is a good rule of thumb [184]. For the purposes of training the BDT only, neutrinos labeled as starting tracks must additionally satisfy several reconstruction quality criteria:

1. The neutrino interaction vertex is contained inside the detector.
2. A muon with energy of more than 100 GeV is produced. This corresponds to an average range of 370m.
3. The reconstructed track direction must be within 5° of the muon direction
4. The path length of the muon must be more than 300 m within the detector.

Any neutrinos not satisfying these conditions are labeled as cascades for BDT training only.

7.1.3 Muon Background Rejection

Because the muon background rate is a factor of 1.6×10^3 times the neutrino rate at this stage, it is necessary to have the largest training sample of atmospheric muons as possible in order to create the most powerful classifier. This is problematic since it is also necessary validate the BDT performance on another large, independent muon sample of sufficient size

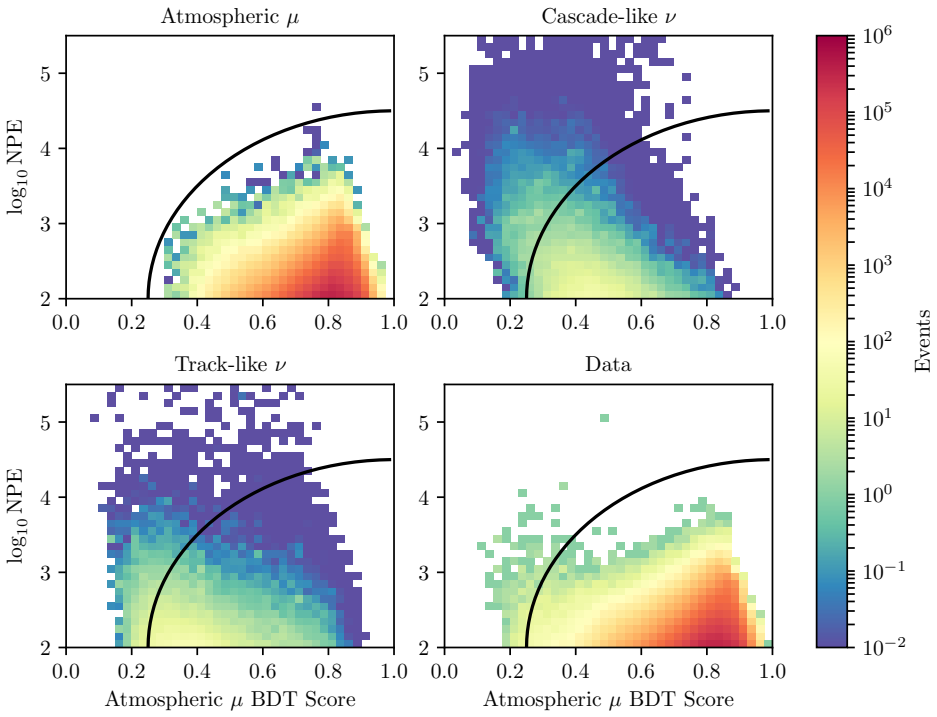

Figure 7.2: Distribution of the total number of photoelectrons and the BDT muon score obtained through k-fold cross validation. Upper left: Distribution for atmospheric muons from CORSIKA simulation. Upper right: Cascade-like neutrinos of all flavors from neutrino-generator simulation. Bottom left: Track-like neutrinos of all flavors from neutrino-generator simulation. Bottom right: Observed data in the 10% burn sample. The elliptical cut to remove atmospheric muon background is shown in black.

to estimate the surviving muon background rate. To overcome this problem, a procedure called k-fold cross validation is performed.

In k-fold cross validation, the training sample is split into k unique testing samples. For each testing sample, the BDT is trained with the remaining data and BDT performance metrics are calculated with the testing sample. The performance metrics are then averaged over all k testing samples in order to estimate how a BDT trained with the entire training sample will perform. The advantage of this approach is that it allows nearly all simulated data to be used for training and validating the BDT if k is large. Here $k = 10$ is used to keep the training time of k BDTs manageable.

To reject muon background, a cut is applied on the muon BDT score, s_μ. Because the muon background tends to fall more rapidly with NPE than neutrinos, the BDT score cut can be relaxed at higher NPE. A two-dimensional cut on NPE and muon BDT score is chosen

that is parameterized by an ellipse,

$$\left(\frac{s_\mu - 1}{a}\right)^2 + \left(\frac{\log_{10}\mathrm{NPE} - 2}{b}\right)^2 > 1. \tag{7.4}$$

Figure 7.2 shows the two-dimensional distribution of muon BDT score and NPE averaged over all 10 testing samples in the k-fold cross validation scheme. Values of $a = 0.75$ and $b = 2.5$ are chosen to exclude all simulated muon events except for one outlier. The remaining neutrino signal rate for this cut is expected to be 326 events in the 10% burn sample.

Since only one simulated muon background event survived, a good estimate of the muon background rate cannot be obtained for the chosen muon rejection cut. To get a better idea of the muon background rate, one can relax the elliptical cut and calculate the background rate as a function of a in Eq. 7.4. In Fig. 7.3, the rate of muon, cascade, and track events is shown as function of a for $b = 2.5$. For a choice of $a = 0.70$, 6.9 ± 4.0 muon background events are expected in the 10% burn sample. Extrapolating to the chosen value of $a = 0.75$, only a lower muon background rate should be found that is negligible compared to the statistical fluctuations in the neutrino signal rate.

An additional estimate of muon background can be obtained from MuonGun simulation of single muons with a much larger sample size than CORSIKA simulation. Of 26,837,295 simulated MuonGun events that passed the outer-layer veto selection, only 55 survived the muon rejection cut with $a = 0.75$. Weighted to the H3a cosmic-ray model, this predicts 0.27 ± 0.10 single muon events in the 10% burn sample. Since bundles are missing in MuonGun simulation, a correction factor can be used to scale the event rate prediction. A factor of 3 is needed to scale MuonGun simulation to match CORSIKA simulation after the outer-layer veto selection due to missing bundles. However, the correction factor after the BDT cut is certainly smaller than 3 since the nearly all CORSIKA events near the BDT cut threshold are single muons. Even with the conservative scaling factor of 3, the MuonGun simulation confirms that the muon background rate is much smaller than statistical fluctuations in the neutrino rate. From now on, muon background can be neglected.

7.1.4 Track and Cascade Classification

After muon rejection, the BDT can also be used to split the resulting neutrino sample into two samples of cascades and tracks. The BDT scores for tracks and cascades, s_{track} and s_{casc} can be used to construct a normalized track score, $\hat{s}_{track} = s_{track}/(s_{track} + s_{casc})$, that uniformly discriminates between cascades and tracks regardless of the value of the muon score, s_μ. A threshold on $\hat{s}_{track}$ to split the neutrino sample into cascade and track samples was chosen to optimize the signal-to-noise ratio of well-reconstructed tracks, $\mathrm{SNR} = S/\sqrt{S+B}$. The signal, S, is the number of well-reconstructed tracks in the track sample meeting the definition described in Sec. 7.1.2, and the noise is the size of Poisson fluctuations in the signal and background rate, $\sqrt{S+B}$, where B is the number of background events in the track sample that are not well-reconstructed tracks.

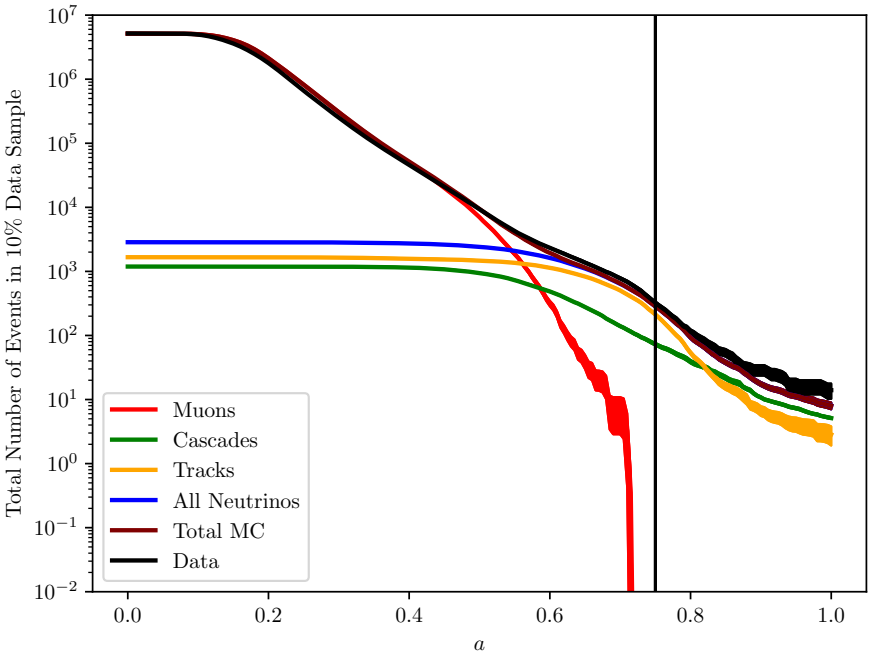

Figure 7.3: The expected rate of atmospheric muons (red), neutrinos (blue), cascades (green), and tracks (orange) in the 10% burn sample as a function of the elliptical BDT cut parameter a in Eq. 7.4 computed through k-fold cross-validation. A value of $a = 0.75$ was chosen for the final cut. The width of the lines corresponds to the statistical uncertainty of the event rate prediction. Though one event passed the final BDT cut, its weight is too small to be visible on the plot.

The SNR as a function of the threshold on normalized BDT track score is shown in Fig. 7.4 for simulated `neutrino-generator` events. Two other performance metrics are shown as well, the purity of the track sample and the efficiency for selecting tracks. The purity is the fraction of events in the track sample that are well-reconstructed. The efficiency is the fraction of all well-reconstructed tracks in the neutrino sample that end up in the track sample. A threshold of $\hat{s}_{track} = 0.52$ was found that optimized the SNR. The corresponding purity and efficiency of the track sample for this choice, are 92% and 98% respectively.

The above performance metrics for cascade and track discrimination are based on the goal of identifying tracks reconstructed well enough that their energy and inelasticity can also be reconstructed. Another approach is to compute the same metrics for identifying charged-current ν_μ tracks regardless of reconstruction quality. For the threshold of $\hat{s}_{track} = 0.52$ obtained here, the purity and efficiency for identifying ν_μ CC events are 99.5% and 87%, respectively. If one optimizes the SNR for all ν_μ CC events instead of only well-reconstructed tracks, a threshold of $\hat{s}_{track} = 0.39$ is obtained and the purity and efficiency are 98% and 95%,

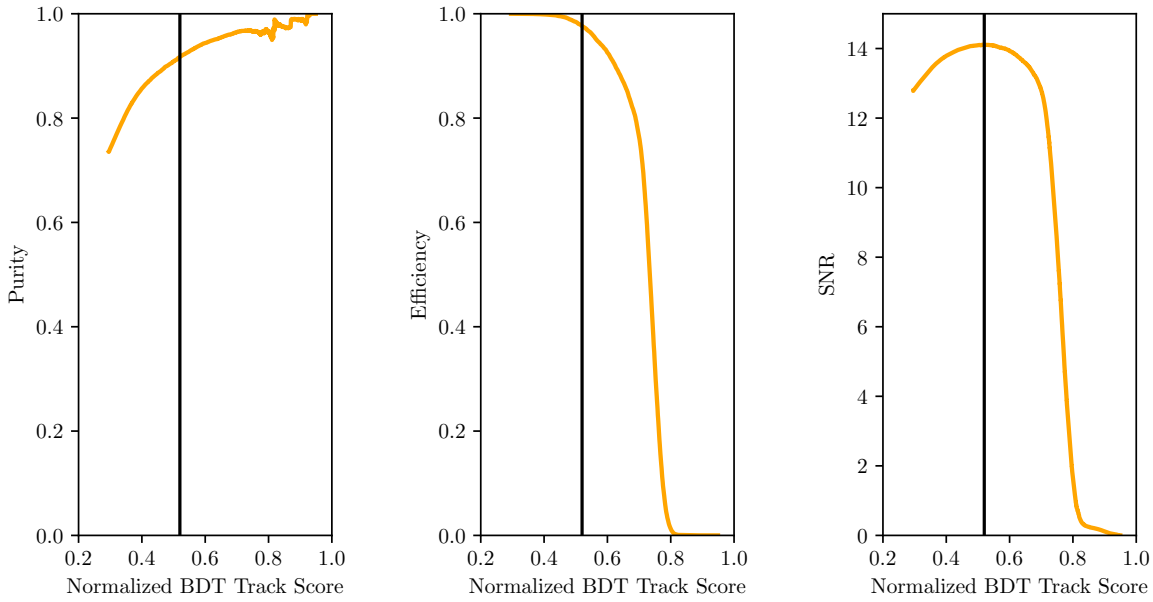

Figure 7.4: Left: The track purity as a function of the BDT track score threshold. Middle: The track efficiency as a function the BDT track score threshold. Right: The track signal-to-noise ratio as a function of BDT track score. A BDT track score threshold of 0.52 (black line) was chosen in order to maximize the signal-to-noise ratio. The purity and efficiency for this threshold are 92% and 98%, respectively. Fluctuations in the purity curve are due to low statistics at a high BDT track cut.

respectively. This out-performs both the purity of 95% and the efficiency of 70% achieved by the likelihood ratio method of Sec. 6.1.4.

7.1.5 Event Sample Properties

After unblinding the full detector live time of 1734 days, 2650 starting tracks and 965 cascades were found. Based on simulations with the baseline flux models used so far, 90% of starting tracks are expected to have a neutrino energy between 740 GeV and 45 TeV and 90% of cascades are expected to have a neutrino energy between 1.1 TeV and 53 TeV. The expected number of events for each neutrino flavor from atmospheric and astrophysical fluxes is shown in Tab. 7.1.

7.2 Reconstructions

In this section, track angular reconstruction will be reviewed, and a new machine learning method will be described that enables the inelasticity of starting tracks to be reconstructed in IceCube. Cascade reconstruction will also be revisited, and a new test of cascade pointing

Tracks							
	$\nu_e + \bar{\nu}_e$		$\nu_\mu + \bar{\nu}_\mu$		$\nu_\tau + \bar{\nu}_\tau$		
	CC	NC	CC	NC	CC	NC	Total
Conv. Atm.	1.0	0.06	2.5×10^3	1.6	0.0	0.0	2.5×10^3
Prompt Atm.	0.10	0.0085	7.1	0.0086	0.0	0.0	7.2
Astrophysical	0.97	0.10	61	0.11	8.1	0.11	70
Total	2.1	0.17	2.5×10^3	1.7	8.1	0.11	2.6×10^3

Cascades							
	$\nu_e + \bar{\nu}_e$		$\nu_\mu + \bar{\nu}_\mu$		$\nu_\tau + \bar{\nu}_\tau$		
	CC	NC	CC	NC	CC	NC	Total
Conv. Atm.	93	6.5	3.1×10^2	1.8×10^2	0.0	0.0	5.9×10^2
Prompt Atm.	11	1.2	1.3	1.0	0.0	0.0	15
Astrophysical	1.3×10^2	17	20	17	83	17	2.9×10^2
Total	2.4×10^2	25	3.3×10^2	1.9×10^2	83	17	8.9×10^2

Table 7.1: The expected number of tracks (top) and cascades (bottom) for the final selection with 1734 days of detector live time. The HKKMS calculation is used for the conventional atmospheric flux, the BERSS calculation is used for the prompt atmospheric flux, and the IceCube global fit is used for the astrophysical flux. 2650 tracks and 965 cascades were observed.

accuracy will be done using a sample of high-inelasticity tracks.

7.2.1 Track Angular Reconstruction

An accurate track angular reconstruction is needed to ensure that later unfolding of the energy loss profile along the track is reliable for later reconstruction of energy and inelasticity. Since accurate angular reconstruction is also vital for IceCube's goal of identifying astrophysical neutrino point sources, much effort has gone into improving the performance of track angular resolution.

Section 6.1.4 described a track reconstruction that used only the first photoelectron time in each DOM. In this approach, the residual time distribution of the first photoelectron, $p_1(\Delta t_i)$, neglected the fact that the more light an event produces, the more likely it is that a photoelectron will be detected earlier. A better approach is to consider the probability distribution describing all photoelectrons detected by a DOM, $p(\Delta t_i)$. If a DOM observes N_i photoelectrons, then the probability distribution describing the residual time of the first photoelectron is [182]

$$p_1(\Delta t_i) = N_i p(\Delta t_i) \left(\int_{t_i}^{\infty} p(\Delta t_i) \right)^{N_i - 1} . \tag{7.5}$$

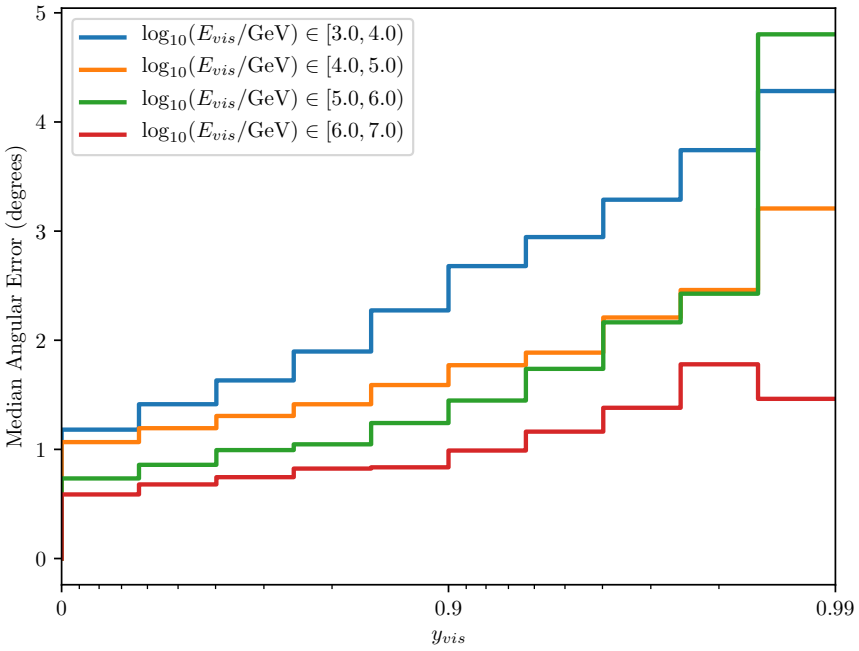

Figure 7.5: The median angle between the reconstructed track and neutrino as a function of visible inelasticity in several bins of total visible energy. Note the inelasticity is on a reverse log scale. For neutrinos of the same energy, the median angular error increases with visible inelasticity but never exceeds 5° up to $y_{vis} = 0.99$. The median angular error also increases at lower energies since fewer photoelectrons are available.

As in Eq. 6.4, this distribution can be used to construct a likelihood that is minimized to find the best estimate of the track's position, time, and direction. The necessary ingredient in this approach is a method to calculate the probability distribution, $p(\Delta t_i)$, describing all photoelectrons per DOM. Analytic parameterizations can be effective but ignore the inhomogeneous optical properties of ice. Instead, $p(\Delta t_i)$ can be obtained from a spline fit of Monte Carlo simulation to light emitted from an infinite, minimum ionizing muon, as described in Sec. 5.3.2.

The spline-based reconstruction described so far is the current state-of-the-art approach used in IceCube's point source searches. However, it is not yet ideal since it only uses the first time and total number of photoelectrons rather than the full photoelectron distribution as done for cascades. Additionally, the light emission profile is assumed to be smooth along the track, which is not a good approximation for the stochastic energy losses of high-energy muons. Nevertheless, it is the most precise direction reconstruction that can be run under current computational constraints. Since the reconstruction was designed for through-going tracks with a smooth energy loss profile, it is natural to question whether it will perform well

for starting tracks, especially for high-inelasticity events where most light is observed from the hadronic cascade rather than the muon. Based on expectations from simulation, the spline-based reconstruction achieves a 1.5° median angular error for the final starting track sample. In Fig. 7.5, the median angular error is shown as a function of visible inelasticity for simulated events in the starting track sample. Even for a visible inelasticity of $y_{vis} = 0.99$, the median angular error does not exceed 5°.

7.2.2 Track Energy and Inelasticity Reconstruction

Having obtained a sample of starting tracks with reliable angular reconstruction, the goal is now to reconstruct their visible energy and inelasticity based on the unfolded energy loss profile of each event. As described in Sec. 6.2.2, the energy loss profile can be unfolded as a sequence of electromagnetic cascades along the reconstructed track. An example energy loss profile for the most energetic starting track observed in full five-year track sample is shown in Fig. 7.6.

Reconstructing the cascade and muon energies from the energy loss profile can be difficult for several reasons. A straightforward approach to identify the cascade energy is to sum the energy loss profile in a distance window after the first nonzero energy loss. However, this can often be fooled by spurious early energy losses that are statistically consistent with zero or result from PMT noise. Furthermore, the optimal distance window over which to sum energy losses can depend strongly on the energy and inelasticity of the event. For low inelasticity events, it is easy to mis-attribute energy losses from the muon to the cascade if the distance window is too large. Similarly for high inelasticity events, if the distance window is too small, some of the cascade energy may be attributed to the muon.

Even if the cascade energy can be reliably identified, reconstructing the energy of the muon faces its own difficulties. Because of the stochastic nature of muon energy loss, large non-Gaussian fluctuations limit the effectiveness of a simple approach based on the finding the average energy loss rate. Truncating the largest energy losses has proven to be an effective solution for through-going tracks [184], but for starting tracks that often leave short tracks in the detector, this throws out valuable information. For example, in events where a large stochastic energy loss took a substantial fraction of the muon's energy, the total deposited energy can be a better energy estimator than the average energy loss rate. All of the above difficulties motivate the application of machine learning to construct the best estimator that will perform well regardless of the energy, inelasticity, track length, or stochasticity of an event.

The problem of reconstructing cascade and muon energies from a set of features for each event is an example of a multi-output regression problem. Ideally, the entire energy loss profile could be used as a set of features, but it has an unequal length depending on the length of the track through the detector. Another approach that can maintain all information about the energy loss profile but ensures an equal length for each event is to use quantiles of the energy loss profile. Taking $n + 1$ total quantiles, the q-th quantile, x_q, is the distance along the track from the entry point where the a fraction q/n of the total deposited energy in

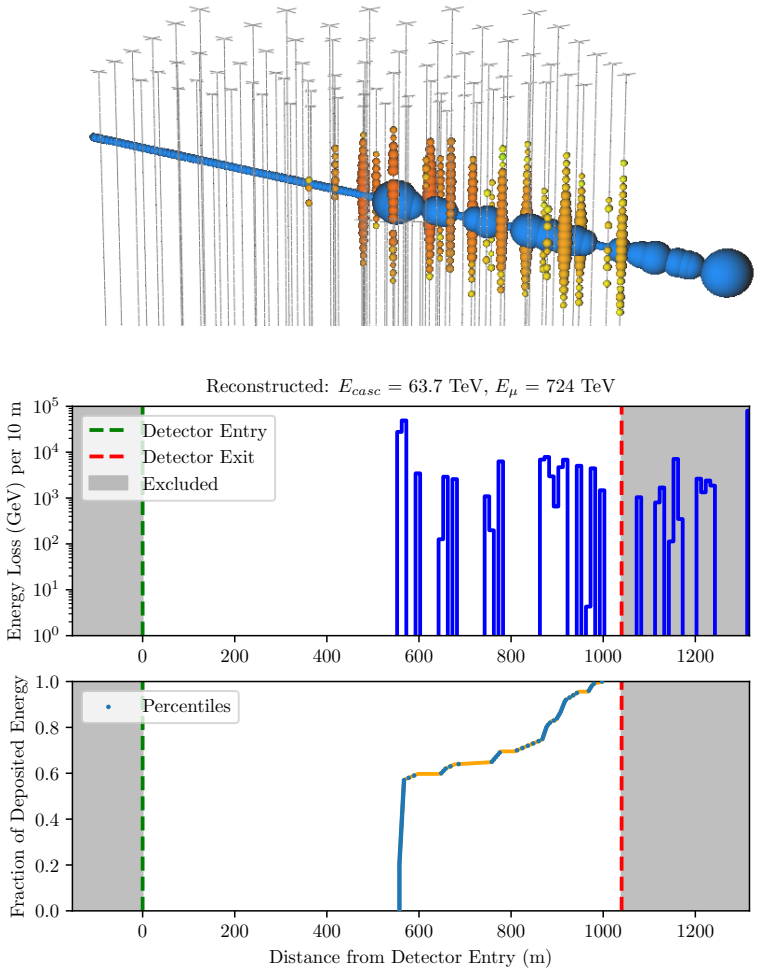

Figure 7.6: Top: A detector view of the most energetic starting track found in the full five-year data sample. The observed number of photoelectrons per DOM is illustrated with colored spheres ranging from red to green, corresponding to the first photoelectron time. The blue spheres correspond to the reconstructed sequence of electromagnetic cascades, and their size is proportional to the reconstructed energy of each cascade on a logarithmic scale. Middle: The reconstructed energy loss profile as a function of distance along the reconstructed track. Entry and exit points of the reconstructed track are shown by green and red dashed lines, respectively. Energy losses outside the detector, shown in grey, are excluded for the purposes of energy and inelasticity reconstruction. Bottom: The cumulative fraction of the total deposited energy with the detector. Percentiles of the energy loss distribution, shown with blue points, are used as features for random forest regression of cascade and muon energy. The cascade and muon energies estimated by random forest regression are $E_{casc} = 63.7\,\text{TeV}$ and $E_\mu = 724\,\text{TeV}$, respectively. This yields values of $E_{vis} = 788\,\text{TeV}$ and $y_{vis} = 0.08$ for the visible energy and inelasticity, respectively. The visible energy is much higher than the total energy of 135 TeV deposited in the detector.

the detector was observed. The first quantile, x_0, corresponds to the distance where the first nonzero energy loss was observed, and the last quantile, x_{n+1}, corresponds to the distance where the last nonzero energy loss was observed. Here, $n = 100$ will be used, and the resulting features are more commonly known as percentiles. Along with the total deposited energy in the detector, E_{dep}, the total track length contained in the detector, L, and the normalized track BDT score, $\hat{s}_{track}$, a set of $n + 4$ features,

$$\mathbf{X}_i = \{\log_{10} E_{dep}, \hat{s}_{track}, L, x_0, x_1, \ldots, x_{n+1}\}, \tag{7.6}$$

is used, which fully characterizes the energy loss profile of an event.

