Study of Seasonal Variation of Cosmic Muons and Neutrino–Nucleus Interaction Modeling in the NOvA Experiment

By

AMIT PAL

PHYS11201904002

National Institute of Science Education and Research, Bhubaneswar

A thesis submitted to the

Board of Studies in Physical Sciences

(as applicable)

In partial fulfillment of requirements

for the Degree of

DOCTOR OF PHILOSOPHY

of

HOMI BHABHA NATIONAL INSTITUTE

April, 2026

Homi Bhabha National Institute

Recommendations of the Viva Voce Committee

As members of the Viva Voce Committee, we certify that we have read the dissertation prepared by **Amit Pal** entitled **"Study of Seasonal Variation of Cosmic Muons and Neutrino–Nucleus Interaction Modeling in the NOvA Experiment"** and recommend that it may be accepted as fulfilling the thesis requirement for the award of Degree of Doctor of Philosophy.

Chairman - Prof. Bedangadas Mohanty	2/9/2026
Guide / Convener - Prof. Sanjay Kumar Swain	2/9/2026
Examiner - Prof. Swagata Mukherjee	Swagata Mukherjee 2/9/2026
Member 1 - Prof. Prolay Kumar Mal	Sept 2, 2026
Member 2 - Prof. Amaresh Kumar Jaiswal	02/09/2026
Member 3 - Prof. Prafulla Kumar Behera	P. K. Behera 02/09/2026

Final approval and acceptance of this thesis is contingent upon the candidate's submission of the final copies of the thesis to HBNI.

I/We hereby certify that I/we have read this thesis prepared under my/our direction and recommend that it may be accepted as fulfilling the thesis requirement.

Date: 2/9/2026

Place: NISER

Signature
Co-guide (if any)

Signature
Guide

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at Homi Bhabha National Institute (HBNI) and is deposited in the Library to be made available to borrowers under rules of the HBNI.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgement of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the Competent Authority of HBNI when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

AMIT PAL
Amit Pal
Name & Signature of the Student

DECLARATION

I hereby declare that I am the sole author of this thesis in partial fulfillment of the requirements for a postgraduate degree from National Institute of Science Education and Research (NISER). I authorize NISER to lend this thesis to other institutions or individuals for the purpose of scholarly research.

AMIT PAL
Amit Pal
Name & Signature of the Student

List of Publications arising from the thesis

Journal

1. ”Explanation of the seasonal variation of cosmic multiple muon events observed with the NOvA Near Detector”, S. Abubakar et al. (The NOvA Collaboration), Phys. Rev. D 113, 012001 (**2026**).

2. ”Precision measurement of neutrino oscillation parameters with 10 years of data from the NOvA experiment”, S. Abubakar et al. (The NOvA Collaboration), Phys. Rev. Lett. 136, 011802 (**2026**).

Conference Proceedings

1. ”Seasonal Variation of Cosmic Muon at the NOvA Near Detector”, A. Pal and S. K. Swain, In Proceedings of the XXV DAE-BRNS HEP Symposium 2022, Mohali, India; Springer Proceedings in Physics, vol 304, 2024, pp. 801-803

Conferences

1. XXV DAE-BRNS High Energy Physics Symposium 2022,
   IISER Mohali, India; Dec. 12 - 16, 2022.
   Poster presentation: “Seasonal variation of cosmic muon at the NOvA near detector”.

2. International Meeting on High Energy Physics (IMHEP-II),
   Institute of Physics (IOP), Bhubaneswar, India; Feb. 16 - 22, 2023.
   Oral presentation: “Seasonal variation of cosmic muon at the NOvA near detector”.

3. 57th Annual Users Meeting at Fermilab,
   Fermilab, USA ; Jul. 9 - 12, 2024.
   Poster presentation: “Impact of HF-CRPA CCQE model on the latest NOvA results”.

4. NuFact 2024 - The 25th International Workshop on Neutrinos from Accelerators,
   Argonne National Laboratory (ANL), USA ; Sep. 16 - 21, 2024.
   Oral presentation: “Seasonal Variation in Cosmic Muon Rate at the NOvA Experiment”.

Poster presentation: "Impact of the HF-CRPA Model on Neutrino Oscillation Parameter Measurements in NOvA".

AMIT PAL

Amit Pal

Name & Signature of the Student

DEDICATIONS

I dedicate this thesis to my মা (Mother) and বাবা (Father) for their unimaginable sacrifices and unwavering support in shaping my life.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Sanjay Kumar Swain, for his constant guidance, encouragement, and unwavering support throughout my doctoral journey. Without his mentorship, this work would not have been possible. I am especially thankful for the countless meaningful discussions and rigorous scientific arguments that have shaped my understanding of physics and strengthened this thesis.

I sincerely thank my Doctoral Committee Members: Prof. Bedangadas Mohanty, Prof. Prolay Kumar Mal, Prof. Amaresh Kumar Jaiswal, and Prof. Prafulla Kumar Behera, for their constructive feedback and valuable suggestions during each annual review. Their insights have consistently helped refine my research and broaden my perspective.

I would like to extend my special gratitude to Prof. Maury Goodman for the many meetings and discussions throughout my PhD. He patiently answered my numerous questions, especially during the cosmic muon analysis, and continuously motivated me to think more deeply about physics. I am also deeply thankful to Prof. Bryan Ramson for his mentorship specially during my stay at Fermilab. His encouragement inspired me to undertake my second major analysis on the HF-CRPA model.

I gratefully acknowledge NISER and the Department of Atomic Energy (DAE) for providing me with the opportunity and financial support to pursue my doctoral research. I am equally thankful to Fermilab for offering computing resources and the opportunity to work on-site for one year. I also sincerely thank the NOvA Collaboration for the collaborative environment and scientific discussions that made this research possible.

I would also like to acknowledge some of my NOvA collaborators: Prof. Alec Habig, Dr. Kathryn Sutton, Dr. Teresa Lackey, Dr. Prabhjot Singh, Dr. Zoya Vallari, Dr. Jon Paley, Dr. Alex Himmel, Dr. Matt Strait, and Dr. Daniel Kaplan, for their collaboration and helpful discussions.

I am indebted to my seniors: Dr. Chandiprasad Kar, Dr. Prafulla Saha, Dr. Aloke Kumar Das, and Dr. Tribeni Mishra, for their guidance, especially in my initial days, in understanding the software framework. Our many discussions on physics and life remain invaluable to me.

I would also like to thank my labmates at NISER: Lipsa, Priyanka, Kuldeep, Sneh, Rahul, Dr. Masudur Rahaman, Sandeep, Divya, Subrajit, Manas, and Vatsal, as well as the engineers from the Clean Room: Ayaskant, Balaram, Devi, Sriniwas, and Shiddarth, for the many discussions on physics, politics, and various aspects of life, and for all the fun moments.

My heartfelt thanks go to my friends at NISER: Subhadip, Debamalya, Sadaf, Soumik, Prottay, Sachin, and Sourav, for the wonderful memories and support. I am equally grateful to my friends outside NISER: Ashish, Anagha, Purnendu, Ritam, Kritiparno, Arko, Shubhadip, Riju, Bivan, Keya, and Sushovon, who have supported me in various ways throughout my journey. I also want to acknowledge some of the friends at Fermilab: Ishwar, Shivam, Sayeed, Prameet, Bhumika, Barnali, Daniel, Maitreyee, and Enrico, for their companionship and support during my time there.

In addition, many friends outside academia have contributed, directly or indirectly, to shaping this thesis. I would especially like to mention Kelly (Susmita) and Niloy, with whom I spent many enjoyable hours playing video games. Their friendship provided a welcome respite from academic pressures and helped me maintain a healthy balance throughout this journey.

Finally, no words are sufficient to express my gratitude to my parents for their unconditional love, relentless support, and sacrifices, especially during the most challenging times. I am deeply thankful to my brother and other family members for always standing beside me.

ABSTRACT

Precision measurements in long-baseline neutrino oscillation experiments require a detailed understanding of both detector-related backgrounds and neutrino–nucleus interaction physics. In experiments such as NOvA, systematic uncertainties arising from cosmic-ray induced backgrounds and nuclear modeling directly impact neutrino energy reconstruction and, consequently, the determination of oscillation parameters. Neutrino detectors such as NOvA also provide opportunities to study cosmic-ray properties and search for particles like dark matter, magnetic monopoles, etc.

This thesis presents two studies based on the data collected by the NOvA Near Detector (ND). The first study investigates the seasonal variation of cosmic ray-induced multiple-muon events observed in the ND. While the seasonal modulation of single-muon events is characterized by a summer maximum and winter minimum and is well understood, previous measurements by the MINOS experiment reported an unexpected inverted seasonal behavior for multiple-muon events, with a winter maximum, for which no conclusive explanation was established. Using nearly four years of NOvA ND data, this work confirms the inverted seasonal behavior of multiple-muon events with high statistical significance. A robust event-rate extraction technique based on Erlang probability distributions is employed to mitigate detector and data-acquisition effects. The observed multiple-muon rate is found to be strongly anti-correlated with the effective atmospheric temperature. To explain this behavior, detailed CORSIKA simulations incorporating realistic atmospheric profiles and detector geometry are performed. The results demonstrate that the interplay between seasonal atmospheric density variations, high-energy air-shower development, and the finite geometry of the detector leads to the observed phase reversal. This study provides the first comprehensive, simulation-backed explanation of the long-standing multiple-muon seasonal anomaly.

The second part of the thesis focuses on neutrino–nucleus interaction modeling, specifically charged-current quasi-elastic (CCQE) interactions, which play a central role in NOvA oscillation analyses. An advanced nuclear model based on the Hartree-Fock Continuum Random Phase Approximation (HF-CRPA) framework is studied and implemented to assess its impact on oscillation measurements. By incorporating realistic nuclear ground-state

descriptions and long-range nuclear correlations, HF-CRPA provides an improved treatment of CCQE interactions compared to traditional models. This work evaluates the effects of HF-CRPA on CCQE kinematics, reconstructed neutrino energy distributions, and oscillation sensitivities within the NOvA analysis framework, demonstrating that nuclear modeling uncertainties constitute a non-negligible systematic effect.

Together, these studies enhance the understanding of both atmospheric backgrounds and neutrino interaction systematics in NOvA, contributing to improved precision in current and future neutrino oscillation measurements.

Contents

Summary

This thesis presents two complementary studies using the NOvA Near Detector (ND), addressing key sources of systematic uncertainty in long-baseline neutrino oscillation experiments. In such experiments, precise determination of oscillation parameters relies on an accurate understanding of both detector-related backgrounds and neutrino–nucleus interaction physics. In addition to their primary role in neutrino measurements, detectors like NOvA also provide valuable opportunities to study cosmic-ray phenomena.

The first part of this thesis focuses on cosmic-ray induced muons and their seasonal variation. Atmospheric muons, produced in cosmic-ray air showers, constitute an important background in neutrino detectors while also serving as probes of high-energy interactions in the atmosphere. Using multi-year NOvA ND data, the behavior of multiple-muon events is investigated to improve the understanding of atmospheric effects influencing their rate. Simulation studies are used alongside data to interpret the observed trends and to explore the role of atmospheric conditions and air-shower development.

The second part of the thesis examines neutrino–nucleus interaction modeling, with a focus on charged-current quasi-elastic interactions that are central to neutrino energy reconstruction. Accurate modeling of these interactions is essential for reducing biases in oscillation analyses. This work studies an advanced nuclear model based on the Hartree-Fock Continuum Random Phase Approximation (HF-CRPA), incorporating more realistic descriptions of nuclear structure and dynamics. The impact of this model is evaluated within the NOvA analysis framework to understand its influence on reconstructed observables and oscillation measurements.

Together, these studies address two important and complementary aspects of neutrino experiments: atmospheric backgrounds and nuclear interaction effects. By improving the understanding of these contributions, this work supports ongoing efforts to enhance the precision and reliability of current and future neutrino oscillation measurements.

List of Figures

List of Tables

Chapter 1

Introduction

Understanding the fundamental constituents of matter and the interactions that govern them has been a central goal of modern physics. Over the past century, a large body of experimental discoveries has led to the development of the Standard Model of particle physics. The Standard Model of particle physics is the theoretical framework that describes the fundamental particles and their interactions through the electromagnetic, weak, and strong forces. In this model, matter is composed of elementary fermions, while interactions between these particles are mediated by gauge bosons associated with the underlying gauge symmetries of the theory.

The fermions of the Standard Model are divided into two categories: quarks and leptons. These particles are arranged in three generations, each containing two quarks and two leptons. The first generation consists of the up and down quarks together with the electron and the electron neutrino. The second generation contains the charm and strange quarks, the muon, and the muon neutrino. The third generation contains the top and bottom quarks, the tau lepton, and the tau neutrino. Particles in higher generations have the same quantum numbers as those in the first generation but possess larger masses.

Interactions among fermions are mediated by gauge bosons. The electromagnetic interaction is carried by the photon, the weak interaction is mediated by the charged $W^{\pm}$ bosons and the neutral Z^0 boson, and the strong interaction is mediated by gluons. In addition to these particles, the Standard Model includes the Higgs boson, which arises from the mechanism responsible for electroweak symmetry breaking and provides masses to the fundamental particles.

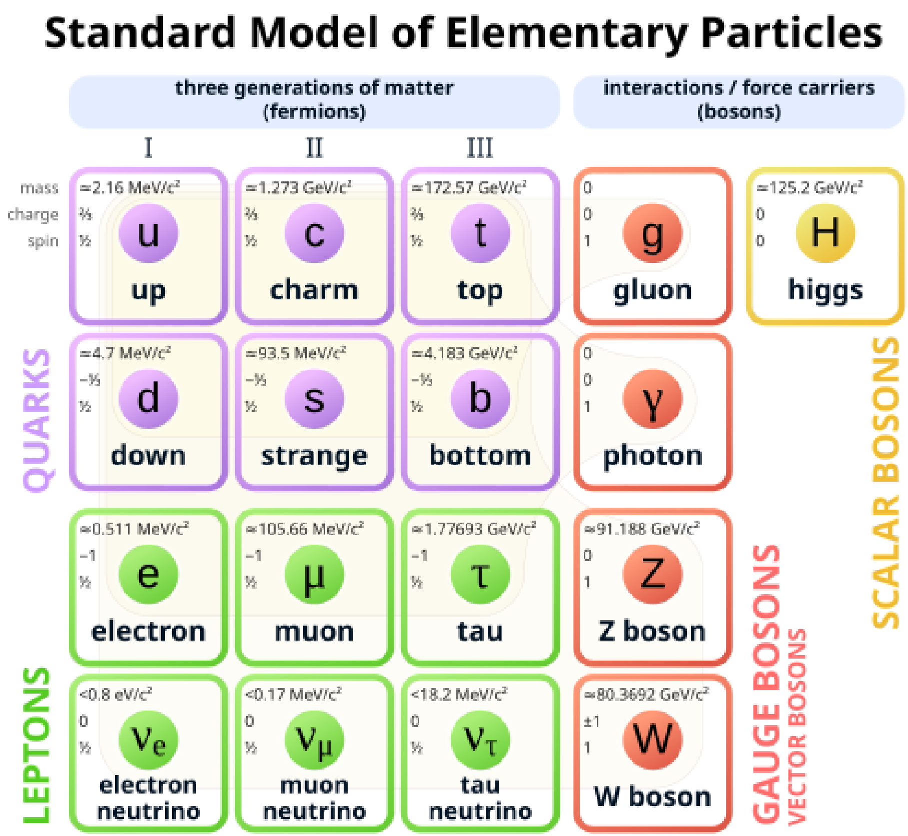

Figure 1.1: The particles of the Standard Model showing the three generations of quarks and leptons together with the gauge bosons and the Higgs boson. Image adapted from [8].

A summary of the particle content of the Standard Model, including the three generations of fermions and the force-mediating gauge bosons, is shown in Fig. 1.1. Although the Standard Model successfully explains a wide range of experimental observations, it does not account for several important phenomena, including neutrino masses, dark matter, and the matter–antimatter asymmetry of the universe.

Among the particles described by the Standard Model, muons and neutrinos play particularly important roles in experimental studies of fundamental physics. Muons are charged leptons with properties similar to electrons but with a mass approximately two hundred times larger. Because of their relatively long lifetime and high penetrating power, muons can traverse large amounts of material and are readily detected in particle detectors. Neutrinos, in contrast, are electrically neutral particles that interact only through the weak interaction and gravity. Their extremely small interaction cross sections make them difficult to detect, but they also allow neutrinos to travel vast distances through matter without significant attenuation. As a result, neutrinos provide unique probes of astrophysical processes and fundamental particle interactions.

1.1 Seasonal variation of cosmic muon

One of the primary sources of muons observed at the Earth's surface and underground detectors is the interaction of high-energy cosmic rays with the Earth's atmosphere. Cosmic rays consist mainly of protons and atomic nuclei that originate from astrophysical sources such as supernova remnants and other energetic environments in the universe. When these particles collide with atmospheric nuclei, they produce cascades of secondary particles known as extensive air showers. Among the particles produced in these interactions are charged mesons, primarily pions and kaons, which subsequently decay into muons and neutrinos. These muons can travel through large amounts of rock and reach detectors located deep underground.

The flux of atmospheric muons observed in underground detectors depends not only on the energy spectrum of the primary cosmic rays but also on the properties of the atmosphere in which the air showers develop. In particular, the production and decay of mesons in the upper atmosphere are sensitive to atmospheric temperature and density profiles. Seasonal variations in atmospheric conditions therefore lead to measurable variations in the rate of cosmic-ray induced muons detected at the Earth's surface or underground. Studying these variations provides valuable insight into the physics of cosmic-ray interactions, the development of air showers, and the propagation of secondary particles through the atmosphere.

Large detectors designed for neutrino experiments are also capable of detecting large numbers of cosmic-ray muons. Their large volumes and fine-grained tracking capabilities allow detailed reconstruction of particle trajectories, making them well suited for studying both neutrino interactions and cosmic-ray induced phenomena. The NOvA (NuMI Off-axis ν_e Appearance) experiment is one such detector system. NOvA is a long-baseline neutrino oscillation experiment that uses an intense neutrino beam produced at Fermilab and two functionally identical detectors located at different distances from the source. The Near Detector measures the properties of the neutrino beam close to its production point, while the Far Detector measures the neutrino spectrum after the beam has traveled approximately 810 km. By comparing observations in the two detectors, NOvA performs precision measurements of neutrino oscillations.

The first part of this thesis focuses on the study of atmospheric muons detected in the NOvA Near Detector. In particular, the seasonal variation of cosmic-ray induced muon rates is investigated using several years of detector data. While the seasonal modulation of single muons has been observed in many underground experiments, the behavior of events containing multiple muons originating from high-energy cosmic-ray air showers exhibits a different seasonal pattern that is not yet fully understood. This thesis presents a detailed study of both single-muon and multiple-muon events observed in the NOvA Near Detector and investigates the physical mechanisms responsible for the observed seasonal variations using comparisons with air-shower simulations.

1.2 Neutrino–nucleus interaction modeling

In addition to studies of cosmic-ray muons, the NOvA experiment is designed to measure neutrino oscillations with high precision. Neutrino oscillations arise from the mixing between neutrino flavor eigenstates and mass eigenstates, causing neutrinos to change flavor as they propagate. Measurements of neutrino oscillation parameters rely on accurate reconstruction of the neutrino energy spectrum. In accelerator-based experiments such as NOvA, neutrinos are detected through their interactions with nuclei in the detector material. However, the interpretation of these interactions is complicated by nuclear effects such as nucleon motion, binding energy, and correlations between nucleons within the nucleus. These effects influence the observed kinematics of the final-state particles and introduce uncertainties in the reconstructed neutrino energy.

Charged-current quasi-elastic scattering is one of the dominant neutrino interaction processes in the energy range relevant for long-baseline neutrino oscillation experiments. Although the fundamental interaction between a neutrino and a nucleon is relatively simple, the nuclear environment in which it occurs introduces significant theoretical challenges. Various models have been developed to describe these nuclear effects, including approaches that incorporate detailed nuclear dynamics such as the Hartree–Fock continuum random phase approximation (HF–CRPA). Understanding the impact of these interaction models on oscillation analyses is essential for reducing systematic uncertainties in neutrino oscillation measurements.

The second part of this thesis investigates the impact of neutrino interaction modeling on oscillation measurements in the NOvA experiment. In particular, the study examines how alternative descriptions of quasi-elastic neutrino scattering based on the HF–CRPA framework affect the predicted event spectra and the extraction of oscillation parameters. By propagating modified interaction predictions through the oscillation analysis framework, the study evaluates the sensitivity of oscillation measurements to uncertainties in neutrino–nucleus interaction modeling.

Together, these studies demonstrate the versatility of the NOvA detectors in addressing a broad range of physics topics. The analysis of cosmic-ray muons provides insight into atmospheric particle production and detector response, while the investigation of neutrino interaction models contributes to reducing systematic uncertainties in precision neutrino oscillation measurements.

The remainder of this thesis is organized as follows. Chapter 2 describes the NOvA experiment, including the NuMI neutrino beamline and the design and operation of the Near and Far detectors. Chapter 3 introduces the physics of cosmic rays and atmospheric muon production. Chapter 4 presents the analysis of seasonal variation of cosmic-ray muons observed in the NOvA Near Detector. Chapter 5 reviews the theoretical framework of neutrino oscillations and neutrino interactions with nuclei. Chapter 6 discusses models of quasi-elastic neutrino scattering and presents a study of the impact of the HF–CRPA interaction model on oscillation measurements. Finally, Chapter 7 summarizes the results and conclusions of this thesis.

Chapter 2

The NOvA Experiment

The NuMI Off-Axis ν_e Appearance (NOvA) experiment is a long-baseline neutrino project based at Fermilab in Illinois, USA. Its primary objective is to investigate neutrino oscillation parameters and determine the neutrino mass ordering by examining both the appearance of electron neutrinos and the disappearance of muon neutrinos in a beam initially composed of muon neutrinos. The neutrino beam used in the experiment originates from Fermilab's Neutrinos at the Main Injector (NuMI) facility. NOvA utilizes a two-detector setup: the Near Detector (ND), situated on the Fermilab site, and the Far Detector (FD), located near Ash River, Minnesota. The FD is positioned approximately 810 kilometers from the NuMI beam source.

Apart from studying neutrino oscillations, the NOvA experiment also explores a wide range of physics topics, including:

- Neutrino–nucleus scattering cross sections
- Searches for sterile neutrinos
- Investigations of non-standard interactions
- Searches for exotic and astrophysical phenomena, such as:
  - Seasonal variations in cosmic ray muon rates
  - Slow and fast magnetic monopoles
  - Dark matter signatures
  - Neutrino magnetic moments

– Multi-messenger astrophysics

2.1 The NuMI Beam

The NOvA experiment is an accelerator-based neutrino experiment that receives its neutrino beam from the *Neutrinos at the Main Injector* (NuMI) facility at Fermi National Accelerator Laboratory (Fermilab). The beam is produced by accelerating protons and directing them onto a graphite target, where their collisions generate a stream of secondary particles that subsequently decay into neutrinos. Several other experiments, such as MINOS, MINERvA, MicroBooNE, ICARUS, and SBND, have also utilized, or previously utilized, the NuMI beam.

The NuMI beam production process begins with the generation of a $35\,\text{keV}$ H^- ion beam. This beam is produced in pulses of $100\,\mu\text{s}$ duration at a repetition rate of $15\,\text{Hz}$. The beam first passes through solenoids for transverse focusing and is then chopped using an Einzel lens. At this initial low-energy stage, the Coulomb repulsion between particles is significant; therefore, the kinetic energy is rapidly increased to minimize this effect. This is achieved using a Radio-Frequency Quadrupole (RFQ), which accelerates the beam from $35\,\text{keV}$ to $750\,\text{keV}$.

The $750\,\text{keV}$ H^- beam is then injected into the linear accelerator (Linac), which increases the energy to $400\,\text{MeV}$ in two stages: the Low-Energy Linac and the High-Energy Linac. The Linac has a total length of approximately $150\,\text{m}$. At its exit, the beam is again chopped using an Einzel lens to select specific portions of the beam pulse. It then passes through a stripping foil that removes the electrons, producing a proton beam.

This proton beam is injected into the Booster synchrotron, where it is further accelerated to $8\,\text{GeV}$ at the same $15\,\text{Hz}$ repetition rate. The Booster has a circumference of $474.2\,\text{m}$. From the Booster, the beam is transferred to the Recycler Ring, located in the same enclosure as the Main Injector (MI). The MI sits on the lower level of the enclosure, while the Recycler is positioned 47 inch above it.

In the Recycler, an RF manipulation process called *slip-stacking* is employed. This involves injecting pairs of proton batches into the Recycler and merging each pair to form double-intensity batches at a fixed energy of $8\,\mathrm{GeV}$. Up to twelve batches can be injected, merged into six high-intensity batches, and then transferred to the Main Injector.

The Main Injector, with a circumference of $3.3\,\mathrm{km}$, accelerates the beam from $8\,\mathrm{GeV}$ to $120\,\mathrm{GeV}$ using a combination of dipole magnets, quadrupole magnets, and RF cavities. The acceleration follows a cycle known as a *ramp*. Initially, the magnetic field is held constant at the $8\,\mathrm{GeV}$ setting during the injection process. Once the injection is complete, the magnetic field gradually increases, initially at a slow rate and then more rapidly, as the beam is accelerated. When the beam reaches $120\,\mathrm{GeV}$, the magnetic field is held constant while the beam is extracted to the NuMI target. Following extraction, the field is reduced back to the $8\,\mathrm{GeV}$ level in preparation for the next cycle. A complete ramp, such as one delivering a $120\,\mathrm{GeV}$ beam to NuMI, takes approximately $1.33\,\mathrm{s}$.

2.2 Neutrino from proton beam

Once the proton beam is accelerated to an energy of $120\,\mathrm{GeV}$, it is directed onto a graphite target at the NuMI beam facility. The interaction of protons with the target produces a cascade of secondary hadrons, which subsequently decay into charged leptons and neutrinos.

Maximizing the neutrino yield requires that a large fraction of the incident proton beam interacts within a compact target volume, thereby reducing additional interactions of the secondary mesons. This constraint motivates the use of a narrow transverse target geometry. The NuMI graphite target is therefore constructed from 47 thin fins, each $20\,\mathrm{mm}$ long in the beam direction and separated by a gap of $0.3\,\mathrm{mm}$, giving a total target length of $95\,\mathrm{cm}$.

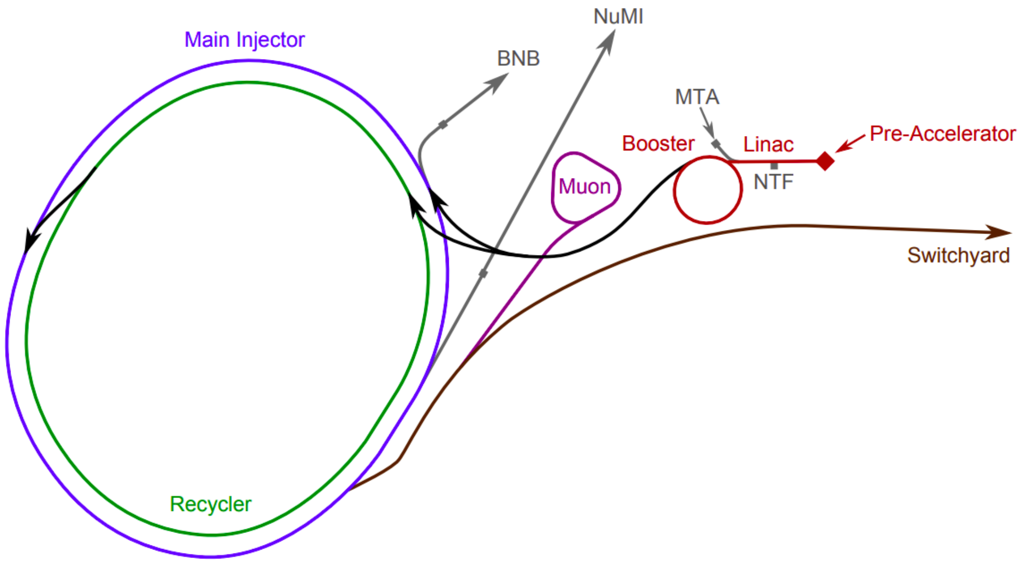

Figure 2.1: Overview of Fermilab Accelerator complex [9]

The hadrons produced in these interactions are subsequently directed using two magnetic focusing horns, which act as pulsed toroidal lenses. These horns increase the flux of mesons in the desired momentum range and allow tuning of the neutrino spectrum. Both the separation between the target and Horn 1, as well as the distance between Horn 1 and Horn 2, can be adjusted to optimize the focusing. Particles emerging close to the beam axis pass essentially unaffected, while those properly focused by Horn 1 generally continue unaltered through Horn 2. However, mesons that are either over- or under-focused by the first horn can be corrected by the second, thereby improving the overall focusing efficiency by nearly 50%. By reversing the polarity of the horn current, the system can be switched between neutrino-dominated and antineutrino-enhanced modes. In the Forward Horn Current (FHC) configuration, positively charged mesons are focused, yielding a beam rich in neutrinos, whereas in the Reverse Horn Current (RHC) mode, negatively charged mesons are selected, producing an antineutrino-enhanced flux. The horn setup is illustrated in Fig. 2.2.