In a multi-output regression problem, the objective is to derive a vector function, $\mathbf{f}(\mathbf{X}_i)$, that predicts the outputs, $\mathbf{y}_i$, for each event. Here the output is taken to be the logarithm of the cascade and muon energies,

$$\mathbf{y}_i = \{\log_{10} E_{casc}, \log_{10} E_{\mu}\}. \tag{7.7}$$

Using a set of training data, one seeks learn the function $\mathbf{f}(\mathbf{X}_i)$ so that it minimizes a metric, here taken to be the mean square error,

$$\text{MSE} = \sum_i w_i \left\|\mathbf{y}_i - \mathbf{f}(\mathbf{X}_i)\right\|^2, \tag{7.8}$$

where w_i is the normalized event weight such that $\sum_i w_i = 1$. The mean square error is an appropriate metric when the residuals, $\mathbf{y}_i - \mathbf{f}(\mathbf{X}_i)$, approximately follow a Gaussian distribution. This is why the logarithm of the cascade and muon energies were chosen as outputs rather than the visible inelasticity, which can have a non-Gaussian residual distribution due to its finite range of $y_{vis} \in [0, 1]$.

To model the function, $\mathbf{f}(\mathbf{X}_i)$, a random forest of decision trees was used. In Sec. 7.1.2, decision trees were detailed as a classification algorithm, but they can also be used for regression problems where the output is continuous rather than a discrete class. The training of decision trees for regression proceeds much as it does for classification. Training data is split at nodes based on the value of the feature that produces the best split into two subsamples according to a metric. In the case of regression, the metric is taken to be the sum of the variance of the outputs for each subsample. Splitting proceeds until certain stopping conditions on the depth of the tree or minimum number of training events are reached. To calculate the output of a new event not in the training sample, the event's features are used to find the leaf node in which it falls, and the mean output of all training data in the leaf node is returned.

A random forest is another ensemble learning technique that attempts to aggregate the results of many decision trees into a single regressor [233]. The strategy of a random forest is to introduce an element of randomness into the training of decision trees and average the output of many trees, which counteracts the tendency of decision trees to over-fit. One way to introduce randomness into training is bootstrapping the training data for every decision tree, i.e. randomly sampling the training data with replacement. Randomness can also be

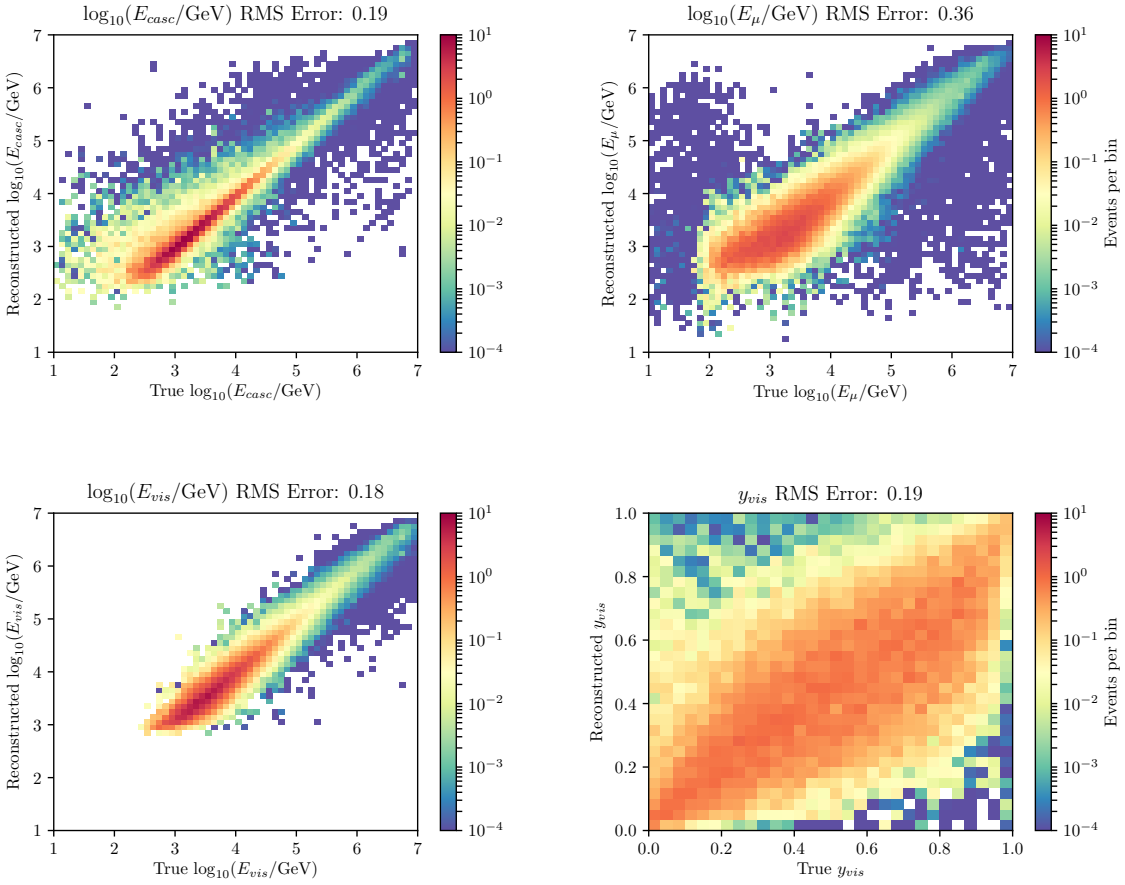

Figure 7.7: From left to right and top to bottom, the joint distribution of reconstructed quantity versus true quantity for cascade energy, E_{casc}, muon energy, E_{μ}, total visible energy, E_{vis}, and visible inelasticity, y_{vis}. The root mean square (RMS) error for each quantity is shown at the top of each panel.

introduced through considering only a random subset of features at each node to determine the best split instead of searching through all features. The software implementation of random forest regression from `scikit-learn` [230] is used here and was found to work well in its default settings.

Training data for the random forest regressor consisted of simulated events from `neutrino-generator` that were weighted to the same atmospheric and astrophysical flux models as used for BDT training. 10% of the simulation, or 105,327 events, was used for training, and another 10% was used for validation. The remaining 80% of the simulated events was reserved for the final data analysis. The training data for the random forest regressor was also independent of the training data used for the BDT.

After training, the 10% validation data was used to quantify the performance of the random forest regression. The root mean square (RMS) error (the square root of Eq. 7.8) is used here to quantify the cascade and muon energy resolution. Though the random forest was trained to learn only the logarithm of the cascade and muon energies, the total visible energy and visible inelasticity are the quantities are of most interest here. Fig. 7.7 shows the two-dimensional joint distribution of the true and reconstructed quantity for cascade energy, muon energy, total visible energy, and visible inelasticity along with an estimate of the RMS error for each. The RMS error is 0.18 for $\log_{10} E_{vis}$, better than the typical energy resolution of 0.22 expected for through-going muons [184]. For reconstructed visible inelasticity, the RMS error is 0.19 and the correlation coefficient with the true visible inelasticity is 0.71.

7.2.3 Cascade Angular Reconstruction

As described in Sec. 6.2.1, cascade direction and energy are reconstructed using the full photoelectron timing distribution recorded by every DOM. For the cascade sample obtained here, the median angle between the true neutrino direction and reconstructed direction is expected to be 15.6° based on simulations with the baseline SPICE Mie ice model described in Sec. 4.3. However, as also shown in Sec. 6.2.1, the photoelectron distributions in individual DOMs from observed events do not always match the best-fit predictions, indicating deficiencies in the model. As a result, verification of the accuracy of cascade angular reconstruction is needed, especially since the cascade zenith angle distribution is a key piece of evidence that enables astrophysical and atmospheric neutrino fluxes to be distinguished. Additionally, neutrino-induced cascades have recently been used to set limits on neutrino point sources in the Southern sky, further increasing the need to verify angular reconstruction [234].

The reconstruction of track inelasticity from the previous section provides an opportunity to study the performance of cascade angular reconstruction. For events that have a high visible inelasticity, the direction of the low energy muon produced in the interaction can be reconstructed accurately, as shown in Fig. 7.5, since it is driven by first photoelectron times that are causally disconnected from the cascade vertex. At the same time, a cascade angular reconstruction applied to high inelasticity events should not be strongly affected by the presence of a low energy muon since it is driven by the photoelectron distribution in DOMs close to the cascade vertex where light has not been strongly scattered away from the Cherenkov cone. Since the hadrons and muon of a ν_μ CC interaction are boosted to nearly the same direction in the laboratory frame at high energies, the distribution of the angle between the cascade and track reconstructions can be used to check the performance of cascade angular reconstruction.

In Fig. 7.8, the distribution of the difference in zenith angle between cascade and track reconstructions is shown for high-inelasticity events with $y_{vis} > 0.75$. The focus is placed on the zenith angle since it is of primary of interest in disentangling astrophysical and atmospheric fluxes. The zenith angle is also more likely to be mis-reconstructed due to imperfect knowledge of ice properties and DOM angular sensitivity, which are both approximately azimuthally symmetric. The most striking feature of Fig. 7.8 is that the median of the observed

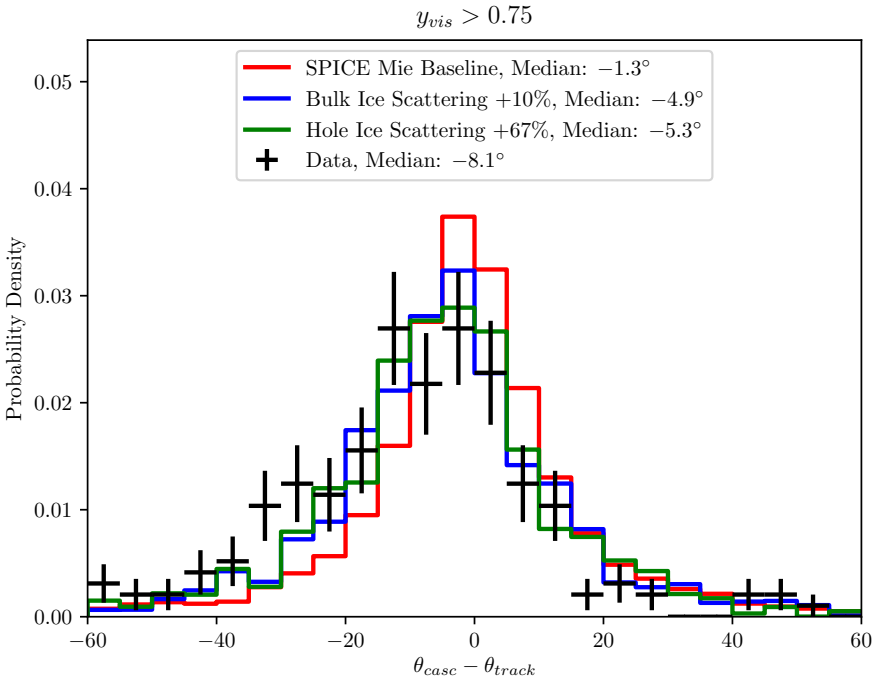

Figure 7.8: The distribution of the difference in zenith angle between spline-based cascade and track reconstructions. Track events with visible inelasticity $y_{vis} > 0.75$ are shown, which are chosen to minimize the effect of the outgoing muon on the cascade reconstruction. 192 events meeting this condition were found in the full 5-year data sample. Expectations from simulation are shown for the baseline SPICE Mie ice model (red), increasing bulk ice scattering globally by +10% (blue), and increasing hole ice scattering by +67% (green).

distribution is −8.1°, which indicates a bias in reconstructing cascades to be more down-going that they actually are. Simulation based on the baseline SPICE Mie ice model only predicts a −1.3° bias, and a number systematic effects were studied that could reproduce the bias seen in data.

Two effects, increased scattering in the bulk glacial ice and increased scattering on air bubbles in the column of "hole ice" surrounding DOMs were found that produce the largest bias. When increasing the scattering coefficients of the bulk ice globally by +10% from the baseline SPICE Mie ice model, a factor consistent with LED flasher studies [175], a bias of −4.9° was found in simulated events. For hole ice, a uniform effective scattering length of 50 cm is assumed inside the bore-hole of every string in order to derive the DOM angular sensitivity correction described in Sec. 5.3.1. However, this model has long been known to be an over-simplification, and the amount of scattering in hole ice is poorly known. Based on calibration studies with LED flashers and down-going muons, an effective scattering length of 30 cm in the hole ice, or a +67% increase in the scattering coefficient, is also

consistent with data. When using a DOM angular sensitivity correction based on a 30 cm effective scattering length, a bias of $-5.3°$ is found. Thus, increased scattering in bulk ice, hole ice, or a combination of both are plausible explanations for the down-going bias in cascade reconstruction seen in high-inelasticity tracks. Accounting for both effects will be an important part of quantifying detector systematic uncertainties in Sec. 7.3.2.

7.3 Model and Likelihood Fit

The data analysis methods used here closely follow those detailed already in Sec. 6.3 and Appendix A. All results will derive from a binned Poisson likelihood fit, where the expected number of counts in a bin, $\mu_i(\boldsymbol{\theta})$, depends on a set of model parameters, $\boldsymbol{\theta}$, that include physics parameters to be measured and nuisance parameters that describe systematic uncertainties. Because of the much larger sample size obtained here compared to the previous chapter, a number of systematic uncertainties that were previously sub-dominant to statistical uncertainties now have an important effect. This section describes the model and all parameters entering a likelihood fit.

Results on the atmospheric neutrino to antineutrino ratio, neutrino charm production, and the astrophysical neutrino flux will all derive from a joint fit on the starting track and cascade samples. For the starting track sample, three dimensions are used: reconstructed total visible energy, visible inelasticity, and cosine zenith. For the cascade sample, only the two dimensions of reconstructed energy and cosine zenith are used. In reconstructed energy, nine half-decade bins are used ranging from $10^{2.5}$ GeV to 10^7 GeV. In cosine zenith, five bins are used ranging from -1 to 1. In reconstructed inelasticity, five bins are used from 0 to 1. Bin sizes are limited by finite simulation statistics, but they are chosen to be comparable to the estimated resolution for each quantity. In the construction of a model and description of systematic uncertainties, the focus will be placed on describing event distributions according to this binning, which are illustrated in Fig. 7.9 and Fig. 7.11.

7.3.1 Neutrino Flux Models

To describe both atmospheric and astrophysical neutrinos, `neutrino-generator` simulation is used. Equal numbers of each neutrino flavor were generated according to a $E^{-1.5}$ spectrum from 10^2 GeV to 10^7 GeV. To increase the efficiency of generating starting events, neutrinos were forced to interact within a cylinder with a height of 1200 m and a radius of 600 m centered on the detector. 3,345,402 simulated events survive the event selection, which are independent of all training data used for BDT classification and random forest regression. Event weights are calculated according to Eq. 5.2 using the combined atmospheric and astrophysical flux as parameterized below. Atmospheric muon background is negligible and is not included.

The conventional atmospheric neutrino flux is described by the HKKMS calculation extrapolated above 10 TeV and corrected for the H3a cosmic-ray model as described in

Sec. 3.3.1. As shown in Tab. 7.1, the sample is dominated by atmospheric neutrinos due to the low energy threshold of the event selection, and a more sophisticated modeling of atmospheric flux uncertainties is necessary. In addition to an overall scaling factor of the conventional flux, uncertainties in the cosmic-ray power-law index and relative contribution of pions and kaons to the atmospheric flux must be taken into account that distort the energy and zenith angle distribution. To describe these uncertainties, the conventional atmospheric flux of each neutrino type can be parameterized as

$$\begin{aligned}
\Phi_{\nu_\mu}(E,\cos\theta) &= \Phi_{\text{conv}} R_{\nu_\mu/\bar{\nu}_\mu} \left(\Phi_{\pi\nu_\mu} + R_{K/\pi}\Phi_{K\nu_\mu}\right) \left(\frac{E}{E_0}\right)^{-\Delta\gamma}, \\
\Phi_{\bar{\nu}_\mu}(E,\cos\theta) &= \Phi_{\text{conv}} \left(\Phi_{\pi\bar{\nu}_\mu} + R_{K/\pi}\Phi_{K\bar{\nu}_\mu}\right) \left(\frac{E}{E_0}\right)^{-\Delta\gamma}, \\
\Phi_{\nu_e}(E,\cos\theta) &= \Phi_{\text{conv}} \left(\Phi_{\pi\nu_e} + R_{K/\pi}\Phi_{K\nu_e}\right) \left(\frac{E}{E_0}\right)^{-\Delta\gamma}, \\
\Phi_{\bar{\nu}_e}(E,\cos\theta) &= \Phi_{\text{conv}} \left(\Phi_{\pi\bar{\nu}_e} + R_{K/\pi}\Phi_{K\bar{\nu}_e}\right) \left(\frac{E}{E_0}\right)^{-\Delta\gamma}.
\end{aligned} \tag{7.9}$$

Here $\Phi_{\pi\nu}$ and $\Phi_{K\nu}$ are flux templates for the contribution of pions and kaons to the atmospheric neutrino flux, respectively, which are obtained from the HKKMS calculation and have the analytic parameterization of Eq. 3.17. An overall flux scaling parameter, Φ_{conv}, acts equally on each neutrino type, and since it is well-constrained from the data, no prior information is needed. Uncertainties on the cosmic-ray power-law index are described through the parameter $\Delta\gamma$, which acts to steepen the neutrino energy spectrum for positive values. A Gaussian prior of $\Delta\gamma = 0 \pm 0.05$ is used, spanning the difference in cosmic-ray spectral indices from the H3a [72] and the Hoerandel [235] cosmic-ray models. A reference energy of $E_0 = 5\,\text{TeV}$ is chosen for the cosmic-ray index correction, which is close to the median neutrino energy of the sample and reduces the correlation with the normalization, Φ_{conv}. The scaling factor $R_{K/\pi}$ describes the enhancement of kaon production in the atmospheric with respect to the HKKMS calculation. Based on a previous measurement of $R_{K/\pi}$ using low-energy cascades in IceCube [121], a Gaussian prior of $R_{K/\pi} = 1 \pm 0.5$ is used. Lastly, a scaling factor on the atmospheric ν_μ to $\bar{\nu}_\mu$ flux ratio of the HKKMS flux calculation, $R_{\nu_\mu/\bar{\nu}_\mu}$, will be measured in Sec. 7.4.3 but is otherwise fixed to be 1.

The prompt atmospheric neutrino flux is described by the BERSS calculation, an update to the ERS calculation used in Ch. 6. The BERSS flux is approximately a factor of 2 smaller than the ERS flux across all energies due to the inclusion of recent LHC data in the calculation of charm production cross sections. An overall scaling factor of the BERSS flux, Φ_{prompt}, is sufficient to describe its uncertainties, and no prior information will be used so as to be as conservative as possible given the large variation in theoretical predictions.

Astrophysical neutrinos are again described by an isotropic flux for each flavor with

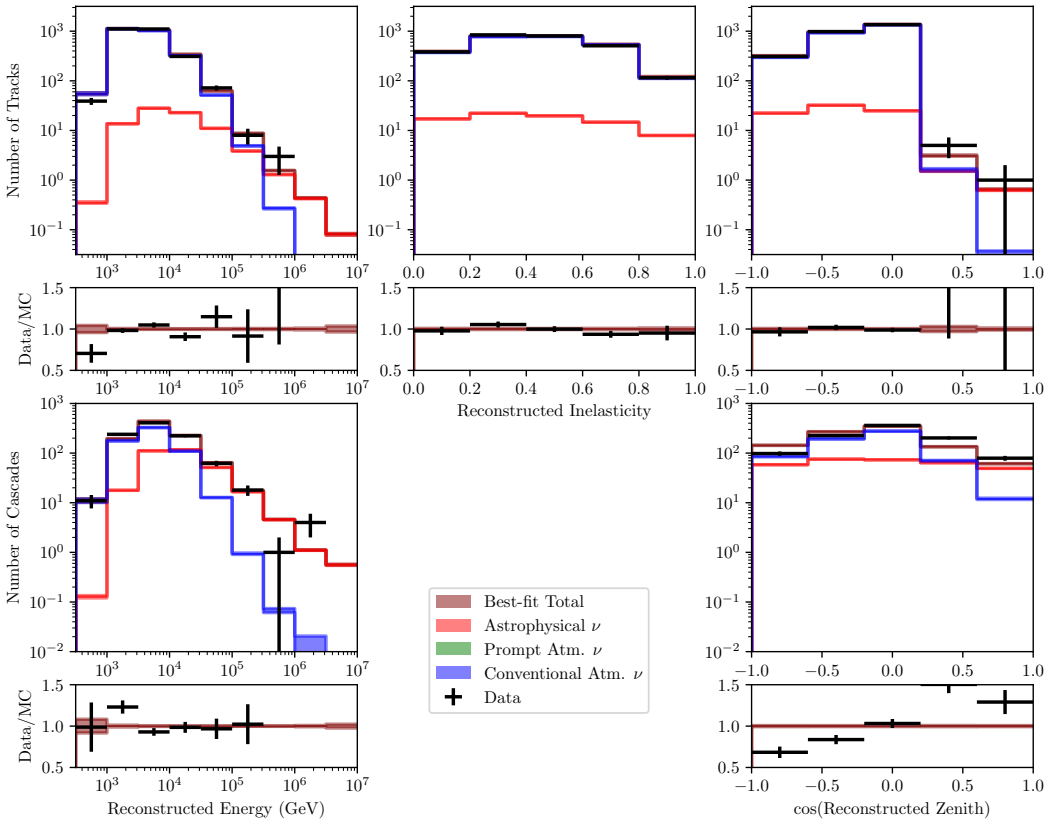

Figure 7.9: Best-fit distributions when only neutrino flux parameters and no detector systematic uncertainties are included. Top: The distribution of reconstructed visible energy, visible inelasticity, and cosine zenith for the sample of starting tracks. Bottom: the distribution of reconstructed energy and cosine zenith for the sample of cascades. Contributions of conventional atmospheric and astrophysical neutrinos are shown in blue and red, respectively, and the total predicted distribution is shown in maroon. The width of colored lines corresponds to the statistical uncertainty in the simulation prediction. The prompt atmospheric neutrino contribution is not shown since its best-fit value is zero. Black error bars indicate observed data from the complete five years of data. The cosine zenith distribution of cascades (bottom right) shows a large excess of down-going cascades and deficit of up-going cascades unaccounted by the model. The best-fit parameters are shown in Tab. 7.2.

Parameter	68% CI	Prior
Φ_{conv}	$1.08^{+0.02}_{-0.02}$	Flat $[0, \infty)$
$\Delta\gamma_{\text{cr}}$	$0.03^{+0.03}_{-0.03}$	Gaussian 0.00 ± 0.05
$R_{K/\pi}$	$0.76^{+0.32}_{-0.21}$	Gaussian 1.00 ± 0.50
Φ_{prompt}	$0.00^{+0.34}_{-0.00}$	Flat $[0, \infty)$
$\Phi_0/10^{-18}\,\text{GeV}^{-1}\text{s}^{-1}\text{cm}^{-2}\text{sr}^{-1}$	$2.05^{+0.21}_{-0.20}$	Flat $[0, \infty)$
γ	$2.64^{+0.07}_{-0.07}$	Flat
Test Statistic $-2\ln\lambda$	235.43	–

Table 7.2: The best-fit neutrino flux parameters when no detector systematic uncertainties are included. 68% confidence intervals are shown as calculated by the profile likelihood method. The prior distribution for each parameter is shown in the last column. The goodness-of-fit statistic is shown in the last row.

power-law energy spectrum,

$$\Phi_\alpha(E) = 3\Phi_0 f_{\alpha,\oplus} \left(\frac{E}{100\,\text{TeV}}\right)^{-\gamma}, \tag{7.10}$$

where $f_{\alpha,\oplus}$ is the fraction of each flavor at Earth, γ is the power-law index, and Φ_0 is the average flux of neutrinos and antineutrinos per flavor at 100 TeV. Unless otherwise noted, the flavor composition of the astrophysical flux will be fixed at $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$, and an equal contribution of neutrinos and antineutrinos is assumed. No prior information on the power-law index, γ, and normalization, Φ_0, will be used. The validity of the power-law model of the astrophysical flux should not be taken for granted given the unknown nature of cosmic sources, and it is important to verify it before moving on to more sophisticated measurements.

Since `neutrino-generator` simulation does not include accompanying muons from cosmic-ray air showers, it is necessary to correct for the atmospheric self-veto probability. Again, the analytic calculation of [141] is used to correct both the conventional and prompt atmospheric fluxes. To approximate the conditions that cause an event to be vetoed, it is required that no muons exist in the air shower that survive to a depth of 1950 m (the center of IceCube) with an energy above 100 GeV. This choice of the muon energy threshold is lower than the 1 TeV threshold used in Sec. 6.3.1 due to the improved muon rejection of the event selection here. The choice of a 100 GeV muon energy threshold was verified with `CORSIKA` simulation of the full detector response to cosmic-ray air showers, analogous to Fig. 6.9.

When the data is fit with the above neutrino flux models using a baseline detector simulation with the SPICE Mie ice model, the best-fit parameters shown in Tab. 7.2 are obtained. As shown in Fig. 7.9, the observable distributions show agreement with the model except for cascade zenith angle distribution. An excess of down-going cascades and a deficit of up-going cascades is seen that cannot be described by the neutrino flux model or the baseline detector simulation. This seemingly indicates a down-going bias in cascade angular reconstruction. In Sec. 7.2.3, a study of high-inelasticity starting tracks confirmed a down-

going bias in cascade angular reconstruction and furthermore indicated increased scattering in bulk and/or hole ice could explain the bias. Given that the determination of astrophysical flux depends strongly on knowledge of the cascade zenith distribution, it is important to understand this discrepancy in the next section.

7.3.2 Detector Systematic Uncertainties

The larger sample size also makes it necessary to improve the modeling of detector systematic uncertainties compared to Ch. 6. Unlike neutrino flux uncertainties, it is not straight-forward to evaluate the change detector systematic uncertainties cause in observable distributions through event re-weighting. Four different detector systematic uncertainties are considered here, scaling factors on the DOM optical effiency, scattering and absorption in bulk glacial ice, and scattering by air bubbles in hole ice. The following section describes the scheme used to incorporate them as additional nuisance parameters in a likelihood fit.

For all of the detector systematic uncertainties considered here, it is possible only to simulate them one at a time by varying each scaling factor by some discrete amount, running a new detector simulation, and calculating the new distribution of observables. For the DOM optical efficiency scaling parameter, ϵ/ϵ_0, four different simulations were run with values $\epsilon/\epsilon_0 \in \{0.9090, 0.9595, 1.10, 1.21\}$. A Gaussian prior of $\epsilon/\epsilon_0 = 1 \pm 0.1$ is used, consistent with calibration studies of minimum ionizing muons [168]. For a global scaling factor on the bulk ice scattering coefficients, $\alpha_{\text{Scat.}}$, a single simulation was run with $\alpha_{\text{Scat.}} = 1.1$. Likewise, for a global scaling factor on the bulk ice absorption coefficients, $\alpha_{\text{Abs.}}$, a single simulation was run with $\alpha_{\text{Abs.}} = 1.1$. Another simulation with lowered bulk ice scattering and absorption, $\alpha_{\text{Scat.}} = 0.929$ and $\alpha_{\text{Abs.}} = 0.929$, was also tested but never found to have an improved likelihood for any fit that will be done later, so it was ignored. Gaussian priors of $\alpha_{\text{Scat.}} = 1 \pm 0.1$ and $\alpha_{\text{Abs.}} = 1 \pm 0.1$ are used, corresponding to uncertainties estimated from LED flasher studies [175]. Lastly, a simulation with a 30 cm effective scattering length in hole ice was done. Using a scaling factor of the hole ice scattering coefficient, $\alpha_{\text{Hole Ice}}$, this corresponds to a value of $\alpha_{\text{Hole Ice}} = 1.67$. Since hole ice scattering is poorly known and both 50 cm and 30 cm effective scattering lengths are consistent with available calibration data at the time of writing, a flat prior in the range $[1, 1.67]$ is used for $\alpha_{\text{Hole Ice}}$.

To treat each scaling factor, θ_j, as a continuous nuisance parameter, an interpolation scheme is necessary so that one can move smoothly between the bin expectation, μ_i^b, of the baseline simulation with baseline nuisance parameter value θ_j^b to the bin expectation for an arbitrary value of the nuisance parameter, $\mu_i(\theta_j)$. For small variations of the nuisance parameter, the bin expectation can be approximated by a first-order Taylor series,

$$\ln \mu_i(\theta_j) = \ln \mu_i^b + a_{ij}(\theta_j - \theta_j^b). \tag{7.11}$$

If a single simulation is run with an altered value of the nuisance parameter, θ_j^s, and yields a bin expectation, $\mu_i^s(\theta_j^s)$, the value of a_{ij} can be found from

$$a_{ij} = \frac{\ln \mu_i^s(\theta_j^s) - \ln \mu_j^b}{\theta_j^s - \theta_j^b} \tag{7.12}$$

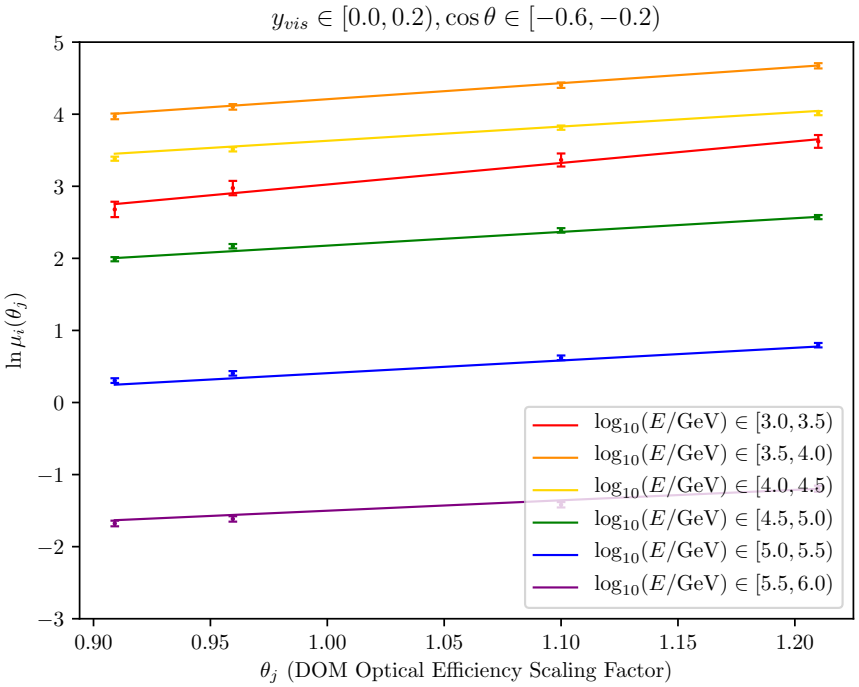

Figure 7.10: The bin expectation as a function of the DOM optical efficiency scaling factor for six half-decade bins with reconstructed energy 1TeV to 1 PeV, reconstructed inelasticity $y_{vis} \in [0, 0.2)$, and reconstructed cosine zenith $\cos\theta \in [-0.6, -0.2)$. Colored error bars show the bin expectation for simulations run with four different values of DOM optical efficiency. The colored curves show the χ^2 fit to the parameterization of Eq. 7.11. The first-order approximation works well to predict the bin expectation from each simulation.

When a simulation is run for multiple values of the nuisance parameter, $\theta_j^{s,k}$, indexed by k and yielding bin expectations, $\mu_i^s(\theta_j^{s,k})$, one can obtain the value of a_{ij} from a χ^2 fit for each bin. The χ^2 statistic for the fit is

$$\chi^2_{ij} = \sum_k w_{ij}^k \left(\ln \mu_i^s(\theta_j^{s,k}) - \ln \mu_j^b - a_{ij}(\theta_j^{s,k} - \theta_j^b) \right)^2 . \tag{7.13}$$

Since simulations of variations in nuisance parameters often have a small sample size, it is necessary to account for the statistical uncertainty in the bin expectation, $\sigma_i^s(\theta_j^{s,k})$. This is reflected in the weights, $w_{ij}^k = \left(\sigma_i^s(\theta_j^{s,k}) / \mu_i^s(\theta_j^{s,k}) \right)^{-2}$, appearing in Eq. 7.13. The χ^2 statistic can be minimized analytically to obtain,

$$a_{ij} = \frac{\sum_k w_{ij}^k \left(\ln \mu_i^s(\theta_j^{s,k}) - \ln \mu_j^b \right) \left(\theta_j^{s,k} - \theta_j^b \right)}{\sum_k w_{ij}^k \left(\theta_j^{s,k} - \theta_j^b \right)^2} . \tag{7.14}$$

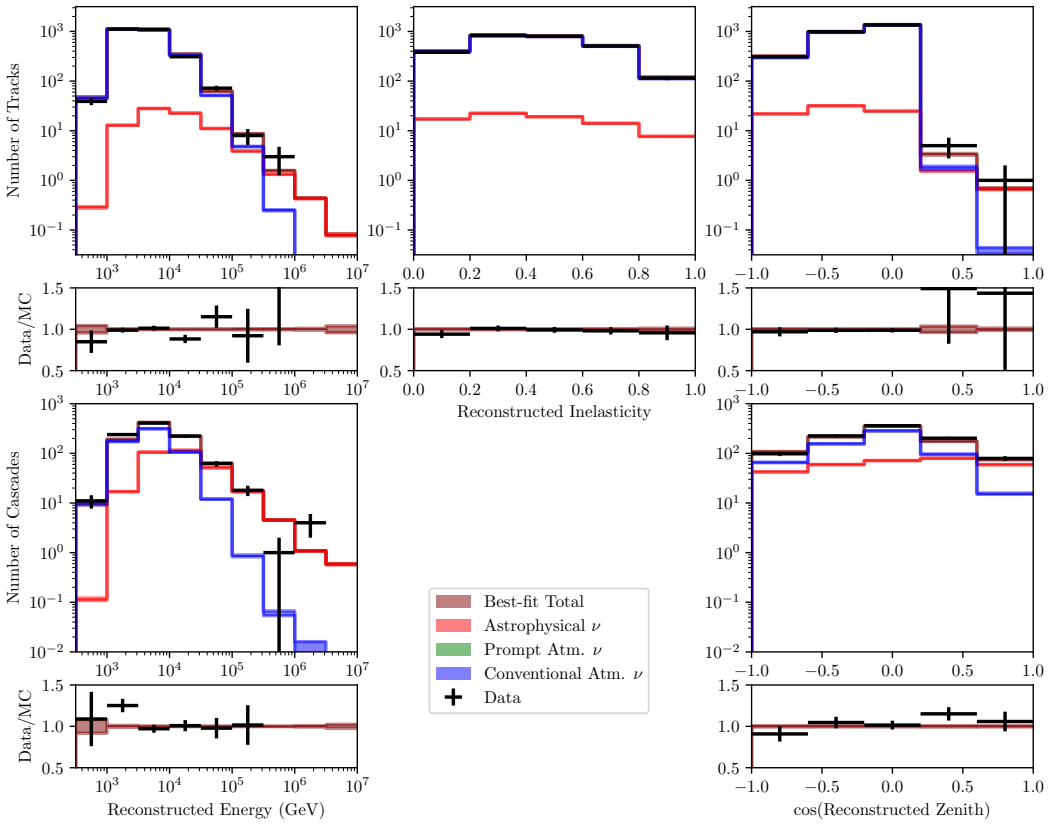

Figure 7.11: Best-fit distributions with all neutrino flux and detector parameters. Top: The distribution of reconstructed visible energy, visible inelasticity, and cosine zenith for the sample of starting tracks. Bottom: the distribution of reconstructed energy and cosine zenith for the sample of cascades. Contributions of conventional atmospheric and astrophysical neutrinos are shown in blue and red, respectively, and the total predicted distribution is shown in maroon. The prompt atmospheric neutrino contribution is not shown since its best-fit value is zero. Black error bars indicate observed data from the complete five years of data. The inclusion of detector parameters describing bulk and hole ice scattering substantially improves the model description of the cascade cosine zenith distribution. The best-fit parameters are shown in Tab. 7.3.

Parameter	68% CI	Prior
Φ_{conv}	$1.05^{+0.07}_{-0.06}$	Flat $[0, \infty)$
$\Delta\gamma_{\text{cr}}$	$0.04^{+0.03}_{-0.03}$	Gaussian 0.00 ± 0.05
$R_{K/\pi}$	$1.11^{+0.35}_{-0.28}$	Gaussian 1.00 ± 0.50
Φ_{prompt}	$0.00^{+1.10}_{-0.00}$	Flat $[0, \infty)$
$\Phi_0/10^{-18}\,\text{GeV}^{-1}\text{s}^{-1}\text{cm}^{-2}\text{sr}^{-1}$	$2.04^{+0.23}_{-0.21}$	Flat $[0, \infty)$
γ	$2.62^{+0.07}_{-0.07}$	Flat
ϵ/ϵ_0	$1.08^{+0.03}_{-0.03}$	Gaussian 1.00 ± 0.10
$\alpha_{\text{Scat.}}$	$1.07^{+0.01}_{-0.01}$	Gaussian 1.00 ± 0.10
$\alpha_{\text{Abs.}}$	$1.02^{+0.02}_{-0.02}$	Gaussian 1.00 ± 0.10
$\alpha_{\text{Hole Ice}}$	$1.46^{+0.12}_{-0.12}$	Flat $[1.00, 1.67]$
Test Statistic $-2\ln\lambda$	175.54	–

Table 7.3: The best-fit parameters including all neutrino flux and detector systematic uncertainties. 68% confidence intervals are shown as calculated by the profile likelihood method. The prior distribution for each parameter is shown in the last column. Though the goodness-of-fit test statistic, $-2\ln\lambda$, shows a substantial improvement when detector parameters are included, the neutrino flux parameters are not largely altered.

To calculate the bin expectation when several nuisance parameters deviate from their baseline value, the generalization of Eq. 7.11 is

$$\ln\mu_i(\theta_1, \theta_2, \ldots, \theta_n) = \ln\mu_i^b + \sum_{j=1}^{n} a_{ij}(\theta_j - \theta_j^b). \tag{7.15}$$

which again reflects a first-order n-dimensional Taylor series expansion. The validity of the first-order approximation is worth checking. The DOM optical efficiency was the only detector nuisance parameter for which simulations of multiple values were available, but as shown in Fig. 7.10, the approximation works quite well for several example bins. Lastly, it is important to note that the values of a_{ij} also depend on the remaining flux nuisance parameters that enter the calculation of μ_i^b, and the above procedure is repeated for every likelihood evaluation where the values of these parameters change.

When the data is fit to this model including both neutrino flux and detector parameters, the best-fit values shown in Tab. 7.3 are obtained. The corresponding best-fit observable distributions are shown in Fig. 7.11. Notably, the description of the cascade cosine zenith distribution is much improved, which is largely due the increased bulk ice scattering and hole ice scattering parameters, $\alpha_{\text{Scat.}} = 1.07^{+0.01}_{-0.01}$ and $\alpha_{\text{Hole Ice}} = 1.46^{+0.12}_{-0.12}$. This confirms the hypothesis from the study of high-inelasticity tracks that increased bulk and hole ice scattering introduce a down-going bias in cascade reconstruction. Additionally, a similar scaling factor of $\alpha_{\text{Scat.}} = 1.07 \pm 0.02$ was also found in an independent analysis of neutrino-induced cascades in IceCube [236].

The goodness-of-fit test statistic also shows a substantial improvement from $-2\ln\lambda = 235.43$ to $-2\ln\lambda = 175.54$ when detector parameters are incorporated, far in excess of the number of degrees of freedom added. Despite this large improvement in the model's description of the data, the remaining astrophysical and atmospheric neutrino flux parameters are not largely affected by the inclusion of detector systematic uncertainties. Most notably, the atmospheric kaon enhancement factor, $R_{K/\pi} = 1.11^{+0.35}_{-0.28}$, shifts due to the improved description of the cascade zenith distribution, and the 68% upper limit on the prompt flux normalization was increased to $\Phi_{\text{prompt}} < 1.10$ and is now compatible with the BERSS calculation. A goodness-of-fit p-value of $p = 0.04$ can be computed using pseudo-experiments generated from the best-fit distributions of Fig. 7.11. This relatively low value might indicate there are additional deficiencies in the detector model, for example ice anisotropy and alternate models of hole ice, but given current computational limitations, this is the best detector model that can be simulated under current computational constraints.

7.4 Results

With a model and likelihood fit in place, a wide range of measurements can be done with inelasticity distributions. To start, a measurement of the energy spectrum of astrophysical neutrinos will be done. Special emphasis will be placed on verifying a power-law energy spectrum, which will be an important assumption for later results on flavor composition, neutrino to antineutrino ratio, and neutrino charm production.

7.4.1 Energy Spectrum of Astrophysical Neutrinos

In Ch. 6, an equal-flavor astrophysical neutrino flux with a power-law energy spectrum was found to be a good description of available data. Since that result, there have been hints that the energy spectrum may be more complex than a power-law, and it is worth investigating the astrophysical flux model again with the improved background rejection, track and cascade classification, and starting track energy resolution obtained here.

In 2015, evidence for an astrophysical flux was found in a sample of through-going ν_μ-induced tracks from the Northern sky in two years of IceCube data [181]. The observed power-law index was $\gamma = 2.2 \pm 0.2$, in moderate tension with the results of Ch. 6 and the value of $\gamma = 2.62 \pm 0.07$ obtained here. An even steeper power-law index of $\gamma = 2.92^{+0.33}_{-0.29}$ was found in an update of the original high-energy starting event analysis to include six years of data [237]. In 2016, an update of the through-going track analysis to include six years of data found a power-law index of $\gamma = 2.13 \pm 0.13$, which only increased the tension with results based on starting events.

Before investigating the tension in power-law indices, it is important to point out that due to the larger atmospheric ν_μ background, through-going track samples are sensitive to the astrophysical flux at higher energies than starting event samples. The six-year analysis was estimated to have sensitivity to an astrophysical flux in the energy range from 194 TeV

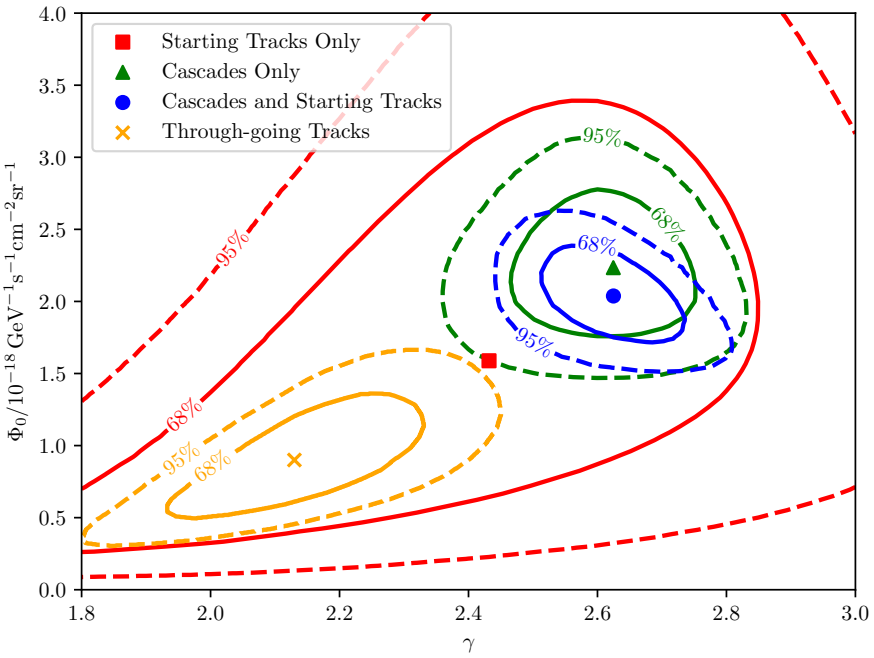

Figure 7.12: Confidence regions for the astrophysical power-law index, γ and flux normalization, Φ_0. The blue contours show the confidence region for the joint fit of the cascade and starting track samples, which is the main result obtained here. The red contours show the confidence region for a fit of starting tracks only, and the green contours show the confidence region for a fit of cascades only, and both are consistent with each other. The confidence region from the IceCube analysis of through-going tracks [136] is shown in orange, which is in significant tension with the cascade confidence region. The contours from the starting track sample are consistent with both cascade and through-going samples. Since the through-going track analysis is sensitive at higher energies, the discrepancy may point to a hardening of the astrophysical power-law index at high energies.

to 7.8 PeV. The method used to construct this sensitive energy range was different than that used previously in Sec. 6.4.1. The method attempts to quantify the central 90% of neutrino energies contributing to the test statistic between the best-fit astrophysical model and an atmospheric background-only model. Since the same method will also be used to quantify to the sensitive energy range for other measurements in this chapter, it is described fully in Appendix B. Applied to the starting event sample here, the sensitive energy range is 3.5 TeV to 2.6 PeV. For completeness sake, the sensitive energy range when using the method of Sec. 6.4.1 is 6.6 TeV to 2.2 PeV, illustrating the considerable improvement at low energies attained by this event selection compared to the previous chapter.

Taken at face value, the two different power-law indices obtained in the two different

energy ranges may indicate that there is a hardening of the astrophysical power-law index above $\sim 200\,\text{TeV}$. The discrepancy in the astrophysical power-law index could also have a more mundane explanation, for example a heretofore unknown systematic bias affecting cascades or tracks, or an even more exotic explanation, for example a highly energy-dependent flavor composition induced by new physics in the cosmic propagation of neutrinos. To explore these possibilities, it is instructive to study the astrophysical flux separately with the starting track sample and cascade sample obtained here.

In Fig. 7.12, confidence regions on the astrophysical flux normalization, Φ_0, and power-law index, γ, are shown for separate fits on the starting track sample, cascade sample, and the combined sample. Confidence regions for the six-year through-going track sample are also shown. For all results, a flavor composition of $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ is assumed. Although the power-law index of $\gamma = 2.43^{+0.28}_{-0.30}$ for starting tracks is harder than the power-law index of $\gamma = 2.62^{+0.08}_{-0.08}$ for cascades, the two results are compatible, and there is no evidence for a discrepancy between the two samples. The starting track results are also compatible with the through-going track results. It will take a much larger starting track sample size to confirm or refute a discrepancy between the results. Additionally, the similarity in the confidence regions for the cascade sample and the combined sample illustrates how measurement of the astrophysical neutrino energy spectrum is dominated by information from cascades.

With no significant discrepancy observed in the power-law index of cascades and starting tracks, one can search for evidence of a more sophisticated structure in the astrophysical neutrino energy spectrum. As in Sec. 6.4.1, the energy spectrum can be measured in a model-independent way through an unfolding. Here the astrophysical flux is unfolded in logarithmic, half-decade bins that are commensurate to the bins in reconstructed energy, and the flux is assumed to follow an E^{-2} spectrum in each bin. The unfolded astrophysical flux is shown in Fig. 7.13 along with the best-fit power-law energy spectrum. From $10^{3.5}\,\text{GeV}$ to $10^{5.5}\,\text{GeV}$, the unfolded flux matches the best-fit power-law flux well. At energies above $10^{5.5}\,\text{GeV}$, all bins are empty except for the bin from $10^{6}\,\text{GeV}$ to $10^{6.5}\,\text{GeV}$, which shows a non-zero flux higher than that expected from the best-fit power law. This bin reflects the four highest energy neutrinos observed in the sample with reconstructed energy around $\sim 1\,\text{PeV}$. The best-fit power-law flux for the six-year through-going track sample is also shown, and notably it better matches the unfolded flux in this bin. This may point to evidence for a hardening in the power-law index within the starting event sample alone, and it worth performing a specific hypothesis test.