After passing the horns, the hadrons enter the decay region. The NuMI decay pipe

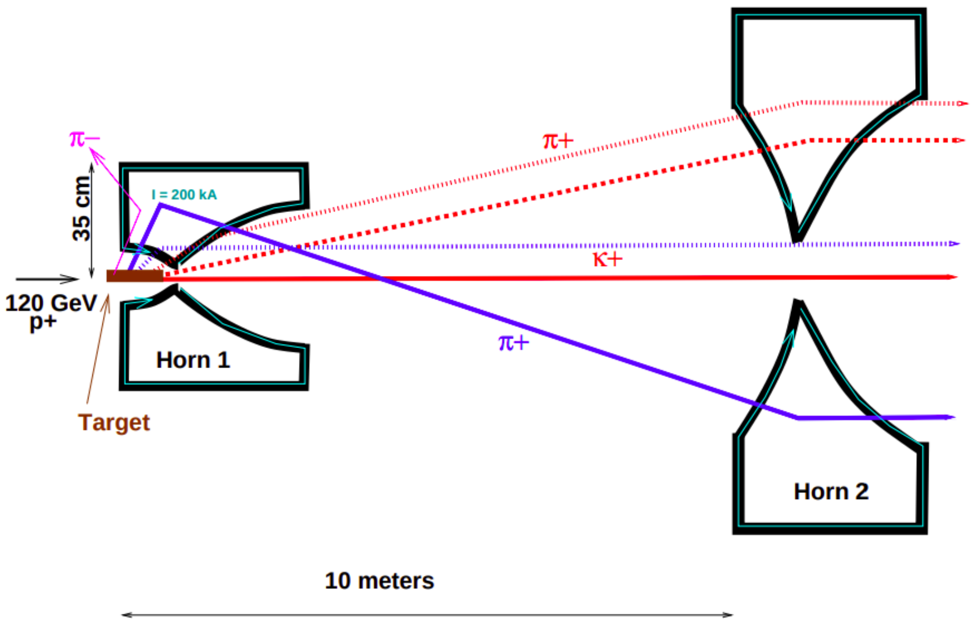

Figure 2.2: Schematic of the magnetic horn focusing system. Mesons are focused to control the energy spectrum and composition of the neutrino beam (FHC or RHC). While the figure shows a $10\,\mathrm{m}$ separation between the two horns, this distance is adjustable. The figure is taken from [10].

begins $46\,\mathrm{m}$ downstream of the target and extends for $675\,\mathrm{m}$ with a diameter of $2\,\mathrm{m}$, enclosed within a larger evacuated tunnel (see Fig. 2.4). Its role is to provide an environment where mesons can decay into tertiary hadrons, charged leptons, and neutrinos with minimal re-interactions. To reduce scattering and absorption, the decay pipe is filled with helium, chosen for its long interaction length.

Horn current	Decay mode	Branching ratio
FHC	$\pi^+ \to \mu^+ + \nu_\mu$	0.9999
	$K^+ \to \mu^+ + \nu_\mu$	0.6356
RHC	$\pi^- \to \mu^- + \bar{\nu}_\mu$	0.9999
	$K^- \to \mu^- + \bar{\nu}_\mu$	0.6356

Table 2.1: Major meson decay modes producing neutrinos and their branching ratios [?].

Figure 2.3: The NuMI Decay pipe tunnel before the decay pipe was installed. The wall is for the concrete shielding. The figure is taken from [10].

The NuMI absorber, located just downstream of the decay pipe, is a large composite structure made of aluminum, steel, and concrete designed to stop the remaining beam. Roughly one-sixth of the original beam power reaches the end of the decay pipe as a mixture of particles. The dominant component, about 80%, consists of primary protons that have not undergone interactions. The rest is mainly composed of undecayed mesons, secondary protons, and a small fraction of neutrons, electrons, and photons. These particles lose energy primarily through ionization as they traverse the absorber, whereas muons and neutrinos contribute little to no energy deposition.

The muons remaining in the NuMI beam after the decay pipe and absorber are ranged out in the muon shield, which consists of approximately $240\,\mathrm{m}$ of solid dolomite rock. Nevertheless, a small fraction of muons pass through the shield and subsequently decay via

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

The resulting electron neutrinos constitute an irreducible background to the ν_e appearance analysis. Additional sources of electron neutrinos arise from kaon decays, such as

$$K_L^0 \to \pi^- + e^+ + \nu_e, \qquad K^+ \to \pi^0 + e^+ + \nu_e.$$

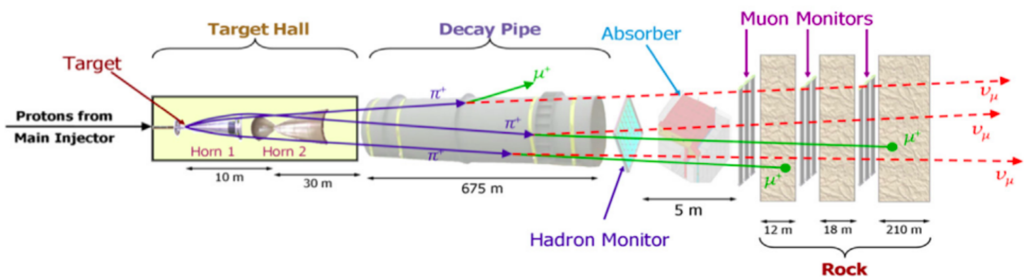

Figure 2.4: Schematic of the NuMI beamline, illustrating the progression from the primary proton beam to the production of neutrinos [10].

2.2.1 Off-axis design

Both NOvA detectors are positioned 14.6 mrad off-axis with respect to the NuMI beam direction. This off-axis angle is optimized to produce a narrow-band neutrino beam centered around the medium-energy region. The optimization can be understood from the kinematics of two-body pion decay.

For the decay $\pi^+ \to \mu^+ + \nu_\mu$, the conservation of four-momentum gives

$$\mathbf{p}_\pi = \mathbf{p}_\mu + \mathbf{p}_\nu \quad \Rightarrow \quad \mathbf{p}_\mu = \mathbf{p}_\pi - \mathbf{p}_\nu, \tag{2.1}$$

where $\mathbf{p}_\pi$, $\mathbf{p}_\mu$, and $\mathbf{p}_\nu$ denote the four-momenta of the pion, muon, and neutrino, respectively. Squaring both sides yields

$$\mathbf{p}_\mu^2 = \mathbf{p}_\pi^2 + \mathbf{p}_\nu^2 - 2\mathbf{p}_\pi.\mathbf{p}_\nu. \tag{2.2}$$

Assuming the neutrino mass $m_\nu \simeq 0$, we obtain

$$m_\mu^2 = m_\pi^2 - 2E_\pi E_\nu + 2\vec{p_\pi} \cdot \vec{p_\nu}, \tag{2.3}$$

where m_π and m_μ are the pion and muon masses, respectively. Rearranging,

$$2E_\pi E_\nu = m_\pi^2 - m_\mu^2 + 2|\vec{p_\pi}||\vec{p_\nu}|\cos\theta, \tag{2.4}$$

with θ being the angle between the pion and neutrino momenta. Since $E_\nu = |\vec{p_\nu}|$, this can be written as

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2(E_\pi - p_\pi \cos\theta)} = \frac{m_\pi^2 - m_\mu^2}{2E_\pi(1 - \beta\cos\theta)}. \tag{2.5}$$

For small angles and in the highly relativistic limit,

$$1 - \beta\cos\theta \approx 1 - \beta\left(1 - \frac{\theta^2}{2}\right) = \frac{1}{2\gamma^2} + \frac{\beta\theta^2}{2} \approx \frac{1}{2\gamma^2} + \frac{\theta^2}{2} \tag{2.6}$$

which gives

$$E_\nu = \frac{m_\pi^2}{2E_\pi\left(\frac{1}{2\gamma^2} + \frac{\theta^2}{2}\right)}\left(1 - \frac{m_\mu^2}{m_\pi^2}\right) = \frac{E_\pi}{2\gamma^2\left(\frac{1}{2\gamma^2} + \frac{\theta^2}{2}\right)}\left(1 - \frac{m_\mu^2}{m_\pi^2}\right). \tag{2.7}$$

Finally, the neutrino energy as a function of pion energy and decay angle can be expressed as

$$E_\nu = E_\pi\left(1 - \frac{m_\mu^2}{m_\pi^2}\right)\frac{1}{1 + \theta^2\gamma^2}. \tag{2.8}$$

Substituting the known pion and muon masses, Eq. 2.8 becomes

$$E_\nu = \frac{0.43\,E_\pi}{1 + \gamma^2\theta^2}. \tag{2.9}$$

Eq. 2.9 is illustrated in Fig. 2.5 for several off-axis angles, including the NOvA configuration at $14.6\,$mrad. At this angle, the medium-energy pion decays yield a narrow-band neutrino energy spectrum, which is ideal for oscillation studies. The narrow band of reconstructed neutrino energy peaking around 2 GeV for 14.6 mrad off-axis is very clear from the simulation of beam as shown in Fig. 2.6.

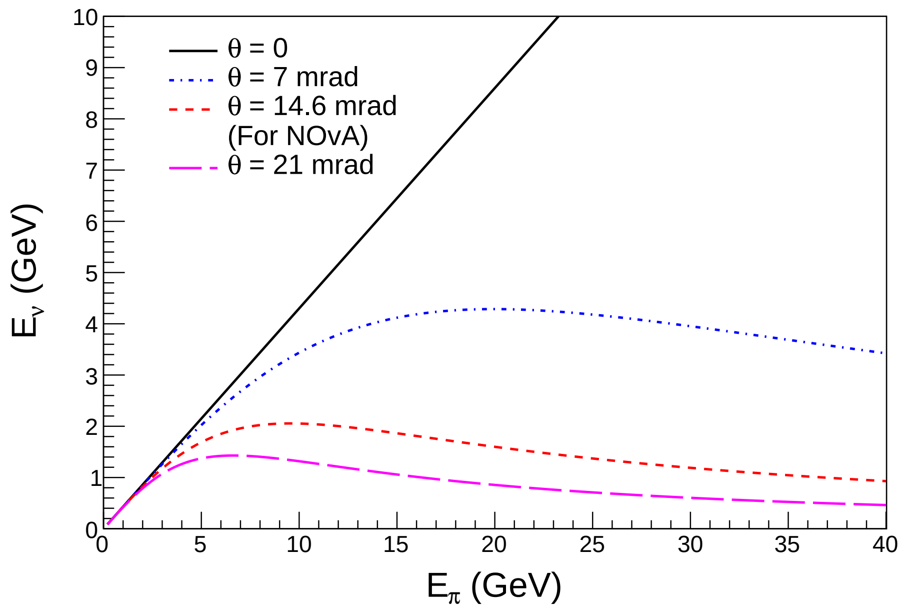

Figure 2.5: Neutrino energy as a function of pion energy for several values of the angle θ between the pion and neutrino directions. The NOvA off-axis angle of $14.6\,\mathrm{mrad}$ is highlighted.

2.3 The NOvA detectors

In order to study neutrino properties and oscillations, two key requirements must be satisfied: (i) the ability to correctly identify the charged leptons (muons and electrons) produced in neutrino interactions and accurately measure their energies for neutrino energy reconstruction, and (ii) the ability to efficiently suppress backgrounds, particularly those from neutral current (NC) interactions.

Both of these goals can be achieved by constructing detectors from low-Z materials. To measure the energies of the outgoing particles, a finely segmented calorimeter is required. In the NOvA detectors, this segmentation is realized using polyvinyl chloride (PVC) cells, which are primarily composed of carbon (a low-Z material). Each PVC cell is filled with liquid scintillator and instrumented with a wavelength-shifting fiber that guides the light

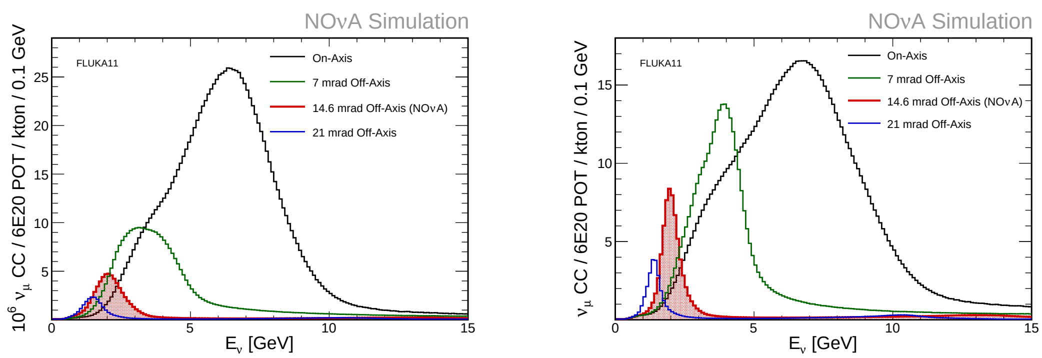

Figure 2.6: Reconstructed neutrino energy flux at four different angles including NOvA's off-axis angle (14.6 mrad) in the ND (left) and FD (Right). Figures are taken from [11]

to photodetectors. Cells are placed side by side to form planes, and alternating planes are oriented orthogonally to provide three-dimensional event reconstruction.

The cell size is chosen to be sufficiently small (high granularity) in order to distinguish the topological features of different particle types. Showering particles such as electrons and photons can then be separated from non-showering particles such as muons and pions. The use of low-Z material is particularly advantageous in distinguishing between electrons and photons, which is essential for reducing NC backgrounds in the ν_e appearance analysis.

In NC interactions, neutral pions (π^0) are often produced. While the π^0 itself does not leave a direct signal in the detector, it decays almost immediately into two photons, which give rise to electromagnetic (EM) showers. In many cases, the two photon showers overlap, mimicking the signature of a single electron. However, a key difference exists: in a true ν_e charged-current (CC) interaction, the electron track begins directly at the primary interaction vertex, whereas for a π^0, the photon-induced shower originates some distance away from the vertex. The ability to resolve this displacement requires both fine segmentation and the use of low-Z materials, which provide a longer radiation length and hence better spatial separation of the shower development.

The only differences are their size and their placement relative to the Earth's surface.

The ND has a mass of about 300 t and is located approximately 100 m underground, whereas the FD has a mass of 14 kt and is located on the surface.

2.3.1 Detector design

The two detectors, the ND and the FD, are functionally identical, both using polyvinyl chloride (PVC) cells as the basic unit of the detector, which is called a cell. Illustration of the NOvA detectors and their basic detection principle is shown in 2.7. Each cell is filled with liquid scintillator and equipped with a looped wavelength-shifting (WLS) fibre. Charged particles traversing the scintillator produce light, which is captured by the fibre and transported to the readout system after undergoing a wavelength shift.

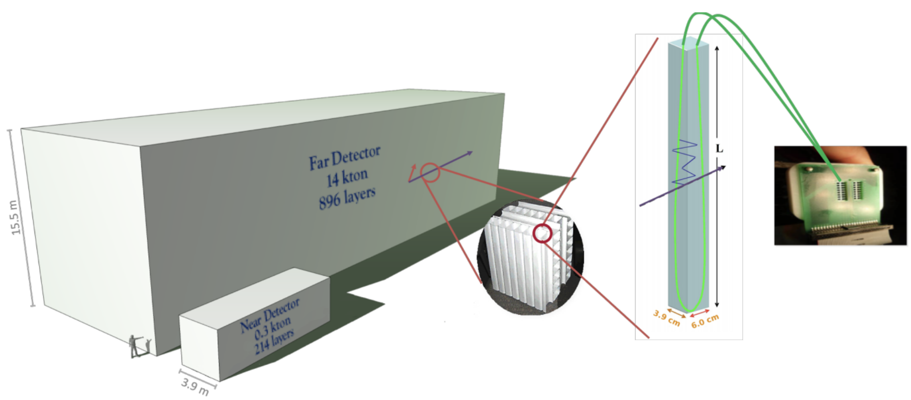

Figure 2.7: The NOvA ND and FD are shown side by side, along with the scale comparison to an average person. The detection principle is also illustrated, showing the PVC cells filled with liquid scintillator and the collection of scintillation light by the Avalanche Photodiode (APD). The picture is taken from [12].

Cells are produced through an extrusion process in groups of sixteen. Each cell has a cross-sectional size of approximately 3.8×5.9 cm. The length of the extrusions differs significantly between the two detectors: the cells in the FD are approximately 15.6 m long, whereas those in the ND are about 3.9 m long. Because of this difference in length, light traveling through the fibre experiences greater attenuation in the FD. To improve light collection efficiency, the PVC material used in the cells is designed to be highly reflective. On average, scintillation photons undergo several reflections before reaching the fibre, meaning that even small improvements in reflectivity significantly enhance the overall light yield.

Two sets of 16-cell extrusions are bonded together to form a 32-cell unit referred to as a module. These modules are arranged into planes that are oriented perpendicular to the direction of the neutrino beam. Successive planes are rotated by $90°$ relative to each other, creating an alternating structure.

Approximately 35% of the detector mass is contributed by the PVC structure, while the remaining 65% consists of the active detection medium: liquid scintillator contained inside the cells. The detailed chemical composition of the scintillator is summarized in Tab. 2.2. Mineral oil serves as the primary solvent, while pseudocumene acts as the main scintillating component. When excited by charged particles, pseudocumene produces ultraviolet light with wavelengths typically between 360 and 390 nm. To better match the optical properties of the WLS fibres, additional wavelength-shifting compounds are included in the mixture to convert this ultraviolet light into the 400–450 nm range.

Componnent	Purpose	Mass Fraction
Mineral Oil	Solvent	95.8%
Pseudocumene	Scintillant	4.1%
2,5-Diphenyloxazole (PPO)	Wavelength shifter	0.091%
bis-MSB	Wavelength Shifter	0.0013%
Stadis-425	Antistatic agent	0.0003%
Tocopherol	Antioxidant	0.0010%

Table 2.2: The chemical composition of the NOvA liquid scintillator [27].

Other additives are also incorporated into the scintillator mixture. A small quantity of tocopherol functions as an antioxidant, helping to preserve the light yield by preventing oxygen contamination from degrading the scintillator performance. Additionally, an anti-static agent is included in small amounts to prevent the accumulation of electrical charge during the filling of the PVC cells. This additive has negligible impact on the scintillation light production. In total, roughly three million gallons of scintillator are required to fill the FD, while the ND requires about thirty thousand gallons.

Each detector cell contains a looped wavelength-shifting fibre responsible for collecting the scintillation light and transporting it to the photodetector. The looped configuration significantly improves the light collection efficiency compared to a single fibre with a reflective end, since photons can propagate toward the readout from either direction along the fibre. The fibre itself consists of a polystyrene core surrounded by two cladding layers with progressively lower refractive indices. This structure enhances total internal reflection and allows efficient transmission of light over long distances. The overall diameter of the fibre is approximately 0.7 mm, with the cladding layers accounting for about 12% of the total thickness.

The scintillation light captured by the fibre is ultimately detected using avalanche photodiodes (APDs). These devices exhibit higher quantum efficiency at longer wavelengths. To better match the spectral response of the APDs, the fibre core contains a fluorescent dye that shifts the absorbed blue light (400–450 nm) to green wavelengths in the range of approximately 490–550 nm. Because the emission and absorption spectra of the dye partially overlap, light with wavelengths below roughly 500 nm is strongly attenuated as it propagates through the fibre.

NOvA employs avalanche photodiodes (APDs), manufactured by Hamamatsu, as the photodetectors for collecting the scintillation light transported by the wavelength-shifting fibres. These devices were selected primarily because of their high quantum efficiency in the wavelength range of approximately 520–550 nm, which corresponds well to the green light emitted from the fibre ends. The APDs are arranged in arrays consisting of 32 pixels, providing a one-to-one correspondence with the 32 detector cells that form a module. Each pixel receives light from both ends of the looped fibre contained within a single cell.

APDs convert incoming photons into an electrical signal through the photoelectric effect. When a photon is absorbed in the semiconductor material, it produces an electron–hole pair. An applied electric field separates these charge carriers and accelerates them through the device. If the electric field is sufficiently strong, the accelerated carriers can generate additional electron–hole pairs through impact ionization. This process repeats, producing an avalanche multiplication that amplifies the initial photocurrent and enables the detection of very small amounts of light.

Thermal generation of charge carriers can introduce noise in the detector signal. To minimize this effect, the APDs are operated at a reduced temperature of approximately $-15°$C using thermoelectric cooling systems. Maintaining this low operating temperature is particularly important for the FD, where scintillation light produced near the far end of a long cell is significantly attenuated before reaching the photodetector.

The electrical signals produced by the APDs are processed by a Front End Board (FEB), which houses several key electronic components including an Application Specific Integrated Circuit (ASIC), an Analogue-to-Digital Converter (ADC), and a Field Programmable Gate Array (FPGA). The FEB forms the first stage of the data acquisition system. Initially, the ASIC performs signal shaping, producing pulses with characteristic rise and fall times. These shaped analogue signals are then digitized by the ADC. The digitized data are subsequently processed by the FPGA, which extracts the timing and amplitude information associated with each signal. Only signals exceeding a predefined threshold are identified as valid “hits” and forwarded for further processing.

Because each module contains 32 cells, a single APD array is used to read out the light from all fibres within that module. At the top of each module, the fibres are guided through dedicated channels known as fibre raceways and routed toward an optical connector. This connector contains 32 precisely aligned openings that couple the fibre ends directly to the corresponding pixels of the APD array. The fibre routing system is enclosed by a manifold cover, while the optical connector itself is positioned within a structure known as the snout. These components are fabricated from black injection-molded plastic in order to suppress stray light reflections and reduce optical cross-talk between adjacent fibres.

Each APD is paired with a dedicated FEB, and the APD–FEB assemblies are housed within an electronics enclosure located on the detector. In addition to signal amplification and digitization, the FEB is responsible for formatting and time-stamping the data before transmitting it to a Data Concentrator Module (DCM). The FEB also hosts the circuitry required to regulate and maintain the operating temperature of the APD.

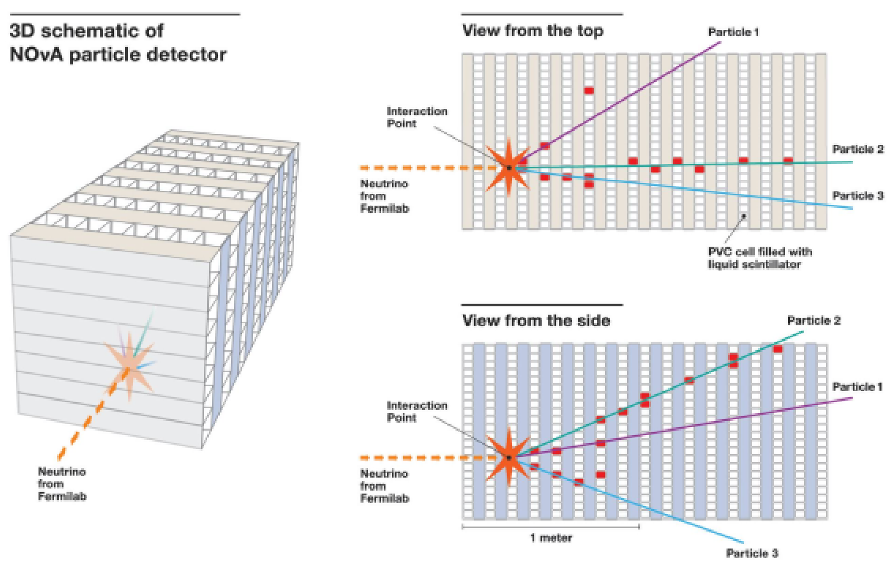

Figure 2.8: Schematic illustration of the particle detection mechanism in the NOvA detector and the reconstruction of a 3D track from two 2D tracks originating from two perpendicular detector planes. The picture is taken from [13].

At a larger scale, modules are arranged side-by-side to form detector planes. Each plane therefore contains 32 parallel cells. Successive planes are rotated by 90° relative to one another, allowing the detector to reconstruct particle trajectories in three dimensions. The full detector is built from structural units known as blocks, where each block consists of 32 planes assembled together in alternating vertical and horizontal orientations, as illustrated in Fig. 2.8.

In the data acquisition chain, every APD–FEB pair reads out 32 cells. A single Data Concentrator Module collects data from 64 FEBs, corresponding to a total of 2,048 detector cells. The DCM aggregates the incoming data streams and forwards them to buffer nodes, where the information is organized into subruns and written to disk for storage and further analysis.

2.3.2 Near Detector

The NOvA ND is located at Fermilab, approximately 1 km downstream from the NuMI neutrino production target. It is situated roughly 100 m underground, close to the location previously occupied by the MINOS ND. The ND became fully operational in August 2014 and serves as a crucial component of the experiment for measuring the unoscillated neutrino beam spectrum and constraining systematic uncertainties in oscillation analyses.

The detector is divided into two main sections: the *Active Region* and the *Muon Catcher*. The Active Region forms the upstream portion of the detector and consists entirely of scintillator planes, which act as the primary interaction and tracking volume for neutrino events. Downstream of this section lies the Muon Catcher, which is designed to enhance the detector's ability to contain and measure high-energy muons produced in neutrino interactions. In this region, planes of scintillator are interleaved with steel absorber plates approximately 10 cm thick. These steel layers increase the stopping power of the detector and improve the probability that long muon tracks are fully contained within the ND.

Within the Active Region, each detector plane is constructed from three PVC modules, leading to a total of 96 scintillator cells per plane. The geometry of the Muon Catcher differs slightly from that of the Active Region. In this section, vertical planes still consist of three modules (corresponding to 96 cells), while horizontal planes contain only two modules, resulting in 64 cells per plane. This asymmetric arrangement reflects the mechanical and structural constraints associated with the integration of steel absorber plates.

Overall, the ND contains 192 scintillator planes in the Active Region and an additional 22 planes in the Muon Catcher. The latter region also includes steel absorber layers placed after every two scintillator planes, giving a total of ten steel plates. Combining all sections, the detector contains 20,192 scintillator cells.

The Active Region is further organized into structural units known as *diblocks*, each consisting of 64 detector planes. This subdivision arises from the structure of the data acquisition system. Specifically, each Data Concentrator Module (DCM) handles the readout from a group of detector modules corresponding to two detector blocks along the beam direction. Although the Muon Catcher contains only 22 planes, it is still treated as a single diblock within the readout architecture, with some DCM input channels remaining unused.

In terms of its physical dimensions, the ND is approximately 4 m wide, 4 m tall, and 15.6 m long. The detector contains about 132 tons of liquid scintillator, 78 tons of steel absorber material, and roughly 83 tons of PVC structure. These components give the ND a total mass of approximately 293 tons [28, 29, 30]. A summary of these detector parameters is provided in Tab. 2.3, and two photographs are shown in Fig. 2.9.

2.3.3 Far Detector

The NOvA FD is located in Ash River, Minnesota, approximately 809 km from the ND at Fermilab. Construction of the detector was completed in July 2014. Similar to the ND, the FD is positioned 14.6 mrad off the central axis of the NuMI beam in order to receive a narrow-band neutrino energy spectrum optimized for oscillation studies. However, unlike the ND, which is located underground, the FD is built near the surface at an altitude of about 380 m. This configuration exposes the detector to a significant flux of cosmic-ray particles, posing additional challenges for both data acquisition and offline event reconstruction.

Figure 2.9: Recent photographs of the NOvA ND. The left panel shows the front view of the detector, while the right panel shows the detector electronics located at the top.

The FD is substantially larger than the ND. It consists of 896 scintillator planes arranged perpendicular to the beam direction. Each plane is composed of 12 PVC modules, resulting in a total of 344,064 readout cells. The detector structure is organized into 28 blocks, which were assembled and positioned using a specialized block pivoter during construction. These blocks form 14 diblocks, each containing 64 detector planes.

The full detector presents an active area of approximately $15.6\,\mathrm{m} \times 15.6\,\mathrm{m}$ to the incoming neutrino beam and extends about $60\,\mathrm{m}$ along the beam direction. The total mass of the detector is roughly 14 kt, with about 65% of the mass consisting of active liquid scintillator and the remaining portion primarily made of PVC structural material [31].

Because the FD operates on the surface, cosmic-ray backgrounds represent an important consideration. To reduce the rate of cosmic particles entering the detector, the building is covered with an overburden consisting of approximately $1.2\,\mathrm{m}$ of concrete together with an

Quantity	**Near Detector**	**Far Detector**
Modules per plane	3	12
Total number of diblocks	4	14
Total number of planes	214 + 10 (steel)	896
Total number of cells	20192	344064
Detector Dimensions (NuMI Coordinate System)		
x (cm)	[-200, 200]	[-758, 768]
y (cm)	[-200, 200]	[-749, 765]
z (cm)	[0, 1560]	[0, 5962]
Mass (metric tons)		
Total	293	14
Scintillator	132	9
PVC	83	5
Steel	78	–

Table 2.3: Summary of the main design parameters and material composition of the NOvA Near Detector and Far Detector [28, 29, 30, 31].

additional 0.15 m layer of barite rock. This shielding corresponds to roughly 14 radiation lengths and significantly reduces the flux of electromagnetic particles reaching the active detector volume.

Despite this shielding, the FD still experiences a high cosmic-ray rate of about 150 kHz. Although most of these events are unrelated to neutrino interactions, a fraction of the collected cosmic-ray data is retained. These events are useful for detector calibration and are also used to estimate the contribution of cosmic-ray backgrounds in neutrino oscillation analyses.

Tab. 2.3 summarizes the main design parameters of the NOvA FD, and photographs are shown in 2.10.

2.4 Data Acquisition and Triggering System

The NOvA experiment employs a data acquisition system (DAQ) designed for continuous readout of detector signals. Instead of recording data only when a trigger occurs, the front-

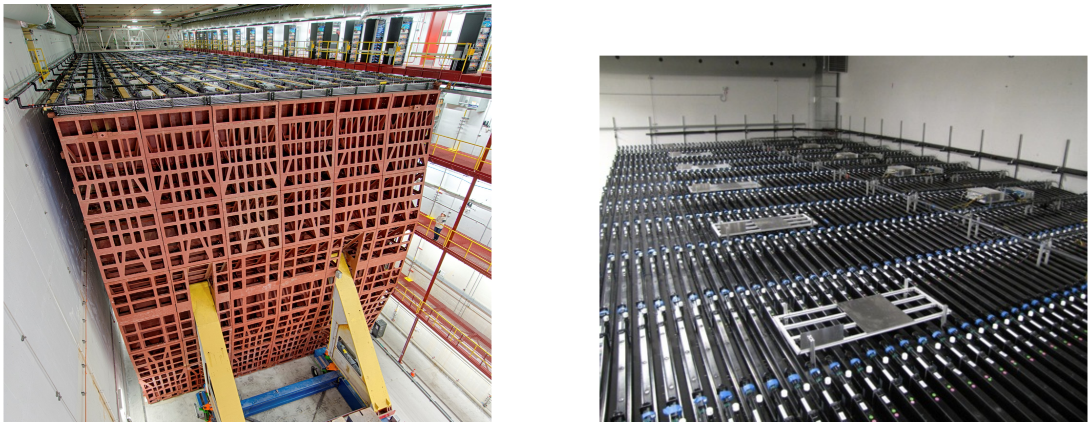

Figure 2.10: Photographs of the NOvA FD. Pictures are taken from [14, 15]

.

end electronics continuously digitize signals from the detector and temporarily store them in intermediate buffers. A global triggering framework subsequently analyzes these buffered data and determines which portions should be written to permanent storage.