In the simplest model that can be tested, a hardening of the power-law index could be explained by two source populations that produce a neutrino energy spectrum with two different power-law indices, γ_1 and γ_2. The total astrophysical neutrino flux can then be parameterized as a double power-law,

$$\Phi_\alpha(E) = 3\Phi_1 f_{\alpha,\oplus} \left(\frac{E}{100\,\text{TeV}}\right)^{-\gamma_1} + 3\Phi_2 f_{\alpha,\oplus} \left(\frac{E}{100\,\text{TeV}}\right)^{-\gamma_2}, \tag{7.16}$$

where Φ_1 and Φ_2 represent the flux at 100 TeV contributed by each source population. An equal flavor composition of $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ will be assumed for each population. In Fig. 7.14,

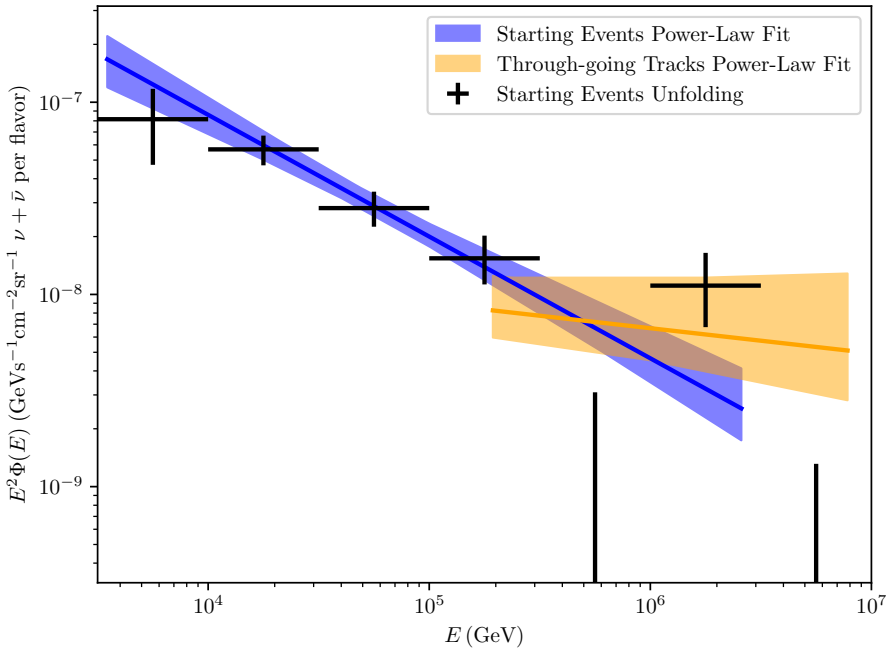

Figure 7.13: The unfolded energy spectrum of astrophysical neutrinos is shown in black. 68% confidence intervals are shown and upper limits are shown for bins that have a best-fit flux of zero. The blue band shows the best-fit power-law astrophysical flux with 68% confidence level uncertainties for the starting track and cascade sample here. The best-fit power-law from the through-going track analysis of [136] is shown in orange. The unfolded spectrum shows hints of a hardening above 1 PeV that is consistent with the results of the power-law through-going track analysis, but it is not possible to state conclusively that a hardening exists with the current starting event sample alone.

confidence regions for the power-law index, γ_1, and flux normalization, Φ_1, are shown, which equally apply to Φ_2 and γ_2 due to symmetry. The best-fit model has equal indices $\gamma_1 = \gamma_2 = 2.62$ or equivalently one of $\Phi_1 = 0$ or $\Phi_2 = 0$ within a numerical tolerance of 10^{-4}, and there is no evidence for a double power-law structure. Nevertheless, the confidence regions do allow for a harder power-law component with $\gamma_1 < 2.62$. Confidence regions for the single power-law fit of through-going tracks are also shown, and there is overlap in the 68% confidence region for Φ_1 and γ_1. Thus, there is no significant incompatibility with the through-going track results under a double power-law model. Although the results of the through-going track analysis are compelling, there is no evidence for a double power law within the starting event sample alone at the focus of this work, and a single power-law model will be used for all remaining results. A future joint fit of starting event and through-going track samples will be able to conclusively determine whether there is evidence for a double power-law or a

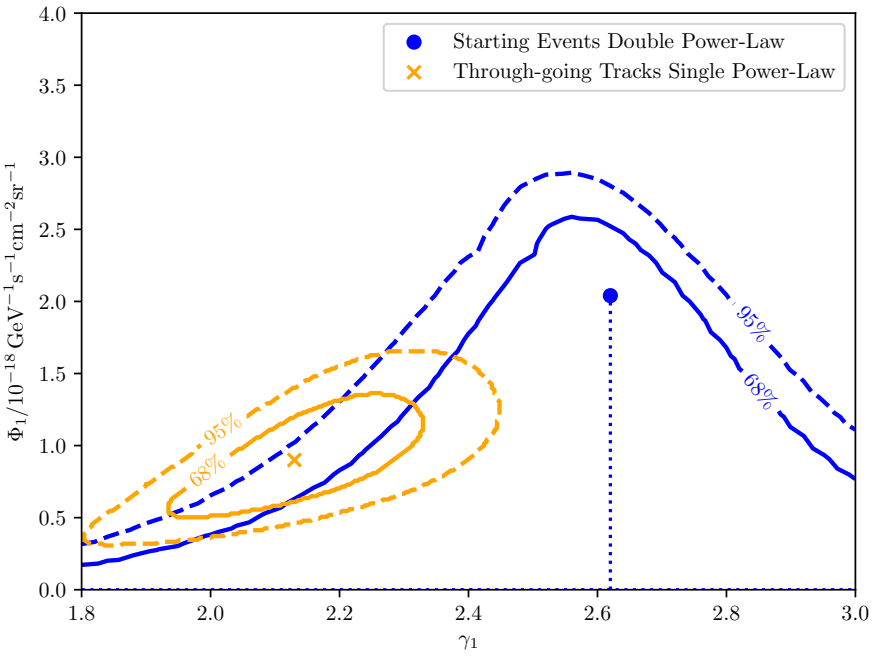

Figure 7.14: Confidence regions on the power-law index, γ_1, and flux normalization, Φ_1, of one of the two components of a double power-law flux model are shown in blue for the starting event sample here. The best-fit model is a single power-law, and as a result there is a degeneracy in Φ_1 at $\gamma_1 = \gamma_2 = 2.62$ and also at $\Phi_1 = 0$ with $\gamma_2 = 2.62$ indicated by the dotted lines. Confidence regions are also shown a single power-law fit of the high-energy through-going track sample [136]. The overlap of the confidence regions suggest that a second component with a harder power-law index suggested by the through-going track results is not incompatible with the starting event results.

more complex structure like a cutoff.

7.4.2 Flavor Composition of Astrophysical Neutrinos

With the improved cascade and track classification as well as track energy and inelasticity reconstructions, limits on the flavor composition of astrophysical neutrinos can be revisited. Updated confidence regions for the flavor composition under the assumption of a single power-law flux are shown in Fig. 7.15, and due to the larger sample size here, Wilks' theorem is found to be an adequate approximation for the test statistic distribution. The limits on non-standard flavor compositions are more constraining than those in Fig. 6.14. For example, a maximally track-like flavor composition of $(0:1:0)_\oplus$ is disfavored at 7.4σ and a maximally cascade-like flavor composition of $(1:0:0)_\oplus$ is disfavored at 5.8σ. Once again, every flavor

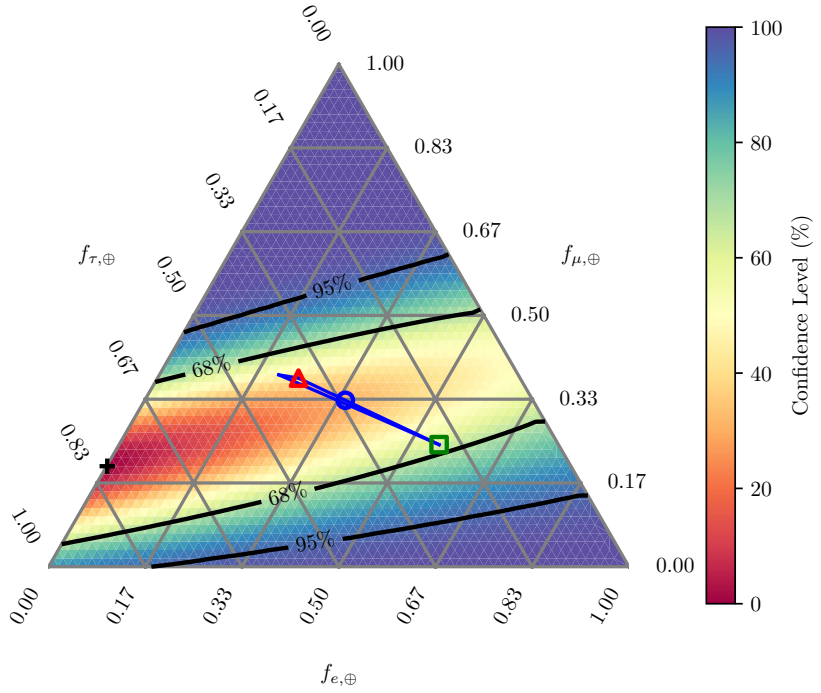

Figure 7.15: Updated confidence regions for astrophysical flavor ratios $(f_e : f_\mu : f_\tau)_\oplus$ at Earth. The labels for each flavor refer to the correspondingly tilted lines of the triangle. Averaged neutrino oscillations map the flavor ratio at sources to points within the extremely narrow blue triangle diagonally across the center. The $\approx \left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ composition at Earth, resulting from a $\left(\frac{1}{3} : \frac{2}{3} : 0\right)_S$ source composition, is marked with a blue circle. The compositions at Earth resulting from source compositions of $(0 : 1 : 0)_S$ and $(1 : 0 : 0)_S$ are marked with a red triangle and a green square, respectively. The updated best-fit neutrino oscillation parameters from [238] are used here. Though the best-fit composition at Earth (black cross) is $(0 : 0.21 : 0.79)_\oplus$, the limits are consistent with all compositions possible under averaged oscillations.

composition allowed by standard neutrino oscillations lies with the 68% confidence region.

The best-fit flavor composition is $(0 : 0.21 : 0.79)_\oplus$, but there remains a strong degeneracy between $f_{e,\oplus}$ and $f_{\tau,\oplus}$ in the flavor composition measurement, which results from the fact that cascade and track classification serve as the primary means to identify neutrino flavor. The distinct visible inelasticity distribution of muons from ν_τ CC events compared to ν_μ CC events does break the degeneracy, but it is not statistically significant presently. The possibility of using visible inelasticity distributions to resolve a ν_τ flux in the future will be studied in Sec. 8.2.1.

The flavor composition limits also do not allow significant limits to be set on the source flavor composition when assuming standard neutrino oscillations. To date, the only means

to set significant limits on the source flavor composition is to incorporate through-going tracks in a joint fit as done for the IceCube global fit result shown in Fig. 6.15. However, it should be noted that that the global fit results assume a single power-law flux model, and if the power-law index hardens at high energies, this would significantly bias the flavor composition measurement since there would be more through-going tracks at high energies than expected from a single power-law model. The result here is more conservative since it relies on cascades and starting tracks in a similar energy range where a single power law can describe both samples.

7.4.3 Neutrino to Antineutrino Ratio

The reconstruction of inelasticity also provides a means to distinguish muon neutrinos and antineutrinos. The composition of neutrinos and antineutrinos in both atmospheric and astrophysical fluxes is a question of great interest. For the atmospheric flux, the neutrino to antineutrino ratio is an energy-dependent quantity that is sensitive to the relative production of positive and negative hadrons in cosmic-ray air showers and is thus a unique diagnostic for hadronic interaction models as discussed in Sec. 3.3.1. As detailed in Sec. 3.2.3, the neutrino to antineutrino ratio of astrophysical neutrinos can distinguish whether neutrinos were produced through $p\gamma$ or pp interactions, adding another piece of information about the environment where cosmic rays are produced.

In Fig. 7.16, the difference in reconstructed inelasticity distributions between muon neutrinos and antineutrinos is illustrated across energy. At low energies and high inelasticity, antineutrinos contribute a smaller fraction of the event rate than neutrinos. At high energies above $\sim 10\,\text{TeV}$, the difference in the inelasticity distribution between neutrinos and antineutrinos vanishes. Unfortunately, this occurs just as the astrophysical flux starts to dominate the total event rate. Thus, there is little hope of measuring the neutrino to antineutrino ratio of the astrophysical flux with the starting track sample here. In the future it may be possible to use a surface cosmic-ray veto to reject atmospheric neutrinos down to $\sim 10\,\text{TeV}$ where there are significant differences in the inelasticity distribution of neutrinos and antineutrinos [239].

From $\sim 1\,\text{TeV}$ to $10\,\text{TeV}$, there are substantial differences in the reconstructed inelasticity distribution that can be exploited to measure the neutrino to antineutrino ratio of the atmospheric flux. It would be desirable to measure the neutrino to antineutrino ratio in several energy bins, but given the limited size of the current event sample, only a scaling factor, $R_{\nu_\mu/\bar{\nu}_\mu}$, of the ν_μ to $\bar{\nu}_\mu$ flux ratio of the HKKMS calculation (parameterized in Eq. 7.9) can be measured. The scaling factor $R_{\nu_\mu/\bar{\nu}_\mu}$ can be added to the list of parameters in Tab. 7.3 and fit using the same maximum likelihood procedure used in the previous section.

The best-fit value of $R_{\nu_\mu/\bar{\nu}_\mu} = 0.77^{+0.44}_{-0.25}$ is compatible with the HKKMS calculation. A corresponding test statistic scan of $R_{\nu_\mu/\bar{\nu}_\mu}$ is shown in Fig. 7.17. Though not shown, the asymptotic value of the test statistic as $R_{\nu_\mu/\bar{\nu}_\mu} \to \infty$ is $-2\Delta \ln L = 14.8$. This implies that a pure neutrino atmospheric flux with no antineutrinos is disfavored with a significance of 3.8σ. By using the central 90% of neutrino energies contributing to this test statistic as

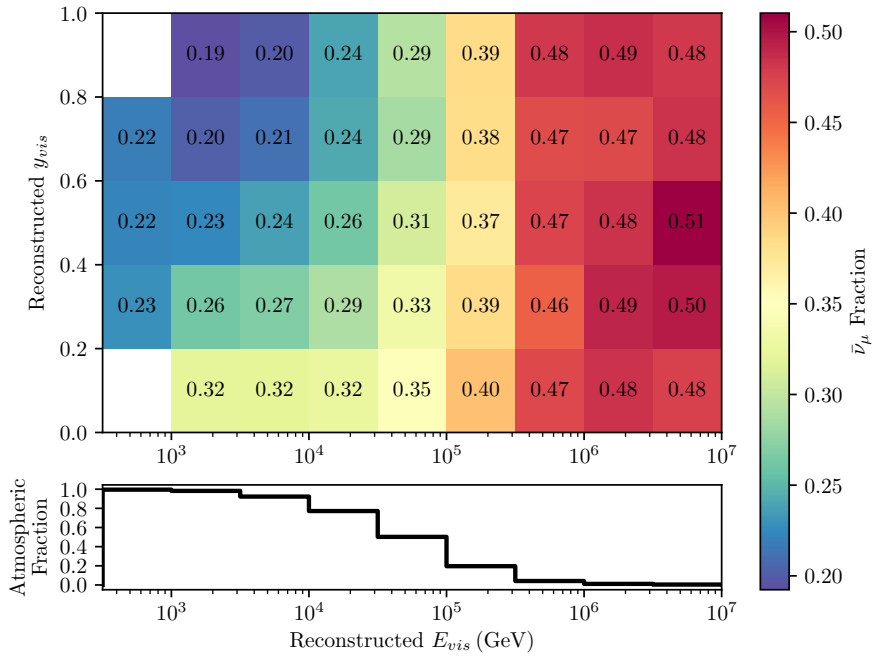

Figure 7.16: The predicted fraction of $\bar{\nu}_\mu$ contributing to the total $\nu_\mu + \bar{\nu}_\mu$ event rate in bins of reconstructed energy and inelasticity. At energies below $\sim 10\,\text{TeV}$, there are substantial differences in the inelasticity distribution that enable the atmospheric neutrino to antineutrino ratio to be measured. The bottom panel shows the fraction of atmospheric neutrinos contributing to the total event rate in bins of reconstructed energy. At energies above $\sim 100\,\text{TeV}$ where astrophysical neutrinos begin to dominate the event rate, differences in the inelasticity distribution vanish and it is not possible to measure the neutrino to antineutrino ratio for the astrophysical flux. An equal neutrino and antineutrino composition is assumed for the astrophysical flux here.

described in Appendix B, one can quantify a sensitive energy range of 770 GeV to 21 TeV for the neutrino to antineutrino ratio measurement.

This is the first measurement of the neutrino to antineutrino ratio ever made in this energy range. It is also in a range comparable to the highest energy measurements of atmospheric muon charge ratio up to 10 TeV [132, 133, 134]. Like the muon charge ratio, it is expected that the neutrino to antineutrino ratio can be used to test the hadronic models used to model cosmic-ray interactions in the atmosphere. However, the utility of the result here is limited by a strong model dependence since one must describe the energy and zenith distribution of the atmospheric flux over the broad energy range where the measurement is sensitive. As a result, the scaling factor on the HKKMS ν_μ to $\bar{\nu}_\mu$ flux ratio should not be interpreted in the context of other hadronic interaction models since they would also produce a different energy

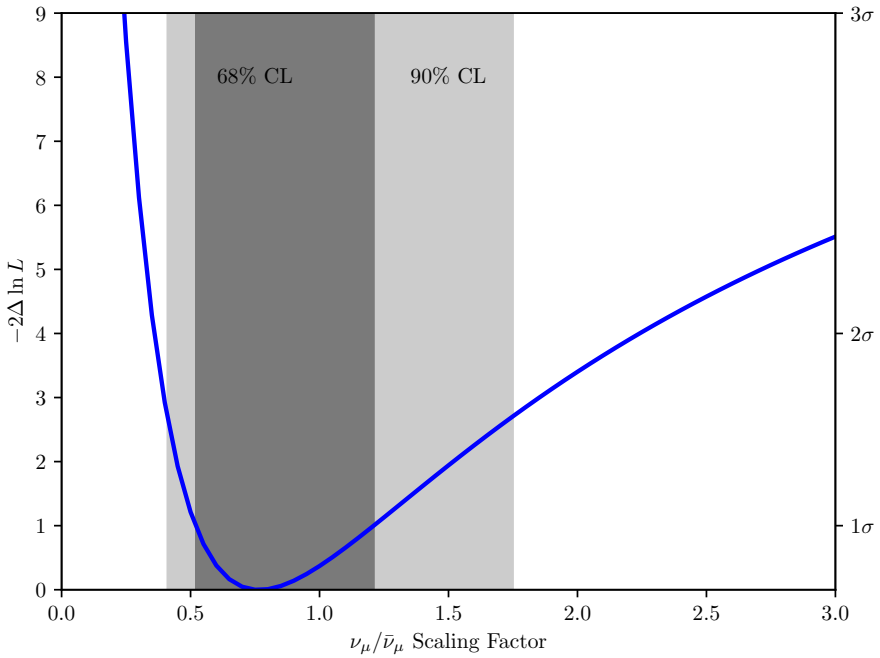

Figure 7.17: A scan of the test statistic for a scaling factor on the ν_μ to $\bar{\nu}_\mu$ flux ratio predicted by HKKMS. 68% and 90% confidence intervals are shown in dark gray and light gray, respectively. The asymptotic value of the test statistic is $2\Delta \ln L = 14.8$ as $R_{\nu_\mu/\bar{\nu}_\mu} \to \infty$.

and zenith angle distribution. In the future, tests of hadronic interaction models could be carried out by fitting the complete energy, zenith, and inelasticity distribution obtained here to a direct calculation of the atmospheric neutrino flux from any hadronic interaction model of interest [240].

7.4.4 Neutrino-Induced Charm Production

The reconstruction of inelasticity can also be used to indirectly identify the contribution of subprocesses contributing to the total neutrino cross section through the imprint that they leave in visible inelasticity distribution. This can be illustrated through the example of charged-current neutrino charm production.

In CC neutrino charm production, charm quarks are primarily produced from scattering off strange sea quarks due to the large value of $|V_{cs}|^2$. As a result, charm production events have a flatter inelasticity distribution that those produced other events that mostly result from scattering on up or down quarks. Additionally, charm hadrons may decay to produce a muon with $\sim$ 10% branching ratio. In a ν_μ CC event, this creates a dimuon, but the small angular separation between the two muons makes them impossible to distinguish at

TeV energies. Instead, the linearity of the mean muon energy loss rate means that the reconstructed muon energy corresponds to roughly the sum of the two muon energies. The visible inelasticity of these events then tends to be lower than in events where the charm hadron did not decay to a muon. If a charm hadron is produced in a ν_e or ν_μ CC interaction and decays to a muon, the visible inelasticity of these events tends to be higher than typical ν_μ CC events. The leading-order cross section calculation of neutrino charm production is used here and has been described in Sec. 2.3.3, and the simulation of charm hadron production, decay, and interaction was detailed in Sec. 5.1.2 and Sec. 5.2.3.

The unique visible inelasticity of charm production events is illustrated in Fig 7.18. On the whole charm production events have a flatter inelasticity distribution and form a larger fraction of the expected rate at high visible inelasticity. Dimuons from charm decay in ν_μ CC events only form $\sim 1\%$ to $\sim 2\%$ of the event rate, but they tend to be reconstructed to lower visible inelasticity than other charm production events due the presence of the second muon. Charm decay to a single muon in ν_e and ν_τ CC interactions can form a significant fraction of the event rate at high inelasticity above $\sim 100\,\text{TeV}$. Though all of these processes contribute to the distinct visible inelasticity distribution of CC neutrino charm production, it is largely the flatter inelasticity distribution caused by scattering off strange sea quarks that dominates the effect.

To identify the contribution of neutrino charm production to CC deep inelastic scattering, one can introduce another parameter, $R_{\text{CC, charm}}$, that scales the event rate for all flavors of CC charm production into the same likelihood fit as done previously. The best-fit value is $R_{\text{CC, charm}} = 0.93^{+0.73}_{-0.59}$, compatible with the expected value of 1. A scan of the test statistic for $R_{\text{CC, charm}}$ is shown in Fig. 7.19. The test statistic for a null hypothesis of zero charm production is $-2\Delta \ln L = 2.8$, and thus zero charm production can be excluded at 91% confidence level.

To estimate the neutrino energy range where there is sensitivity to neutrino charm production, the central 90% of neutrino energies contributing to the test statistic is found to extend from 1.5 TeV to 340 TeV as described in Appendix B. Since the central 90% of neutrino energies contributing to the total charm production event rate is 1.3 TeV to 44 TeV, the upper end of 340 TeV may seem counter-intuitive. However, Fig. 7.18 illustrates that in the region above $\sim 100\,\text{TeV}$ charm production forms the largest fraction of the event rate at high visible inelasticity and thus contributes strongly to the statistical evidence for charm production.

Although the present limits on charm production are unlikely to add useful constraints that will improve knowledge of nuclear structure or neutrino DIS, this analysis serves as a blueprint for how inelasticity distributions can be used to identify the contribution of individual processes to neutrino scattering. In Sec. 8.2.2, the possibility of using inelasticity information to search for additional neutrino interactions both in and beyond the Standard Model will be explored.

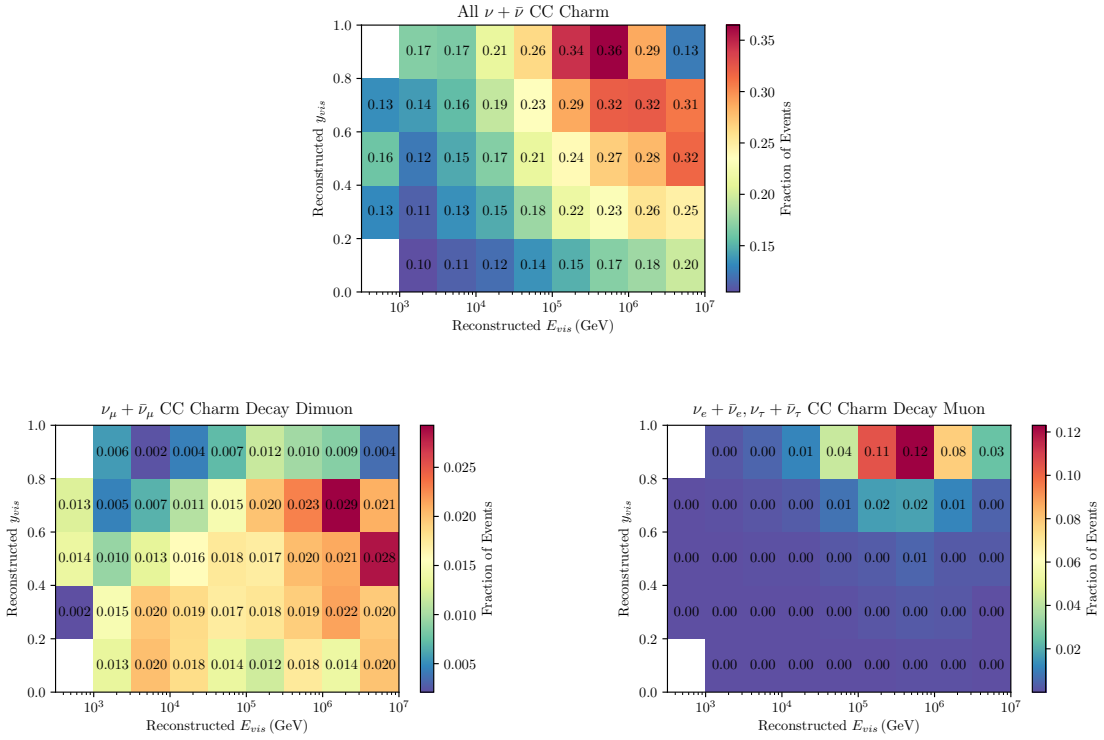

Figure 7.18: Top: The predicted fractional contribution of all-flavor neutrino CC charm production events to the total event rate in bins of reconstructed visible energy and inelasticity. The increased charm production fraction at high visible inelasticity and high energy (up to 36%) provides a shape difference that allows the presence of charm production to be identified in a likelihood fit of observed data. Bottom left: The fractional contribution of dimuons from charm decay in ν_μ CC interactions to the total event rate. This only provides a $\sim$ 1% contribution to the event rate, but dimuons to tend to have lower reconstructed inelasticity than other events. Bottom right: The fractional contribution of single muons from the decay of charm hadrons in ν_e or ν_τ CC interactions to the total event rate. At high energies above 100 TeV, these events can be as much as 12% of the total event rate. Most of the power to identify neutrino charm production comes from the distinct inelasticity distribution of single muons in ν_μ CC charm production events rather than muons from the decay of charm hadrons.