Each detector module, consisting of 32 scintillator cells, is read out using an Avalanche Photodiode (APD) coupled to Front End Boards (FEBs). The APD converts scintillation light collected by the wavelength-shifting fibers into an electrical signal. The FEB then amplifies, digitizes, formats, and assigns a timestamp to the signal before forwarding the processed data to the higher-level readout system.

The digitized signals are transmitted to Data Concentrator Modules (DCMs). A single DCM receives input from up to 64 FEBs. The DCM aggregates incoming data over short time intervals and forms data packets known as *microslices*, each representing approximately 50 μs of detector activity. These microslices are subsequently grouped into larger time segments. Multiple microslices are combined into a *millislice* representing 5 ms of data, which is then transferred through an Ethernet network to a set of buffer nodes.

The buffer nodes form a distributed memory pool where data are temporarily stored while trigger decisions are evaluated. The system operates in a round-robin fashion: each

5 ms block of data from all DCMs is sent to a particular buffer node, ensuring that each node holds a complete snapshot of the detector activity during that time window. This unit of data is often referred to as a *milliblock*. Once stored, the milliblock is examined by the triggering system to determine whether the data contain information relevant for physics analyses or detector monitoring. A schematic flowchart of NOvA DAQ system is shown in Fig. 2.11.

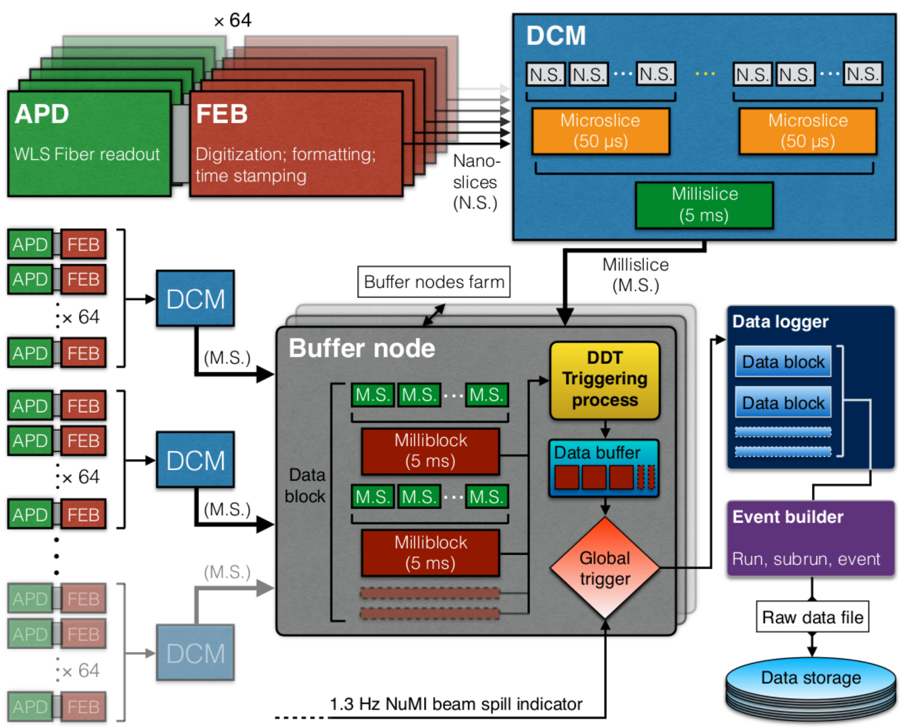

Figure 2.11: Schematic diagram of the NOvA data acquisition system showing the flow of information from the APDs and front-end electronics to the buffer nodes and long-term storage. The diagram is taken from [13].

.

The readout architecture is closely connected to the detector segmentation. Each DCM reads out modules belonging to planes of a single view and spans a region corresponding to two consecutive detector blocks, commonly referred to as a *diblock*. Because of this arrangement, the readout system is organized in terms of diblocks rather than the physical block structure of the detector. In total, the ND is instrumented with 16 DCMs, while the much larger FD requires 168 DCMs to handle its readout.

2.4.1 Timing Synchronization

Precise timing synchronization across the entire experiment is essential for identifying neutrino interactions associated with the NuMI beam spills. To achieve this, NOvA employs a dedicated timing distribution system linked to a Global Positioning System (GPS) reference. A Master Timing Distribution Unit (MTDU) receives the GPS signal and establishes a global reference clock, commonly referred to as "NOvA time". The MTDU distributes this clock to multiple Slave Timing Distribution Units (STDUs), which in turn synchronize the timing of individual DCMs and front-end electronics.

The MTDU is also connected to the accelerator timing network. Whenever a NuMI beam spill occurs, the accelerator timestamp is translated into the NOvA timing framework, allowing the experiment to identify and record the detector activity corresponding to the neutrino beam window.

2.4.2 Triggering System

Given the continuous readout of the detector and the high data rate, it is not feasible to store all recorded information permanently. The triggering system therefore plays a crucial role in selecting relevant portions of the data stream for long-term storage.

Two main categories of triggers are used in the experiment:

- **Clock-Based Triggers:** These triggers select data around predefined time intervals. A primary example is the *NuMI trigger*, which records detector activity around the beam spill. For each spill, approximately $550\ \mu$s of data are stored, centered on the $\sim\ 10\ \mu$s duration of the beam pulse. Another example is the *cosmic pulser* *trigger*, which periodically records detector data (typically at a frequency of about 10 Hz). These cosmic events are widely used for detector calibration and for estimating cosmic-ray backgrounds in oscillation analyses.

- **Data-Driven Triggers (DDT):** Unlike clock-based triggers, data-driven triggers analyze the detector activity itself to determine whether a physics signature of interest is present. These triggers operate on the buffered data and search for specific patterns or clusters of hits. Examples include triggers designed to identify rare phenomena such as supernova neutrino bursts or magnetic monopole candidates. In such cases, the duration of the stored data window is determined by the observed time structure of the detected signals.

Whenever a trigger condition is satisfied, the selected data segment is validated by the global trigger system and forwarded to the data logger. The data logger organizes the information into structured files divided into runs, subruns, and events before transferring them to the permanent storage system for further offline processing and analysis.

Chapter 3

Cosmic Ray Physics

Cosmic rays [32, 33] are highly energetic particles of extraterrestrial origin that continuously bombard the Earth. The primary cosmic-ray flux is dominated by fully ionized atomic nuclei, with protons (hydrogen nuclei) forming the largest component. This predominance reflects the primordial abundance of hydrogen produced during Big Bang nucleosynthesis. Helium nuclei constitute the second most abundant component, followed by heavier nuclei such as carbon, oxygen, and other intermediate-mass elements synthesized through stellar nucleosynthesis. Iron, being the heaviest element that can be efficiently produced in stellar fusion processes, represents a significant contribution to the high-mass end of the cosmic-ray composition. The approximate elemental composition of primary cosmic rays at GeV–TeV energies is summarized in Tab. 3.1.

Table 3.1: Composition of primary cosmic rays at GeV–TeV energies [16].

Component	**Fraction of total flux (%)**
Protons (H)	$\sim$79
Helium (He)	$\sim$20
Other (C, N, O, Fe etc.)	$\sim$1

The experimental discovery of cosmic rays is credited to the Austrian physicist Victor Hess in 1912 [34, 35]. At that time, the origin of the ionizing radiation detected by ground-based electroscopes was unclear. To investigate this, Hess performed a series of balloon flights to study the variation of ionization with altitude. Early flights in 1911, reaching altitudes of about 1100 m, showed no substantial decrease in radiation relative to ground level.

On 17 April 1912, Hess carried out a balloon ascent to approximately 5300 m during a near-total solar eclipse. The absence of any reduction in ionization during the eclipse, together with a clear increase in radiation intensity at higher altitudes, ruled out both terrestrial and solar origins. Hess concluded that the radiation must be entering the atmosphere from outer space, providing the first direct evidence for cosmic rays as a natural source of highly penetrating, high-energy particles. In recognition of this fundamental discovery, Victor Hess was awarded the 1936 Nobel Prize in Physics, shared with Carl D. Anderson [36] [37].

3.1 Acceleration of Cosmic Rays

Understanding the origin of cosmic rays requires identifying both the astrophysical sources that supply the necessary energy and the physical mechanisms responsible for accelerating particles to relativistic energies. Candidate sources must satisfy stringent energetic constraints, as the power required to sustain the observed Galactic cosmic-ray population is of the order

$$L_{\mathrm{CR}} \sim 10^{41}\,\mathrm{erg\,s^{-1}}. \tag{3.1}$$

This requirement restricts viable accelerators to highly energetic astrophysical environments such as supernova remnants, active galactic nuclei, and gamma-ray bursts.

The most widely accepted acceleration mechanism is Fermi acceleration, in which charged particles gain energy through repeated, collisionless interactions with moving magnetized plasma. The energy change experienced by a particle depends on the relative velocity between the particle and the scattering centers embedded in the plasma.

Fermi's original idea of stochastic acceleration was proposed in the late 1940s [38]. The modern formulation of diffusive shock acceleration, which naturally leads to a power–law spectrum, is reviewed by Blandford and Eichler [39]. The Hillas criterion, which constrains the maximum energy reachable in an astrophysical accelerator, is discussed in [40].

In second-order Fermi acceleration, particles scatter off randomly moving magnetic irregularities, leading to a mean fractional energy gain per interaction given by

$$\left\langle \frac{\Delta E}{E} \right\rangle \propto \left(\frac{V}{c}\right)^2, \tag{3.2}$$

where V is the characteristic velocity of the scattering centers and c is the speed of light. Since $V \ll c$ for typical interstellar turbulence, this process is relatively inefficient and contributes primarily at lower energies.

A more efficient mechanism is first-order Fermi acceleration, also known as diffusive shock acceleration, which occurs at astrophysical shock fronts such as those produced by supernova explosions. Particles gain energy by repeatedly crossing the shock boundary and scattering off magnetic irregularities on either side. In this case, the average fractional energy gain per acceleration cycle scales linearly with the shock velocity,

$$\left\langle \frac{\Delta E}{E} \right\rangle \propto \frac{V_{\text{sh}}}{c}, \tag{3.3}$$

where V_{sh} denotes the shock velocity. This first-order dependence makes shock acceleration significantly more efficient than stochastic acceleration in turbulent media.

Diffusive shock acceleration naturally leads to a power-law energy spectrum for the accelerated particles,

$$\frac{dN}{dE} \propto E^{-\gamma}, \tag{3.4}$$

where the spectral index γ depends on the shock compression ratio. For strong, non-relativistic shocks, the theory predicts $\gamma \simeq 2$, which is consistent with the observed cosmic-ray spectrum after accounting for propagation effects in the interstellar medium.

The maximum energy attainable in an astrophysical accelerator is limited by the size R of the acceleration region and the strength B of the magnetic field. This constraint is commonly expressed by the Hillas criterion,

$$E_{\text{max}} \simeq ZeBR\beta, \tag{3.5}$$

where Z is the particle charge number, e is the elementary charge, and $\beta = V/c$ characterizes the velocity of the scattering region. This criterion provides a useful framework for evaluating the capability of different astrophysical objects to accelerate cosmic rays to the highest observed energies.

3.2 Energy spectra

The intensity of cosmic rays observed at Earth decreases rapidly with increasing particle energy. Over a wide range of energies, the all-particle cosmic-ray spectrum is well described by an approximate power-law behavior, extending from a few GeV up to beyond 10^{20} eV. As a result, low-energy cosmic rays are abundant, whereas ultra–high-energy cosmic rays are extremely rare.

The differential energy spectrum of primary cosmic rays can be parameterized as [32]

$$\frac{dN}{dE} = k\,E^{-\gamma}, \tag{3.6}$$

where k is a normalization constant and γ is the spectral index, which varies with energy. The flux is typically expressed in units of $\mathrm{m^{-2}\,sr^{-1}\,s^{-1}\,GeV^{-1}}$.

At energies of order $\sim 10\,\mathrm{GeV}$, the cosmic-ray flux at the top of the atmosphere is approximately 1 particle $\mathrm{cm^{-2}\,s^{-1}}$ [32]. In contrast, at energies approaching 10^{20} eV, the flux decreases dramatically to roughly one particle per square kilometer per century [16]. This enormous dynamic range poses a significant experimental challenge, requiring detectors with very large effective areas and long exposure times in order to study the highest-energy cosmic rays.

The all-particle spectrum exhibits several characteristic features that reflect changes in the dominant acceleration mechanisms and/or propagation effects. Up to energies of about $E \sim 10^{15}$ eV, the spectrum follows a power law with spectral index $\gamma \simeq 2.7$. Near this energy, known as the *knee*, the spectrum steepens to $\gamma \simeq 3.0$ [41, 16], which is commonly interpreted as evidence for the maximum energy reachable by typical Galactic accelerators, such as supernova remnants.

A second steepening, referred to as the *second knee*, occurs around $E \sim 10^{17}$ eV. Beyond this region, the spectrum gradually hardens again, forming a feature known as the *ankle* at approximately $E \sim 10^{19}$ eV. The ankle is often interpreted as marking the transition from cosmic rays of predominantly Galactic origin to those of extragalactic origin.

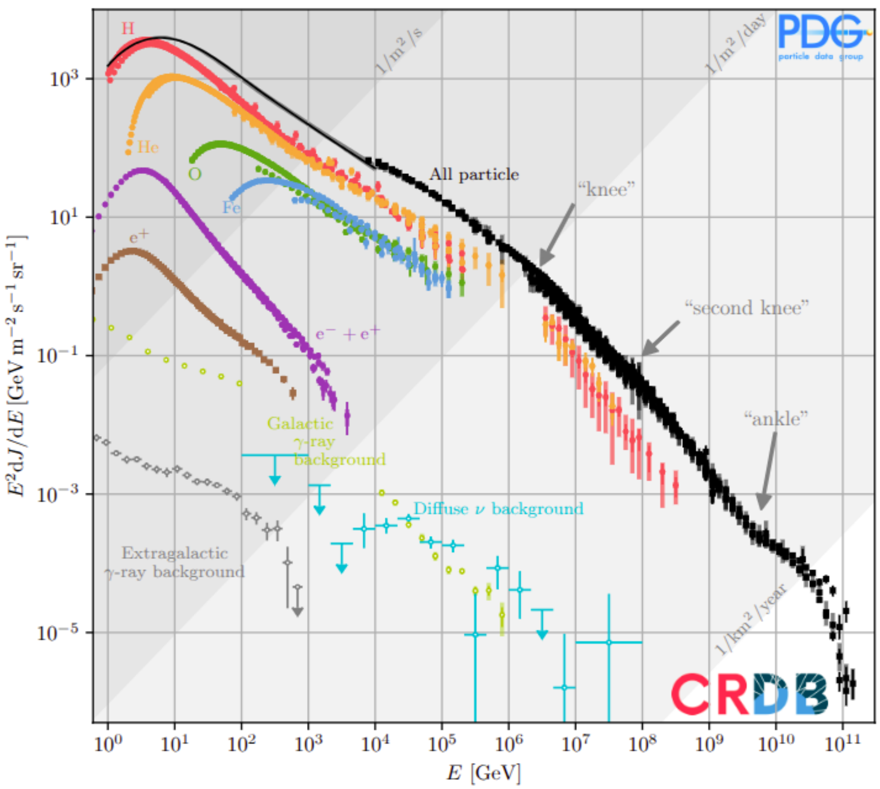

Figure 3.1: All-particle cosmic-ray energy spectrum compiled by the Particle Data Group [16]. It shows measurements of the cosmic-ray intensity (including both charged and neutral components) multiplied by the square of the particle kinetic energy.

At the highest energies, cosmic rays are expected to lose energy through interactions with the cosmic microwave background (CMB). This phenomenon, known as the Greisen–Zatsepin–Kuz'min (GZK) suppression, arises primarily from photopion production and electron–positron pair production during cosmic-ray propagation over cosmological distances. The characteristic energy scale of this suppression is given by

$$E_{\text{GZK}} \sim (5\text{–}6) \times 10^{19}\,\text{eV}. \tag{3.7}$$

As a consequence, the Universe becomes increasingly opaque to cosmic rays with energies above this threshold unless they originate from nearby sources.

Measurements performed by large air-shower experiments, such as the HiRes experiment and the Pierre Auger Observatory [42, 43], have observed a suppression of the cosmic-ray flux consistent with the GZK prediction. Nevertheless, rare events with energies exceeding 10^{19} eV have been reported, motivating continued experimental and theoretical efforts to understand the nature of ultra–high-energy cosmic-ray sources and their propagation through intergalactic space.

Fig. 3.1 shows the all-particle cosmic-ray energy spectrum compiled by the Particle Data Group, illustrating the power-law behavior over many decades in energy and highlighting the knee, ankle, and GZK suppression features.

3.3 Cosmic Air Shower

When a high-energy cosmic-ray particle enters the Earth's atmosphere, it undergoes an inelastic collision with an atmospheric nucleus at altitudes of several tens of kilometers. This initial interaction initiates a cascade of secondary particles that propagates toward the ground and is collectively known as an extensive air shower (EAS). Through successive interactions and decays, the shower multiplies rapidly, distributing the primary particle's energy among a large number of secondaries over a wide lateral area that can extend to several square kilometers.

Hadronic interactions dominate the early stages of the shower development. The primary cosmic ray produces secondary hadrons, mainly charged and neutral pions and kaons, which in turn interact with air nuclei to generate further hadrons as well as leptons and photons. This hadronic component drives the longitudinal development of the shower until the particle energies fall below the threshold for further particle production and the cascade is eventually absorbed in the lower atmosphere or at the Earth's surface.

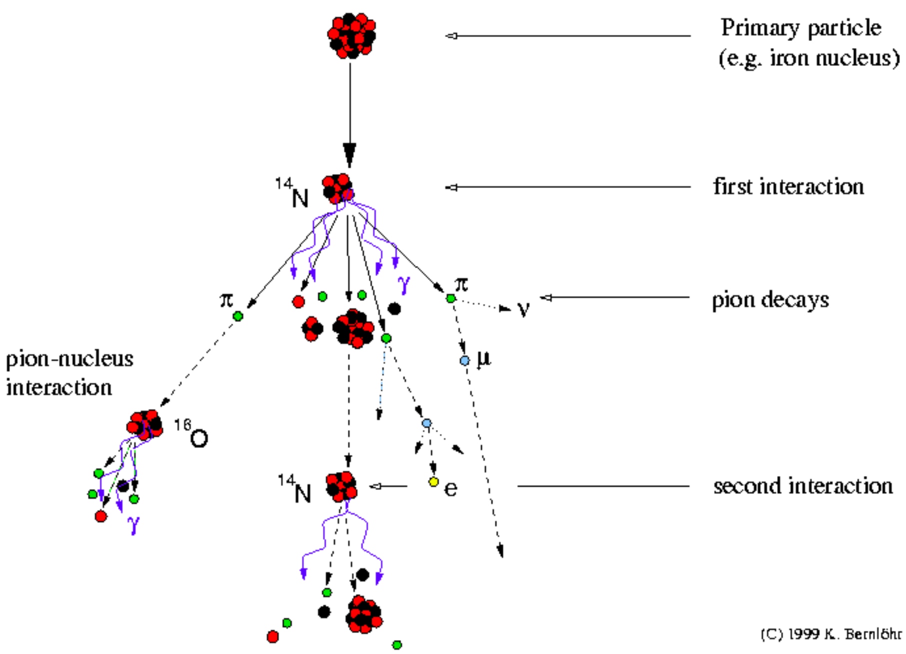

Figure 3.2: Developement of cosmic air showers produced from a interaction of a primary particle with the atmosphere [17].

In addition to the hadronic component, electromagnetic sub-showers arise primarily from the decay of neutral pions into photon pairs. These photons initiate electromagnetic cascades through pair production and bremsstrahlung processes, leading to large numbers of electrons, positrons, and photons. Although the microscopic origins differ, the particle distributions observed at ground level from hadron- and gamma-induced showers can appear similar to surface detector arrays. The development of cosmic air showers after the first interaction is shown in Fig. 3.2.

Extensive air showers play a crucial role in cosmic-ray studies because the flux of ultra-high-energy primary particles is extremely low. Direct detection of such primaries is impractical due to their small interaction probability, whereas the large spatial footprint of an EAS provides an observable signature that can be sampled by ground-based detectors. This amplification mechanism enables indirect measurements of primary cosmic rays at energies far exceeding those accessible by man-made accelerators.

The secondary pions and kaons produced in hadronic showers are unstable and typically decay on short timescales if they do not undergo further interactions. The dominant decay channels are

$$\pi^{\pm} \to \mu^{\pm} + \nu_{\mu}\,(\bar{\nu}_{\mu}), \tag{3.8}$$

$$\pi^{0} \to \gamma + \gamma, \tag{3.9}$$

$$K^{\pm} \to \mu^{\pm} + \nu_{\mu}\,(\bar{\nu}_{\mu}), \tag{3.10}$$

where photons from neutral pion decay feed the electromagnetic component of the shower, while charged meson decays produce muons and neutrinos. At sufficiently low energies, muons further decay according to

$$\mu^{\pm} \to e^{\pm} + \nu_{e}\,(\bar{\nu}_{e}) + \bar{\nu}_{\mu}\,(\nu_{\mu}). \tag{3.11}$$

At higher energies, however, time dilation suppresses muon decay in flight, and neutrino production is dominated by the decay of kaons rather than muons.

Muons constitute a particularly important component of extensive air showers. Due to their relatively large mass and weak interaction cross section, muons lose energy primarily through ionization and radiative processes and are far less likely than electrons to initiate secondary electromagnetic cascades. As a result, muons can traverse the atmosphere with minimal energy loss and reach the Earth's surface with a substantial fraction of their initial energy.

High-energy muons are capable of penetrating significant depths underground, where they continue to lose energy through interactions with matter. This penetrating power makes them an effective probe of high-energy cosmic-ray interactions and provides a nearly background-free signal for underground detectors. Since the overwhelming majority of muons observed at the Earth's surface originate from cosmic-ray interactions in the atmosphere, muon measurements offer a powerful indirect method for studying the properties of primary cosmic rays.

3.4 Seasonal variation of cosmic muon flux

Primary cosmic rays interacting with nuclei in the upper atmosphere initiate extensive air showers (EAS), producing cascades of hadronic and electromagnetic particles. The hadronic component of these showers is dominated by secondary mesons, primarily charged pions and kaons, which play a central role in the generation of atmospheric muons. Muons arise mainly from the weak decays of these mesons, with the channels $\pi^{\pm} \rightarrow \mu^{\pm}\nu_{\mu}$ and $K^{\pm} \rightarrow \mu^{\pm}\nu_{\mu}$ providing the dominant contribution to the observable muon flux.

The production and survival of muons in the atmosphere are strongly influenced by the conditions under which their parent mesons propagate. In particular, the competition between meson decay and hadronic interaction depends on both the meson energy and the density of the surrounding atmosphere. At higher altitudes, where the atmospheric density is lower, mesons are more likely to decay before interacting, leading to an enhanced production of high-energy muons. Conversely, in denser atmospheric layers, secondary mesons have a higher probability of interacting before decaying, which suppresses the resulting muon flux.

Atmospheric density profiles are not static but vary with time, primarily due to temperature variations. While short-term temperature fluctuations are most pronounced in the troposphere and stratosphere, the temperature structure of the upper atmosphere evolves more gradually over seasonal timescales. Since temperature and density are inversely related, seasonal changes in atmospheric temperature effectively modulate the decay probability of mesons produced in cosmic-ray interactions.

As a result, the flux of atmospheric muons observed at the underground detector is expected to exhibit a seasonal dependence. During warmer periods, the atmosphere expands, reducing the density at a given altitude and increasing the likelihood that mesons decay into muons. During colder periods, the opposite effect occurs, leading to a reduced muon flux. Because the primary cosmic-ray flux is approximately isotropic and stable over these timescales, the observed seasonal modulation of muons can be attributed primarily to atmospheric effects rather than changes in the incident cosmic-ray intensity.

This seasonal variation of the cosmic muon flux provides a valuable probe of the interplay between cosmic-ray interactions and atmospheric dynamics.

3.4.1 Muon intensity produced in atmosphere

The determination of the muon intensity at an underground detector begins with the calculation of the differential muon intensity at the Earth's surface. Muons that reach the underground detector originate predominantly from the decay of secondary mesons produced by the interaction of primary cosmic rays with atmospheric nuclei.

The production rate of these mesons decreases with atmospheric slant depth X approximately as an exponential attenuation factor e^{-X/Λ_N}, where Λ_N represents the effective interaction length of the primary nucleon in air. The total differential muon intensity at ground level is obtained by integrating the muon production spectrum over the entire atmospheric column:

$$\frac{dI_\mu}{dE_\mu} = \int_0^\infty P'_\mu(X, E_\mu)\, dX. \tag{3.12}$$

Here, $P'_\mu(X, E_\mu)$ accounts for muons originating from the two-body decays $M \to \mu\nu_\mu$, with $M = \pi, K$. Explicitly,

$$P'_\mu(X, E_\mu) = \sum_{M=\pi,K} \frac{\epsilon_M}{X\cos\theta(1-r_M)} \int_{E_\mu}^{E_\mu/r_M} \frac{M'(X,E)}{E}\frac{dE}{E}, \tag{3.13}$$

where $r_M = m_\mu^2/m_M^2$ and $M'(X, E)$ denotes the meson production spectrum.

The meson spectrum can be expressed as

$$M'(X,E) = e^{-X/\Lambda_M} \frac{Z_{NM}}{\lambda_N} E^{-(\gamma+1)} \int_0^\infty \left(\frac{X'}{X}\right)^{\epsilon_M/(E\cos\theta)} \exp\left(\frac{X'}{\Lambda_M} - \frac{X'}{\Lambda_N}\right) dX'. \tag{3.14}$$

In this expression: Λ_M and Λ_N are the attenuation lengths for mesons and nucleons, λ_N is the nucleon interaction length, Z_{NM} is the spectrum-weighted moment of the inclusive production cross section, γ is the spectral index of the muon energy spectrum, ϵ_M is the critical energy of meson M.

The critical energy ϵ_M represents the energy at which decay and interaction probabilities are equal, and is given by

$$\epsilon_M = \frac{m_M c^2 H(T)}{c\tau_M}. \tag{3.15}$$

For standard atmospheric conditions, one obtains approximately $\epsilon_\pi \simeq 115$ GeV and $\epsilon_K \simeq 850$ GeV.

Solving Eq. 3.12,

$$\frac{dI_\mu}{dE_\mu} \approx A E_\mu^{-(\gamma+1)} \left(\frac{A_\pi^1}{1+1.1E_\mu\cos\theta/\epsilon_\pi} + \frac{0.635\, A_K^1}{1+1.1E_\mu\cos\theta/\epsilon_K}\right), \tag{3.16}$$

where

$$A_M^1 = \frac{Z_{NM}(1-r_M^{\gamma+1})}{(1-r_M)(\gamma+1)}. \tag{3.17}$$

Using standard parameter values, this reduces numerically to

$$\frac{dI_\mu}{dE_\mu} \approx 0.14\, E_\mu^{-(\gamma+1)} \left(\frac{1}{1+1.1E_\mu\cos\theta/\epsilon_\pi} + \frac{0.054}{1+1.1E_\mu\cos\theta/\epsilon_K}\right) \frac{\text{muons}}{\text{cm}^2\,\text{sr}\,\text{GeV}}. \tag{3.18}$$

The total surface muon intensity above a threshold energy E_{th} is then

$$I_\mu(E_{\text{th}}) = \int_{E_{\text{th}}}^{\infty} \frac{dI_\mu}{dE_\mu} dE_\mu. \tag{3.19}$$

Substituting Eq. 3.16 and integrating yields

$$I_\mu \approx B\, E_{\text{th}}^{-\gamma} \left(\frac{1}{\gamma + (\gamma+1) 1.1 E_{\text{th}} \cos\theta / \epsilon_\pi} + \frac{0.054}{\gamma + (\gamma+1) 1.1 E_{\text{th}} \cos\theta / \epsilon_K} \right), \tag{3.20}$$

where B is a normalization constant. The quantity $E_{\text{th}} \cos\theta$ corresponds to the vertical surface energy required for a muon to reach a given underground depth.

3.4.2 Relation of muon intensity with atmospheric temperatures

The dependence of the muon intensity on atmospheric temperature arises through the critical energies ϵ_M, which are proportional to the atmospheric scale height $H(T)$. Since the real atmosphere is not strictly isothermal and meson production occurs over a broad range of altitudes, a direct treatment using $T(X)$ is complex.

To address this, an effective atmospheric temperature T_{eff} is introduced. This parameter represents the temperature of an equivalent isothermal atmosphere that would reproduce the observed muon intensity I_μ^0.

Define the fractional temperature deviation as

$$\eta(X) = \frac{T(X) - T_{\text{eff}}}{T_{\text{eff}}}, \tag{3.21}$$

so that the critical energy becomes

$$\epsilon_M(X) = \epsilon_M^0 [1 + \eta(X)], \tag{3.22}$$

with ϵ_M^0 corresponding to $T = T_{\text{eff}}$.

A perturbation in temperature modifies the meson production spectrum and therefore the differential muon intensity. The resulting variation can be written as

$$\Delta\left(\frac{dI_\mu}{dE_\mu}\right) = \sum_{M=\pi,K} \alpha_M(X)\frac{\Delta T(X)}{T_{\text{eff}}}, \tag{3.23}$$

where the temperature correlation coefficient is

$$\alpha_M(X) = \int_{E_{\text{th}}}^{\infty} \frac{A_M^1 E_\mu^{-(\gamma+1)} e^{-X/\Lambda_M}}{1 + B_M^1 K(X)(E_\mu \cos\theta/\epsilon_M^0)^2} dE_\mu, \tag{3.24}$$

with

$$B_M^1 = \frac{(\gamma+3)(1-r_M^{\gamma+1})}{(\gamma+1)(1-r_M^{\gamma+3})}, \tag{3.25}$$

$$K(X) = \frac{(1-X/\Lambda_M')^2}{(1-e^{-X/\Lambda_M'})\Lambda_M'/X}, \tag{3.26}$$

and

$$\frac{1}{\Lambda_M'} = \frac{1}{\Lambda_N} - \frac{1}{\Lambda_M}. \tag{3.27}$$

Integrating over atmospheric depth gives the total variation in intensity:

$$\Delta I_\mu = \sum_{M=\pi,K} \int_0^\infty \alpha_M(X)\frac{\Delta T(X)}{T_{\text{eff}}} dX. \tag{3.28}$$

Requiring $\Delta I_\mu = 0$ when $T(X) = T_{\text{eff}}$ leads to the effective temperature definition

$$T_{\text{eff}} = \frac{\sum_M \int_0^\infty T(X)\alpha_M(X)dX}{\sum_M \int_0^\infty \alpha_M(X)dX}. \tag{3.29}$$

For discrete atmospheric data,

$$T_{\text{eff}} = \frac{\sum_i T(X_i)[W_\pi(X_i) + W_K(X_i)]\Delta X_i}{\sum_i [W_\pi(X_i) + W_K(X_i)]\Delta X_i}, \tag{3.30}$$

where

$$W_M(X) = \frac{A_M^1(1-X/\Lambda_M')^2 e^{-X/\Lambda_M}}{\gamma + (\gamma+1)B_M^1 K(X)(E_{\text{th}}\cos\theta/\epsilon_M)^2}. \tag{3.31}$$

An effective temperature coefficient is defined as

$$\alpha_T = \frac{1}{I_\mu^0} \sum_M \int_0^\infty \alpha_M(X) dX. \tag{3.32}$$

Relation Between Muon Intensity and Measured Muon Rate

In practice, underground experiments do not measure the muon intensity directly. Instead, what is recorded is the muon counting rate. The intensity must therefore be related to measurable detector quantities.

The muon intensity I_μ is defined as the flux per unit area, per unit solid angle, per unit time. For a detector with effective area A_{eff}, solid angle acceptance Ω, and detection efficiency ξ, the measured muon rate R_μ is related to the intensity through

$$I_\mu = \frac{R_\mu}{\xi A_{\text{eff}} \Omega}. \tag{3.33}$$

Since the rate is defined as the number of detected muons per unit time,

$$R_\mu = \frac{N_\mu}{\Delta t}, \tag{3.34}$$

we can rewrite Eq. 3.33 as

$$I_\mu = \frac{N_\mu / \Delta t}{\xi A_{\text{eff}} \Omega}. \tag{3.35}$$

To study seasonal variations, we consider relative changes with respect to the long-term average values. Let the deviation of the intensity from its average value be

$$\Delta I_\mu = I_\mu - I_\mu^0, \tag{3.36}$$

and similarly for the rate,

$$\Delta R_\mu = R_\mu - \langle R_\mu \rangle. \tag{3.37}$$

Using Eq. 3.35, the fractional variation of the intensity becomes

$$\frac{\Delta I_\mu}{I_\mu} = \frac{\dfrac{N_\mu^i}{\delta t} - \dfrac{N_\mu}{\Delta t}}{\dfrac{N_\mu}{\Delta t}}, \tag{3.38}$$

where N_μ^i is the number of muons counted during a shorter time interval δt within the total exposure period Δt.

Recognizing that $R_\mu^i = N_\mu^i/\delta t$ and $\langle R_\mu \rangle = N_\mu/\Delta t$, we obtain

$$\frac{\Delta I_\mu}{I_\mu} = \frac{\Delta R_\mu}{\langle R_\mu \rangle}. \tag{3.39}$$

This result shows that the fractional variation of the measured muon rate is identical to the fractional variation of the muon intensity. Detector-specific quantities such as efficiency, effective area, and solid angle acceptance cancel in the ratio.

Combining Eq. 3.39 with the temperature–intensity relation derived previously,

$$\frac{\Delta I_\mu}{I_\mu^0} = \alpha_T \frac{\Delta T_{\text{eff}}}{\langle T_{\text{eff}} \rangle}, \tag{3.40}$$

leads to the experimentally useful expression

$$\frac{\Delta R_\mu}{\langle R_\mu \rangle} = \alpha_T \frac{\Delta T_{\text{eff}}}{\langle T_{\text{eff}} \rangle}. \tag{3.41}$$

Eq. 3.41 establishes the direct proportionality between seasonal variations in the underground muon rate and variations in the effective atmospheric temperature. The proportionality constant α_T is referred to as the effective temperature coefficient and encodes the sensitivity of the detector to atmospheric changes.

3.4.3 Theoretical Prediction of the Effective Temperature Coefficient

A theoretical estimate of the effective temperature coefficient α_T can be obtained from the definition of the relative variation of the muon intensity with respect to temperature.

Starting from the general relation between intensity variations and temperature variations, one may write

$$\alpha_T = \frac{T}{I_\mu^0}\frac{\partial I_\mu}{\partial T}. \tag{3.42}$$

At first glance, Eq. 3.42 suggests that the muon intensity must be expressed explicitly as a function of temperature. However, in the semi-analytic treatment developed in the previous subsection, the temperature dependence does not appear directly in I_μ. Instead, it enters implicitly through the meson critical energies ϵ_M, which depend on the atmospheric scale height and therefore on temperature.

To obtain a more practical expression, one exploits the fact that a variation in temperature effectively shifts the competition between meson decay and interaction, which in turn modifies the energy dependence of the muon intensity. Using this property, the temperature derivative may be recast as an energy derivative, yielding the equivalent form [44]

$$\alpha_T = -\frac{E_{th}}{I_\mu^0}\frac{\partial I_\mu}{\partial E_\mu} - \gamma, \tag{3.43}$$

where E_{th} is the vertical threshold energy required for a muon to reach the detector depth and γ is the muon spectral index.

Substituting the approximate expression for the integral muon intensity (Eq. 3.20) together with the differential spectrum (Eq. 3.16) into Eq. 3.43, one obtains

$$\alpha_T = \frac{1}{D_\pi}\frac{\dfrac{1}{\epsilon_K} + A_K^1\left(\dfrac{D_\pi}{D_K}\right)^2\dfrac{1}{\epsilon_\pi}}{\dfrac{1}{\epsilon_K} + A_K^1\left(\dfrac{D_\pi}{D_K}\right)\dfrac{1}{\epsilon_\pi}}, \tag{3.44}$$

where the quantities D_π and D_K are defined as

$$D_{\pi,K} = \frac{\gamma}{\gamma+1}\frac{\epsilon_{\pi,K}}{1.1E_{th}\cos\theta} + 1. \tag{3.45}$$

The parameter A_K^1 encodes the relative contribution of kaon decays to the total muon intensity. If one neglects the kaon component and considers only pion-induced muons, the above expression simplifies considerably. Setting $A_K^1 = 0$ yields

$$(\alpha_T)_\pi = \left[\frac{\gamma}{\gamma + 1} \frac{\epsilon_\pi}{1.1 E_{th} \cos\theta} + 1 \right]^{-1} . \tag{3.46}$$

This limiting case illustrates that the effective temperature coefficient increases as the detector threshold energy becomes large compared to the pion critical energy. In deep underground experiments, where $E_{th} \cos\theta \gg \epsilon_\pi$, α_T approaches unity, reflecting the dominance of meson decay over interaction at high energies.

3.5 Overview and history of measurements of cosmic muon flux modulation

The variation of muon flux as a function of month of the year is quite well measured phenomenon and it has been reported by several experiments placed underground and on the Earth surface. In this section, I will briefly describe historical measurements of seasonal variation with the brief description of the experiment. Results are separated based on the number of muons detected in the detector from same primary: single muon measurements (if a single muon is detected from a single primary) and multiple muon measurements (if more than one muon is detected from same primary).

3.5.1 Single muon measurements

One of the earliest measurements was performed by the MACRO detector at the Gran Sasso National Laboratory in Italy [18]. The detector was located under an overburden of about 3100 mwe and consisted of a large liquid scintillator and tracking system with approximate dimensions of $77 \times 12 \times 9$ m^3. Using data collected between 1991 and 1994, the experiment observed a clear annual modulation of the underground muon rate correlated with

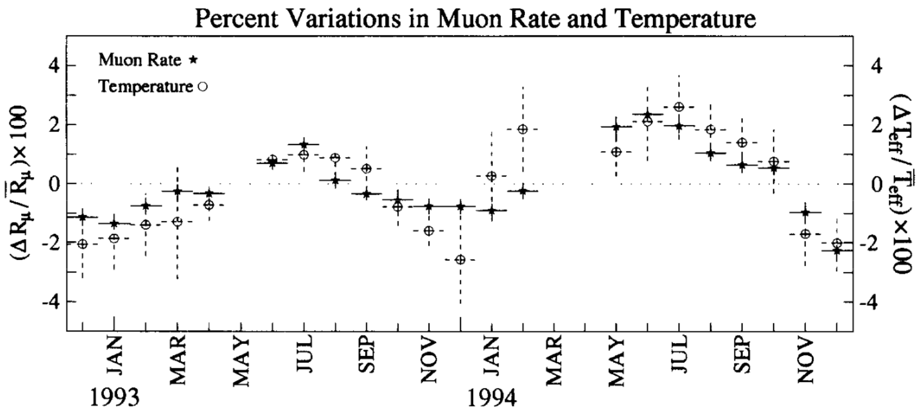

Figure 3.3: Seasonal variation of the underground muon intensity measured by the MACRO experiment together with the corresponding effective atmospheric temperature variation [18].

atmospheric temperature variations, as shown in Fig. 3.3. This study reported one of the first measurements of the effective temperature coefficient.

The AMANDA detector at the South Pole also studied the seasonal behavior of cosmic muons [19]. The detector consisted of optical modules embedded deep in Antarctic ice at depths between approximately 1500 m and 2000 m. Using about 225 days of data collected in 1997, the experiment observed a strong correlation between the muon rate and atmospheric temperature, yielding a temperature coefficient of $\alpha_T = 0.86 \pm 0.05$. The correlation between the muon rate and the effective atmospheric temperature is shown in Fig. 3.4.

The Large Volume Detector (LVD), also located at the Gran Sasso underground laboratory, is a 1 kton liquid scintillator detector designed primarily to detect neutrinos from supernova explosions [20]. The detector consists of 840 scintillation counters arranged in three modular towers. Muon data collected between January 2001 and December 2008, corresponding to about 2907 live days, showed a clear seasonal modulation in the underground muon intensity consistent with the expected atmospheric temperature dependence,

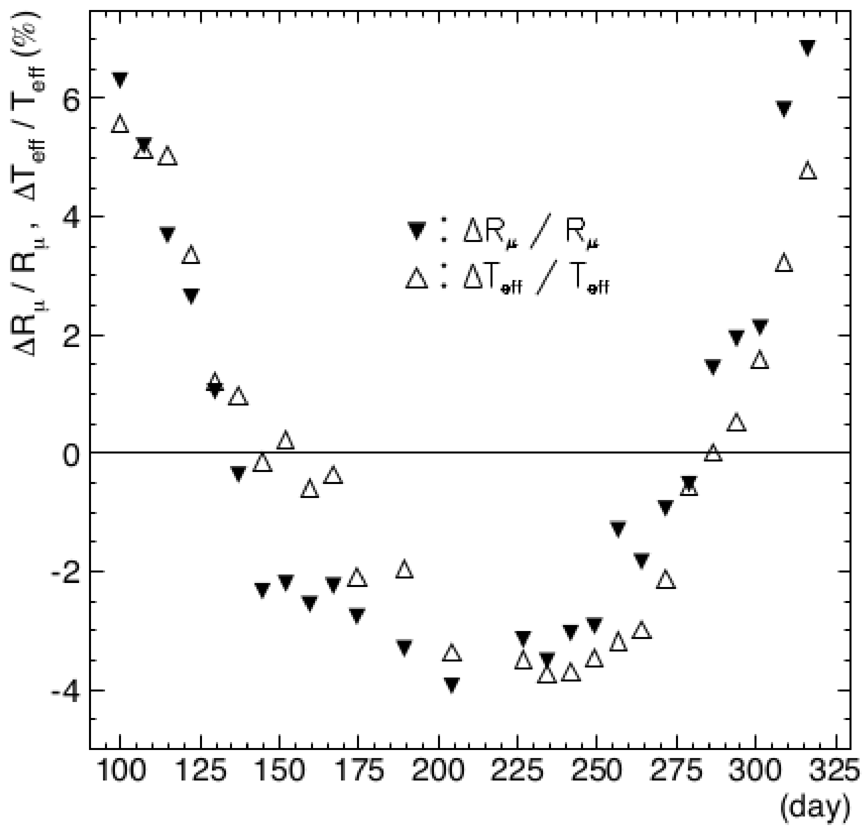

Figure 3.4: Seasonal variation of the underground muon intensity and the corresponding atmospheric temperature variation measured by the AMANDA experiment [19].

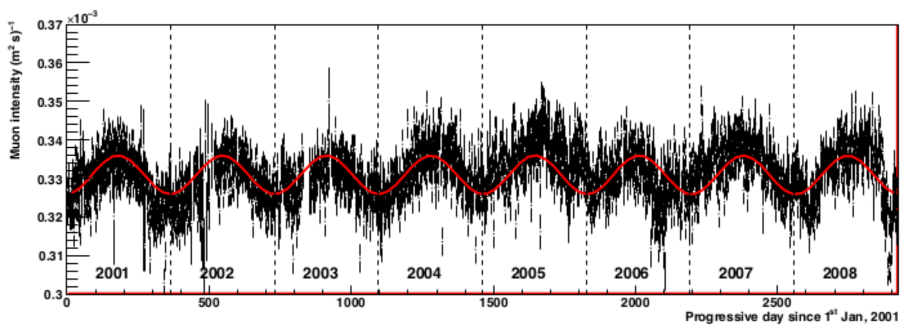

Figure 3.5: Seasonal variation of the underground muon intensity measured by the LVD experiment together with the corresponding atmospheric temperature variation [20].

as illustrated in Fig. 3.5.

Seasonal variations of muon rates have also been studied using the IceCube neutrino observatory at the South Pole [21]. IceCube instruments a cubic kilometer of Antarctic ice with arrays of photomultiplier tubes deployed on vertical strings. Using data collected

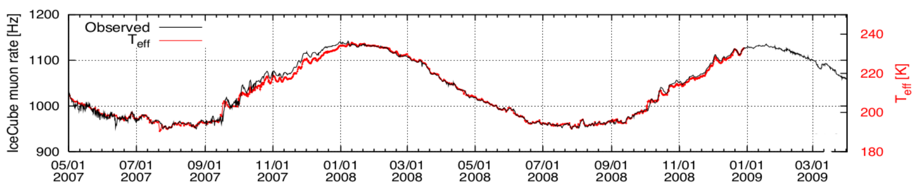

Figure 3.6: Seasonal variation of the underground muon intensity and the corresponding effective atmospheric temperature measured by the IceCube experiment [21].

between 2007 and 2009, the experiment observed a strong correlation between the muon rate and the effective atmospheric temperature, as shown in Fig. 3.6, reporting a temperature coefficient of approximately $\alpha_T \approx 0.90$.

The Borexino experiment at the Gran Sasso National Laboratory has also measured seasonal variations of cosmic-ray muons [22]. Although the detector was primarily designed to study low-energy solar neutrinos, its large liquid scintillator target and surrounding water Cherenkov veto system provide excellent sensitivity to cosmic muons. Data collected between 2007 and 2011 showed a clear annual modulation in the muon flux with a measured temperature coefficient of $\alpha_T = 0.93 \pm 0.04$. The variation of the muon rate and its correlation with atmospheric temperature are shown in Fig. 3.7.

Seasonal variation measurements have also been performed in reactor neutrino experiments such as the Double Chooz experiment in France [23]. The experiment operates two liquid scintillator detectors located approximately 400 m and 1050 m from the reactor cores, with overburdens of about 120 mwe and 300 mwe, respectively. Analysis of several years of data revealed a measurable seasonal modulation of the muon flux, with

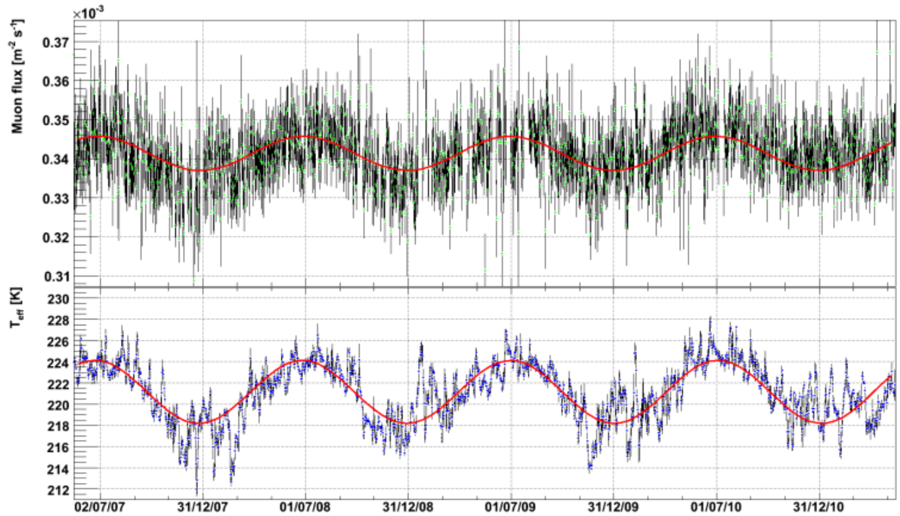

Figure 3.7: Seasonal variation of the underground muon intensity and the corresponding atmospheric temperature measured by the Borexino experiment [22].

temperature coefficients of $\alpha_T = 0.212 \pm 0.013$ (stat) $\pm$ 0.011 (syst) for the ND and $\alpha_T = 0.355 \pm 0.002$ (stat) $\pm$ 0.017 (syst) for the FD. The comparison between the muon rate variation and the effective atmospheric temperature is shown in Fig. 3.8.

The MINOS experiment, designed to study neutrino oscillations, operated two detectors: a ND at Fermilab and a FD located 731 km away in the Soudan Underground Mine in Minnesota [24]. The FD, situated under an overburden of about 2100 mwe, was used to study the seasonal variation of cosmic-ray muons. Using data collected between 2003 and 2008, the experiment observed a clear annual modulation of the underground muon rate that was strongly correlated with variations in the effective atmospheric temperature. The seasonal variation of the muon rate and the corresponding temperature variation are shown in Fig. 3.9.

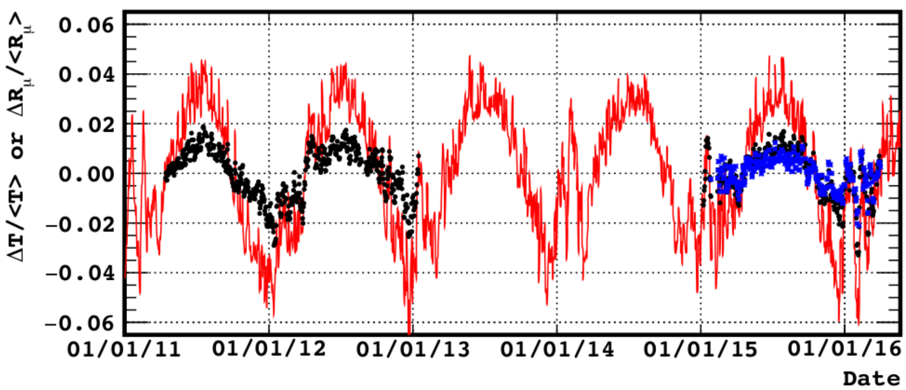

Figure 3.8: Seasonal variation of the underground muon intensity and the corresponding atmospheric temperature measured by the Double Chooz experiment [23].

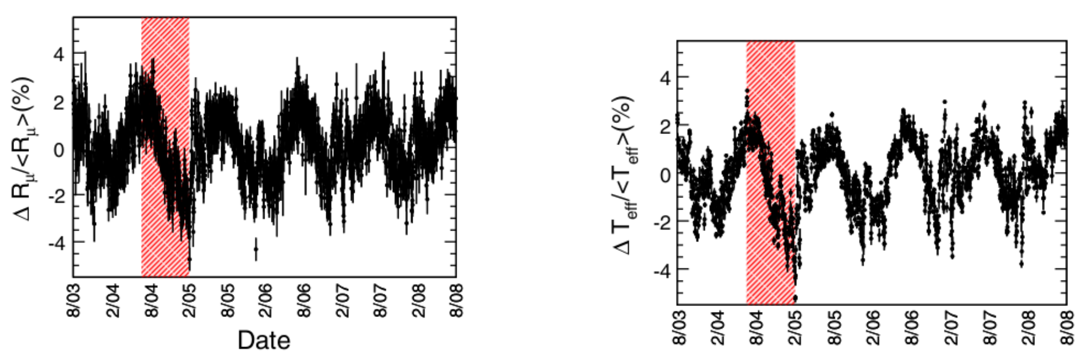

Figure 3.9: Seasonal variation of the underground muon intensity measured by the MINOS experiment [24]. The left panel shows the variation of the muon rate with time, while the right panel shows the corresponding effective atmospheric temperature variation.

3.5.2 Multiple muon measurements

The MINOS experiment first reported an unexpected winter maximum in multiple-muon events, contrasting with the well-established summer peak observed for single muons. This effect was observed in both the ND [45], located at 225 meters water equivalent (mwe), and the FD [45], located at 2100 mwe underground. The strength of the seasonal variation was found to depend on the spatial separation of muons in an event, with closely spaced muons exhibiting a winter maximum while widely separated muons followed the expected summer maximum behavior. The MINOS study proposed that secondary hadronic interactions in the atmosphere played a key role in this effect, though a definitive explanation remained unresolved.

Based on this, the NOvA ND [46], located at the same depth as the MINOS ND but with a different detector design, confirmed the winter peak in multiple-muon events in an independent dataset spanning 2015–2017. The NOvA FD [47], situated at the surface with minimal overburden, further reinforced this anomaly by detecting a strong seasonal dependence in multiple-muon showers, with a 30% stronger winter maximum at higher multiplicities and a 20% stronger effect for showers arriving from near the horizon.

However, none of these studies explains the inconsistency between single- and multiple-muon seasonal trends and it remained an open question. The MINOS study explored four possible explanations for this effect [45], including the altitude geometry effect, but ultimately discarded it based on the spatial separation among muons at the detector level. However, this conclusion was influenced by the size of the detector itself. According to MINOS, a $\pm 2\%$ variation in atmospheric temperature leads to a $\pm 2\%$ variation in the altitude of muon production, which was deemed insufficient to significantly affect the number of multiple-muon events. However, when considering the true separation of muons at the Earth's surface, which has a mean distance of ≈ 130 m, this effect becomes more significant. A $\pm 2\%$ change in separation at the surface could cause many multiple-muon events to be misidentified as single-muon events by the time they reach the detector, particularly in surface-based or shallow underground experiments.

In this thesis, I extended the analysis of reference [46] with four additional years of NOvA ND data and, for the first time, present a CORSIKA-based simulation of the seasonal effects in the NOvA ND.

Chapter 4

Analysis of Seasonal Variation of Cosmic Muon

This chapter presents a analysis of the seasonal modulation of cosmic-ray muons observed in the NOvA ND. The primary objective is to quantify the correlation between atmospheric temperature variations and the underground muon flux, and to interpret the results within the theoretical framework of meson production and decay in the atmosphere.

The first step of the analysis is the simulation of extensive air showers using CORSIKA, which provides a description of primary interactions and secondary particle production in the atmosphere. The simulated muons are then propagated to detector depth, randomized over the detector surface, and processed through the detector simulation.

Subsequently, the reconstruction procedure is applied to both simulated and experimental data. This includes hit clustering, track reconstruction using the Multi-Hough algorithm, removal of spurious tracks, and merging of two-dimensional projections into three-dimensional muon tracks. Event selection criteria are then defined to distinguish between single-muon and multimuon events.

An important component of the seasonal analysis is the determination of the effective atmospheric temperature. Temperature profiles are obtained from atmospheric data, and the effective temperature is calculated by weighting different atmospheric layers according to their contribution to muon production. This allows a comparison between temporal variations in atmospheric temperature and fluctuations in the observed muon rate.

The seasonal behavior of the muon flux is first presented using experimental data. The modulation amplitude and phase are extracted and compared with theoretical expectations.

The same analysis is then performed using Monte Carlo simulations to validate the interpretation and to assess systematic effects. Finally, the observed behavior is discussed in terms of the competition between meson decay and interaction processes in the atmosphere, providing a physical explanation for the measured correlation.

Through this approach, the chapter establishes a connection between atmospheric dynamics, air-shower physics, and the underground muon flux measured in the NOvA ND.

4.1 Cosmic Ray Simulation

A realistic modeling of the cosmic-ray muon flux is essential for interpreting the seasonal modulation observed in the NOvA ND. In the standard NOvA framework, cosmic-ray backgrounds are frequently simulated using the CRY [48] Monte Carlo generator. CRY produces secondary particles at a predefined altitude according to parameterized distributions derived from experimental measurements. While computationally efficient, this approach does not simulate the full development of an extensive air shower initiated by a primary cosmic ray.

In particular, CRY generates particles independently according to empirical distributions and does not model the initial primary interaction nor the correlated hadronic cascade that follows. As a consequence, correlations among secondary particles originating from the same primary are not preserved. This limitation becomes critical when studying multiple-muon events, since such events arise from a single primary cosmic ray producing several correlated secondaries in the same air shower. Because CRY lacks this capability, it is not suitable for analyses involving muon multiplicity and inter-particle correlations.

For this reason, we employ the CORSIKA [49] (COsmic Ray SImulations for KAscade) Monte Carlo framework. CORSIKA provides a detailed simulation of extensive air showers in the Earth's atmosphere, beginning with the interaction of primary cosmic rays and tracking all relevant secondary particles through subsequent hadronic and electromagnetic processes. The program includes multiple interaction models appropriate for different energy ranges and provides a realistic description of meson production, decay, and propagation.

In this work, Corsika-v77410 was used with the default high-energy hadronic interaction model QGSJET and the low-energy model GHEISHA. These models together describe particle production and interactions over the full energy range relevant to this analysis.

```
RUNNR   1                               run number
EVTNR   1                               number of first shower event
NSHOW   2000000                         number of showers to generate
PRMPAR  14                              particle type of prim. particle
ESLOPE  -2.7                            slope of primary energy spectrum
                                        (E^(-gamma))
ERANGE  200  1.E6                       energy range of primary particle (GeV)
THETAP  0.  90.                         range of zenith angle (degree)
PHIP    -180.  180.                     range of azimuth angle (degree)
SEED    1   0   0                       seed for 1. random number sequence
SEED    2   0   0                       seed for 2. random number sequence
QGSJET  T   0
QGSSIG  T
OBSLEV  190.E2                          altitude of observation level (cm)
MAGNET  19.2733  49.7109                earth magnetic field at FermiLab
                                        (micro tesla)
HADFLG  0  0  0  0  0  2                flags hadr.interact.&fragmentation
ECUTS   50 50 1.E3 1.E3                 energy cuts (GeV) for particles
                                        (for hadrons, muons, electrons,
```

```
                                               photons)
MUADDI  T                                      additional info for muons
MUMULT  T                                      muon multiple scattering angle
ELMFLG  T   T                                  em. interaction flags (NKG,EGS)
STEPFC  1.0                                    mult. scattering step length fact.
RADNKG  200.E2                                 outer radius for NKG lat.dens.distr.
ARRANG  0                                      rotation of array to north
MAXPRT  1                                      max. number of printed events
ATMFILE /home/amit/corsika/Ana/ATM012017.txt    atmospheric file for Fermilab
DIRECT  /home/amit/corsika/data/               output directory
DATDIR  ./
USER    apal                                   user
DEBUG   F  6  F  1000000                       debug flag and log.unit for out
EXIT                                           terminates input
```

Table A: Input parameters for cosmic ray simulation

4.1.1 Primary Composition and Energy Spectrum

The cosmic-ray flux incident on the Earth's atmosphere consists predominantly of protons, which account for approximately 79% of primaries. Helium nuclei contribute roughly 20%, while heavier nuclei such as carbon, silicon, and iron form the remaining component. Since heavier nuclei produce air showers with different multiplicities and energy distributions, it is important to include representative species in the simulation.

Accordingly, primary showers were generated for protons, ^{4}He, ^{12}C, and ^{56}Fe nuclei. The differential energy spectrum of primaries was modeled as a power law:

$$\frac{dN}{dE_0} \propto E_0^{-\gamma}, \tag{4.1}$$

where E_0 denotes the primary energy and γ is the spectral index. Experimental measurements indicate $\gamma \approx 2.68 \pm 0.02$ below 10^{15} eV [13]. For simplicity and consistency with previous studies, a value of $\gamma = 2.7$ was adopted.

The primary energy range was chosen between 200 GeV and 10^6 GeV. The lower bound is motivated by the depth of the NOvA ND, which is located approximately 110 m underground. A muon must possess a minimum energy of about 54 GeV to penetrate the rock overburden and reach the detector [46]. Simulations indicate that primary energies below 200 GeV do not produce multimuon events at detector depth.

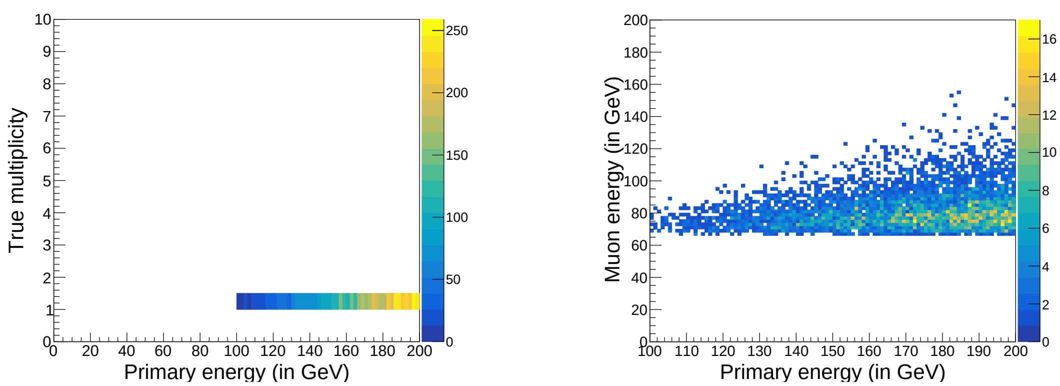

Figure 4.1: Muon multiplicity (left) and muon energy (right) as a function of primary energy for 100–200 GeV primaries. No multimuon events are observed in this range.

To verify this threshold, five million proton showers were simulated in the energy range 100–200 GeV with a muon cutoff of 50 GeV. Although 5539 muons satisfied the energy requirement, no multimuon events were observed. This confirms that multimuon production becomes significant only above 200 GeV. The upper bound of 10^6 GeV ensures coverage of high-energy primaries capable of generating large muon multiplicities, even though such events are rare due to the steep $E^{-2.7}$ spectrum. The flux falls by approximately a factor of

10^{10} in the energy range of 200 to 10^6 GeV.