7.4.5 Inelasticity Distribution across Energy

For the previous measurements of neutrino to antineutrino ratio and neutrino-induced charm production, it was necessary to fit the reconstructed energy and zenith angle distributions in addition to the inelasticity distribution in order to constrain uncertainties in the neutrino flux model and detector model. Although this is the best approach that should

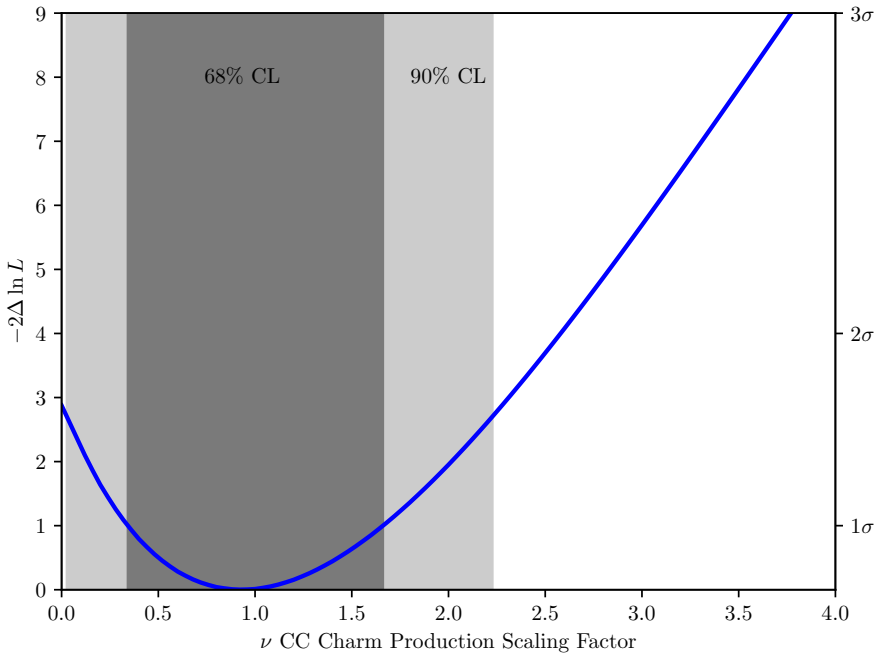

Figure 7.19

be taken whenever one wants to measure a quantity derived from inelasticity distributions, it is also desirable to have a model-independent characterization of inelasticity distribution so that other physicists can explore other measurements that could be made, for example searches for physics beyond the Standard Model or other Standard Model effects not present in the CSMS cross section calculation.

Unfolding of the differential cross section $d\sigma/dy$ could provide such a model-independent characterization of the inelasticity distribution. The goal of an unfolding would be to measure $d\sigma/dy$ in several bins of energy and inelasticity by deconvolving the detector resolution and then correcting for the detection efficiency and neutrino flux in each bin. However, unfolding faces several challenges with the event sample here. The first issue is that the limited resolution and sample size greatly limit the number of bins for which there is a meaningful constraint on $d\sigma/dy$. Additionally, the detection efficiency is a strong function of energy and inelasticity, and for any choice of a bin size large enough to have a sufficient number of events, the efficiency varies strongly enough within a bin that it is impossible to correct for it without a model of how $d\sigma/dy$ and the neutrino flux vary within that bin.

While more data may alleviate the problems with an unfolding in the future, the present approach will be to describe the inelasticity distribution across energy with a physically motivated parameterization. If a physicist has an alternate model of an inelasticity distribution, it can be fit to the same parameterization, and the parameters can be compared to those

determined from the data here. To motivate the parameterization, recall from Eq. 2.26 that the leading order differential CC cross section can be written in the form

$$\frac{d\sigma}{dxdy} = \frac{2G_F^2 M E_\nu}{\pi}\left(\frac{M_W^2}{Q^2+M_W^2}\right)^2 \left[xq(x,Q^2) + (1-y)^2 x\bar{q}(x,Q^2)\right], \tag{7.17}$$

where q and $\bar{q}$ represent sums of quark PDFs. At high-energies where the behavior of the PDFs is being probed at low Bjorken-x, the PDFs will have a power-law behavior, $xq(x,Q^2) \sim A(Q^2)x^{-\lambda}$ with $\lambda \sim 0.4$. When transforming variables from (x,y) to (Q^2,y), the Q^2-dependence of Eq. 7.17 can be separated from the y-dependence and integrated out to give a two-parameter function,

$$\frac{d\sigma}{dy} \propto \left(1+\epsilon(1-y)^2\right) y^{\lambda-1}. \tag{7.18}$$

The parameter ϵ represents the relative importance of the term proportional to $(1-y)^2$ in Eq. 7.17. Its value will be different for neutrinos and antineutrinos. The normalized inelasticity distribution can then be written as

$$\frac{dp}{dy} = N\left(1+\epsilon(1-y)^2\right) y^{\lambda-1}, \tag{7.19}$$

where N is a normalization constant

$$N = \frac{\lambda(\lambda+1)(\lambda+2)}{2\epsilon+(\lambda+1)(\lambda+2)}. \tag{7.20}$$

Though it only contains two degrees of freedom, this simple parameterization can accurately represent sophisticated calculations of inelasticity distributions. For example, a fit of Eq. 7.18 to the CSMS calculation produces no more than a 1% root mean square error for neutrino energies from 1 TeV to 10 PeV.

In practice, the parameters ϵ and λ tend to be highly correlated when fitting Eq. 7.19 to realistic inelasticity distributions. To alleviate this issue, it is convenient to fit for the mean of the distribution, $\langle y \rangle$, and λ instead, which show far less correlation. The mean inelasticity can be found analytically,

$$\langle y \rangle = \int_0^1 y\frac{dp}{dy}dy = \frac{\lambda(2\epsilon+(\lambda+2)(\lambda+3))}{(\lambda+3)(2\epsilon+(\lambda+1)(\lambda+2))} \tag{7.21}$$

It is then straightforward to substitute

$$\begin{aligned} \epsilon &= -\frac{(\lambda+2)(\lambda+3)}{2}\frac{\langle y\rangle(\lambda+1)-\lambda}{\langle y\rangle(\lambda+3)-\lambda}, \\ N &= \frac{1}{2}(\lambda+1)(\langle y\rangle(\lambda+3)-\lambda) \end{aligned} \tag{7.22}$$

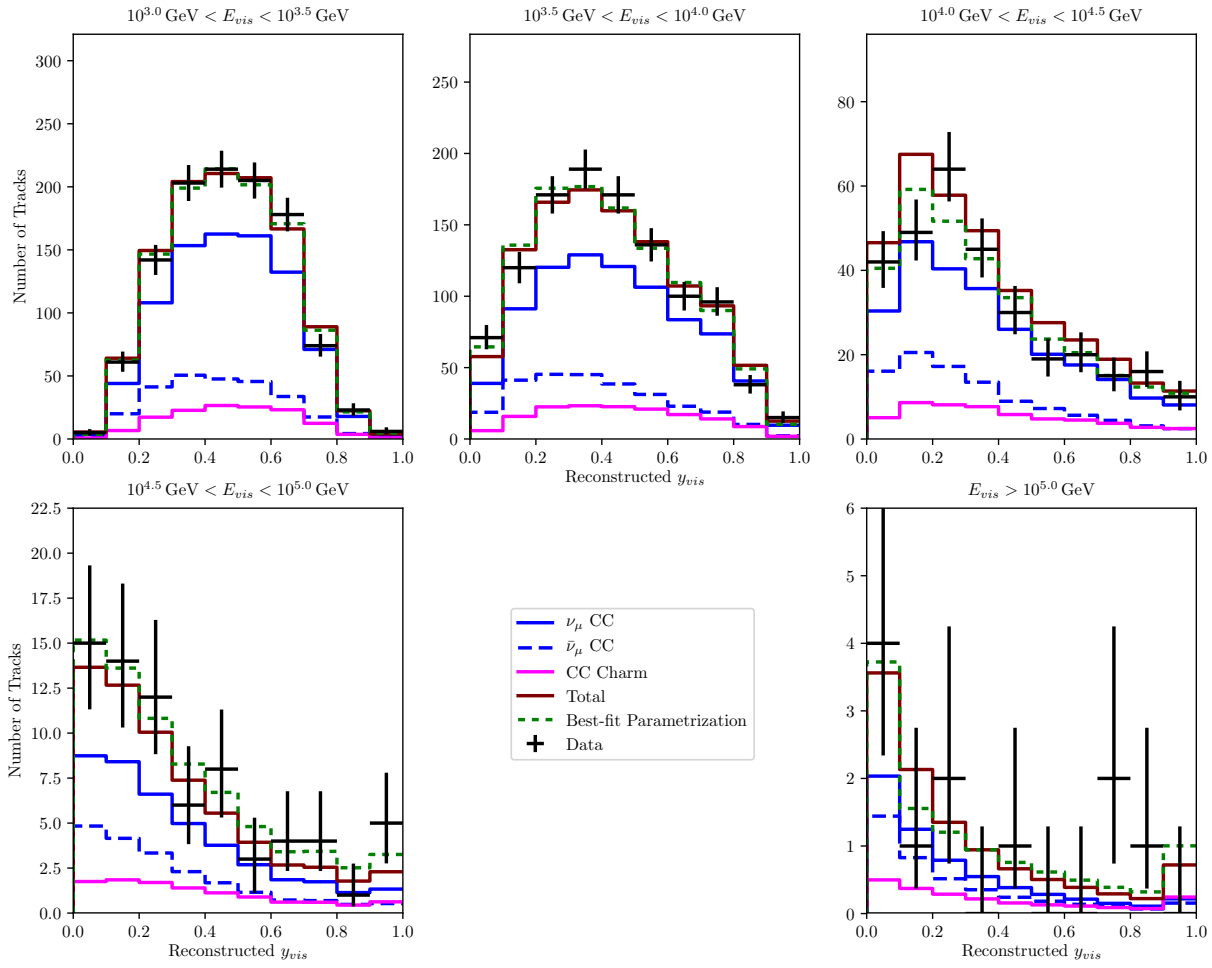

Figure 7.20: The reconstructed visible inelasticity distribution in five different bins of reconstructed energy. Observed data is shown in black and the result of fitting the distribution to the parameterization of Eq. 7.18 is shown with dashed green lines. The prediction of the CSMS differential CC cross section are shown for neutrinos with solid blue lines and antineutrinos with dashed blue lines. The total CC charm contribution is shown in pink, illustrating its flatter inelasticity distribution. The best-fit flux models of Tab. 7.3 are assumed for all predictions.

into Eq. 7.19 so that dp/dy can easily be found as a function of $\langle y \rangle$ and λ only.

To characterize the inelasticity distribution across energy, the starting track sample is split into four half decade bins in reconstructed total visible energy from 1TeV to 100 TeV. A fifth bin containing all events above 100 TeV is also used. For events in each bin of reconstructed energy, the reconstructed visible inelasticity distribution is binned into a histogram containing 10 bins from 0 to 1 as illustrated in Fig. 7.20.

A binned Poisson likelihood fit is then used to fit the observed data to bin expectations

$\log_{10}(E_{vis}/\text{GeV})$ Range	68% $\log_{10}(E_\nu/\text{GeV})$ Range	Events	$\langle y \rangle$	λ
$[3.0, 3.5)$	$3.33^{+0.21}_{-0.22}$	1111	$0.41^{+0.06}_{-0.06}$	$0.72^{+0.86}_{-0.72}$
$[3.5, 4.0)$	$3.73^{+0.22}_{-0.22}$	1107	$0.42^{+0.02}_{-0.02}$	$1.09^{+0.27}_{-0.48}$
$[4.0, 4.5)$	$4.18^{+0.22}_{-0.22}$	310	$0.38^{+0.03}_{-0.03}$	$0.96^{+0.28}_{-0.35}$
$[4.5, 5.0)$	$4.65^{+0.22}_{-0.22}$	72	$0.38^{+0.05}_{-0.07}$	$0.76^{+0.45}_{-0.76}$
$[5.0, \infty)$	$5.22^{+0.50}_{-0.31}$	11	$0.31^{+0.13}_{-0.20}$	$0.11^{+0.82}_{-0.11}$

Table 7.4: The best-fit parameters when reconstructed inelasticity distributions are fit to Eq. 7.18 in five bins of reconstructed energy. The energy range containing the central 68% of simulated neutrino energies for each bin is shown in the second column.

using `neutrino-generator` simulation. To determine the best-fit parameters $\langle y \rangle$ and λ in each energy bin, it is necessary to re-weight simulated events, which are generated with an inelasticity distribution according to the CSMS differential cross section as described in Sec. 5.1.1. If w_i^0 is the nominal event weight with energy E_i and inelasticity y_i and $\frac{dp}{dy}(E_i, y_i)_{\text{CSMS}}$ is the value of the normalized inelasticity distribution of the CSMS calculation found from precomputed tables, then the new event weight is

$$w_i = w_i^0 \frac{\frac{dp}{dy}(y_i; \langle y \rangle, \lambda)}{\frac{dp}{dy}(E_i, y_i)_{\text{CSMS}}}, \tag{7.23}$$

where $\frac{dp}{dy}(y; \langle y \rangle, \lambda)$ is the parameterization of Eq. 7.19. A neutrino flux is needed to calculate the nominal weight, w_i^0, so the best-fit fluxes from Tab. 7.3 are used. Because the bin width in reconstructed energy is not significantly larger than the energy resolution, the choice of neutrino flux model has very little influence. However, detector systematic uncertainties can influence the reconstructed inelasticity distribution, so the four detector parameters describing DOM optical efficiency, bulk ice scattering and absorption, and hole ice scattering are also used in the fit by the same method described in Sec. 7.3.2 with the same prior information. Lastly, a single overall scaling factor of the inelasticity distribution predicted from simulation is included without any prior information, which accounts for a normalization uncertainty in the neutrino flux.

The best-fit parameters describing the inelasticity distribution in each of the five reconstructed energy bins are shown in Tab. 7.4. As shown in Fig. 7.20, the best-fit distributions all show good agreement with the observed distributions, confirming that the two-parameter description of inelasticity distributions is sufficient to describe the data. To illustrate the energy-dependence of the inelasticity distributions, one can plot the mean inelasticity obtained from the fit versus neutrino energy as shown in Fig. 7.21. One can compare the mean inelasticity to the CSMS calculation of $\langle y \rangle$ for both neutrinos and antineutrinos. The observed mean inelasticity is expected to fall between the predictions for neutrinos and antineutrinos, and it is better compare to the mean inelasticity weighted by the atmospheric HKKMS flux. In all bins, the mean inelasticity agrees well with the flux-averaged CSMS prediction and follows the downward trend caused by the W-boson propagator.

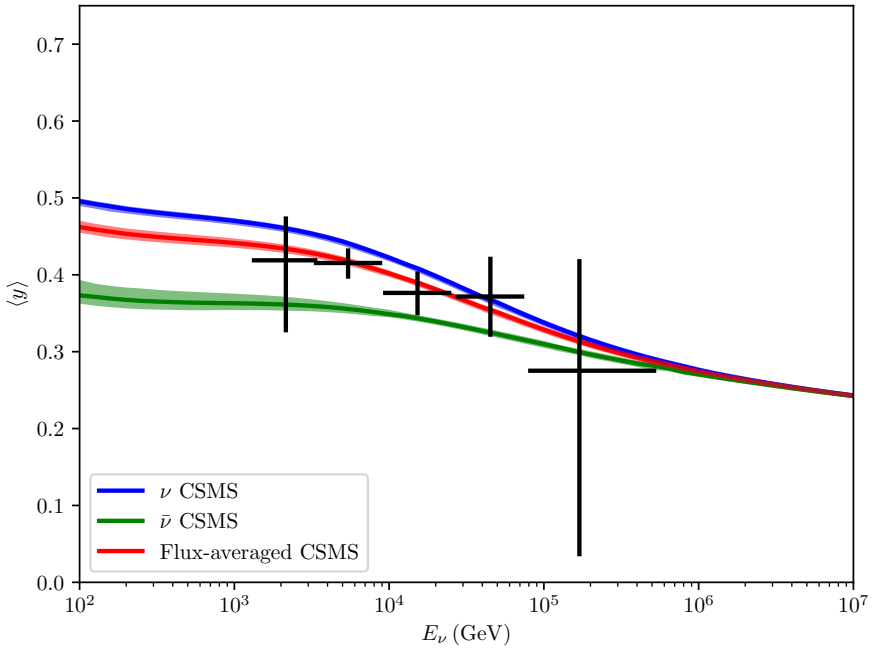

Figure 7.21: The mean inelasticity obtained from the fit to Eq. 7.18 in five bins of reconstructed energy. Vertical error bars indicate the 68% confidence interval for the mean inelasticity, and horizontal error bars indicate the expected central 68% of neutrino energies in each bin. The predicted mean inelasticity from CSMS is shown in blue for neutrinos and in green for antineutrinos. A flux-averaged mean inelasticity according the HKKMS atmospheric flux calculation is shown in red.

Chapter 8

Conclusion

8.1 Summary of Results

In Ch. 6, the first IceCube measurement of the flavor composition of astrophysical neutrinos was presented [213]. A sample of 137 neutrino candidates was obtained using an outer-layer veto to reject atmospheric muon background in 974 days of IceCube data. A likelihood ratio approach using first recorded photoelectron times was used to identify 129 cascades and 8 tracks. A joint maximum likelihood fit was performed on the energy and zenith distribution of the cascade and track samples in order to determine the energy spectrum and flavor composition of the astrophysical neutrino flux. The fit also included the contribution of backgrounds from atmospheric neutrinos and muons. The energy spectrum of the astrophysical flux was found to be consistent with a power law and an index of $\gamma = 2.6 \pm 0.15$ was measured in an energy range extending from 35 TeV to 1.9 PeV. The best-fit flavor composition was found to be $(0 : 0.2 : 0.8)_\oplus$, but the 68% confidence region includes the flavor composition of $\approx \left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ typically expected from complete pion decay in astrophysical sources and standard neutrino oscillations. The confidence regions also do not enable constraints on the flavor composition of astrophysical neutrinos at their sources, but they can be used to place limits on non-standard flavor compositions that arise in exotic physics scenarios. The most extreme maximally cascade-like and maximally track-like flavor compositions of $(1 : 0 : 0)_\oplus$ and $(0 : 1 : 0)_\oplus$ were disfavored at 2.3σ and 3.3σ, respectively.

In Ch. 7, the inelasticity distribution of contained neutrino interactions in IceCube was studied for the first time using improved machine learning techniques. A sample of 3615 neutrino candidates was found in 1734 days of IceCube data. The use of a boosted decision tree enabled neutrino interactions down to 1 TeV to be identified with negligible atmospheric muon background. The boosted decision tree was also used to classify events into samples of 2650 tracks and 965 cascades. By using a random forest regressor, the energy and inelasticity of tracks were reconstructed. A joint maximum likelihood fit of the distribution of energy, inelasticity, and zenith distribution for tracks and the energy and zenith distribution for cascades was developed that accounts for systematic uncertainties in neutrino flux and detector models. The likelihood fit was used to perform several measurements.

First of all, the measurements of the astrophysical neutrino energy spectrum and flavor composition were updated to reflect the enhanced event classification and energy reconstruction techniques. The energy spectrum is compatible with a single power law and an index of $\gamma = 2.62 \pm 0.07$ was measured. Due to the improved background rejection at low energies, the measurement was sensitive to the astrophysical flux from 6.6 TeV to 2.2 PeV. Since recent results using through-going tracks at IceCube indicate a harder astrophysical power-law index of 2.13 ± 0.13 above 200 TeV, a double power-law model was fit to the data, but no evidence for a second power-law component with a harder index was found. However, the confidence regions for the second power-law component are compatible with the through-going track results.

Confidence regions on the astrophysical flavor composition were found that are compatible with the expected flavor composition of $\approx \left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_{\oplus}$. Although the constraints on non-standard flavor compositions from exotic physics are improved, there are not significant constraints on the flavor composition at astrophysical sources. A joint fit of the sample with through-going tracks could improve the results on flavor composition but should be treated cautiously due to the possibility of an energy spectrum more complex than a single power law above 200 TeV. All measurements of flavor composition so far show a strong degeneracy between ν_τ and ν_e. In principle, the inclusion of inelasticity information allows this degeneracy to be broken, although a much larger dataset will be required for the effect to be statistically significant as demonstrated in the next section.

Secondly, the neutrino to antineutrino ratio of the atmospheric flux was measured. A scaling factor on the HKKMS ν_μ to $\bar{\nu}_\mu$ flux ratio was measured to be $0.77^{+0.44}_{-0.25}$. The sensitive energy range for this measurement is estimated to be from 770 GeV to 21 GeV. This result is complementary to the measurements of the atmospheric muon charge ratio up to 10 TeV and can be used to test hadronic interaction models used for cosmic-ray air shower simulations.

Thirdly, neutrino-induced charm production leaves an imprint in the reconstructed inelasticity distribution that can be used to identify the process. A scaling factor on the charged-current charm production cross section was measured to be $0.93^{+0.73}_{-0.59}$ when assuming a fixed inelasticity distribution. Zero charm production is excluded at 91% confidence level, and the sensitive energy range is estimated to be from 1.5 TeV to 340 TeV. The result illustrates how the inelasticity distribution can be used to identify the contribution of different scattering processes to the total neutrino cross section provided the inelasticity distribution can be calculated. As described in Sec. 8.2.2, the same techniques could be used to search for many other processes both in and beyond the Standard Model.

Last of all, the reconstructed inelasticity distribution was studied in bins of energy from 1 TeV to 100 TeV. In every bin, the reconstructed inelasticity distribution shows good agreement with the CSMS calculation of the differential cross section, $d\sigma/dy$. Additionally, a two-parameter function, $dp/dy \propto (1 + \epsilon(1-y)^2)y^{\lambda-1}$, can be fit to the inelasticity distribution in every bin and describes the data well. The mean of the best-fit distribution in every energy bin matches the mean inelasticity $\langle y \rangle$ expected from the CSMS calculation. In the future, an unfolding could enable a better model-independent characterization of the inelasticity distribution across energy.

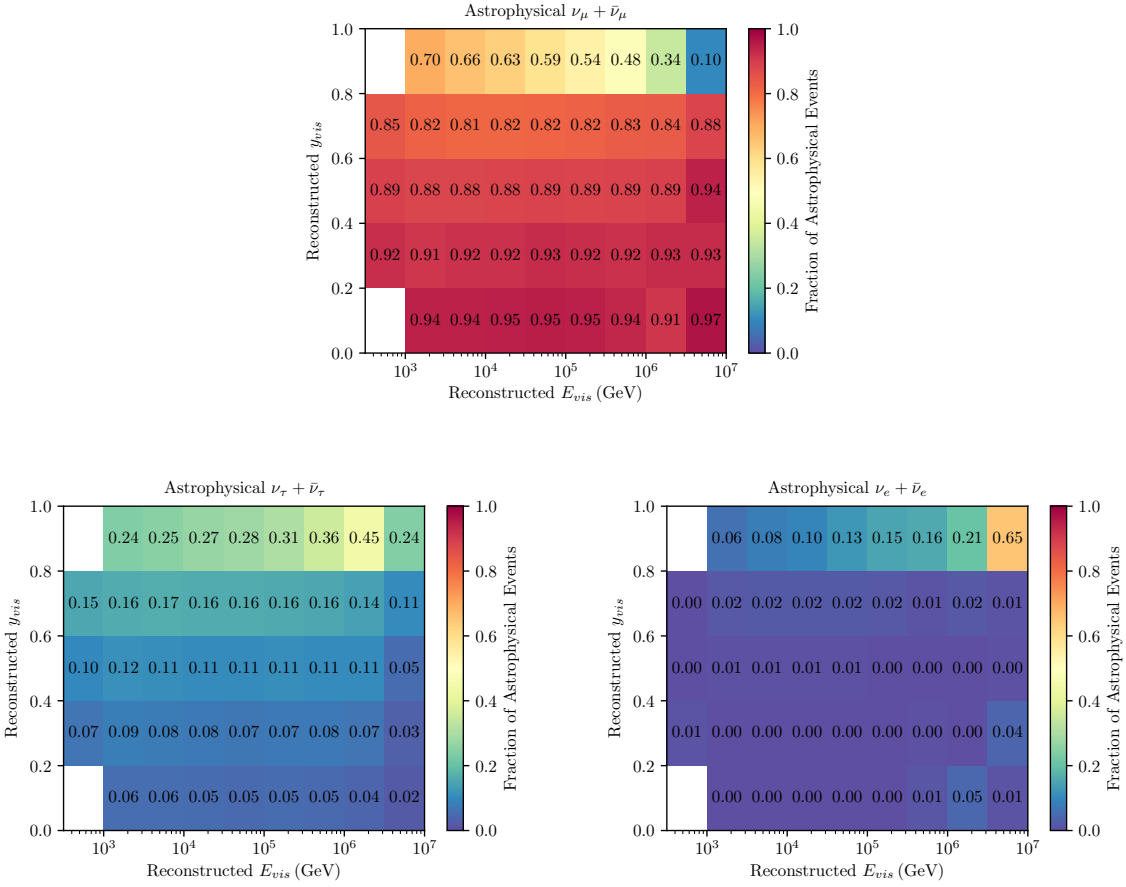

Figure 8.1: The fractional contribution to the event rate of astrophysical neutrino-induced tracks from ν_e (bottom right), ν_μ (top), and ν_τ (bottom left) in bins of reconstructed visible energy and inelasticity. A flavor composition of $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ is assumed. The unique distribution for each neutrino flavor can help break the degeneracy in a flavor composition measurement.

8.2 Future Topics with Inelasticity Distributions

8.2.1 Tau Neutrinos

In Sec. 7.4.2, limits on the flavor composition of astrophysical neutrinos were presented that showed a strong degeneracy between the contributions from ν_e an ν_τ. In this section, the possibility of using the unique visible inelasticity distribution of starting tracks for each flavor will be explored as a method to break the degeneracy and definitively identify a ν_τ contribution to the astrophysical flux.

The unique inelasticity distribution for each neutrino flavor is illustrated in Fig. 8.1, where the fractional contribution of each neutrino flavor to the rate of astrophysical starting tracks is shown in bins of reconstructed visible energy and inelasticity. Muon neutrinos generally dominate the event rate of starting tracks, but they tend to have a lower contribution at high inelasticity. Since muons from tau decay events generally have lower energy than muons in ν_μ CC events as discussed in Sec. 4.1.5, tau neutrinos contribute more to the event rate at high inelasticity. Low-energy muons from pion, kaon, and charm hadron decays in the accompanying hadronic cascade of a ν_τ interaction also lead to a large contribution to the event rate in the highest inelasticity bin at energies above $\sim 100\,\text{TeV}$. Very few tracks are created by ν_e interactions, but the same mechanism of decays in hadronic cascades can lead to a substantial contribution of ν_e in the highest inelasticity bin above $\sim 100\,\text{TeV}$. The contribution of ν_e from hadronic decays is generally lower than the corresponding contribution from ν_τ since a tau decay initiates a hadronic cascade with a 68% branching fraction, whereas an electron creates an electromagnetic cascade that produces few muons. Contributions from the Glashow resonance are also visible in the highest inelasticity bin from $10^{6.5}\,\text{GeV}$ to $10^7\,\text{GeV}$ where low energy muons are produced in the cascade from the dominant hadronic mode of W-boson decay.

To estimate how inelasticity information can aid a future flavor composition measurement, a sensitivity study can be done by scaling the number of events from the current five-year analysis and estimating the median significance for discovering a nonzero ν_τ flux. A test statistic can be formed from a hypothesis test between a null hypothesis with $f_{\tau,\oplus} = 0$ but both $f_{e,\oplus}$ and $f_{\mu,\oplus}$ free and an alternate hypothesis where the flavor composition is completely free. To estimate the median value of the test statistic for a larger sample size, the "Asimov" method described in [241] is used and the two hypotheses are fit to an observed distribution that taken to be exactly equal to to the scaled best-fit distribution with $\left(\frac{1}{3} : \frac{1}{3} : \frac{1}{3}\right)_\oplus$ flavor composition in Tab. 7.3. The same model and parameters described in Sec. 7.3 are used.

In addition to projecting the sensitivity for more years of IceCube data, it is also possible to project the sensitivity for a detector volume upgrade like the proposed IceCube-Gen2 [221]. If one assumes that the energy threshold of the event selection and the energy resolution are not significantly degraded by the sparser instrumentation, e.g. through improved optical sensor design or data analysis techniques, then scaling the event rate of the current IceCube analysis can represent the improvement gained by a larger instrumented volume. The median significance can then be estimated not only as a function of time, but also as a function of the exposure of the detector in $\text{km}^3\,\text{yr}$. For the present flavor composition analysis, the exposure is $4\,\text{km}^3\,\text{yr}$ corresponding to the five-year live time and the IceCube instrumented volume of $0.8\,\text{km}^3$. A twelve-year live time and an instrumented volume of $10\,\text{km}^3$ proposed for IceCube-Gen2 corresponds to an exposure of $120\,\text{km}^3\,\text{yr}$, or 30 times as many events.

In Fig. 8.2, the median significance for discovering a ν_τ flux is shown as a function of exposure. For the IceCube-Gen2 scenario considered here, a ν_τ flux may be identified with a significance of 2.5σ in twelve years of operation. This is likely an optimistic estimate due to the fact that no attempt is made to account for sparser instrumentation. On the other

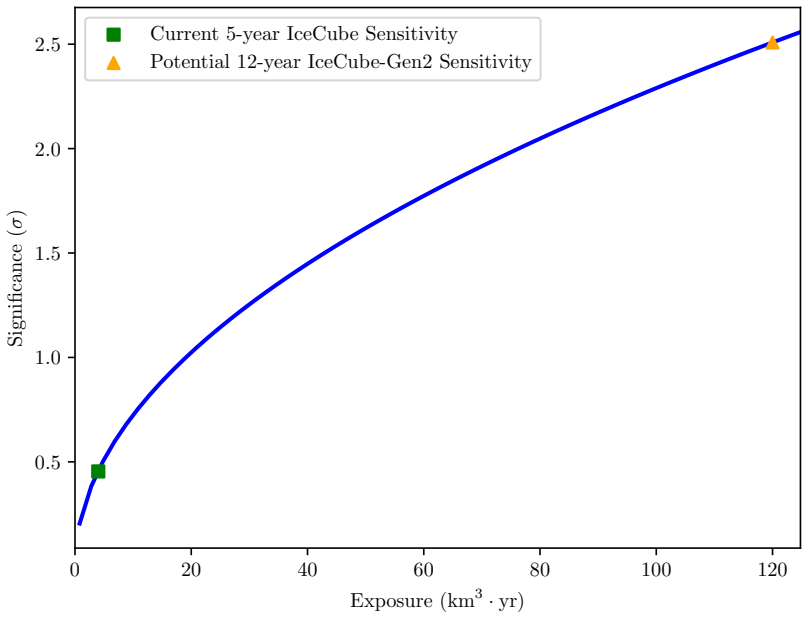

Figure 8.2: The median significance for identifying a ν_τ flux as a function of exposure using the visible inelasticity distribution of starting tracks. The significance for the current analysis with five-years of IceCube is marked by a green square. The potential significance for IceCube-Gen2 with an instrumented volume of 10 km^3 and observation time of 12 years is marked by an orange triangle.

hand, these results depend on the assumption of a single power-law astrophysical flux with $\gamma = 2.62$. If a hardening of the astrophysical power-law index exists as hinted by through-going tracks, then the enhanced rate of high energy events could reduce the time needed to resolve a ν_τ component. Lastly, other signatures of ν_τ interactions not considered here such as double cascades [242] and/or double pulses [243] may identify a ν_τ flux on a smaller time scale.

8.2.2 Additional Neutrino Interactions in and beyond the Standard Model

In Sec. 7.4.4, a strategy was developed for identifying the contribution of charm production to the total CC neutrino cross section. In the future, this same strategy could also be applied to identify top quark production, a process that forms $\sim 5\%$ of the CC neutrino cross section at 10 PeV as shown in Fig. 2.10. As first pointed out in [51], top production events tend to have higher inelasticity than charm production events due to the top quark's much heavier mass of 173 GeV. The decay of the top quark could also lead to dimuon or

trimuon signatures that also leave a distinct visible inelasticity distribution.

Visible inelasticity distributions could also explore neutrino interactions in the Standard Model beyond deep inelastic scattering. One processes that has been neglected so far in IceCube analyses is the production of a single W-boson in the Coulomb field of the nucleus,

$$\nu_\ell + N \to \ell^- + W^+ + X. \tag{8.1}$$

The cross section for this process on an oxygen target reaches $\sim 8\%$ of the CC DIS cross section at 1 PeV [244, 245]. Since muons can be created either in ν_μ interactions or in the subsequent decay of the W-boson, this process will leave a unique visible inelasticity distribution. Both top production and single W-boson production are examples of Standard Model processes that leave distinct signatures, but it will likely require a volume upgrade like IceCube-Gen2 or KM3NET [222] to observe them due to their small cross section.

There is also promise to use inelasticity distributions to search for neutrino interactions beyond the Standard Model. Inelasticity distributions were first recognized as a method to search for the production of supersymmetric particles in 1998 [246]. Another possibility is to search for leptoquarks, particles that directly couple leptons and quarks in Grand Unified Theories [247]. Muons from the decay of scalar second-generation leptoquarks are expected to have a flat inelasticity distribution whereas the DIS inelasticity distribution is strongly peaked near $y = 0$. A last example is searching for models of low-scale quantum gravity, where black holes can be produced in high-energy neutrino interactions and evaporate to produce muons [248]. The resulting black hole production events should have a higher visible inelasticity than DIS events. All of these processes also produce enhancements in the total neutrino cross section that may be detectable by observing increased Earth absorption in a sample of through-going tracks [48]. In the future, improved sensitivity to models of physics beyond the Standard Model could be obtained from a joint fit of through-going tracks and starting events.

For several of the above scenarios, recent results from the Large Hadron Collider (LHC) can rule out regions of parameters space that would produce an observable IceCube signal. To access regions completely inaccessible to man-made particle accelerators, new technologies are needed to detect ultra-high-energy neutrinos above 10^{17} eV where the center-of-mass energy exceeds that of the LHC. Detection of coherent radio Cherenkov emission provides the best prospects for observing ultra-high-energy neutrinos, and the ARA or ARIANNA detectors have been proposed to accomplish this goal [89]. For radio detection, the inelasticity of ν_e CC interactions could be observable since the electromagnetic cascade initiated by the electron is highly elongated by the Landau-Pomeranchuk-Migdal effect while the hadronic cascade is not, a possibility first pointed out in [249]. Along with measuring the absorption of Earth skimming neutrinos [250, 251], measuring inelasticity distributions is a promising technique to search for physics beyond the Standard Model with the highest energy neutrinos.

Bibliography

[1] K. Greisen, Annu. Rev. Nucl. Sci. **10**, 63 (1960), ISSN 0066-4243, URL http://www.annualreviews.org/doi/10.1146/annurev.ns.10.120160.000431.

[2] F. Reines, Annu. Rev. Nucl. Sci. **10**, 1 (1960), ISSN 0066-4243, URL http://www.annualreviews.org/doi/10.1146/annurev.ns.10.120160.000245.

[3] M. A. Markov, in *Proceedings, 10th Int. Conf. High-Energy Phys. (ICHEP 60) Rochester, NY, USA, 25 Aug - 1 Sep 1960* (1960), pp. 578–581, URL http://inspirehep.net/record/1341439/files/C60-08-25-p578.pdf.

[4] C. L. Cowan, F. Reines, F. B. Harrison, et al., Science **124**, 103 (1956), URL http://science.sciencemag.org/content/124/3212/103.

[5] T. Kajita, Rev. Mod. Phys. **88**, 030501 (2016), ISSN 0034-6861, URL https://link.aps.org/doi/10.1103/RevModPhys.88.030501.

[6] A. B. McDonald, Rev. Mod. Phys. **88**, 030502 (2016), ISSN 0034-6861, URL https://link.aps.org/doi/10.1103/RevModPhys.88.030502.

[7] M. G. Aartsen, R. Abbasi, Y. Abdou, et al., Science **342**, 1242856 (2013), ISSN 1095-9203, URL http://www.ncbi.nlm.nih.gov/pubmed/24264993.

[8] W. Pauli, Phys. Today **31N9**, 27 (1979), URL http://inspirehep.net/record/45177/files/meitner_0393.pdf.

[9] E. Fermi, Zeitschrift für Phys. **88**, 161 (1934), ISSN 1434-6001, URL http://link.springer.com/10.1007/BF01351864.

[10] S. L. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev. D **2**, 1285 (1970), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.2.1285.

[11] A. Salam, in *Conf. Proc. C680519* (1968), pp. 367–377, URL http://physics.princeton.edu/~mcdonald/examples/EP/salam_ns_367_68.pdf/53083.

[12] S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.19.1264.

[13] F. Hasert, S. Kabe, W. Krenz, et al., Phys. Lett. B **46**, 138 (1973), ISSN 03702693, URL http://linkinghub.elsevier.com/retrieve/pii/0370269373904991.

[14] N. Cabibbo, Phys. Rev. Lett. **10**, 531 (1963), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.10.531.

[15] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49**, 652 (1973), URL +http://dx.doi.org/10.1143/PTP.49.652.

[16] G. Danby, J.-M. Gaillard, K. Goulianos, et al., Phys. Rev. Lett. **9**, 36 (1962), ISSN

0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[17] K. Kodama, N. Ushida, C. Andreopoulos, et al., Phys. Lett. B **504**, 218 (2001), ISSN 03702693, URL http://www.sciencedirect.com/science/article/pii/S0370269301003070.

[18] N. Agafonova, A. Aleksandrov, A. Anokhina, et al., Phys. Rev. Lett. **115**, 121802 (2015), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.115.121802.

[19] S. Dell'Oro, S. Marcocci, M. Viel, et al., Adv. High Energy Phys. **2016**, 1 (2016), ISSN 1687-7357, URL http://www.hindawi.com/journals/ahep/2016/2162659/.

[20] C. S. Wu, E. Ambler, R. W. Hayward, et al., Phys. Rev. **105**, 1413 (1957), ISSN 0031-899X, URL https://link.aps.org/doi/10.1103/PhysRev.105.1413.

[21] M. Goldhaber, L. Grodzins, and A. W. Sunyar, Phys. Rev. **109**, 1015 (1958), ISSN 0031-899X, URL https://link.aps.org/doi/10.1103/PhysRev.109.1015.

[22] M. Aartsen, K. Abraham, M. Ackermann, et al., Phys. Rev. Lett. **117**, 071801 (2016), ISSN 0031-9007, URL http://link.aps.org/doi/10.1103/PhysRevLett.117.071801.

[23] B. Pontecorvo, Zh. Eksp. Teor. Fiz. **34**, 247 (1958), URL https://www.osti.gov/scitech/biblio/4349231.

[24] B. T. Cleveland, T. Daily, R. Davis, Jr., et al., Astrophys. J. **496**, 505 (1998), ISSN 0004-637X, URL http://stacks.iop.org/0004-637X/496/i=1/a=505.

[25] Q. R. Ahmad, R. C. Allen, T. C. Andersen, et al., Phys. Rev. Lett. **87**, 071301 (2001), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.87.071301.

[26] Y. Fukuda, T. Hayakawa, E. Ichihara, et al., Phys. Rev. Lett. **81**, 1562 (1998), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.81.1562.

[27] G. Drexlin, V. Hannen, S. Mertens, et al., Adv. High Energy Phys. **2013**, 1 (2013), ISSN 1687-7357, URL http://www.hindawi.com/journals/ahep/2013/293986/.

[28] K. Abazajian, E. Calabrese, A. Cooray, et al., Astropart. Phys. **35**, 177 (2011), ISSN 09276505, URL http://www.sciencedirect.com/science/article/pii/S0927650511001344.

[29] K. Olive, D. Weinberg, C. Woody, et al., Chinese Phys. C **40**, 100001 (2016), ISSN 1674-1137, URL http://stacks.iop.org/1674-1137/40/i=10/a=100001?key=crossref.5ca5f9c166adc59f41e03a802b715c19.

[30] M. C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz, J. High Energy Phys. **2014**, 52 (2014), ISSN 1029-8479, URL http://dx.doi.org/10.1007/JHEP11(2014)052.

[31] M. G. Aartsen, K. Abraham, M. Ackermann, et al., J. Phys. G Nucl. Part. Phys. **44**, 054006 (2017), ISSN 0954-3899, URL http://stacks.iop.org/0954-3899/44/i=5/a=054006?key=crossref.7805601c31f929b21f483e709b5d48ea.

[32] J. M. Conrad, M. H. Shaevitz, and T. Bolton, Rev. Mod. Phys. **70**, 1341 (1998), ISSN 0034-6861, URL https://link.aps.org/doi/10.1103/RevModPhys.70.1341.

[33] R. Devenish and A. Cooper-Sarkar, *Deep inelastic scattering* (Oxford University Press, 2004), ISBN 9780198506713.

[34] J. D. Bjorken, Phys. Rev. **179**, 1547 (1969), ISSN 0031-899X, URL https://link.aps.org/doi/10.1103/PhysRev.179.1547.

[35] V. Radescu, H1, Z. Collaborations, et al., in *Proceedings, 35th Int. Conf. High energy Phys. (ICHEP 2010) Paris, Fr. July 22-28, 2010* (Paris, France, 2010), pp. 22–25, URL https://pos.sissa.it/120/168/pdf.

[36] J. C. Collins, D. E. Soper, and G. Sterman, in *Adv. Ser. Direct. High Energy Phys.* (1989), vol. 5, pp. 1–91, URL http://www.worldscientific.com/doi/abs/10.1142/9789814503266_0001.

[37] Y. L. Dokshitzer, Zh. Eksp. Teor. Fiz. **46**, 641 (1977), URL http://inspirehep.net/record/126153?ln=en.

[38] V. Gribov and L. Lipatov, Sov. J. Nucl. Phys. **15**, 781 (1972), URL https://inspirehep.net/record/73449?ln=en.

[39] G. Altarelli and G. Parisi, Nucl. Phys. B **126**, 298 (1977), ISSN 05503213, URL http://linkinghub.elsevier.com/retrieve/pii/0550321377903844.

[40] J. Mousseau, M. Wospakrik, L. Aliaga, et al., Phys. Rev. D **93**, 071101 (2016), ISSN 2470-0010, URL https://link.aps.org/doi/10.1103/PhysRevD.93.071101.

[41] K. J. Eskola, P. Paakkinen, H. Paukkunen, et al., Eur. Phys. J. C **77**, 163 (2017), ISSN 1434-6052, URL http://dx.doi.org/10.1140/epjc/s10052-017-4725-9.

[42] A. Cooper-Sarkar, P. Mertsch, and S. Sarkar, J. High Energy Phys. **2011**, 42 (2011), ISSN 1029-8479, URL http://link.springer.com/10.1007/JHEP08(2011)042.

[43] L. A. Anchordoqui, A. M. Cooper-Sarkar, D. Hooper, et al., Phys. Rev. D **74**, 043008 (2006), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.74.043008.

[44] Y. S. Jeong and M. H. Reno, Phys. Rev. D **81**, 114012 (2010), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.81.114012.

[45] Y. S. Jeong and M. H. Reno (2010), URL http://arxiv.org/abs/1007.1966http://dx.doi.org/10.1103/PhysRevD.82.033010.

[46] M. Tzanov, D. Naples, S. Boyd, et al., Phys. Rev. D **74**, 12008 (2006), URL https://link.aps.org/doi/10.1103/PhysRevD.74.012008.

[47] M. G. Aartsen, G. C. Hill, A. Kyriacou, et al., Nature (2017), ISSN 0028-0836, URL http://www.nature.com/doifinder/10.1038/nature24459.

[48] S. Miarecki, Ph. D. thesis, University of California, Berkeley (2016), URL https://escholarship.org/uc/item/7q09d51t.

[49] H. Georgi and H. D. Politzer, Phys. Rev. D **14**, 1829 (1976), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.14.1829.

[50] S. A. Rabinowitz, C. Arroyo, K. T. Bachmann, et al., Phys. Rev. Lett. **70**, 134 (1993), ISSN 0031-9007, URL http://link.aps.org/doi/10.1103/PhysRevLett.70.134.

[51] V. Barger, E. Basso, Y. Gao, et al., Phys. Rev. D **95**, 093002 (2017), ISSN 2470-0010, URL http://link.aps.org/doi/10.1103/PhysRevD.95.093002.

[52] J. A. Formaggio and G. P. Zeller, Rev. Mod. Phys. **84**, 1307 (2012), ISSN 0034-6861, URL http://link.aps.org/doi/10.1103/RevModPhys.84.1307.

[53] S. L. Glashow, Phys. Rev. **118**, 316 (1960), ISSN 0031-899X, URL http://link.aps.

org/doi/10.1103/PhysRev.118.316.
[54] V. Berezinsky and A. Gazizov, JETP Lett. **25**, 254 (1977), URL http://inspirehep.net/record/124058?ln=en.
[55] J. Matthews and C. Jui, Nucl. Phys. B - Proc. Suppl. **87**, 411 (2000), ISSN 09205632, URL http://linkinghub.elsevier.com/retrieve/pii/S0920563200007064.
[56] A. Hillas and V. F. Hess, in *Cosm. Rays* (1972), pp. 139–147, ISBN 9780080167244, URL http://www.sciencedirect.com/science/article/pii/B9780080167244500136.
[57] K. Greisen, Phys. Rev. Lett. **16**, 748 (1966), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.16.748.
[58] G. T. Atsepin and V. A. Kuz'min, J. Exp. Theor. Phys. Lett. **4**, 78 (1966), ISSN 0021-3640.
[59] G. A. Medina-Tanco, Astrophys. J. **510**, L91 (1999), ISSN 0004637X, URL http://stacks.iop.org/1538-4357/510/i=2/a=L91.
[60] J. Abraham, P. Abreu, M. Aglietta, et al., Phys. Rev. Lett. **101**, 061101 (2008), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.101.061101.
[61] R. U. Abbasi, T. Abu-Zayyad, M. Allen, et al., Phys. Rev. Lett. **100**, 101101 (2008), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.100.101101.
[62] R. U. Abbasi, M. Abe, T. Abu-Zayyad, et al., Astrophys. J. **790**, L21 (2014), ISSN 2041-8205, URL http://stacks.iop.org/2041-8205/790/i=2/a=L21?key=crossref.7925149da575c297568066043f0ac173.
[63] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Astrophys. J. **826**, 220 (2016), ISSN 1538-4357, URL http://stacks.iop.org/0004-637X/826/i=2/a=220?key=crossref.80edeb8cd150bd107512bad85b132bce.
[64] E. Fermi, Phys. Rev. **75**, 1169 (1949), ISSN 0031-899X, URL https://link.aps.org/doi/10.1103/PhysRev.75.1169.
[65] A. R. Bell, Mon. Not. R. Astron. Soc. **182**, 147 (1978), ISSN 0035-8711, URL https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/182.2.147.
[66] T. K. Gaisser, *Cosmic rays and particle physics* (Cambridge University Press, 1990), ISBN 9780521339315.
[67] A. Obermeier, P. Boyle, J. Hörandel, et al., Astrophys. J. **752**, 69 (2012), ISSN 0004-637X, URL http://stacks.iop.org/0004-637X/752/i=1/a=69?key=crossref.c92c3112e3dc30b487d42d62f0c09b72.
[68] P. Chen, T. Tajima, and Y. Takahashi, Phys. Rev. Lett. **89**, 161101 (2002), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.89.161101.
[69] D. Giannios, Mon. Not. R. Astron. Soc. Lett. **408**, L46 (2010), URL +http://dx.doi.org/10.1111/j.1745-3933.2010.00925.x.
[70] M. G. Aartsen, M. Ackermann, J. Adams, et al., Adv. Sp. Res. (2017), ISSN 0273-1177, URL http://www.sciencedirect.com/science/article/pii/S0273117717303757.
[71] A. M. Hillas, Annu. Rev. Astron. Astrophys. **22**, 425 (1984), ISSN 0066-4146, URL http://www.annualreviews.org/doi/10.1146/annurev.aa.22.090184.002233.

[72] T. K. Gaisser, Astropart. Phys. **35**, 801 (2012), ISSN 0927-6505, URL http://www.sciencedirect.com/science/article/pii/S0927650512000497.

[73] A. M. Hillas, J. Phys. G Nucl. Part. Phys. **31**, R95 (2005), ISSN 0954-3899, URL http://stacks.iop.org/0954-3899/31/i=5/a=R02?key=crossref.fe40e031a4035007189b0da8706b8db9.

[74] T. K. Gaisser, T. Stanev, and S. Tilav, Front. Phys. **8**, 748 (2013), ISSN 2095-0462, URL http://link.springer.com/10.1007/s11467-013-0319-7.

[75] M. Ackermann, M. Ajello, A. Allafort, et al., Science **339** (2013), URL http://science.sciencemag.org/content/339/6121/807.

[76] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Astrophys. J. **835**, 151 (2017), ISSN 1538-4357, URL http://stacks.iop.org/0004-637X/835/i=2/a=151?key=crossref.8a01080309e6e445250568d23589b0c0.

[77] W. D. Arnett, J. N. Bahcall, R. P. Kirshner, et al., Annu. Rev. Astron. Astrophys. **27**, 629 (1989), ISSN 0066-4146, URL http://www.annualreviews.org/doi/10.1146/annurev.aa.27.090189.003213.

[78] M. Aartsen, M. Ackermann, J. Adams, et al., Phys. Rev. Lett. **113**, 101101 (2014), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.113.101101.

[79] C. Spiering, in *Rev. Mod. Astron.* (Wiley-VCH Verlag GmbH and Co. KGaA, Weinheim, Germany, 2008), pp. 375–405, ISBN 9783527622993, URL http://doi.wiley.com/10.1002/9783527622993.ch16.

[80] S. H. Margolis, D. N. Schramm, and R. Silberberg, Astrophys. J. **221**, 990 (1978), ISSN 0004-637X, URL http://adsabs.harvard.edu/doi/10.1086/156104.