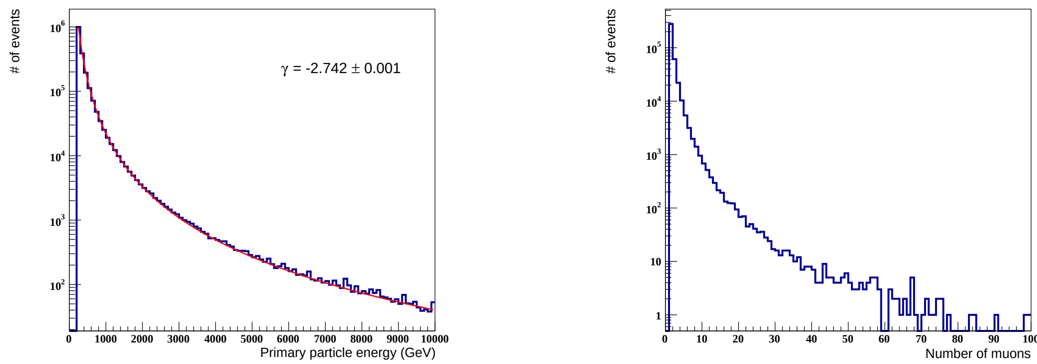

Figure 4.2: Primary energy distribution (left) and corresponding muon multiplicity (right) from simulated proton showers.

Primary directions were generated isotropically, with azimuthal angles in the range $[-180°, 180°]$ and zenith angles between $0°$ and $90°$.

4.1.2 Energy Thresholds and Observation Conditions

Since the primary focus of this analysis is muons, tracking thresholds were chosen to optimize computational efficiency while preserving physical accuracy. Muons and hadrons were tracked down to 50 GeV to allow hadronic decays into muons to be properly simulated. In contrast, electrons and photons were assigned a high cutoff energy of 10^3 GeV to suppress electromagnetic shower development, which is not relevant for the present study.

The Earth's magnetic field influences the propagation of charged particles in the atmosphere. The geomagnetic field at Fermilab (41.84°N, 88.25°W) was included using NOAA values [50], with horizontal and vertical components of 19.27 μT and 49.71 μT, respectively. The observation level in CORSIKA was set to the Fermilab altitude of approximately 190 m above sea level.

The CORSIKA binary output files contain detailed information about run parameters, event headers, and all particles reaching the observation level. These files were converted

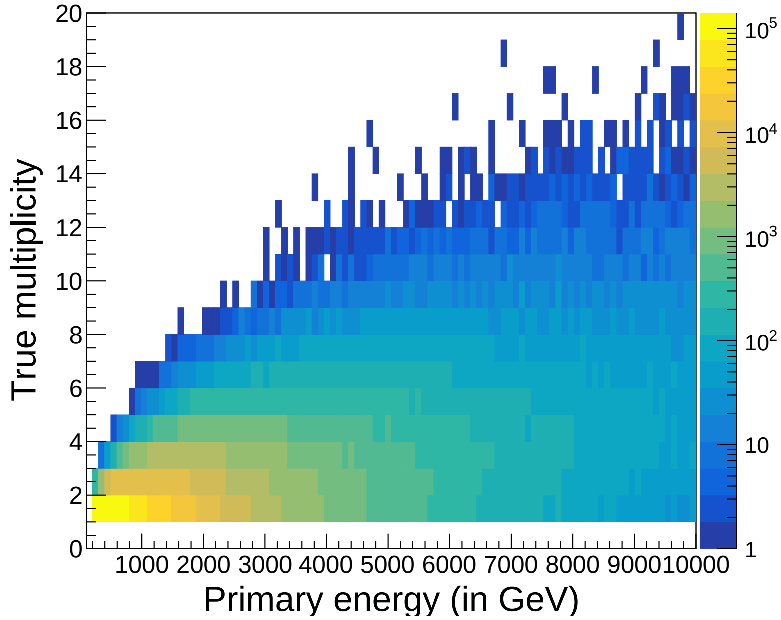

Figure 4.3: True muon multiplicity as a function of primary energy.

into ROOT ntuples to facilitate further analysis.

The fitted primary spectrum yields an index of -2.742 ± 0.001, consistent with the input value. As expected, the average muon multiplicity increases with primary energy (Fig. 4.3), reflecting the larger number of secondary mesons produced in high-energy interactions.

Fig. 4.4 shows that multimuon events originate from significantly higher-energy primaries compared to single-muon events.

4.1.3 Muon Kinematics and Production Mechanisms

The kinematic properties of simulated muons are shown in Figs. 4.5 and 4.6. The CORSIKA coordinate system was transformed into the NOvA convention, where the Y-axis is vertically upward. Since only downward-going muons are relevant for seasonal studies, the Y-component of the momentum is negative, while X and Z components are symmetric.

Approximately 94.7% of muons originate from pion decays, with a smaller fraction

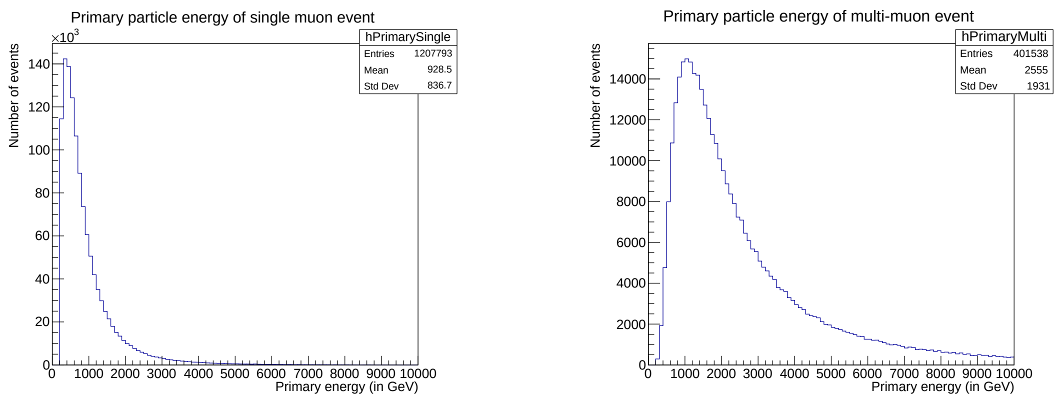

Figure 4.4: Primary energy distributions for single- and multimuon events.

from kaon decays. Most muons are produced at altitudes around 20 km, corresponding to the typical decay height of mesons in the upper atmosphere.

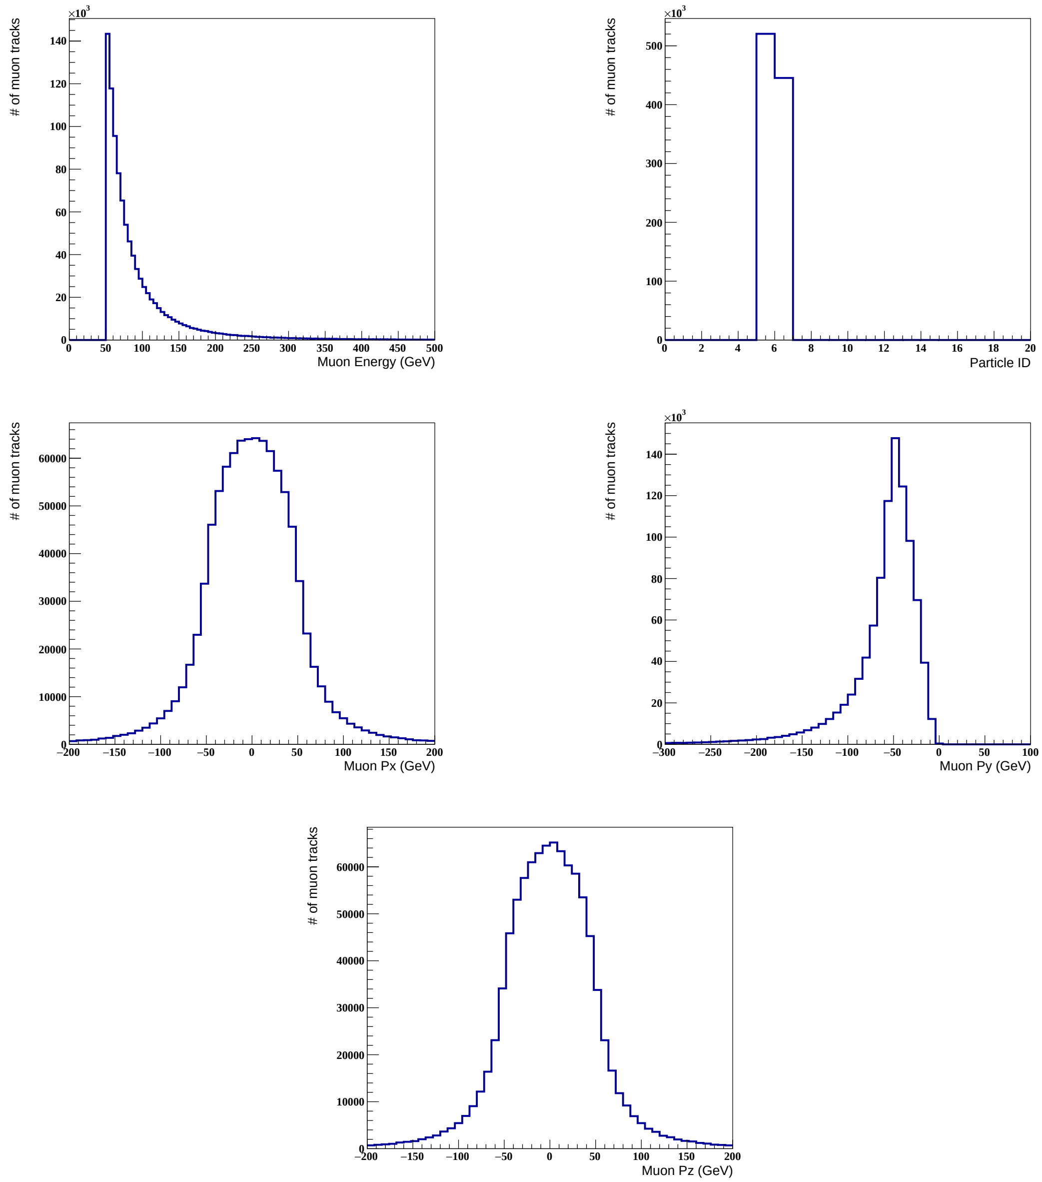

Figure 4.5: Muon energy, particle identification, and momentum components in NOvA coordinates.

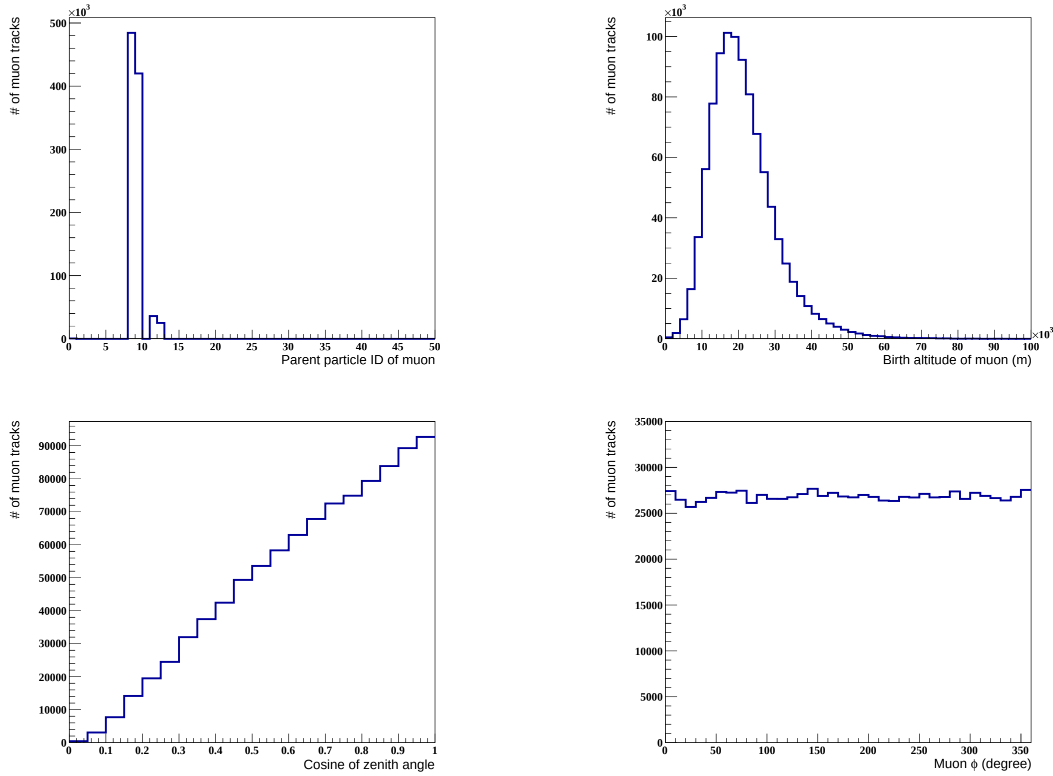

Figure 4.6: Muon parent particle, production altitude, zenith angle, and azimuthal angle distributions.

4.1.4 Spatial Randomization of Muons

CORSIKA simulates extensive air showers at a fixed reference position and provides the positions and momenta of particles at the observation level. However, in reality, cosmic-ray showers occur uniformly over the Earth's surface. To reproduce the correct geometric acceptance of the NOvA ND, it is therefore necessary to randomize the spatial positions of simulated muons before passing them to the detector simulation.

The simplest approach would be to distribute muons uniformly across the top surface of the detector. However, this would underestimate the contribution from inclined tracks entering through the sides of the detector. Muons with sufficiently large zenith angles can intersect the detector volume even if their vertical projection lies outside the physical footprint of the top surface. Consequently, the randomization region must extend beyond the detector boundaries.

To accommodate such trajectories, an enlarged rectangular region in the XZ plane was defined. The coordinates were chosen as

$$-3200 \text{ cm} < X < 3200 \text{ cm}, \qquad -3000 \text{ cm} < Z < 4600 \text{ cm}.$$

This extension corresponds to approximately 3000 cm beyond the detector dimensions in each horizontal direction. The asymmetry in the Z-range arises from the NOvA coordinate definition, where the Z-axis spans 0 to 1600 cm within the detector while the X-axis is centered symmetrically about zero.

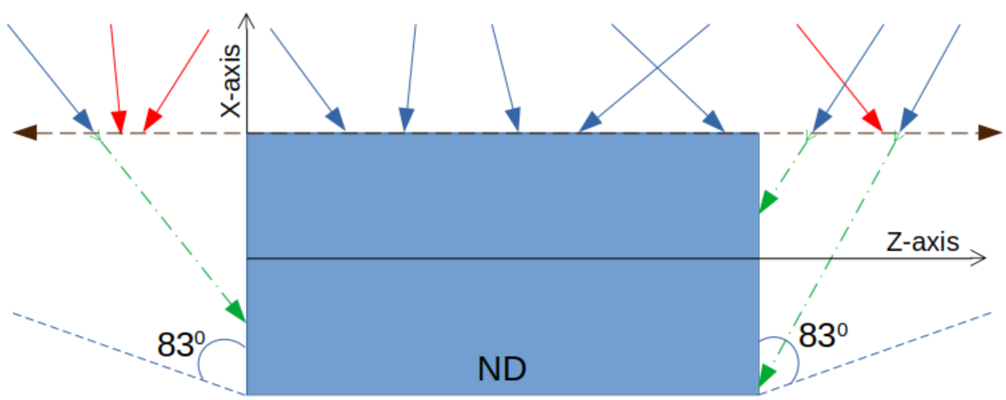

Figure 4.7: Illustration of the extended randomization region relative to the detector geometry. Inclined muons entering through the detector sides are accounted for by enlarging the injection area.

The maximum zenith angle for which a muon can still enter the lowest portion of the detector from the side is approximately

$$\theta_{\text{max}} = \tan^{-1}\left(\frac{3000}{400}\right) \approx 82.4^{\circ}.$$

By extending the randomization region accordingly, we ensure that muons with zenith angles up to roughly $83°$ are properly included in the acceptance.

For single-muon events, the muon position is drawn uniformly within this extended region. For multimuon events, the first muon is randomized uniformly, while subsequent muons retain their relative spatial separation and angular configuration as determined by CORSIKA. This procedure preserves the intrinsic shower geometry while embedding the event into a realistic detector acceptance.

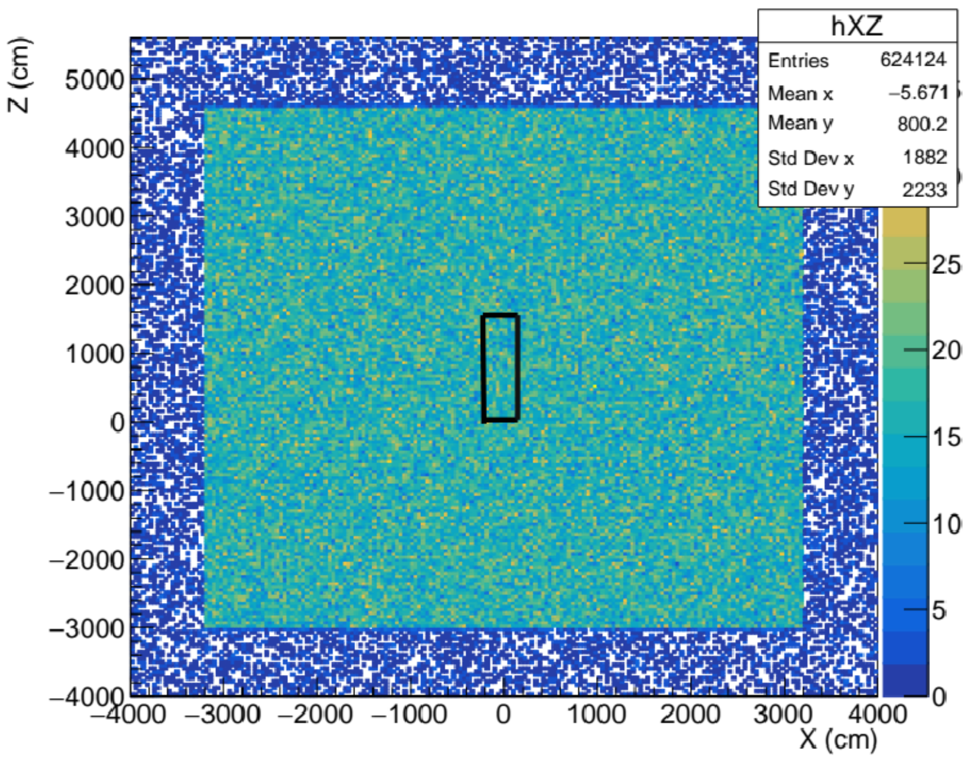

Figure 4.8: Muon impact positions in the XZ plane after randomization. The black rectangle denotes the detector footprint, while the shaded region represents the extended injection area.

The resulting spatial distribution of muons in the XZ plane is shown in Fig. 4.8. While most muons fall within the defined randomization region, some extend beyond it in multimuon events due to preserved shower geometry.

This randomization procedure ensures that the simulated flux accurately reproduces the angular acceptance and geometrical exposure of the detector.

4.1.5 Energy Scaling Through Rock Overburden

The NOvA ND is located approximately 110 m underground. Muons reaching the detector must therefore traverse a substantial thickness of rock before entering the active volume. Simulating this overburden directly within the GEANT [51, 52, 53] framework would significantly increase computational cost. Instead, a semi-analytic treatment of muon energy loss in rock is employed.

The average rate of energy loss of a muon traversing matter is given by [**?**]

$$-\frac{dE_\mu}{dl} = a + bE_\mu, \tag{4.2}$$

where l denotes the thickness of material (expressed in g/cm^2), a represents ionization energy loss, and bE_μ accounts for radiative processes including bremsstrahlung, pair production, and photonuclear interactions.

This first-order differential equation can be integrated to relate the muon energy at the surface ($E_{\mu,0}$) to the energy after traversing a thickness l:

$$E_\mu = (E_{\mu,0} + \epsilon)e^{-bl} - \epsilon, \tag{4.3}$$

where

$$\epsilon = \frac{a}{b}$$

is the critical energy below which ionization losses dominate over radiative processes.

For standard rock, $\epsilon \approx 500$ GeV and $b^{-1} \approx 2.5\times10^5$ g/cm^2 [16]. Using these parameters and the known overburden thickness, the residual muon energy at detector depth can be calculated analytically.

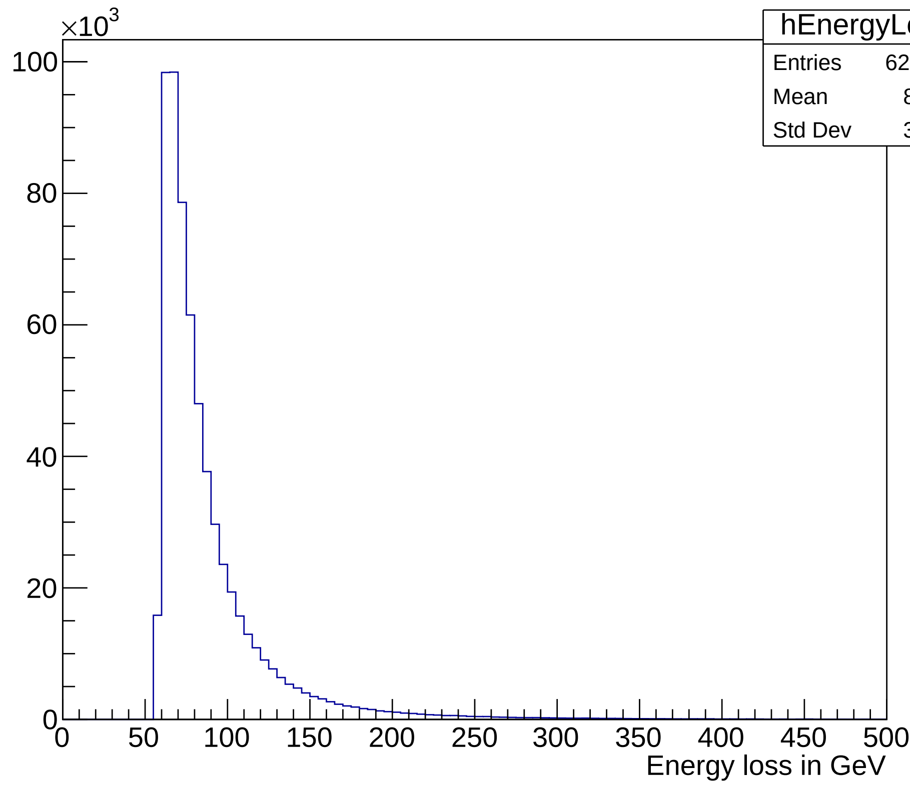

Figure 4.9: Muon energy loss after traversing 110 m of rock. Only muons above approximately 50 GeV at the surface can reach the detector.

The energy loss depends on the effective path length through rock, which increases with zenith angle as $l \propto 1/\cos\theta$. Consequently, highly inclined muons experience larger attenuation. However, as shown in the zenith distribution (Fig. 4.6), the flux decreases toward large angles because of their large path length and the rapidly falling of flux ($E^{-3.7}$), limiting the overall impact of extreme path lengths.

Since the seasonal analysis focuses on flux variations rather than detailed angular deflections, multiple Coulomb scattering within the rock is neglected. This approximation is justified because scattering effects are not expected to introduce seasonal biases.

This energy scaling procedure provides an efficient and physically well-motivated approximation of underground muon propagation.

4.1.6 Simulation of Event Timing

After spatial randomization and energy scaling, the final muon information is stored in HEPEVT format for input to the detector simulation. In real data, the classification of events into single- and multimuon categories depends on the temporal separation between reconstructed tracks. Specifically, if an additional muon is detected within a 50 μs window following the first, the event is classified as a multimuon event.

In CORSIKA simulations, the multiplicity of each event is intrinsically known, since all secondaries from a given primary are generated together. However, the GEANT-based detector simulation requires explicit time stamps for each particle in order to model detector response and readout windows correctly.

To emulate realistic timing conditions, each muon is assigned a random time uniformly distributed between 0 and 50 μs. This interval matches the DDActivity1 readout window used in data acquisition. By assigning random times within this window, we ensure compatibility with detector simulation requirements while preserving the physical multiplicity information provided by CORSIKA.

The resulting HEPEVT file is processed through the TextFileGen module (configured via `prodtextgen.fcl`), which injects the muons into the full GEANT detector simulation. This step models the response of both the active detector volume and the muon catcher system, producing simulated raw data comparable to experimental observations.

An example of a simulated two-muon event is shown in Fig. 4.10, where the color scale corresponds to the recorded ADC values.

This completes the chain from primary cosmic-ray generation in the atmosphere to a fully simulated detector response, providing a consistent framework for studying seasonal variations in muon flux.

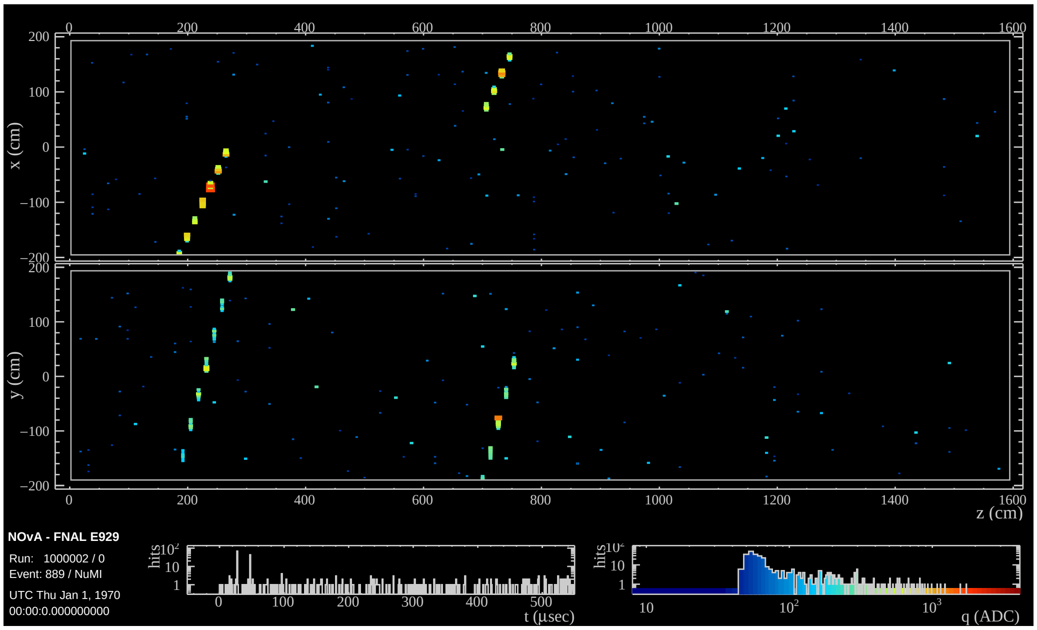

Figure 4.10: Event display of a simulated two-muon event. Color indicates the recorded ADC signal.

4.2 Modeling of Seasonal Variation in CORSIKA

Seasonal modulation of the underground muon flux originates from temperature-dependent changes in the atmospheric density profile. Since the development of an extensive air shower (EAS) depends sensitively on the density of the medium through which secondary particles propagate, an accurate modeling of the atmosphere is essential for any quantitative study of seasonal effects.

When a primary cosmic ray interacts with an atmospheric nucleus, typically in the upper stratosphere, it initiates a cascade of secondary hadrons and electromagnetic particles. The subsequent evolution of the shower is governed by the amount of matter traversed by the secondaries before decaying or interacting. This quantity is characterized by the *atmospheric depth*, which represents the integrated column density of air above a given altitude [54].

The atmospheric depth $X(h)$ at altitude h is defined as

$$X(h) = \int_h^\infty \rho(h')\, dh', \tag{4.4}$$

where $\rho(h')$ is the atmospheric density at height h'. Physically, $X(h)$ corresponds to the total mass per unit area above altitude h. It is typically expressed in units of g/cm^2.

The probability that a secondary meson decays before interacting depends on the competition between its decay length and interaction length, both of which depend on the local density. Therefore, seasonal changes in atmospheric temperature—which modify the density profile—directly affect the meson decay probability and, consequently, the resulting muon flux at detector depth.

Atmospheric Parameterization in CORSIKA

CORSIKA models the atmosphere as a layered medium composed primarily of nitrogen (N_2), oxygen (O_2), and argon (Ar), with volume fractions of 78.1%, 21.0%, and 0.9%, respectively. The vertical density profile is approximated using five atmospheric layers.

For the first four layers ($h < 100$ km), the mass overburden is parameterized by an exponential form:

$$X(h) = a_i + b_i e^{-h/c_i}, \tag{4.5}$$

where a_i, b_i, and c_i are layer-dependent parameters.

For the fifth layer ($h > 100$ km), the mass overburden decreases linearly with height:

$$X(h) = a_5 - b_5 \frac{h}{c_5}. \tag{4.6}$$

The parameter c_i has a direct physical interpretation: it is related to the atmospheric scale height and depends on temperature through

$$c_i = \frac{kT}{Mg}, \tag{4.7}$$

where k is the Boltzmann constant, T is the atmospheric temperature, M is the mean molecular mass of air, and g is the acceleration due to gravity.

Thus, c_i increases with temperature, reflecting the expansion of the atmosphere during warmer periods. As temperature rises, the atmosphere becomes less dense at a given altitude but extends to greater heights. The parameters a_i and b_i are chosen such that $X(h)$ remains continuous across the boundaries between adjacent layers. Because temperature varies seasonally, these parameters change accordingly, leading to a seasonal modulation of the atmospheric depth profile.

Most high-energy muons detected underground originate from the decay of charged pions and kaons produced in the upper atmosphere, typically between 10 and 40 km above sea level. In this region, the competition between meson interaction and decay is particularly sensitive to density variations.

Accurately reproducing this density variation in simulation is therefore essential for modeling the seasonal behavior observed in data.

Atmospheric Models in CORSIKA

CORSIKA includes several predefined atmospheric models, such as the U.S. Standard Atmosphere, monthly models for Central Europe, and monthly models for the South Pole. However, these default parameterizations do not include monthly atmospheric profiles specific to the Fermilab location.

To obtain realistic atmospheric profiles for Fermilab (41.84°N, 88.25°W), we use data from the Global Data Assimilation System (GDAS) [55]. GDAS provides global atmospheric profiles as a function of time, including temperature, pressure, and density at multiple altitudes.

The `gdastool` script included in CORSIKA (located in `src/utils/`) allows automatic extraction of atmospheric data from the GDAS database and converts it into CORSIKA-compatible format. The output file contains:

- ATMLAY: atmospheric layer boundaries,
- ATMA: parameters a_i,
- ATMB: parameters b_i,
- ATMC: parameters c_i,
- refractive index profile as a function of altitude.

An example atmospheric file generated using `gdastool` is shown in Fig. 4.11. The first row specifies the altitude ranges of the five layers, followed by rows containing a_i, b_i, and c_i. The lower section provides the refractive index profile.

```
# atmospheric parameters ATMLAY, A, B, C respectively
 0.00000000E+00  3.57100049E+05  9.18401848E+05  2.62664147E+06  1.00000000E+07
-1.76494084E+01 -1.19258722E+02  8.32049076E-01  5.76602088E-04  1.12829200E-02
 1.07057266E+03  1.16100574E+03  1.22133018E+03  1.11578929E+03  1.00000000E+00
 8.16709396E+05  9.36735836E+05  6.78458212E+05  7.00635217E+05  1.00000000E+09
# atmospheric height [m] and refractive index columns
-1.00025332E+03  1.000349261805015E+00
-9.99253315E+02  1.000349219043119E+00
-9.98253315E+02  1.000349176286458E+00
-9.97253315E+02  1.000349133535033E+00
-9.96253315E+02  1.000349090788841E+00
-9.95253315E+02  1.000349048047884E+00
-9.94253315E+02  1.000349005312159E+00
-9.93253315E+02  1.000348962581667E+00
-9.92253315E+02  1.000348919856406E+00
-9.91253315E+02  1.000348877136376E+00
-9.90253315E+02  1.000348834421577E+00
-9.89253315E+02  1.000348791712008E+00
-9.88253315E+02  1.000348749007668E+00
-9.87253315E+02  1.000348706308556E+00
-9.86253315E+02  1.000348663614672E+00
-9.85253315E+02  1.000348620926015E+00
-9.84253315E+02  1.000348578242585E+00
```

Figure 4.11: One example of the atmospheric file: The first row provides a range of atmospheric layers in cm, and the next three give a_i, b_i, and c_i parameters for those five layers. After that, it shows atmospheric height and its corresponding refractive index.

Since GDAS uses longitudes defined between $0°$ and $360°$, the Fermilab longitude of $88.25°$W must be converted to $360° - 88.25° = 271.75°$.

Atmospheric files for each day of 2017 were extracted using the command:

python gdastool -t 1483282800 -c 41.84 271.75 -o ATMOS012017.DAT

Daily atmospheric parameters within each month were averaged to produce a single representative monthly atmospheric file. These files were then supplied to CORSIKA using the `ATMFILE` keyword, allowing shower simulations to incorporate realistic seasonal atmospheric conditions.

4.3 Track Reconstruction

In the NOvA detector, a recorded event consists of a collection of cell hits stored following a trigger issued by the data acquisition (DAQ) system. These hits represent localized energy depositions in the detector and form the basic input for event reconstruction. The first task in the reconstruction chain is to identify geometrical patterns within this collection of hits, particularly straight-line trajectories characteristic of cosmic-ray muons.

Because a through-going muon produces approximately linear patterns in both orthogonal detector views (XZ and YZ), a line-finding technique is well suited for this purpose. For this analysis, a modified Hough Transform [56, 57] algorithm is employed. The Hough Transform is a well-established method in image analysis for identifying parametric shapes, especially straight lines, in a set of discrete points. Its robustness against noise makes it particularly effective for cosmic-ray reconstruction in a scintillator-based detector.

Hough Transform Formalism

Instead of describing a line in slope-intercept form, which cannot represent vertical lines, the Hough Transform uses a polar parameterization. In this representation, a straight line in the (x, y) plane is expressed as

$$\rho(\theta) = x \cos\theta + y \sin\theta, \tag{4.8}$$

where ρ is the perpendicular distance from the origin to the line, and θ is the angle between this perpendicular direction and the x-axis.

Each detector hit at position (x, y) corresponds to a curve in the (ρ, θ) parameter space. If multiple hits lie on the same physical line, the corresponding curves intersect at a common point (ρ_0, θ_0). Therefore, identifying regions of high intersection density in the Hough parameter space allows one to determine the best candidate lines in the original coordinate system.

Modified Hough Implementation in NOvA

For the ND, the Hough Transform is applied independently to each detector view [58, 59]. Instead of mapping single hits, the implementation considers all possible hit pairs and determines the parameters (ρ_0, θ_0) corresponding to the line connecting each pair.

Each candidate line contributes to the Hough accumulator using a Gaussian-weighted smearing:

$$w = \exp\left(-\frac{(\rho - \rho_0)^2}{2\sigma_\rho^2}\right) \exp\left(-\frac{(\theta - \theta_0)^2}{2\sigma_\theta^2}\right), \tag{4.9}$$

with resolutions defined as

$$\sigma_\rho = \frac{3\ \text{cm}}{\sqrt{12}}, \tag{4.10}$$

$$\sigma_\theta = \frac{3\ \text{cm}}{d\sqrt{6}}, \tag{4.11}$$

where d denotes the spatial separation between the two hits forming the pair, and the factor of 3 cm corresponds to the detector cell width.

To reduce computational overhead and suppress unphysical solutions, two geometric constraints are imposed:

- Only hit pairs separated by $d \leq \sqrt{15000}$ cm are considered.
- For hit pairs aligned horizontally within a view, a minimum separation of $0.25\sqrt{15000}$ cm is required.

The latter condition mitigates artificial accumulation of horizontal Hough solutions arising from detector geometry rather than true particle trajectories.

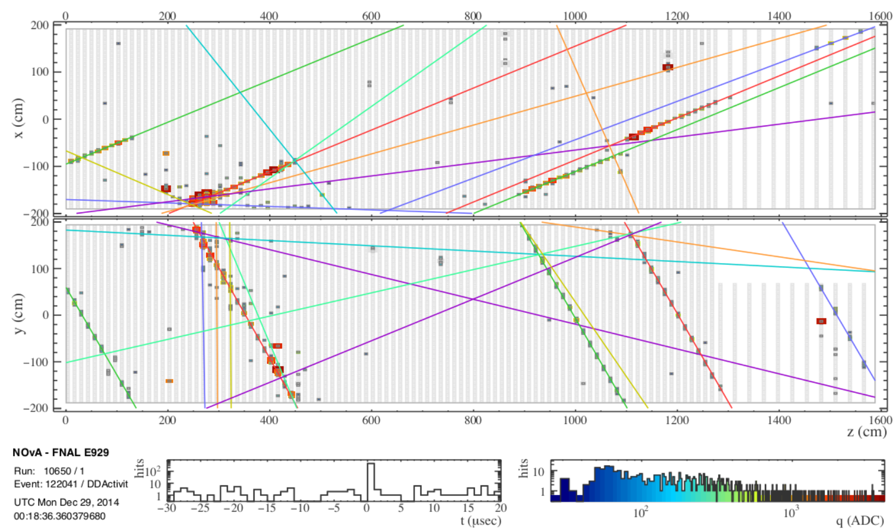

Figure 4.12: Five-muon event recorded in the NOvA ND. The colored lines represent track candidates identified by the Multi-Hough algorithm in each detector view [13].

Fig. 4.12 illustrates an example multi-muon event together with the reconstructed Hough lines in both views.

Iterative Multi-Hough Procedure

To reconstruct multiple tracks within the same event, the Hough procedure is applied iteratively. After identifying the strongest peak in the Hough map, the hits associated with that line are removed from the event. A new Hough map is then constructed using the remaining hits. This process continues until either no peak exceeds a predefined threshold or a maximum number of iterations is reached.

2D Track Formation

Following the iterative Hough procedure, final 2D tracks are constructed by selecting hits whose perpendicular distance from a given Hough line does not exceed 10 cm. The resulting 2D tracks in each detector view are ordered by decreasing number of associated hits.

Removal of Spatially Isolated Hits

A genuine charged-particle trajectory produces hits that are contiguous in space. Spurious noise hits may appear far from the main cluster and must be rejected. The maximum allowed gap between consecutive hits depends on the track inclination and is defined as

$$\Delta r_{\text{max}} = d_{\text{gap}} + d_{\text{gap}} \sin\theta, \tag{4.12}$$

where $d_{\text{gap}} = 70$ cm and θ represents the reconstructed track angle. This parameterization allows larger tolerances for shallow tracks, with

$$70 \text{ cm} \leq \Delta r_{\text{max}} \leq 140 \text{ cm}.$$

Suppression of Rogue 2D Tracks

Spurious 2D track candidates may arise due to noise or partial overlaps with true trajectories. To suppress such solutions, tracks within a given view are examined pairwise.

A boolean matrix H of dimension $k \times k$ is constructed, where k is the number of reconstructed 2D tracks in that view. Tracks are ordered by decreasing hit multiplicity. For each pair (n, m) with $n < m$, overlap is tested by comparing all hit combinations. If

$$\Delta Z = |Z_n^i - Z_m^j| < 10 \text{ cm}, \tag{4.13}$$

for any pair of hits, the shorter track is flagged as overlapping and considered non-physical.

The longest track is always retained. Additional tracks are accepted only if they exhibit no overlap with previously accepted tracks. This ensures that:

- At least one reconstructable track exists per triggered event,
- The most significant track is preserved,
- Multiple tracks in genuine multi-muon events remain spatially separated.

Any residual false 2D tracks typically fail the subsequent 3D merging stage.

2D Track Merging into 3D Tracks

After cleaning both views independently, remaining 2D tracks are compared between XZ and YZ views. Since hit timing is not sufficiently reliable to determine directionality, hits are ordered by increasing Z coordinate.

For a candidate pair $(track_n^{XZ}, track_m^{YZ})$, merging occurs if the following conditions are simultaneously satisfied:

$$\Delta startZ = |Z_0^{XZ}(n) - Z_0^{YZ}(m)| < 50 \text{ cm}, \tag{4.14}$$

$$\Delta endZ = |Z_{\text{last}}^{XZ}(n) - Z_{\text{last}}^{YZ}(m)| < 50 \text{ cm}, \tag{4.15}$$

$$N_{\text{hits}}^{XZ}(n) + N_{\text{hits}}^{YZ}(m) \geq 6. \tag{4.16}$$

When these criteria are met, the two 2D tracks are combined to form a 3D track, and neither can participate in further combinations. The procedure terminates when no valid pairs remain.

Finally, the reconstructed 3D track endpoints are defined from the ordered hit list. The starting point is assigned to the hit with the larger Y coordinate, corresponding to a downward-going muon assumption. Upward-going muons are not considered in this analysis.

4.4 Selection Cuts

After reconstruction, additional quality requirements are imposed to define a clean and well-understood sample of single- and multi-muon events. Although the reconstruction algorithm identifies candidate tracks, not all reconstructed objects correspond to well-measured through-going muons. Detector edge effects, limited geometric sampling, and accidental temporal overlaps can introduce ambiguities in classification.

To suppress these effects and improve the purity and angular resolution of the selected sample, three additional selection criteria are applied. These are summarized in Tab. 4.1.

4.4.1 Fiducial Volume Cut

The first requirement is a fiducial volume constraint. Cosmic muons traversing the rock above the detector may produce secondary low-energy electrons via bremsstrahlung. These particles can enter the detector near its boundaries and create short tracks parallel to the muon that contaminate the muon sample.

To reduce such edge-related backgrounds, each reconstructed track is required to begin and end at least 50 cm inside the detector boundaries. This ensures that only well-contained tracks are selected and removes ambiguities associated with partially contained or external activity.

The effect of this cut on the spatial distribution of tracks is shown in Fig. 4.13. The XY and YZ projections demonstrate the removal of tracks close to detector edges.

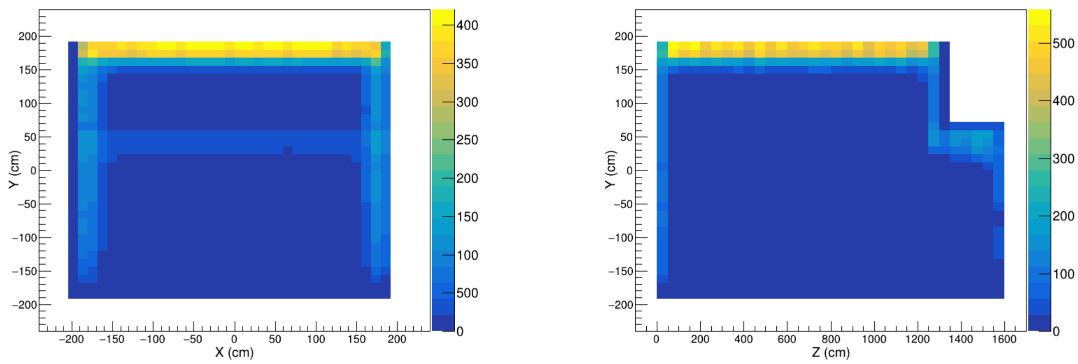

Figure 4.13: XY (left) and YZ (right) projections of muon tracks after applying the fiducial volume cut. The sharp truncation ND boundaries reflects the removal of edge-associated tracks.

The fiducial volume cut removes 16.91% of the data and 15.6% of the MC sample, indicating consistent modeling of boundary effects.

4.4.2 Number of Planes Cut

The second selection addresses reconstruction quality. The NOvA detector consists of alternating planes oriented along orthogonal directions. Reliable track reconstruction requires sufficient sampling across multiple planes.

Tracks traveling nearly parallel to the Y-axis or horizontally along the X-axis intersect only a small number of planes. As illustrated in Fig. 4.14, such tracks produce very few hits and are difficult to reconstruct accurately.

This insufficient geometric constraint leads to discrepancies between true and reconstructed direction cosines, particularly near zero values, as shown in Fig. 4.15.

These poorly reconstructed tracks also cause noticeable differences between data and MC before quality cuts, as shown in Figs. 4.19 and 4.20.

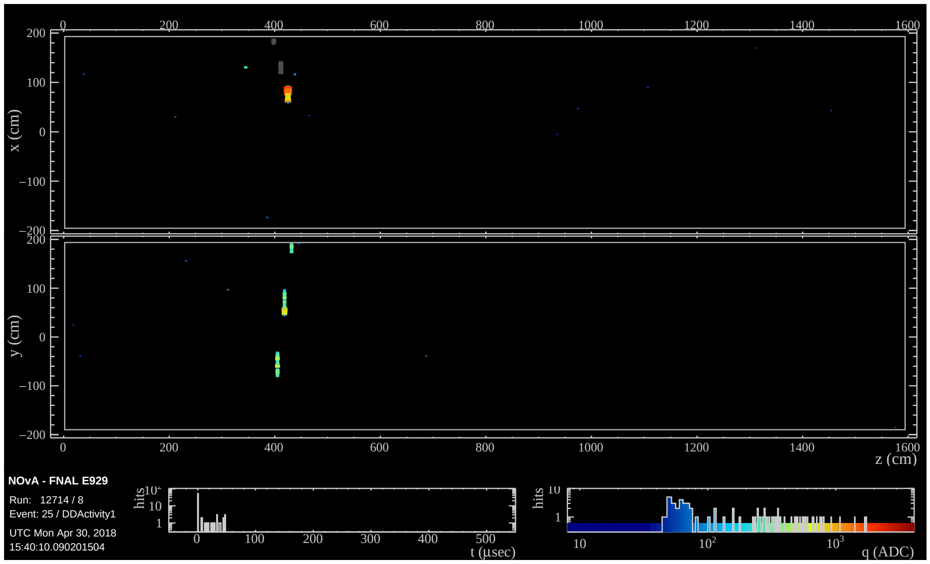

Figure 4.14: Event display for run 12714, event 25. The track passes through only four planes, making reconstruction unreliable.

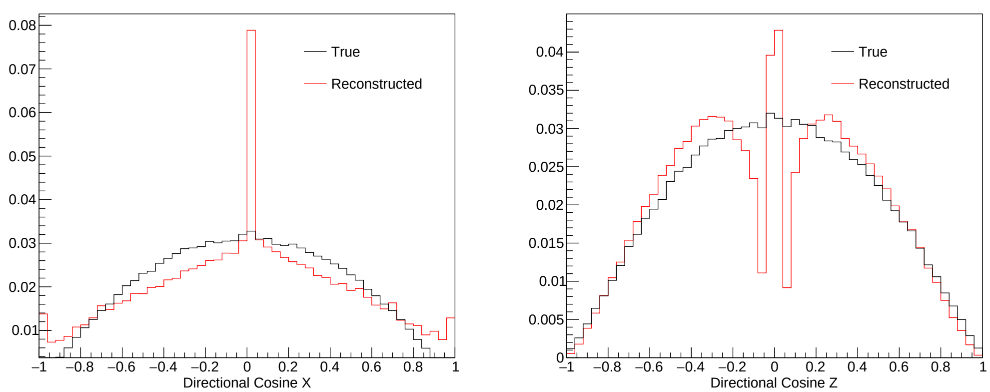

Figure 4.15: True vs reconstructed direction cosine in X (left) and Z (right). The discrepancy near zero arises from tracks crossing too few planes.

To improve angular resolution and remove reconstruction artifacts, a minimum requirement of

$$n_{\text{planes}} \geq 10$$

is imposed. After applying this cut, the mismatches near zero in the direction cosine distributions are significantly reduced.

The number-of-planes distribution and applied threshold are shown in Fig. 4.16.

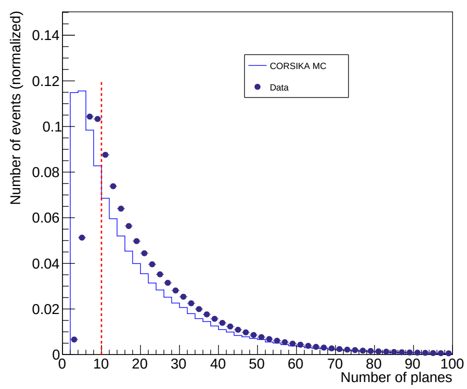

Figure 4.16: Distribution of the number of crossed planes for data and MC. The dashed line indicates the applied threshold at $n_{\text{planes}} = 10$.

This cut removes 19.5% of data and 22.6% of MC events.

4.4.3 Time Cut

The final selection ensures proper classification of multi-muon events. In rare cases, two independent single-muon events occurring close in time may be recorded within the same trigger window and incorrectly reconstructed as a multi-muon event.

An example is shown in Fig. 4.17, where two sequential single-muon events are stored as a double-muon topology.

To prevent such misclassification, the time difference between the first reconstructed track and any subsequent track in the same event is required to satisfy:

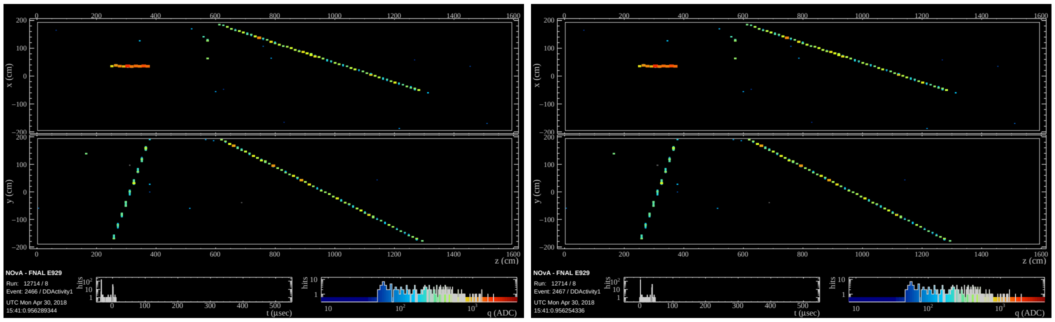

Figure 4.17: Two sequential events reconstructed as a double-muon topology but originating from distinct primaries.

$$\Delta t \leq 100 \text{ ns}.$$

The distribution of this time difference is shown in Fig. 4.18. The threshold removes accidental coincidences at large time separation.

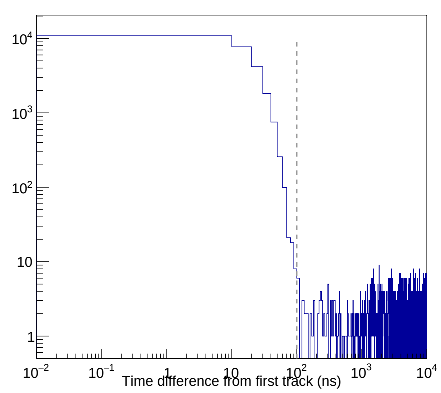

Figure 4.18: Distribution of time difference between the first track and subsequent tracks within the same event. The dashed line indicates the applied cut at 100 ns.

The time cut removes only 0.06% of data and does not affect MC events, as simulated events do not contain overlapping independent primaries.

Criteria	Value
Number of planes	$n_{\text{planes}} \geq 10$
Fiducial volume	$d \geq 50$ cm
Time separation	$\Delta t \leq 100$ ns

Table 4.1: Analysis selection criteria

4.4.4 Data–MC Comparison

A critical validation step in the analysis is verifying that the Monte Carlo simulation accurately reproduces detector observables. Data–MC comparisons are therefore performed both before and after applying the selection cuts.

Fig. 4.19 and 4.20 show comparisons prior to cuts. Noticeable discrepancies appear in direction cosines, number of planes, and related reconstruction-dependent variables. These differences are primarily driven by poorly reconstructed tracks and edge effects.

After applying all selection criteria, the comparisons shown in Fig. 4.21 and Fig. 4.22 demonstrate significantly improved agreement across spatial distributions, angular variables, track length, number of hits, and deposited energy. Both data and MC are normalized to unit area to emphasize shape agreement.

The excellent post-cut agreement confirms that the simulation reliably reproduces the detector response and that the selected muon sample is well understood. This validated sample is then used for the extraction of seasonal modulation parameters.

4.5 Multiple Muon Rate Calculation

Once multiple-muon events are identified and selected, the next step is to determine their rate as a function of time. A straightforward approach would be to divide the number of selected events in a given time interval by the elapsed live time of that interval. However, this naïve method introduces several subtle complications.

First, the numerator (event count) may include residual background events or misre-

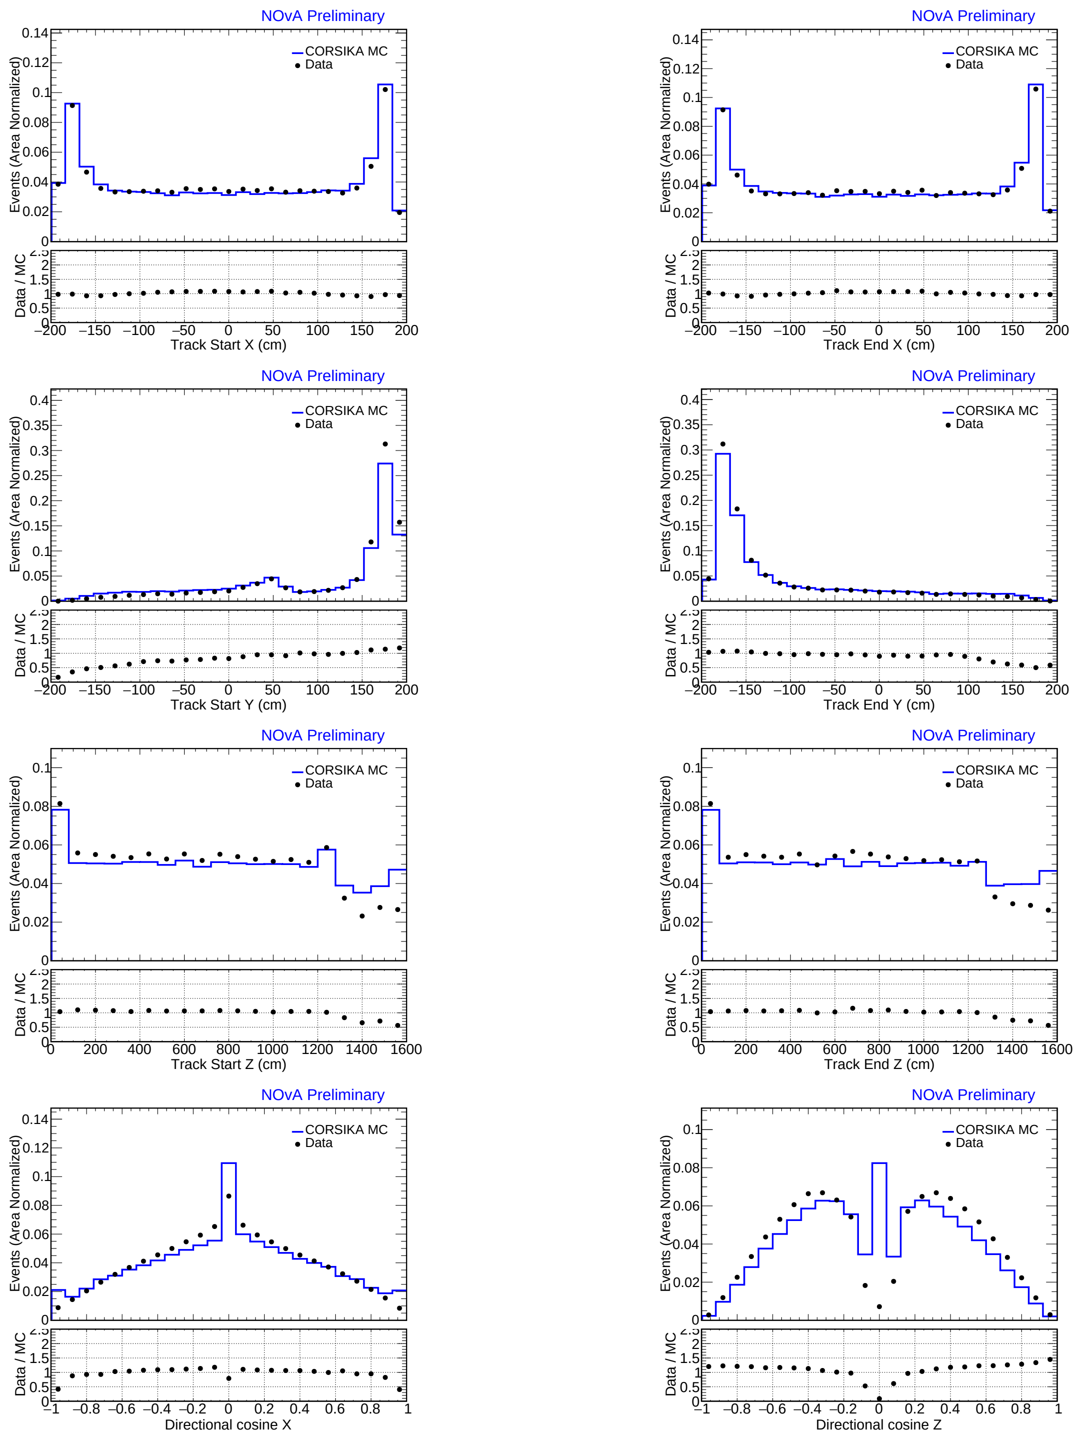

Figure 4.19: Data-MC comparison before cuts - Top three rows: the start (left) and end point (right) of muon tracks in the X, Y, and Z directions. The bottom rows: direction cosines in X and Z directions.

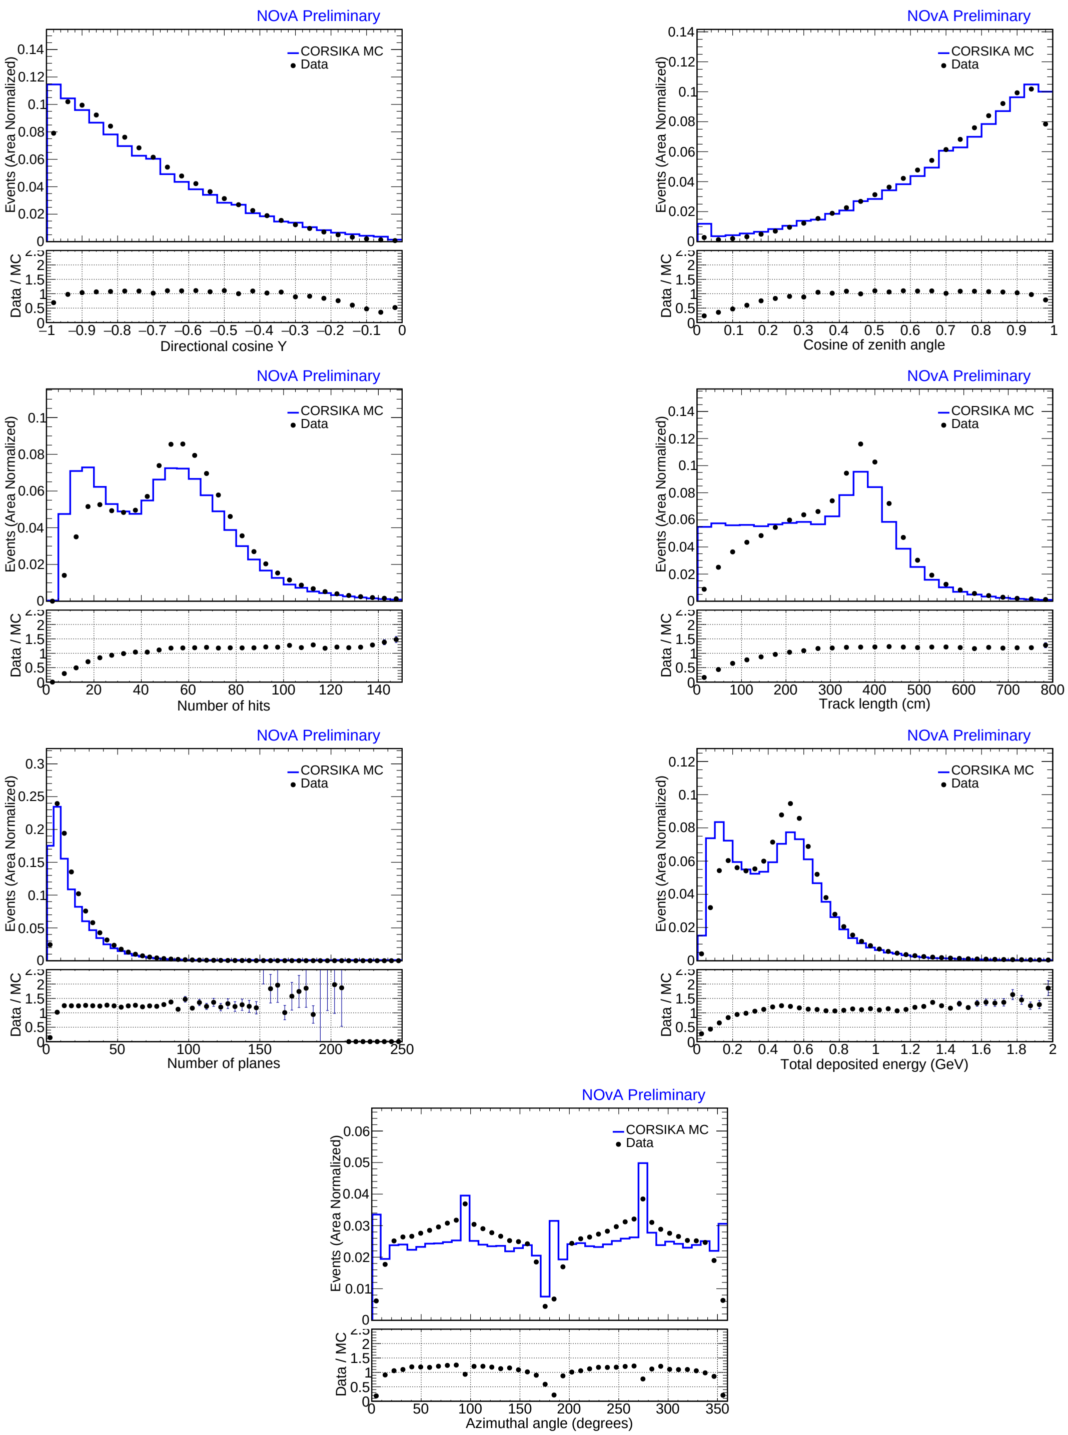

Figure 4.20: Data-MC comparison before cut - for every reconstructed muon, the top row: direction cosine in Y direction (left) and cosine of zenith angle (right), next row: the number of hits (left) and track length (right), then in the next row: its number of planes (left) and total energy deposited by the track (right), and the bottom row: azimuthal angle in the XZ plane with $0°$ corresponding to the +X direction.