[81] F. W. Stecker, Astrophys. J. **228**, 919 (1979), ISSN 0004-637X, URL http://adsabs.harvard.edu/doi/10.1086/156919.

[82] E. Waxman and J. Bahcall, Phys. Rev. Lett. **78**, 2292 (1997), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.78.2292.

[83] F. W. Stecker, C. Done, M. H. Salamon, et al., Phys. Rev. Lett. **66**, 2697 (1991), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.66.2697.

[84] M. Mandelartz and J. B. Tjus, Astropart. Phys. **65**, 80 (2015), ISSN 0927-6505, URL http://www.sciencedirect.com/science/article/pii/S0927650514001881.

[85] J. K. Becker, Phys. Rep. **458**, 173 (2008), ISSN 0370-1573, URL http://www.sciencedirect.com/science/article/pii/S037015730800029X.

[86] A. Loeb and E. Waxman, J. Cosmol. Astropart. Phys. **2006**, 003 (2006), ISSN 1475-7516, URL http://stacks.iop.org/1475-7516/2006/i=05/a=003?key=crossref.59902d06c6b3effb9bb25e1ebaa3e6b0.

[87] V. Beresinsky and G. Zatsepin, Phys. Lett. B **28**, 423 (1969), ISSN 03702693, URL http://linkinghub.elsevier.com/retrieve/pii/0370269369903414.

[88] M. Aartsen, K. Abraham, M. Ackermann, et al., Phys. Rev. Lett. **117**, 241101 (2016), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.117.241101.

[89] A. Connolly and A. Vieregg, in *Neutrino Astron.* (WORLD SCIENTIFIC, 2017),

pp. 217–241, ISBN 978-981-4759-40-3, URL http://www.worldscientific.com/worldscibooks/10.1142/9964.

[90] E. Waxman and J. Bahcall, Phys. Rev. D **59**, 023002 (1998), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.59.023002.

[91] M. Bustamante, J. F. Beacom, and W. Winter, Phys. Rev. Lett. **115**, 161302 (2015), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.115.161302.

[92] P. Lipari, M. Lusignoli, and D. Meloni, Phys. Rev. D **75**, 123005 (2007), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.75.123005.

[93] S. R. Klein, R. E. Mikkelsen, and J. Becker Tjus, Astrophys. J. **779**, 106 (2013), ISSN 0004-637X, URL http://stacks.iop.org/0004-637X/779/i=2/a=106?key=crossref.2e7509ef86bf245c4b4bee135bef1fed.

[94] M. Kachelrieß and R. Tomàs, Phys. Rev. D **74**, 063009 (2006), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.74.063009.

[95] L. A. Anchordoqui, H. Goldberg, F. Halzen, et al., Phys. Lett. B **593**, 42 (2004), ISSN 03702693, URL http://linkinghub.elsevier.com/retrieve/pii/S0370269304006847.

[96] L. A. Anchordoqui, J. F. Beacom, H. Goldberg, et al., Phys. Rev. D **75**, 063001 (2007), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.75.063001.

[97] T. Kashti and E. Waxman, Phys. Rev. Lett. **95**, 181101 (2005), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.95.181101.

[98] S. Hümmer, M. Maltoni, W. Winter, et al., Astropart. Phys. **34**, 205 (2010), ISSN 0927-6505, URL http://www.sciencedirect.com/science/article/pii/S092765051000126X.

[99] A. Bhattacharya, R. Gandhi, W. Rodejohann, et al., J. Cosmol. Astropart. Phys. **2011**, 017 (2011), ISSN 1475-7516, URL http://stacks.iop.org/1475-7516/2011/i=10/a=017?key=crossref.b212dcb51b19ae4da01fa2d6358ae1ee.

[100] L. A. Anchordoqui, H. Goldberg, F. Halzen, et al., Phys. Lett. B **621**, 18 (2005), ISSN 03702693, URL http://linkinghub.elsevier.com/retrieve/pii/S0370269305008658.

[101] V. Barger, L. Fu, J. Learned, et al., Phys. Rev. D **90**, 121301 (2014), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.90.121301.

[102] G. Bertone, D. Hooper, and J. Silk, Phys. Rep. **405**, 279 (2005), ISSN 03701573, URL http://linkinghub.elsevier.com/retrieve/pii/S0370157304003515.

[103] J. A. Frieman, M. S. Turner, and D. Huterer, Annu. Rev. Astron. Astrophys. **46**, 385 (2008), ISSN 0066-4146, URL http://www.annualreviews.org/doi/10.1146/annurev.astro.46.060407.145243.

[104] J. L. Feng, Annu. Rev. Astron. Astrophys. **48**, 495 (2010), ISSN 0066-4146, URL http://www.annualreviews.org/doi/10.1146/annurev-astro-082708-101659.

[105] M. G. Aartsen, M. Ackermann, J. Adams, et al., Eur. Phys. J. C **77**, 146 (2017), ISSN 1434-6052, URL https://doi.org/10.1140/epjc/s10052-017-4689-9.

[106] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Eur. Phys. J. C **77**, 82 (2017),

ISSN 1434-6052, URL https://doi.org/10.1140/epjc/s10052-016-4582-y.
[107] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Eur. Phys. J. C **76**, 531 (2016), ISSN 1434-6052, URL https://doi.org/10.1140/epjc/s10052-016-4375-3.
[108] M. G. Aartsen, R. Abbasi, Y. Abdou, et al., Phys. Rev. D **88**, 122001 (2013), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.88.122001.
[109] R. Allahverdi and K. Richardson, Phys. Rev. D **85**, 113012 (2012), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.85.113012.
[110] J. F. Beacom, N. F. Bell, D. Hooper, et al., Phys. Rev. D **68**, 093005 (2003), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.68.093005.
[111] J. F. Beacom, N. F. Bell, D. Hooper, et al., Phys. Rev. Lett. **90**, 181301 (2003), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.90.181301.
[112] M. Bustamante, J. F. Beacom, and K. Murase, Phys. Rev. D **95**, 063013 (2017), ISSN 2470-0010, URL https://link.aps.org/doi/10.1103/PhysRevD.95.063013.
[113] J. F. Beacom, N. F. Bell, D. Hooper, et al., Phys. Rev. Lett. **92**, 011101 (2004), URL https://link.aps.org/doi/10.1103/PhysRevLett.92.011101.
[114] D. Morgan, E. Winstanley, J. Brunner, et al., Astropart. Phys. **25**, 311 (2006), ISSN 09276505, URL http://linkinghub.elsevier.com/retrieve/pii/S0927650506000338.
[115] L. A. Anchordoqui, H. Goldberg, M. C. Gonzalez-Garcia, et al., Phys. Rev. D **72**, 065019 (2005), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.72.065019.
[116] V. Brdar, J. Kopp, and X.-P. Wang, J. Cosmol. Astropart. Phys. **2017**, 026 (2017), ISSN 1475-7516, URL http://stacks.iop.org/1475-7516/2017/i=01/a=026?key=crossref.9e31419ba883b4b46d2b124242be3228.
[117] D. Hooper, D. Morgan, and E. Winstanley, Phys. Rev. D **72**, 065009 (2005), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.72.065009.
[118] V. A. Kostelecký and M. Mewes, Phys. Rev. D **70**, 031902 (2004), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.70.031902.
[119] M. Gasperini, Phys. Rev. D **39**, 3606 (1989), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.39.3606.
[120] C. A. Argüelles, T. Katori, and J. Salvado, Phys. Rev. Lett. **115**, 161303 (2015), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.115.161303.
[121] M. Aartsen, M. Ackermann, J. Adams, et al., Phys. Rev. D **91**, 122004 (2015), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.91.122004.
[122] M. Honda, T. Kajita, K. Kasahara, et al., Phys. Rev. D **75**, 043006 (2007), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.75.043006.
[123] G. D. Barr, T. K. Gaisser, P. Lipari, et al., Phys. Rev. D **70**, 023006 (2004), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.70.023006.
[124] M. G. Aartsen, M. Ackermann, J. Adams, et al., Eur. Phys. J. C **75**, 116 (2015), ISSN 1434-6044, URL http://link.springer.com/10.1140/epjc/s10052-015-3330-z.
[125] M. G. Aartsen, R. Abbasi, Y. Abdou, et al., Phys. Rev. Lett. **110**, 151105 (2013), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.110.151105.

[126] R. Enberg, M. H. Reno, and I. Sarcevic, Phys. Rev. D **78**, 043005 (2008), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.78.043005.
[127] T. K. Gaisser and S. R. Klein, Astropart. Phys. **64**, 13 (2015), ISSN 0927-6505, URL http://www.sciencedirect.com/science/article/pii/S0927650514001613.
[128] Tech. Rep., Washington, D.C. (1976), URL https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19770009539.pdf.
[129] S. Roesler, R. Engel, and J. Ranft, in *Adv. Monte Carlo Radiat. Physics, Part. Transp. Simul. Appl.* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2001), pp. 1033–1038, URL http://link.springer.com/10.1007/978-3-642-18211-2_166.
[130] D. Chirkin (2004), URL http://arxiv.org/abs/hep-ph/0407078.
[131] M. G. G. Aartsen, R. Abbasi, M. Ackermann, et al., Phys. Rev. D **89**, 062007 (2013), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.89.062007.
[132] N. Agafonova, A. Aleksandrov, A. Anokhina, et al., Eur. Phys. J. C **74**, 2933 (2014), ISSN 1434-6044, URL http://link.springer.com/10.1140/epjc/s10052-014-2933-0.
[133] V. Khachatryan, A. M. Sirunyan, A. Tumasyan, et al., Phys. Lett. B **692**, 83 (2010), ISSN 0370-2693, URL http://www.sciencedirect.com/science/article/pii/S0370269310008725.
[134] P. Adamson, C. Andreopoulos, K. E. Arms, et al., Phys. Rev. D **76**, 052003 (2007), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.76.052003.
[135] A. Bulmahn and M. H. Reno, Phys. Rev. D **82**, 057302 (2010), URL https://journals.aps.org/prd/abstract/10.1103/PhysRevD.82.057302.
[136] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Astrophys. J. **833**, 3 (2016), ISSN 0004-637X, URL http://dx.doi.org/10.3847/0004-637X/833/1/3.
[137] A. Bhattacharya, R. Enberg, M. H. Reno, et al., J. High Energy Phys. **2015**, 110 (2015), ISSN 1029-8479, URL https://doi.org/10.1007/JHEP06(2015)110.
[138] R. Gauld, J. Rojo, L. Rottoli, et al., J. High Energy Phys. **2016**, 130 (2016), ISSN 1029-8479, URL https://doi.org/10.1007/JHEP02(2016)130.
[139] M. Aartsen, M. Ackermann, J. Adams, et al., Phys. Rev. D **91**, 022001 (2015), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.91.022001.
[140] S. Schönert, T. K. Gaisser, E. Resconi, et al., Phys. Rev. D **79**, 043009 (2009), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.79.043009.
[141] T. K. Gaisser, K. Jero, A. Karle, et al., Phys. Rev. D **90**, 023009 (2014), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.90.023009.
[142] P. Cherenkov, Dokl. Akad. Nauk Ser. Fiz. **2**, 451 (1934).
[143] I. Frank and I. Tamm, C.R.Acad.Sci.URSS **14**, 109 (1937).
[144] P. Price and K. Woschnagg, Astropart. Phys. **15**, 97 (2001), ISSN 09276505, URL http://www.sciencedirect.com/science/article/pii/S0927650500001420.
[145] J. D. Jackson, *Classical electrodynamics* (Wiley, 1999), ISBN 9780471309321.
[146] S. Agostinelli, J. Allison, K. Amako, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **506**, 250 (2003), ISSN 01689002.
[147] C. Wiebusch, Ph. D. thesis, RWTH Aachen (1995), URL https://web.physik.rwth-

aachen.de/~wiebusch/Publications/Various/phd.pdf.
[148] L. Landau and I. Pomeranchuk, Dokl. Akad. Nauk Ser. Fiz. **92**, 535 (1953).
[149] A. B. Migdal, Phys. Rev. **103**, 1811 (1956), ISSN 0031-899X, URL http://link.aps.org/doi/10.1103/PhysRev.103.1811.
[150] S. Klein, Rev. Mod. Phys. **71**, 1501 (1999), ISSN 0034-6861, URL http://link.aps.org/doi/10.1103/RevModPhys.71.1501.
[151] L. Rädel and C. Wiebusch, Astropart. Phys. **44**, 102 (2013), ISSN 09276505, URL http://www.sciencedirect.com/science/article/pii/S092765051300025X?via%3Dihub.
[152] D. Chirkin and W. Rhode (2004), URL http://arxiv.org/abs/hep-ph/0407075.
[153] T. Sjöstrand, S. Mrenna, and P. Skands, J. High Energy Phys. **2006**, 026 (2006), ISSN 1029-8479, URL http://stacks.iop.org/1126-6708/2006/i=05/a=026?key=crossref.7fbc8fa1a47a48f7565bead655446685.
[154] M. Kowalski, Tech. Rep. (2002), URL https://internal.icecube.wisc.edu/reports/amanda/data/20020803-track.pdf.
[155] D. Heck, J. Knapp, J. N. Capdevielle, et al. (1998), URL http://inspirehep.net/record/469835/files/FZKA6019.pdf.
[156] S. Panknin, Ph. D. thesis, Humboldt-Universität zu Berlin (2011), URL http://edoc.hu-berlin.de/dissertationen/panknin-sebastian-2011-09-15/PDF/panknin.pdf.
[157] L. Rädel, Master thesis, RWTH Aachen (2012), URL https://internal.icecube.wisc.edu/reports/data/icecube/2012/10/001/icecube_201210001_v1.pdf.
[158] D. E. Groom, N. V. Mokhov, and S. I. Striganov, At. Data Nucl. Data Tables **78**, 183 (2001), ISSN 0092640X.
[159] T. DeYoung, S. Razzaque, and D. Cowen, Astropart. Phys. **27**, 238 (2007), ISSN 09276505.
[160] P. Lipari, Astropart. Phys. **1**, 195 (1993), ISSN 09276505.
[161] A. Achterberg, M. Ackermann, J. Adams, et al., Astropart. Phys. **26**, 155 (2006), ISSN 09276505.
[162] J. Yang, *IceCube Array* (2012), URL https://gallery.icecube.wisc.edu/internal/v/graphics/arraygraphics2011/blueTopArrayWLabels.jpg.html.
[163] R. Abbasi, Y. Abdou, T. Abu-Zayyad, et al., Astropart. Physics, Vol. 35, Issue 10, p. 615-624. **35**, 615 (2011), ISSN 0927-6505, URL http://dx.doi.org/10.1016/j.astropartphys.2012.01.004.
[164] R. Abbasi, Y. Abdou, M. Ackermann, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **700**, 188 (2013), ISSN 01689002, URL http://www.sciencedirect.com/science/article/pii/S016890021201217X.
[165] R. Abbasi, M. Ackermann, J. Adams, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **601**, 294 (2009), ISSN 01689002, URL http://www.sciencedirect.com/science/article/pii/S0168900209000084?via%3Dihub.
[166] M. Aartsen, M. Ackermann, J. Adams, et al., J. Instrum. **12**, P03012 (2017),

ISSN 1748-0221, URL http://stacks.iop.org/1748-0221/12/i=03/a=P03012?key=crossref.70875ec02751d59944b8f24f1551de41.

[167] R. Abbasi, Y. Abdou, T. Abu-Zayyad, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **618**, 139 (2010), ISSN 01689002, URL http://www.sciencedirect.com/science/article/pii/S0168900210006662?via%3Dihub.

[168] M. G. Aartsen, R. Abbasi, M. Ackermann, et al., J. Instrum. **9**, P03009 (2013), ISSN 1748-0221, URL http://dx.doi.org/10.1088/1748-0221/9/03/P03009.

[169] *Data Movement*, URL https://icecube.wisc.edu/science/data/datamovement.

[170] M. Ackermann, J. Ahrens, X. Bai, et al., J. Geophys. Res. **111**, D13203 (2006), ISSN 0148-0227, URL http://doi.wiley.com/10.1029/2005JD006687.

[171] P. B. Price, K. Woschnagg, and D. Chirkin, Geophys. Res. Lett. **27**, 2129 (2000), ISSN 00948276, URL http://doi.wiley.com/10.1029/2000GL011351.

[172] K. Casey, T. Fudge, T. Neumann, et al., Ann. Glaciol. **55**, 137 (2014), ISSN 02603055, URL http://openurl.ingenta.com/content/xref?genre=article&issn=0260-3055&volume=55&issue=68&spage=137.

[173] P. Askebjer, S. W. Barwick, L. Bergström, et al., Appl. Opt. **36**, 4168 (1997), ISSN 0003-6935, URL https://www.osapublishing.org/abstract.cfm?URI=ao-36-18-4168.

[174] P. B. Price and L. Bergström, Appl. Opt. **36**, 4181 (1997), ISSN 0003-6935, URL https://www.osapublishing.org/abstract.cfm?URI=ao-36-18-4181.

[175] M. Aartsen, R. Abbasi, Y. Abdou, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **711**, 73 (2013), ISSN 01689002, URL http://www.sciencedirect.com/science/article/pii/S0168900213001460?via%3Dihub.

[176] D. Chirkin, in *Proceeding 33rd Int. Cosm. Ray Conf.* (The Astroparticle Physics Conference, Rio de Janeiro, Brazil, 2013), p. 0580, URL http://www.cbpf.br/~icrc2013/papers/icrc2013-0580.pdf.

[177] Y. D. He and P. B. Price, J. Geophys. Res. Atmos. **103**, 17041 (1998), ISSN 01480227, URL http://doi.wiley.com/10.1029/98JD01643.

[178] N. E. Bramall, R. C. Bay, K. Woschnagg, et al., Geophys. Res. Lett. **32**, L21815 (2005), ISSN 0094-8276, URL http://doi.wiley.com/10.1029/2005GL024236.

[179] R. C. Bay, R. A. Rohde, P. B. Price, et al., J. Geophys. Res. **115**, D14126 (2010), ISSN 0148-0227, URL http://doi.wiley.com/10.1029/2009JD013741.

[180] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Astropart. Physics, Vol. 78, p. 1-27. **78**, 1 (2015), ISSN 0927-6505, URL http://dx.doi.org/10.1016/j.astropartphys.2016.01.006.

[181] M. Aartsen, K. Abraham, M. Ackermann, et al., Phys. Rev. Lett. **115**, 081102 (2015), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.115.081102.

[182] T. A. The AMANDA Collaboration and J. Ahrens, Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **524**, 169 (2004), URL

http://www.sciencedirect.com/science/article/pii/S0168900204001871?via%3Dihub.

[183] M. Aartsen, R. Abbasi, Y. Abdou, et al., Phys. Rev. D **89**, 102004 (2014), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.89.102004.

[184] R. Abbasi, Y. Abdou, M. Ackermann, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **703**, 190 (2013), ISSN 01689002, URL http://www.sciencedirect.com/science/article/pii/S0168900212014234.

[185] D. F. Cowen, J. Phys. Conf. Ser. **60**, 227 (2007), ISSN 1742-6588, URL http://stacks.iop.org/1742-6596/60/i=1/a=048?key=crossref.d3808261195728e49c01cbcc1554a223.

[186] J. G. Learned and S. Pakvasa, Astropart. Phys. **3**, 267 (1995), ISSN 09276505, URL http://linkinghub.elsevier.com/retrieve/pii/0927650594000433.

[187] M. G. Aartsen, K. Abraham, M. Ackermann, et al. (2015), URL http://dx.doi.org/10.1103/PhysRevD.93.022001.

[188] A. Ishihara and K. Hoshina, *neutrino-generator*, URL http://code.icecube.wisc.edu/svn/projects/neutrino-generator.

[189] A. Gazizov and M. Kowalski, Comput. Phys. Commun. **172**, 203 (2005), ISSN 00104655.

[190] G. Prézeau, Astrophys. J. **814**, 122 (2015), ISSN 1538-4357, URL http://stacks.iop.org/0004-637X/814/i=2/a=122?key=crossref.2bdabffa24637b9d6bbfd126bd51fd99.

[191] A. M. Dziewonski and D. L. Anderson, Phys. Earth Planet. Inter. **25**, 297 (1981), ISSN 00319201.

[192] C. Andreopoulous, C. Barry, S. Dytman, et al. (2015), URL http://arxiv.org/abs/1510.05494.

[193] G. De Lellis, F. Di Capua, and P. Migliozzi, Phys. Lett. B **550**, 16 (2002), ISSN 03702693, URL http://www.sciencedirect.com/science/article/pii/S0370269302029544?via%3Dihub.

[194] A. E. Hedin, J. Geophys. Res. Sp. Phys. **96**, 1159 (1991), ISSN 01480227, URL http://doi.wiley.com/10.1029/90JA02125.

[195] E.-J. Ahn, R. Engel, T. K. Gaisser, et al., Phys. Rev. D **80**, 094003 (2009), ISSN 1550-7998, URL http://link.aps.org/doi/10.1103/PhysRevD.80.094003.

[196] J. van Santen, Ph. D. thesis, University of Wisconsin–Madison (2014), URL https://na02.alma.exlibrisgroup.com/view/action/uresolver.do?operation=resolveService&package_service_id=13405786150002122&institutionId=2122&customerId=2120.

[197] J.-H. Koehne, K. Frantzen, M. Schmitz, et al., *PROPOSAL*, URL code.icecube.wisc.edu/svn/projects/PROPOSAL/.

[198] A. Bueno and A. Gascón, Comput. Phys. Commun. **185**, 638 (2014), ISSN 00104655.

[199] A. Bueno, A. Gascón, J. Illana, et al., J. Cosmol. Astropart. Phys. **2012**, 028 (2012), ISSN 1475-7516, URL http://stacks.iop.org/1475-7516/2012/i=02/a=028?key=crossref.7c630381873ec21c1e14b2068fc2de5a.

[200] L. Rädel and C. Wiebusch, Astropart. Phys. **38**, 53 (2012), ISSN 09276505, URL http://www.sciencedirect.com/science/article/pii/S0927650512001831?via%3Dihub.

[201] C. Kopper, *Photons from IceCube Muon*, URL http://news.psu.edu/photo/451969/2017/02/21/photons-icecube-muon.

[202] J. Lundberg, P. Miočinović, K. Woschnagg, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **581**, 619 (2007), ISSN 01689002, URL http://www.sciencedirect.com/science/article/pii/S0168900207015161?via%3Dihub.

[203] D. Chirkin, *PPC*, URL http://code.icecube.wisc.edu/svn/projects/ppc/.

[204] C. Kopper, *clsim*, URL http://code.icecube.wisc.edu/svn/projects/clsim/.

[205] N. Whitehorn, J. van Santen, and S. Lafebre, Comput. Phys. Commun. **184**, 2214 (2013), ISSN 00104655, URL http://www.sciencedirect.com/science/article/pii/S0010465513001434?via%3Dihub.

[206] C. Weaver, *PMTResponseSimulator*, URL http://code.icecube.wisc.edu/svn/projects/DOMLauncher/trunk/private/DOMLauncher/PMTResponseSimulator.cxx.

[207] S. Flis, *DOMLauncher*, URL http://code.icecube.wisc.edu/svn/projects/DOMLauncher/.

[208] A. Bouchta, A. Olivas, and G. Wikstrom, *trigger-sim*, URL code.icecube.wisc.edu/svn/projects/trigger-sim.

[209] M. G. Aartsen, R. Abbasi, Y. Abdou, et al., Phys. Rev. Lett. **111**, 021103 (2013), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.111.021103.

[210] O. Mena, S. Palomares-Ruiz, and A. C. Vincent, Phys. Rev. Lett. **113**, 091103 (2014), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.113.091103.

[211] S. Palomares-Ruiz, A. C. Vincent, and O. Mena, Phys. Rev. D **91**, 103008 (2015), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.91.103008.

[212] A. Palladino, G. Pagliaroli, F. Villante, et al., Phys. Rev. Lett. **114**, 171101 (2015), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.114.171101.

[213] M. G. Aartsen, M. Ackermann, J. Adams, et al., Phys. Rev. Lett. **114**, 171102 (2015), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.114.171102.

[214] D. Pandel, Diploma thesis, Humboldt-Universität zu Berlin (1996).

[215] G. Japaridze and M. Ribordy (2005), URL http://arxiv.org/abs/astro-ph/0506136.

[216] N. van Eijndhoven, O. Fadiran, and G. Japaridze, Astropart. Phys. **28**, 456 (2007), ISSN 0927-6505, URL http://www.sciencedirect.com/science/article/pii/S0927650507001260.

[217] G. J. Feldman and R. D. Cousins, Phys. Rev. D **57**, 3873 (1998), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.57.3873.