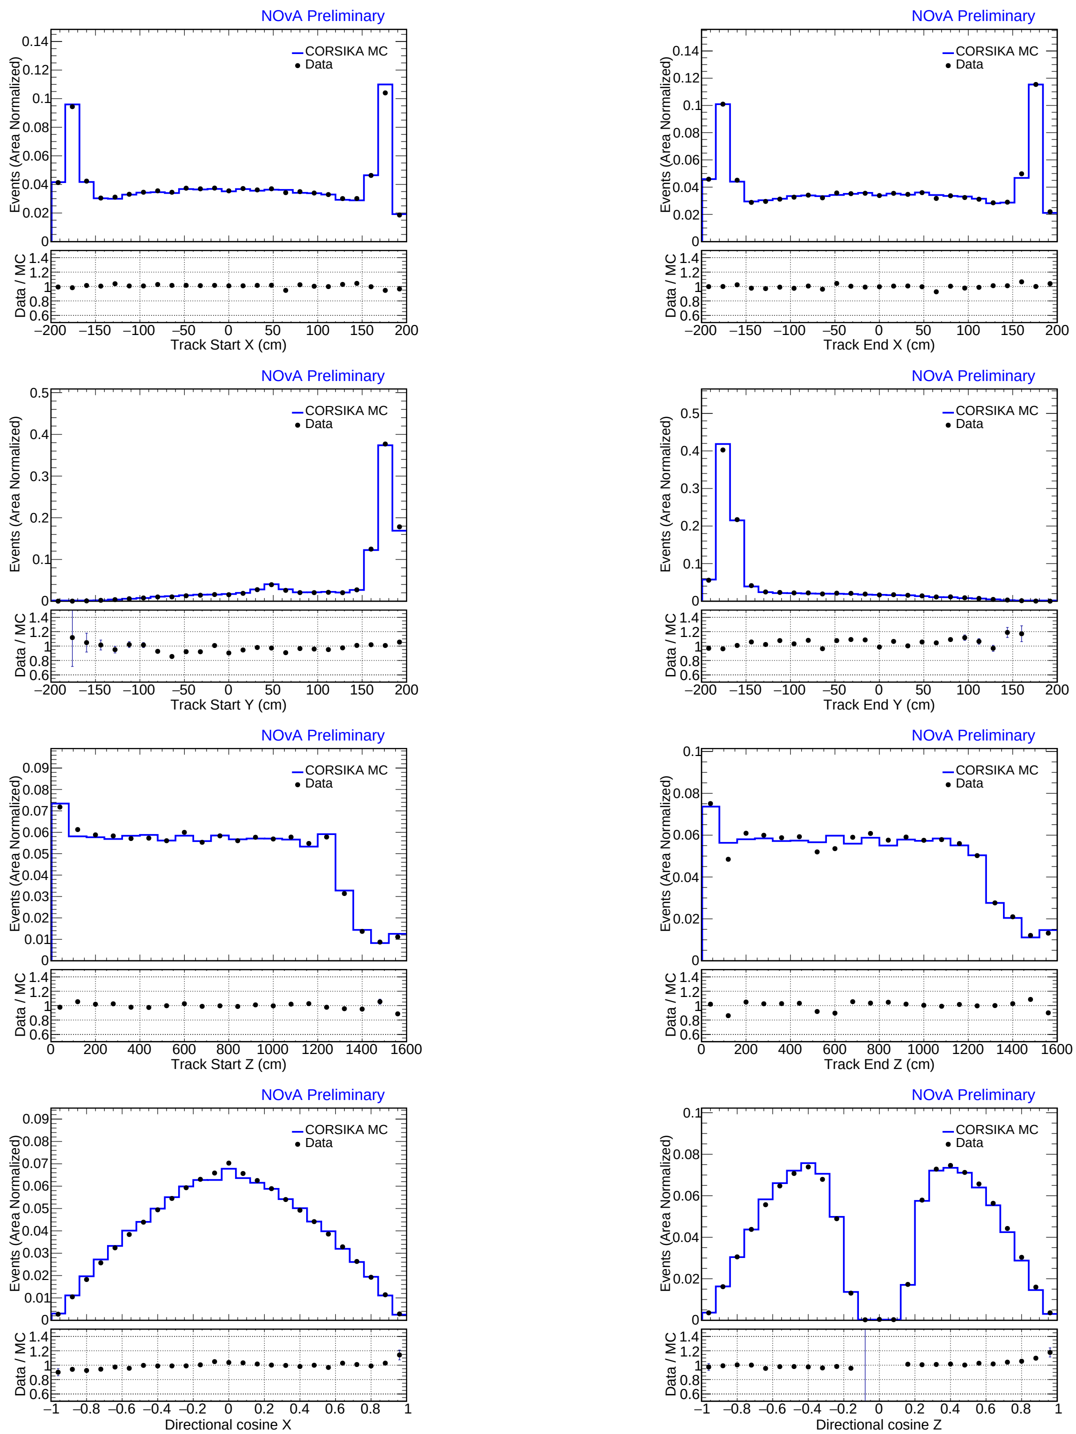

Figure 4.21: Data-MC comparison after cut - Top three rows: the start (left) and end point (right) of muon track in the X, Y, and Z directions. The bottom row: direction cosines in X and Z directions.

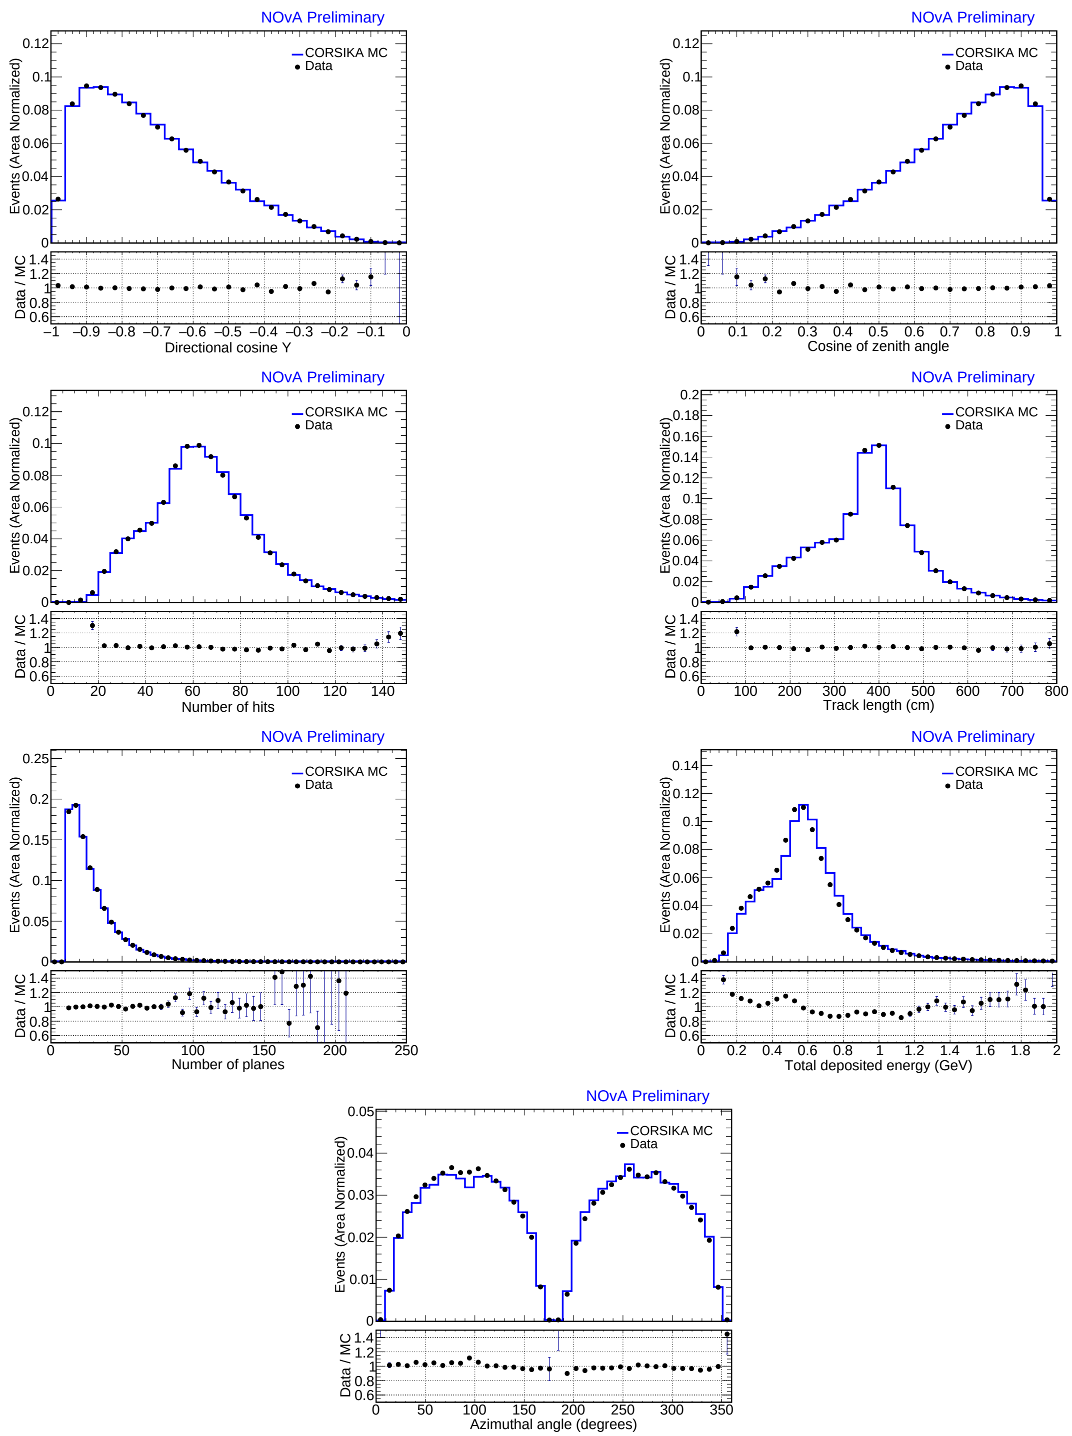

Figure 4.22: Data-MC comparison after cut - for every reconstructed muon, the top row: direction cosine in Y direction (left) and cosine of zenith angle (right), next row: the number of hits (left) and track length (right), then in the next row its number of planes (left) and total energy deposited by the track (right), and the bottom row: azimuthal angle in the XZ plane with 0° corresponding to the +X direction.

constructed tracks. Second, the denominator (live time) can be affected by unaccounted detector dead time or data acquisition (DAQ) interruptions. Furthermore, NOvA data files are not stored strictly in chronological order on disk. As a result, events that physically occurred in one time interval may be written into neighboring files, leading to distortions in simple count-based rate estimates.

To avoid these issues, a method based on interarrival time statistics is employed. This approach relies on the fundamental property that independent cosmic-ray showers arrive according to a Poisson process. For a Poisson process with rate λ, the time difference between two consecutive events follows an exponential distribution. This provides a statistically robust way to extract the event rate without explicit dependence on file boundaries or live-time bookkeeping.

4.5.1 Erlang Distribution in the Time Domain

For a Poisson process with rate λ, the waiting time until the k-th event follows an Erlang distribution [60] with integer shape parameter k and rate parameter λ. The probability density function (PDF) is

$$f_T(t; k, \lambda) = \frac{\lambda^k}{(k-1)!} t^{k-1} e^{-\lambda t}, \qquad t \geq 0. \tag{4.17}$$

In the present analysis, we are interested in the time difference Δt between two *consecutive* multiple-muon events. This corresponds to the $k = 1$ case, reducing the Erlang distribution to a simple exponential law:

$$f_T(\Delta t) = \lambda e^{-\lambda \Delta t}, \qquad \Delta t \geq 0. \tag{4.18}$$

Therefore, extracting the rate λ reduces to fitting the interarrival time distribution to an exponential function.

The interarrival times span several orders of magnitude due to occasional dead time and statistical fluctuations. To improve visualization and fitting stability, the distribution is expressed in terms of

$$x = \log_{10}(\Delta t), \tag{4.19}$$

so that

$$\Delta t = 10^x. \tag{4.20}$$

The Jacobian of the transformation is

$$\frac{d(\Delta t)}{dx} = \ln(10)\, 10^x. \tag{4.21}$$

The probability density in x is then obtained using the change-of-variable formula:

$$f_X(x) = f_T(\Delta t) \left| \frac{d(\Delta t)}{dx} \right|. \tag{4.22}$$

Substituting Eq. 4.18 gives

$$\begin{aligned} f_X(x) &= \lambda e^{-\lambda 10^x} \cdot (\ln(10)\, 10^x) && (4.23) \\ &= (\ln 10)\, \lambda\, 10^x\, e^{-\lambda 10^x}. && (4.24) \end{aligned}$$

Defining a normalization constant

$$A \equiv \ln(10), \tag{4.25}$$

and absorbing possible bin-width factors into A, the fitting function becomes

$$f_X(x) = A\, \lambda\, 10^x\, e^{-\lambda 10^x}. \tag{4.26}$$

This is the functional form used to extract the multiple-muon event rate from the $\log_{10}(\Delta t)$ distribution.

Since Δt is computed in milliseconds, the fitted λ has units of ms^{-1} and is subsequently converted to Hz.

Prior to constructing the Δt distribution, all selected multiple-muon events are reordered chronologically based on their arrival times. The time difference between event i and event $(i+1)$ is then calculated.

A representative fit for one day of data is shown in Fig. 4.23. The central region of the distribution follows the expected exponential behavior, while deviations appear at very small and very large Δt.

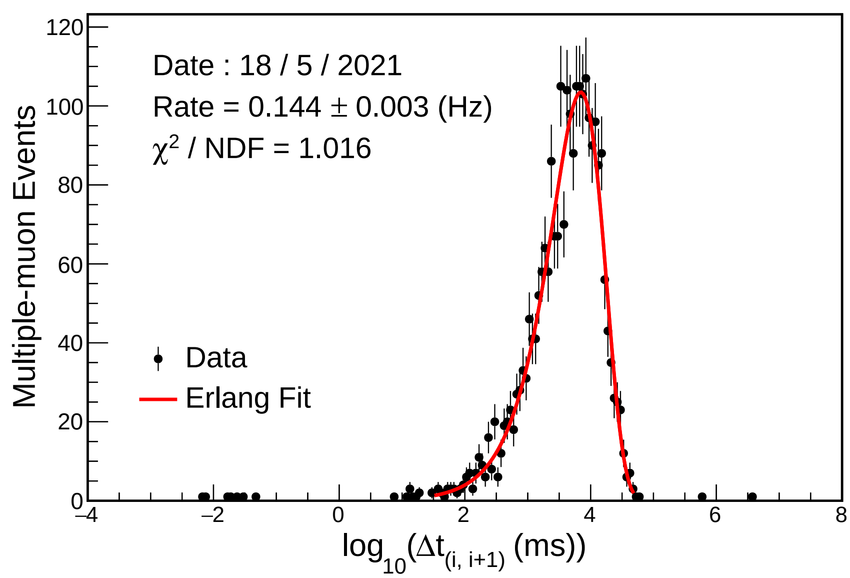

Figure 4.23: Logarithm (base 10) of the time difference (in ms) between consecutive multiple-muon events from one day of data. The red curve shows the Erlang ($k = 1$) fit in the range 1.5 to 4.7. The extracted rate for this day is 0.144 ± 0.003 Hz.

Very small Δt values (left side of the plot) may arise from split tracks or DAQ double triggers. Very large Δt values (right side) can result from detector dead time or temporary data acquisition interruptions. Because the fit is performed only in the central region where the distribution follows the exponential hypothesis, these pathological regions do not bias the extracted rate.

The goodness-of-fit for weekly binned data is summarized in Fig. 4.24. The χ^2/NDF distribution demonstrates that the exponential model provides an excellent description of the data over time, confirming that multiple-muon arrivals are consistent with a Poisson process.

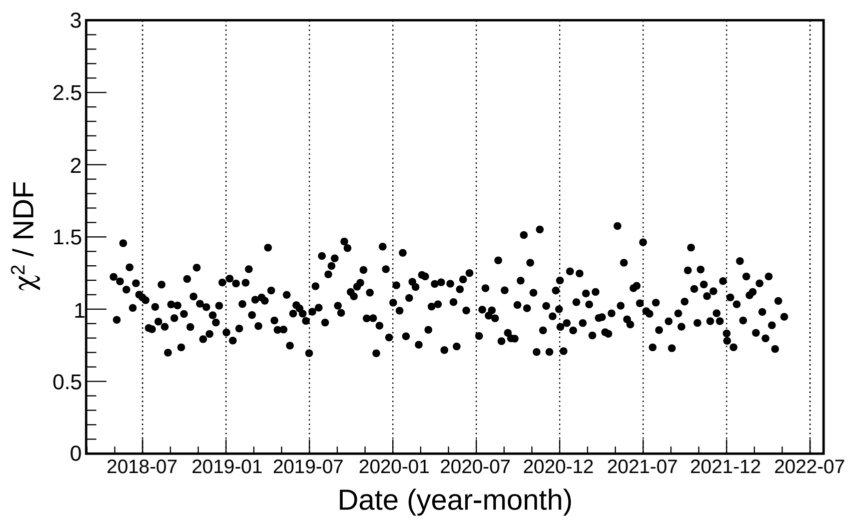

Figure 4.24: χ^2/NDF distribution of Erlang fits for weekly multiple-muon data. The good agreement indicates that the Poisson hypothesis remains valid over the entire dataset.

The principal advantage of this method is its robustness against irregularities in the data stream. Occasional dead time or file-boundary effects only distort the extreme tails of the Δt distribution. Since the rate is extracted from the central region where the exponential law holds, the result is insensitive to such distortions.

In contrast to the naïve “event count divided by live time” approach, this method does not rely explicitly on accurate live-time accounting within individual files. Instead, it uses only the statistical structure of the arrival process itself. Consequently, the uncertainty on λ arises purely from counting statistics in the fit, minimizing systematic biases in both numerator and denominator.

This statistically robust extraction of the multiple-muon rate forms the basis for the subsequent seasonal modulation analysis.

4.6 Atmospheric Temperature Extraction and Effective Temperature Calculation

As discussed in previous sections, the seasonal modulation of the underground muon rate is driven by temperature-dependent changes in the atmospheric density profile. Variations in temperature alter the competition between meson decay and interaction in the upper atmosphere, thereby modifying the flux of muons that reach detector depth.

However, muons are produced over a wide range of altitudes, primarily between 10 and 40 km. Since the exact production altitude of each muon is unknown, it is not sufficient to correlate the muon rate with temperature at a single atmospheric level. Instead, one must consider the full vertical temperature profile and construct a suitably weighted average that reflects the altitude-dependent probability of meson decay. This motivates the introduction of the *effective temperature*, T_{eff}, which encapsulates the relevant atmospheric information into a single parameter suitable for correlation studies.

Temperature data used in this analysis are obtained from the European Centre for Medium-Range Weather Forecasts (ECMWF) [61]. The details of the data extraction procedure are described below.

4.6.1 ECMWF and ERA5 Reanalysis Data

The ECMWF is an intergovernmental organization that provides global weather forecasts and climate reanalysis products. Forecasts are generated using ensemble-based numerical weather prediction models, in which multiple simulations are performed with slightly varied initial conditions to account for uncertainties in atmospheric measurements. This ensemble framework provides probabilistic predictions across various timescales.

For long-term atmospheric studies, ECMWF provides reanalysis datasets. Reanalysis combines historical observational data with modern numerical weather models to produce a temporally consistent and spatially complete dataset. The resulting product incorporates the laws of atmospheric physics to ensure internal consistency across time.

In earlier NOvA analyses [46], temperature data from the ERA-Interim reanalysis [62] were used. In the present work, the newer ERA5 reanalysis [63] is adopted. ERA5 is the fifth-generation ECMWF reanalysis product and provides improved spatial resolution, better model physics, and reduced systematic uncertainties. It covers the period from 1940 onward and supersedes ERA-Interim.

4.6.2 Extraction of Temperature Profiles

For this analysis, ERA5 temperature data were extracted for the NOvA ND location. The extraction parameters are summarized in Tab. 4.2.

Table 4.2: Parameters and options used to extract ECMWF temperature data at the NOvA ND location

Parameter	Selection
Product type	Reanalysis
Variable	Temperature
Pressure levels	1–1000 hPa (all 37 levels)
Years	2017–2022
Months	January–December
Days	1–31
Times	00, 06, 12, 18 UTC
Geographical area	(41°N, 87°W) to (42°N, 88°W)

The pressure levels available in ERA5 are non-uniformly spaced and include 37 discrete levels between 1 and 1000 hPa. These levels provide fine resolution in the lower atmosphere and adequate coverage of the stratosphere, where most meson production occurs.

ERA5 provides temperature data on a $0.25° \times 0.25°$ latitude–longitude grid. For the NOvA ND region, four grid points surrounding Fermilab are selected. At each pressure level and time, the four corresponding temperature values are averaged to obtain a representative temperature for the detector location.

Temperature data are stored in GRIdded Binary (GRIB) format [64], the standard format adopted by the World Meteorological Organization for meteorological datasets. Decoding of GRIB files is performed using the ECMWF `ecCodes` interface [65]. Temperatures are extracted using:

grib_get -P dataDate,hour,level -l 'latitude,longitude' file.grib

The resulting temperature profile as a function of pressure level and time is shown in Fig. 4.25 for the years 2017–2022.

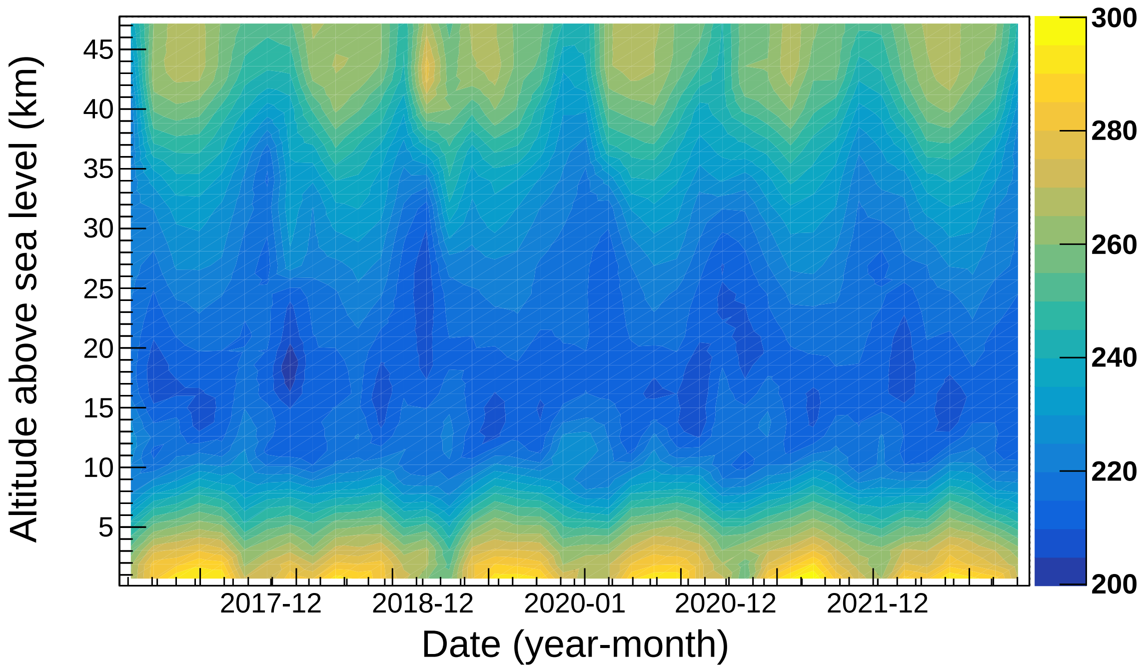

Figure 4.25: Temperature profile of the Fermilab region at different pressure levels. The color scale represents temperature in Kelvin.

Since meson production and decay are typically discussed in terms of altitude or atmospheric depth, pressure levels are converted to altitude using the barometric relation:

$$P(h) = P_0 \exp\left(-\frac{hMg}{k_B T}\right), \tag{4.27}$$

where $P(h)$ is the pressure at altitude h, $P_0 = 1010$ hPa is the surface pressure, M is the molar mass of dry air, g is the acceleration due to gravity, and k_B is the Boltzmann constant. This relation follows from hydrostatic equilibrium and the ideal gas law.

4.6.3 Effective Temperature

Although temperature varies significantly with altitude, not all atmospheric layers contribute equally to the production of underground muons. The dominant contribution arises from regions where pions and kaons are produced and subsequently decay before interacting, typically between 10 and 40 km.

To account for this altitude-dependent sensitivity, an effective temperature T_{eff} is defined as a weighted average of the atmospheric temperature profile:

$$T_{\text{eff}} = \frac{\sum_{i=1}^{n} T(X_i)\left[W_\pi(X_i) + W_K(X_i)\right]\Delta X_i}{\sum_{i=1}^{n}\left[W_\pi(X_i) + W_K(X_i)\right]\Delta X_i}. \tag{4.28}$$

Here:

- $T(X_i)$ is the temperature at atmospheric depth (or pressure level) X_i,
- W_π and W_K are weighting functions describing the contribution of pion and kaon decays to the muon flux,
- ΔX_i represents the non-uniform spacing between pressure levels.

The weighting functions incorporate the probability of meson production, interaction, and decay at each atmospheric layer, and depend on parameters such as attenuation lengths, critical energies, and the muon energy threshold. These parameters follow the formalism described in [45, 66].

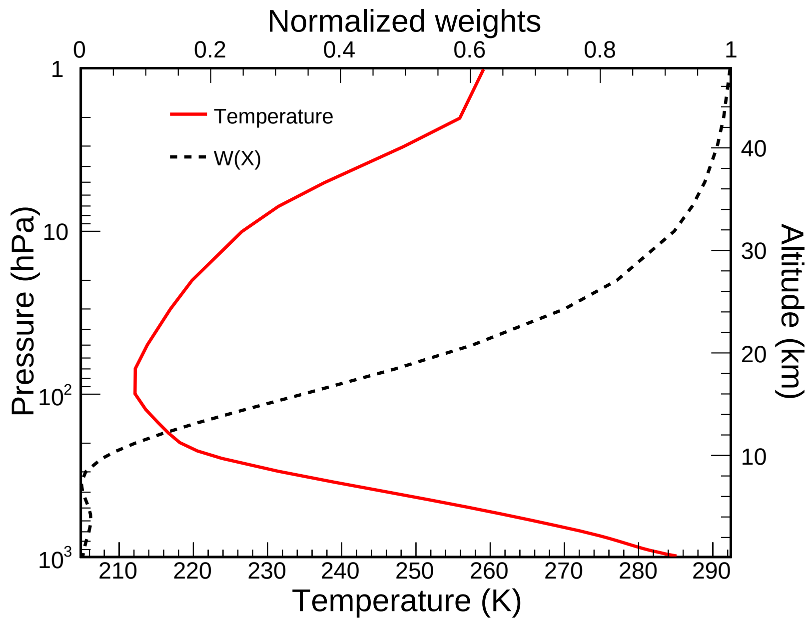

Figure 4.26: Average temperature profile (solid line) and normalized weighting function (dashed line) used in the calculation of T_{eff}.

The total weight $W_{\mathrm{tot}} = W_\pi + W_K$ is shown in Fig. 4.26 together with the mean temperature profile. The weighting function peaks in the stratospheric region, reflecting the altitude range most relevant for muon production.

The effective temperature is calculated at 00, 06, 12, and 18 UTC for each day in the dataset. Daily and monthly effective temperatures are then obtained by averaging the corresponding six-hour values.

Fig. 4.27 shows the resulting effective temperature time series in both daily and monthly binning. This T_{eff} will be used in the next section to quantify the correlation between atmospheric temperature and multiple-muon rate.