[218] F. James, *MINUIT Function Minimization and Error Analysis: Reference Manual Version 94.1* (1994), URL https://inspirehep.net/record/1258343/.
[219] H. L. Lai, J. Huston, S. Kuhlmann, et al. (1999), URL http://dx.doi.org/10.1007/s100529900196.
[220] M. G. Aartsen, K. Abraham, M. Ackermann, et al., Astrophys. J. **809**, 98 (2015), ISSN 1538-4357, URL http://stacks.iop.org/0004-637X/809/i=1/a=98?key=crossref.2f347e3557d5fe39880a0cb51b6d6dad.
[221] M. G. Aartsen, M. Ackermann, J. Adams, et al. (2014), URL http://arxiv.org/abs/1412.5106.
[222] A. Kappes (2007), URL http://arxiv.org/abs/0711.0563.
[223] S. Seunarine, Tech. Rep. (2011), URL https://docushare.icecube.wisc.edu/dsweb/Get/Document-56859/TFT_Filter_2010_11_Cascade.pdf.
[224] C. Kopper and L. Schulte, Tech. Rep. (2012), URL https://docushare.icecube.wisc.edu/dsweb/Get/Document-59396/cascade_filter_proposal_2012_v0.4.pdf.
[225] L. Schulte and M. Lesiak-Bzdak, Tech. Rep. (2012), URL http://icecube.wisc.edu/~mlesiak-bzdak/Documents/Proposals/cscd-filter-l2-ic86-2013-v05.pdf.
[226] N. Kurahashi-Neilson, Tech. Rep. (2010), URL https://docushare.icecube.wisc.edu/dsweb/Get/Document-56581/MuonFilterProposal.pdf.
[227] T. Glusenkamp, Tech. Rep. (2012), URL https://docushare.icecube.wisc.edu/dsweb/Get/Document-59906/MuonFilter2012_v3_with_addendum.pdf.
[228] D. Chirkin and C. Weaver, *TopologicalSplitter*, URL http://code.icecube.wisc.edu/svn/projects/TopologicalSplitter/.
[229] Y. Freund and R. E. Schapire, J. Comput. Syst. Sci. **55**, 119 (1997), ISSN 00220000, URL http://linkinghub.elsevier.com/retrieve/pii/S002200009791504X.
[230] F. Pedregosa, G. Varoquaux, A. Gramfort, et al., J. Mach. Learn. Res. **12**, 2825 (2011), URL https://dl.acm.org/citation.cfm?id=2078195.
[231] K. M. Gorski, E. Hivon, A. J. Banday, et al., Astrophys. J. **622**, 759 (2005), ISSN 0004-637X, URL http://stacks.iop.org/0004-637X/622/i=2/a=759.
[232] M. Aartsen, R. Abbasi, Y. Abdou, et al., Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip. **736**, 143 (2014), ISSN 01689002, URL http://linkinghub.elsevier.com/retrieve/pii/S0168900213014654.
[233] L. Breiman, Mach. Learn. **45**, 5 (2001), ISSN 08856125, URL http://link.springer.com/10.1023/A:1010933404324.
[234] M. G. Aartsen, M. Ackermann, J. Adams, et al., Astrophys. J. **846**, 136 (2017), ISSN 1538-4357, URL http://stacks.iop.org/0004-637X/846/i=2/a=136?key=crossref.a6419e182e069e85db5823ef6d966289.
[235] J. R. Hörandel, Astropart. Phys. **19**, 193 (2003), ISSN 0927-6505, URL http://www.sciencedirect.com/science/article/pii/S0927650502001986?via%3Dihub.
[236] H. Niederhausen and Y. Xu, in *Proceeding 35th Int. Cosm. Ray Conf.* (Proceedings of Science (ICRC2017) 968, Bexco, Busan, Korea, 2017), URL https://pos.sissa.it/301/968/pdf.

[237] C. Kopper, in *Proceeding 35th Int. Cosm. Ray Conf.* (Proceedings of Science (ICRC2017) 981, Bexco, Busan, Korea, 2017), URL http://icecube.wisc.edu/collaboration/authors/icrc17_icecube.

[238] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, et al., J. High Energy Phys. **2017**, 87 (2017), ISSN 1029-8479, URL http://link.springer.com/10.1007/JHEP01(2017)087.

[239] J. Auffenberg, in *34th Int. Cosm. Ray Conf.* (Proceedings of Science (ICRC2015) 1156, The Hague, The Netherlands, 2015), URL http://icecube.wisc.edu/collaboration/authors/icrc15_gen2http://pos.sissa.it/.

[240] A. Fedynitch, R. Engel, T. K. Gaisser, et al., EPJ Web Conf. **99**, 08001 (2015), ISSN 2100-014X, URL http://www.epj-conferences.org/10.1051/epjconf/20159908001.

[241] G. Cowan, K. Cranmer, E. Gross, et al., Eur. Phys. J. C **71**, 1554 (2011), ISSN 1434-6052, URL https://doi.org/10.1140/epjc/s10052-011-1554-0.

[242] M. Usner, in *Proceeding 35th Int. Cosm. Ray Conf.* (Proceedings of Science (ICRC2017) 974, Bexco, Busan, Korea, 2017), URL https://pos.sissa.it/301/974/pdf.

[243] S. Kopper, M. Meier, L. Wille, et al., in *Proceeding 35th Int. Cosm. Ray Conf.* (Proceedings of Science (ICRC2017) 1009, Bexco, Busan, Korea, 2017), URL http://icecube.wisc.edu/collaboration/authors/icrc17_icecubehttp://pos.sissa.it/.

[244] D. Seckel, Phys. Rev. Lett. **80**, 900 (1998), ISSN 0031-9007, URL https://link.aps.org/doi/10.1103/PhysRevLett.80.900.

[245] I. Alikhanov, Phys. Lett. B **756**, 247 (2016), ISSN 0370-2693, URL http://www.sciencedirect.com/science/article/pii/S0370269316001805?via%3Dihub.

[246] M. Carena, D. Choudhury, S. Lola, et al., Phys. Rev. D **58**, 095003 (1998), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.58.095003.

[247] L. A. Anchordoqui, C. A. García Canal, H. Goldberg, et al., Phys. Rev. D **74**, 125021 (2006), ISSN 1550-7998, URL http://link.aps.org/doi/10.1103/PhysRevD.74.125021.

[248] L. A. Anchordoqui, M. M. Glenz, and L. Parker, Phys. Rev. D **75**, 024011 (2007), ISSN 1550-7998, URL http://link.aps.org/doi/10.1103/PhysRevD.75.024011.

[249] J. Alvarez-Muñiz, R. A. Vázquez, and E. Zas, Phys. Rev. D **61**, 023001 (1999), ISSN 0556-2821, URL https://link.aps.org/doi/10.1103/PhysRevD.61.023001.

[250] A. Connolly, R. S. Thorne, and D. Waters, Phys. Rev. D **83**, 113009 (2011), ISSN 1550-7998, URL https://link.aps.org/doi/10.1103/PhysRevD.83.113009.

[251] S. R. Klein and A. Connolly (2013), URL http://arxiv.org/abs/1304.4891.

[252] S. Baker and R. D. Cousins, Nucl. Instruments Methods Phys. Res. **221**, 437 (1984), ISSN 01675087, URL http://linkinghub.elsevier.com/retrieve/pii/0167508784900164.

[253] S. S. Wilks, Ann. Math. Stat. **9**, 60 (1938), ISSN 0003-4851, URL http://projecteuclid.org/euclid.aoms/1177732360.

[254] S. Schoenen, Ph. D. thesis, RWTH Aachen University (2017), URL http://publications.rwth-aachen.de/record/696221.

Appendix A

Statistical Methods

The following section is a general review of binned maximum likelihood fitting and the frequentist techniques used to quantify goodness-of-fit, calculate confidence regions, and perform hypothesis testing that are used throughout this work.

A.1 Likelihood

In a binned maximum likelihood fit, the distribution of observable quantities is binned to obtain the number of observed events per bin, n_i, and the expected number of events per bin, $\mu_i(\boldsymbol{\theta})$, which depends on a set of parameters $\boldsymbol{\theta}$. Here, $\mu_i(\boldsymbol{\theta})$ can be calculated by summing the weights of events from Monte Carlo simulation. The likelihood, L, can be found from

$$-2\ln L(\boldsymbol{\theta}) = 2\sum_i \mu_i(\boldsymbol{\theta}) - n_i \ln \mu_i(\boldsymbol{\theta}) + \ln n_i! + \sum_j \left(\frac{\theta_j - \theta_j^*}{\sigma_j}\right)^2 . \tag{A.1}$$

The first term describes the Poisson probability for observing n_i counts in each bin. The second term describes prior knowledge on any of the parameters that may exist from previous measurements or theoretical considerations. The prior distribution on each parameter θ_j is usually assumed to be Gaussian with mean θ_j^* and standard deviation σ_j. Maximum likelihood estimates of all parameters, $\hat{\boldsymbol{\theta}}$, can be found by minimizing Eq. A.1. It is conventional to always work with $-2\ln L$ rather than L in order to facilitate comparisons to the more common χ^2 test statistic. In the limit of large n_i, $-2\ln L$ approaches χ^2.

A.2 Goodness-of-Fit

The test statistic used to evaluate goodness-of-fit is [252, 29]

$$-2\ln \lambda(\hat{\boldsymbol{\theta}}) = 2\sum_i \mu_i(\hat{\boldsymbol{\theta}}) - n_i + n_i \ln\left(\frac{n_i}{\mu_i(\hat{\boldsymbol{\theta}})}\right) + \sum_j \left(\frac{\theta_j - \theta_j^*}{\sigma_j}\right)^2 . \tag{A.2}$$

A p-value quantifying the goodness-of-fit of a model can be obtained from the cumulative distribution of the test statistic $-2\ln\lambda$ over an ensemble of experiments. According to Wilks' theorem [253], the distribution of $-2\ln\lambda$ asymptotically approaches a χ^2 distribution in the large sample limit, and this is often a useful approximation for calculating p-values. However, because of many bins commonly have small counts here, it will be necessary to determine the distribution of $-2\ln\lambda$ through randomly generating pseudo-experiments – repeatedly sampling synthetic data, n_i, from the best-fit model expectations, $\mu_i(\hat{\boldsymbol{\theta}})$, according to a Poisson distribution and minimizing Eq. A.1 to find the test statistic for each experiment.

A.3 Confidence Regions

To construct confidence regions for parameters, the profile likelihood method is used [241]. Parameters can be divided into physics parameters, $\boldsymbol{\theta}_p$, that are of interest to be measured, and nuisance parameters, $\boldsymbol{\theta}_n$, that are needed to describe systematic uncertainties. A profile likelihood can be calculated by minimizing Eq. A.1 over $\boldsymbol{\theta}_n$ while $\boldsymbol{\theta}_p$ is held fixed, yielding the best-fit nuisance parameters, $\hat{\boldsymbol{\theta}}_n(\boldsymbol{\theta}_p)$. A test statistic based on the profile likelihood ratio to the maximum likelihood with all parameters free can be formed,

$$-2\Delta\ln L(\boldsymbol{\theta}_p) = -2\ln\left(\frac{L(\boldsymbol{\theta}_p, \hat{\boldsymbol{\theta}}_n(\boldsymbol{\theta}_p))}{L(\hat{\boldsymbol{\theta}})}\right). \tag{A.3}$$

If the cumulative distribution of $-2\Delta\ln L$ over an ensemble of experiments, $P(-2\Delta\ln L)$, is known, one can construct a confidence region with confidence level $1-\alpha$ by finding all values of $\boldsymbol{\theta}_p$ that satisfy $P(-2\Delta\ln L(\boldsymbol{\theta}_p)) < 1-\alpha$. Usually no analytic form for the cumulative distribution $P(-2\Delta\ln L)$ is known, but it can be evaluated by randomly generating pseudo-experiments sampled from the best-fit model expectations, $\mu_i(\boldsymbol{\theta}_p, \hat{\boldsymbol{\theta}}_n(\boldsymbol{\theta}_p))$. However, according Wilks' theorem, in many cases the distribution of $-2\Delta\ln L$ can be well approximated by a χ^2 distribution with a number of degrees of freedom corresponding to the size of $\boldsymbol{\theta}_p$, and here all confidence regions quoted will make this approximation unless otherwise noted. Wilks' theorem makes it straightforward to translate the test statistic to a confidence level through

$$-2\Delta\ln L(\boldsymbol{\theta}_p) = P^{-1}_{\chi^2_n}(1-\alpha), \tag{A.4}$$

where $P^{-1}_{\chi^2_n}$ is the inverse cumulative χ^2 distribution for n degrees of freedom. For example, to obtain a one-dimensional 68% confidence interval, one can scan the test statistic to find where $-2\Delta\ln L(\boldsymbol{\theta}_p) = 1$.

A.4 Hypothesis Testing

The same profile likelihood method can also be applied to hypothesis testing. If the null hypothesis is described by parameters $\boldsymbol{\theta}_0$ and the alternate hypothesis is described by

parameters $\boldsymbol{\theta}_1$, then the maximum likelihood for each hypothesis can be found by minimizing Eq. A.1, yielding parameter estimates $\hat{\boldsymbol{\theta}}_0$ and $\hat{\boldsymbol{\theta}}_1$ and corresponding likelihood values. A test statistic based on the likelihood ratio can be formed,

$$-2\Delta \ln L(\hat{\boldsymbol{\theta}}_0, \hat{\boldsymbol{\theta}}_1) = -2 \ln \left(\frac{L(\hat{\boldsymbol{\theta}}_0)}{L(\hat{\boldsymbol{\theta}}_1)} \right). \tag{A.5}$$

A p-value for the rejection of the null hypothesis can be calculated from the cumulative distribution of $-2\Delta \ln L$. As a result of Wilks' theorem, often the χ^2 distribution is a suitable approximation when the number of degrees of freedom is taken to be the difference in the size of $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$. When this is not a good approximation, the distribution of $-2 \ln L$ may also be calculated from randomly generated pseudo-experiments sampled from the best-fit model expectations of the null hypothesis, $\mu_i(\hat{\boldsymbol{\theta}}_0)$. Again, the Wilks' theorem approximation will be used unless otherwise noted.

Appendix B

Sensitive Energy Range Construction

Whenever a scaling factor on an energy-dependent signal is measured, for example the astrophysical neutrino flux, the atmospheric neutrino/antineutrino ratio, or the neutrino CC charm production cross section, it is important to characterize the energy range in which the measurement is valid. The simplest approach to characterize the energy range is to calculate the central 90% of simulated neutrino energies that are weighted to the best-fit signal model. However, this approach neglects the fact that in certain energy ranges, there may be a large background rate or insufficient signal rate to have any sensitivity to the quantity of interest.

A generalization of this approach is to weight neutrino energies not by their contribution to the event rate, but by their contribution to a test-statistic describing a hypothesis test between two models, one containing the signal being measured and background and another containing only background. Specifically, a simulated neutrino event in bin i of the observable distribution can be assigned a weight

$$w_{ij} = -2 \ln \left(\frac{L_i^b}{L_i^{s+b}} \right) \times \frac{w_{ij}^s}{\mu_i^s} \tag{B.1}$$

where $L_i^b = \left(\mu_i^b\right)^{n_i} e^{-\mu_i^b}/n_i!$ is the Poisson likelihood of the best-fit background-only model in bin i, $L_i^{s+b} = \left(\mu_i^{s+b}\right)^{n_i} e^{-\mu_i^{s+b}}/n_i!$ is the Poisson likelihood of the best-fit signal and background model in bin i, w_{ij}^s is the weight of each event in bin i for the best-fit signal model, and μ_i^s is the total expected signal count in bin i. The weights, w_{ij}, have the property that they sum to the test statistic computed over the entire observerable distribution, $-2\Delta \ln L$, neglecting any contribution from prior information on nuisance parameters. The weights, w_{ij}, can then be used to calculate a central 90% energy range that better illustrates the neutrino energies that contribute to the statistical evidence for the signal being measured. This energy range construction was first devised for the measurement of the astrophysical flux with the six-year through-going track sample [136], and its statistical properties are investigated thoroughly in [254].

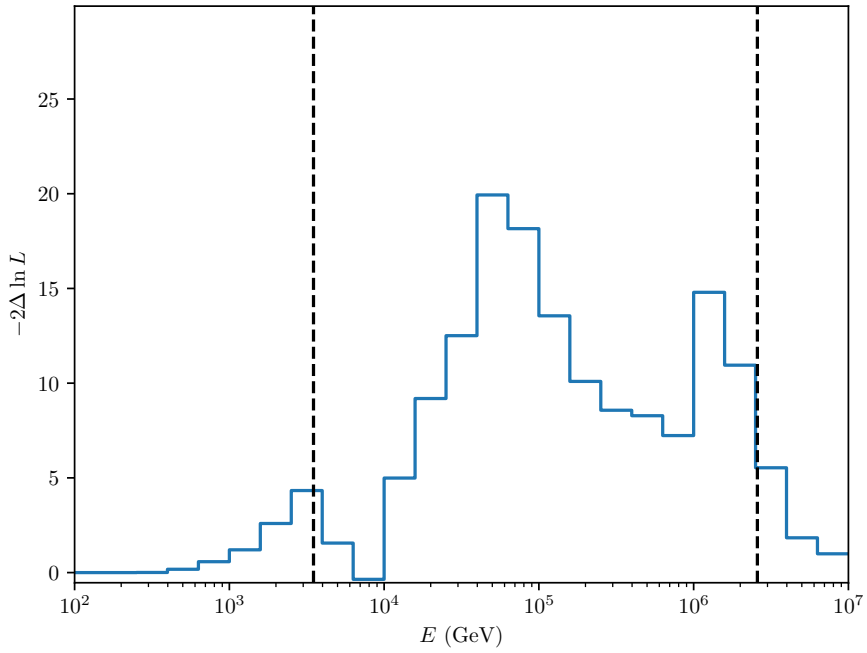

Figure B.1: The blue histogram shows the contribution to the test statistic from different bins of neutrino energy for the best-fit astrophysical neutrino power-law model. The central energy range contributing 90% of the total test statistic is shown with black dashed lines and extends from 3.5 TeV to 2.6 PeV.

B.1 Astrophysical Neutrinos

For the measurement of the single power-law astrophysical flux presented in Sec. 7.4.1, it is straight forward to also fit the data to a background-only model containing only conventional and prompt atmospheric neutrinos. This yields a test statistic of $-2\Delta \ln L = 157.3$ between the atmospheric-only model and atmospheric plus astrophysical model, corresponding to a significance of $> 12\sigma$ by Wilks' theorem. Weighting simulated events with the best-fit astrophysical flux according to Eq. B.1, one can plot the differential contribution of neutrino energies to the test-statistic as shown in Fig. B.1. The central 90% of neutrino energies contributing to the test statistic extends from 3.5 TeV to 2.6 PeV. The central 90% range of neutrino energies contributing the total event rate of astrophysical neutrinos is 3.3 TeV to 220 TeV, illustrating how the highest energy PeV neutrinos contribute very strongly to the statistical evidence for astrophysical neutrinos.

In Fig. B.1, one may be wary of the bimodal shape of the distribution, but this results from statistical fluctuations in data. Similar behavior was seen in pseudo-experiments with synthetic data generated from the best-fit distribution, and the calculated energy range for

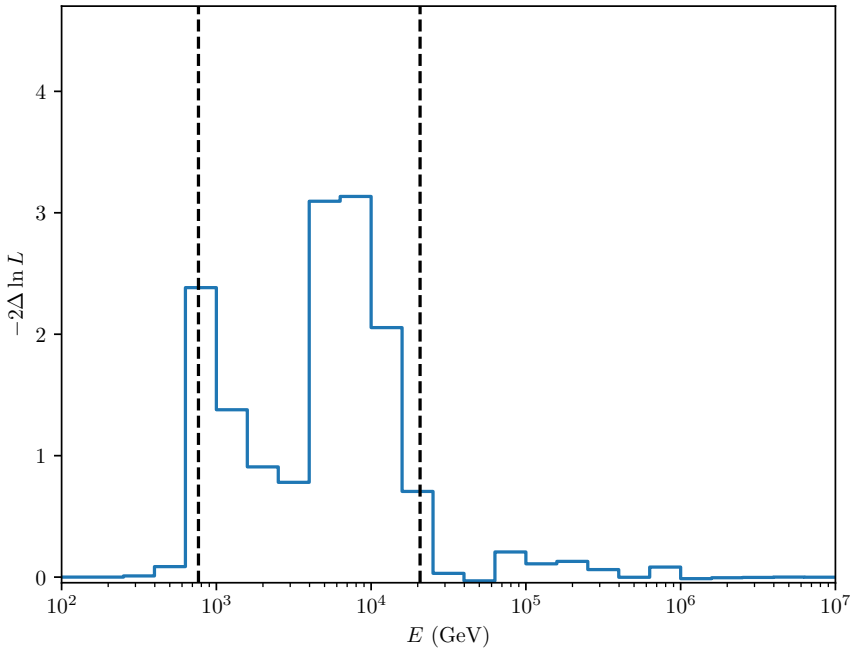

Figure B.2: The blue histogram shows the contribution to the test statistic from different bins of antineutrino energy for the best-fit conventional atmospheric flux model. The central energy range contributing 90% of the total test statistic is shown with black dashed lines and extends from 770 GeV to 21 TeV.

all pseudo-experiments did not differ substantially from the observed energy range for all cases here.

B.2 Atmospheric Neutrino to Antineutrino Ratio

For measurement of the atmospheric neutrino to antineutrino ratio, there is some ambiguity in the definition of signal and background. Here signal is chosen to be atmospheric antineutrinos rather than neutrinos since they form a smaller contribution to the atmospheric flux. The background-only model is then the best-fit neutrino-only HKKMS flux with $R_{\nu_\mu/\bar{\nu}_\mu} \to \infty$ in the parametrization of Eq. 7.9. The best-fit signal and background model has mixed neutrinos and antineutrinos with $R_{\nu_\mu/\bar{\nu}_\mu} = 0.77$. The total test statistic between neutrino-only and the best-fit mix of neutrinos and antineutrinos is $-2\Delta \ln L = 14.8$. Weighting simulated antineutrinos with the best-fit value of $R_{\nu_\mu/\bar{\nu}_\mu}$ according to Eq. B.1, a plot of the differential contribution of antineutrino energies to the test-statistic is shown in Fig. B.2. The central 90% of antineutrino energies extends from 770 GeV to 21 TeV. This is not very different for the central 90% range of neutrino energies contributing to the event

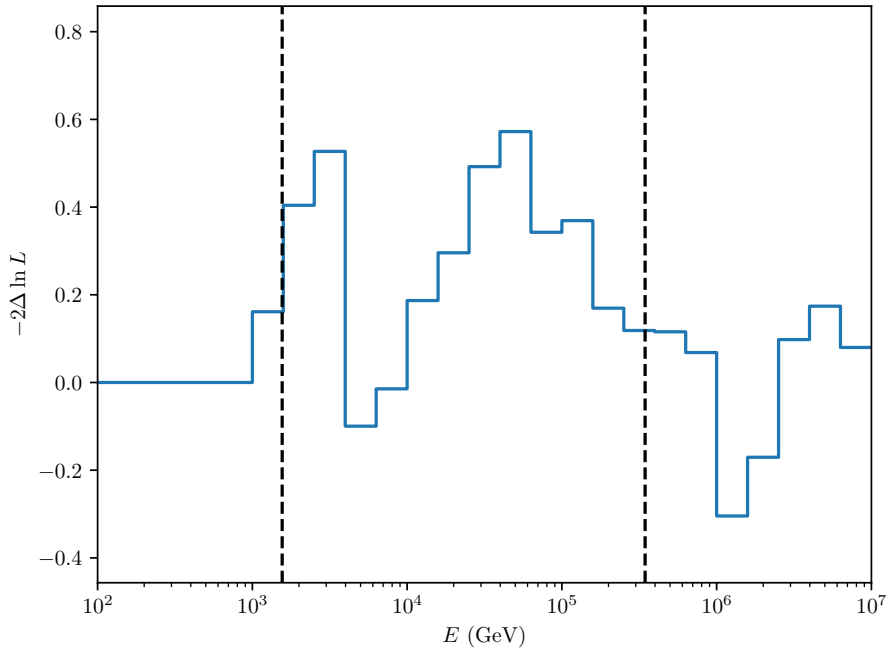

Figure B.3: The blue histogram shows the contribution to the test statistic from different bins of neutrino energy for CC charm production events. The central energy range contributing 90% of the total test statistic is shown with black dashed lines and extends from 1.6 TeV to 340 TeV.

rate, 1.1 TeV to 25 TeV. Again the distribution in Fig. B.2 is bimodal, but this is a result of statistical fluctuations in data and a similar structure is seen in randomly generated pseudo-experiments.

B.3 Neutrino Charm Production

For measurement of a scaling factor on the CC neutrino charm production cross section, the background-only model has the weight of all simulated CC events producing a charm hadron set to zero. The test statistic between the best-fit models with and without charm production is $-2\Delta \ln L = 2.8$. Weighting simulated neutrinos to the best-fit scaling of the CC charm production cross section of $R_{\mathrm{CC,\,charm}} = 0.93$ according to Eq. B.1, a plot of the differential contribution of neutrino energies to the test-statistic is shown in Fig. B.2. The central 90% of neutrino energies extends from 1.6 TeV to 340 TeV. This extends substantially higher than the central 90% of neutrino energies contributing to the charm production event rate, 1.3 TeV to 44 TeV, and it is a consequence of the fact that charm production events form an increasingly large fraction of high visible inelasticity events at higher energy as shown

in Fig. 7.18. The distribution in Fig. B.3 shows a bimodal structure and even becomes negative in places, but this is a consequence of statistical fluctuations in data. The effect is more pronounced compared to Fig. B.2 and Fig. B.3 due to the weaker statistical significance, and similar structure is seen in random pseudo-experiments.