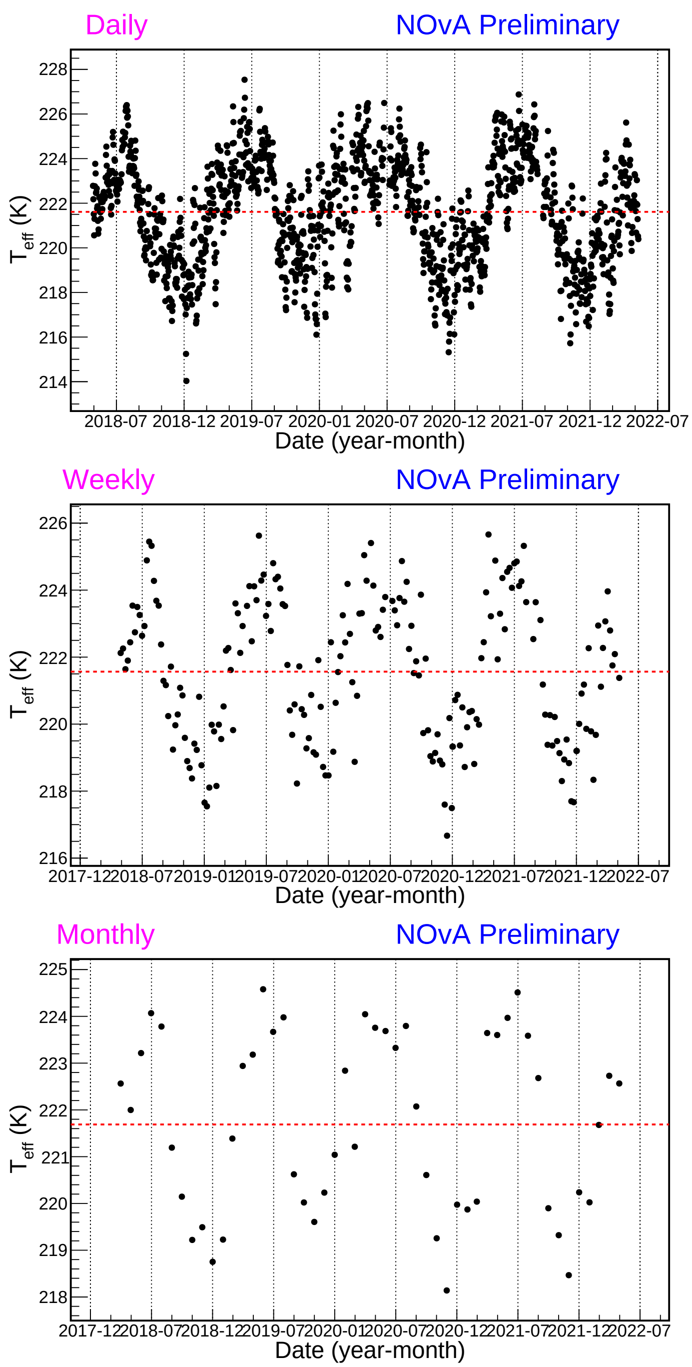

Figure 4.27: Effective temperature as a function of time: for three different binnings: daily (top), weekly (middle), and monthly (bottom). The dashed line in every plot shows the average effective temperature for the full temperature dataset.

4.6.4 Seasonal variation in ND data

The seasonal modulation of the multimuon rate observed in NOvA ND data is shown in Fig. 4.28, 4.29, and 4.30 for daily, weekly, and monthly binnings, respectively. In each case, the percentage variation is computed relative to the mean rate over the full data-taking period.

A clear annual modulation is observed. Most importantly, the variation of the multimuon rate is found to be opposite in phase to that of the effective atmospheric temperature. While the effective temperature peaks during summer, the multimuon rate reaches its maximum during winter.

To quantify this periodic behavior, both the multimuon rate variation and the effective temperature variation are fitted with a sinusoidal function of the form

$$f(t) = V_0 + V \cos\left[\frac{2\pi}{T}(t-\phi)\right], \tag{4.29}$$

where V_0 is the mean offset, V is the amplitude of modulation, T is the period, and ϕ is the phase.

	Daily	Weekly	Monthly
V_0(%)	-0.06 $\pm$ 0.002	-0.04 $\pm$ 0.002	-0.02 $\pm$ 0.003
Amplitude	2.15 $\pm$ 0.03	2.143 $\pm$ 0.04	2.11 $\pm$ 0.04
Time period (day)	355.9 $\pm$ 0.77	352 $\pm$ 0.02	355.8 $\pm$ 0.8
Phase	-1.21 π $\pm$ 0.69	-1.33π $\pm$ 0.22 π	-1.27π $\pm$ 0.23 π

Table 4.3: Best fitted parameters for the sinusoidal fit of percentage variation of multimuon rate

The best-fit parameters for the effective temperature are summarized in Tab. 4.4, while the corresponding parameters for the multimuon rate are listed in Tab. 4.3. In all binnings, the fitted period is consistent with one year, confirming that the modulation is seasonal in origin.

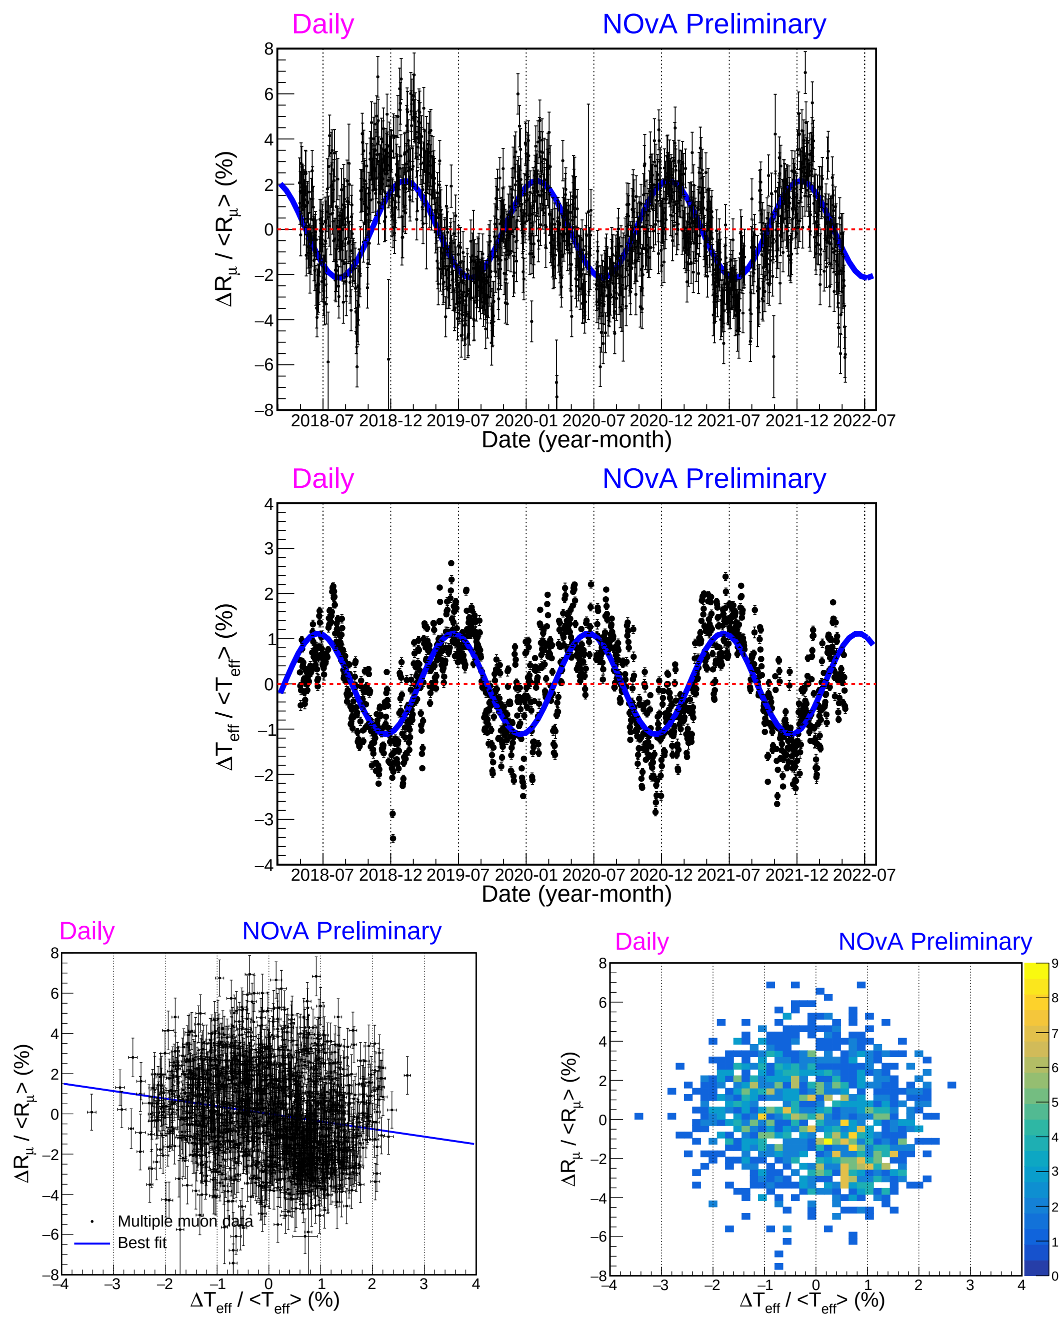

Figure 4.28: Seasonal variation for the daily bin: (top) the percentage variation of rate of multimuon data and a sinusoidal best fit (solid blue line), (middle) the percentage variation of effective temperature and a sinusoidal best fit (solid blue line), (bottom left) the multimuon rate percentage variation as a function of the percentage variation effective and a linear best fit (solid blue line), and (bottom right) the same as the bottom left but in the form of 2D histogram without error bar

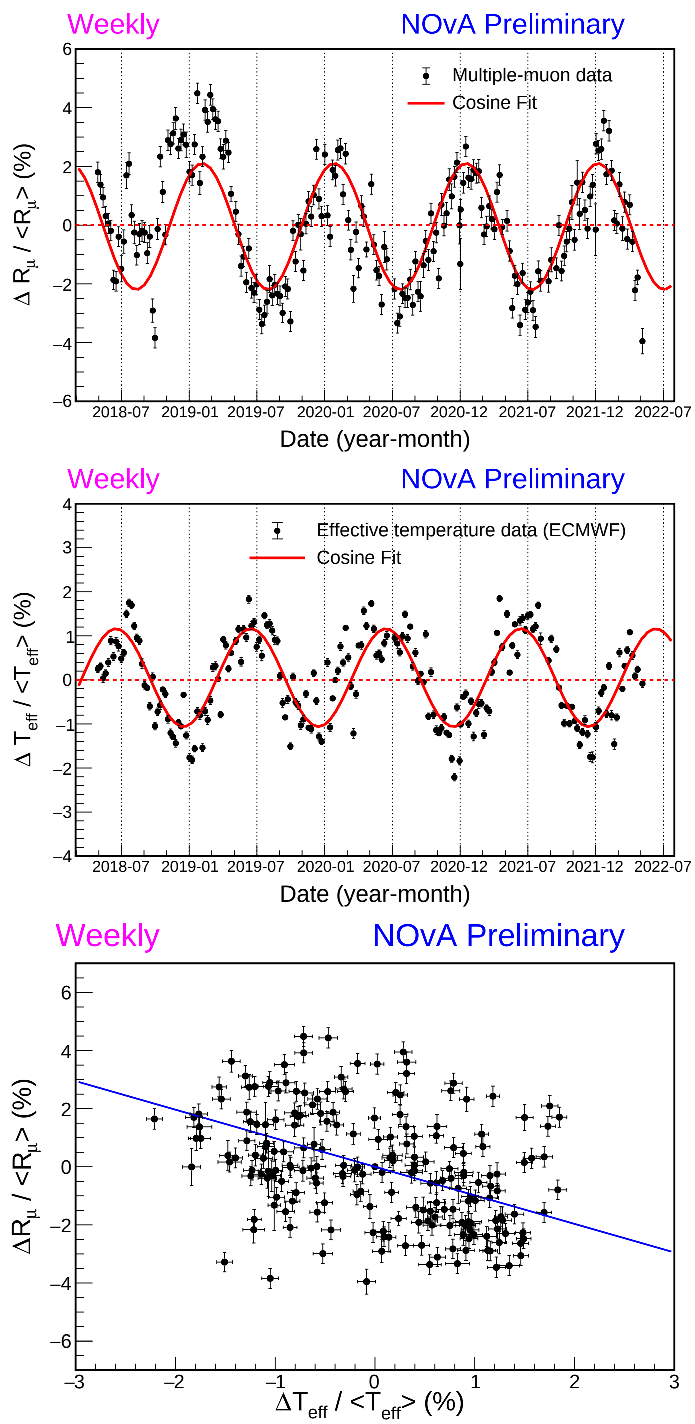

Figure 4.29: Seasonal variation for the weekly bin: (Top) the percentage variation of rate of multimuon data and a sinusoidal best fit (solid blue line), (middle) the percentage variation of effective temperature and a sinusoidal best fit (solid blue line), and (bottom) the multimuon rate percentage variation as a function of the percentage variation effective and a linear best fit (solid blue line)

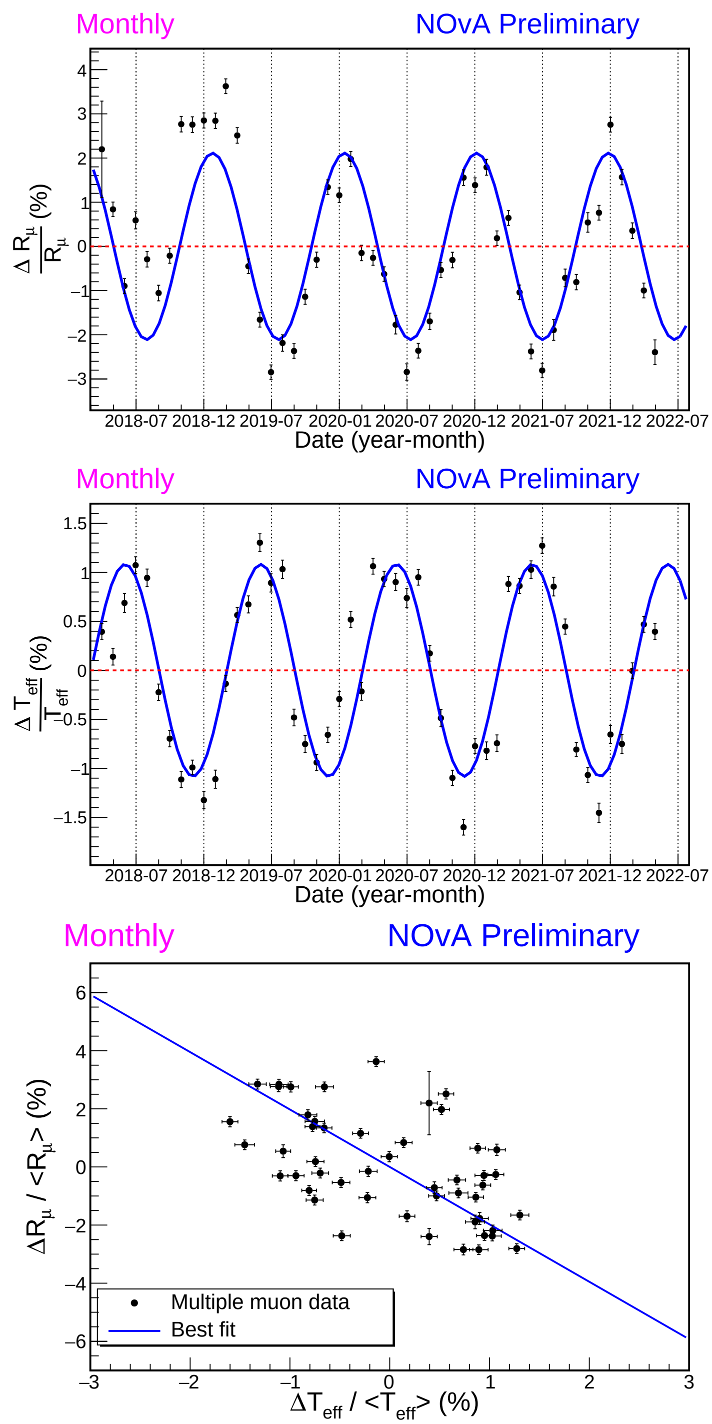

Figure 4.30: Seasonal variation for the monthly bin: (Top) the percentage variation of rate of multimuon data and a sinusoidal best fit (solid blue line), (middle) the percentage variation of effective temperature and a sinusoidal best fit (solid blue line), and (bottom) the multimuon rate percentage variation as a function of the percentage variation effective and a linear best fit (solid blue line)

	Daily	Weekly	Monthly
V_0(%)	0.003 $\pm$ 0.0002	0.003 $\pm$ 0.0006	-0.01 $\pm$ 0.001
Amplitude	1.12 $\pm$ 0.003	1.089 $\pm$ 0.01	1.082 $\pm$ 0.02
Time period (day)	364.52 $\pm$ 0.16	354.7 $\pm$ 0.01	365.8 $\pm$ 0.85
Phase	-2.58 π $\pm$ 0.05 π	-2.7 π $\pm$ 0.12π	-2.61π $\pm$ 0.24π

Table 4.4: Best fitted parameters for the sinusoidal fit of percentage variation of the effective temperature of the atmosphere

The amplitude of the multimuon modulation is approximately 2.1%, which is larger than the $\sim 1.1\%$ modulation observed in the effective temperature. The phase difference between the two confirms their anticorrelated behavior.

Correlation with Effective Temperature

The strength of the correlation between the multimuon rate and the effective temperature is quantified using the temperature correlation coefficient α_T, defined through a linear fit of the form

$$\frac{\Delta R}{R} = \alpha_T \frac{\Delta T_{eff}}{T_{eff}}. \tag{4.30}$$

The fitted values of α_T for different time binnings are presented in Tab. 4.5. The negative values of α_T clearly demonstrate the anticorrelation between the multimuon rate and the effective temperature.

Bin size	Best fitted α_T	χ^2/NDF
Daily	-0.376 $\pm$ 0.024	6.02
Weekly	-0.983 $\pm$ 0.03	24.5
Monthly	-2.495 $\pm$ 0.076	43.4

Table 4.5: Multimuon α_T fitted values according to different time binnings. The error comes from the straight-line fitting.

Although a linear trend is visible (see lower panels of Fig. 4.28, 4.29, and 4.30), the large χ^2/NDF values indicate that a simple linear model does not fully capture the underlying behavior. This suggests that additional physical effects beyond a purely temperature-driven mechanism are contributing to the observed modulation.

Overall, the data establish three key observations:

- A clear annual modulation of the multimuon rate.
- A phase opposite to that of the effective temperature.
- A strong but non-ideal linear anticorrelation with temperature.

These features motivate a detailed comparison with Monte Carlo simulations.

4.6.5 Seasonal variation from CORSIKA

To understand the origin of the observed behavior, dedicated CORSIKA simulations were performed using monthly atmospheric profiles obtained from the GDAS database for the year 2017. Cosmic-ray showers were generated for proton, helium, carbon, and iron primary particles using the configuration described in Section 2.

Surface-level behavior

The seasonal variation of single- and multimuon rates at the Earth's surface, assuming an infinitely large detector, is shown in Fig. 4.31.

For proton primaries, both single and multimuon rates peak in summer and reach a minimum in winter.

This indicates that primary composition influences the seasonal modulation at the surface, but it does not yet explain the detector-level behavior observed in data.

Detector-level behavior

When the simulated surface muons are propagated through the NOvA ND geometry, the seasonal behavior changes significantly. As shown in Fig. 4.32, the detector-level multi-

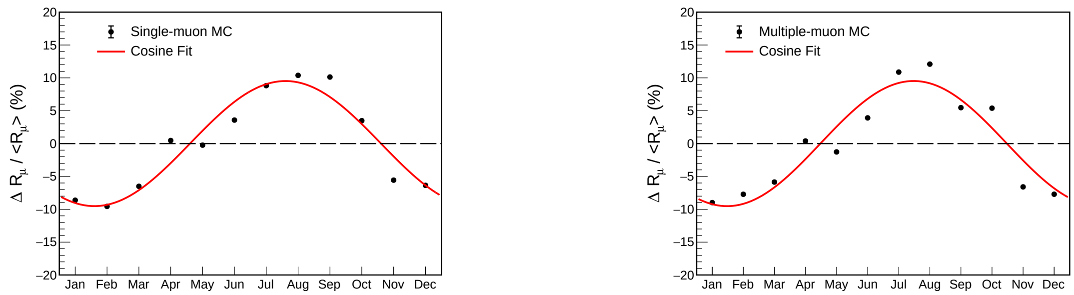

Figure 4.31: Seasonal variation of single (left) and multimuon (right) event rates for infinitely sized detector at the Earth's surface from proton (top row), ^{4}He (next row), ^{12}C (next row) primaries

muon rate now peaks in winter and reaches a minimum in summer, reproducing the trend observed in data.

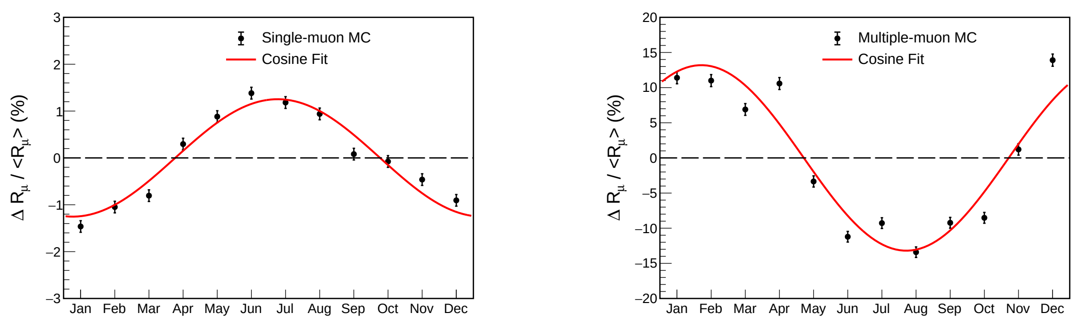

Figure 4.32: Seasonal variation of single and multimuon event rates at the NOvA ND from proton primary

This inversion demonstrates that the effect is not solely determined by the production rate of muons in the atmosphere, but is strongly influenced by detector geometry and event topology.

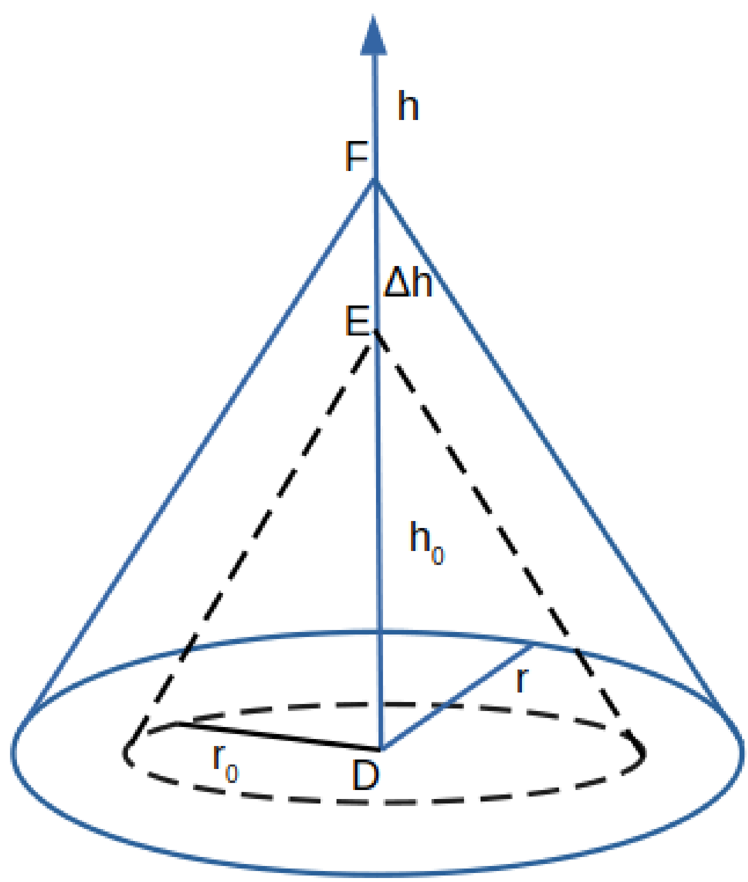

Figure 4.33: This shows how the radial spread of muons increases with increasing temperature. In summer, the atmosphere expands, so the muon birth altitude (h = DF) increases compared to that in winter (h_0 = DE), and hence the radial spread also increases in summer (r) compared to winter (r_0) (figure from Ref. [25]).

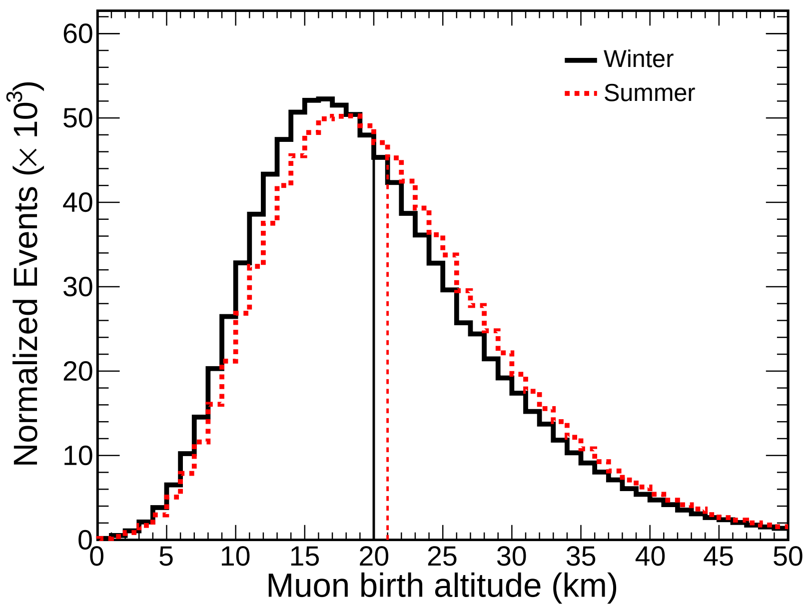

Figure 4.34: Production altitude distribution of multiple-muon events from CORSIKA with muon threshold energy of $54\,\mathrm{GeV}$. Three winter months (December–February) and three summer months (June–August) are compared. Vertical lines represent the mean muon birth altitude for summer (red dashed line) and winter (black solid line).

4.7 Explanation

4.7.1 The altitude-geometry explanation

The total mass of the atmosphere remains approximately constant throughout the year, but seasonal temperature differences cause vertical expansion of atmosphere in summer and contraction in winter. Under the assumption of an isothermal atmosphere, the ideal gas law $PV = nRT$ implies that a $\pm 2\%$ change in absolute temperature results in a corresponding $\pm 2\%$ shift in altitude for a given pressure level. As shown in Fig. 4.33 [25], for a shower origin at fixed pressure, the increased altitude (point F) in summer leads to a broader lateral spread of muons at the detector. In contrast, during winter, the lower altitude (point E) results in a more concentrated footprint.

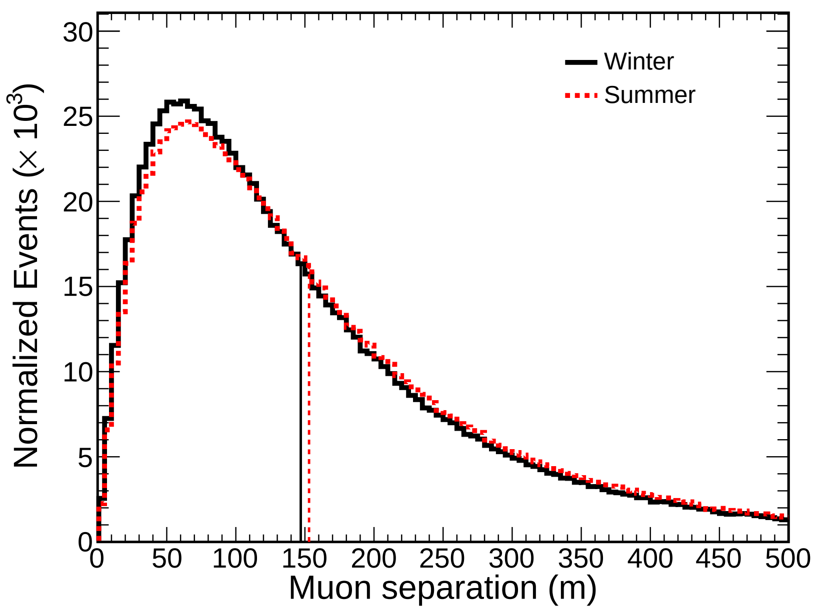

Figure 4.35: Muon separation distribution of multiple-muon events from CORSIKA with muon threshold energy of $54\,\mathrm{GeV}$. Three winter months (December–February) and three summer months (June–August) are compared. Vertical lines represent the mean muon separation for summer (red dashed line) and winter (black solid line).

We use the events generated with the CORSIKA simulation to study the altitude and muon separation variation during summer and winter. Fig. 4.34 shows the mean birth altitude of muons in the upper atmospehere from multiple muon events is higher in summer than in winter. As a result, the separation among muons for those events also increases in summer compared to winter, as shown in the Fig. 4.35. It is instructive to consider the typical transverse momentum of mesons in hadronic showers, and how that impacts the separation distribution of muons reaching a detector. The opening angle between a secondary pion and the primary cosmic ray direction in an air shower is approximately proportional to p_t/p_l, where p_t is the pion's transverse momentum and $p_l \sim E_\pi \sim E^\mu$ is its longitudinal momentum. Consequently, the typical lateral separation between the two muons originating from different pions in the same shower can be expressed as

$$\sqrt{2} \times h \times (p_t/E^\mu), \tag{4.31}$$

where h is the altitude at which the pions decay. Since only higher-energy muons can reach deeper underground detectors, the required E^μ increases with depth, leading to smaller lateral separations between muons for such underground detectors. To illustrate how the altitude-geometry effect depends on detector depth and muon energy, Table 4.6 shows the typical lateral separation of muons at various underground depths, calculated using Eq. 4.31 for pions decaying at an altitude of 20 km and a typical pion p_t of 300 MeV/c. The corresponding minimum muon energies to reach the detector are calculated assuming standard rock with a density of 2.8 g/cm^3 and a minimum ionizing energy loss of $2\ \mathrm{MeV}/(\mathrm{g\,cm}^{-2})$. The muon separation decreases with increasing the energy (required to reach deeper detectors) and with increasing depth due to energy-dependent selection of muons closer to the shower core.

For the NOvA or MINOS underground near detectors, with an energy threshold close to $50\,\mathrm{GeV}$, the typical muon separation is $170\,\mathrm{m}$ which is much larger than the detector size.

Depth (m)	E^{μ}_{min} (GeV)	Muon Separation (m)
18	10	850
80	50	170
180	100	85
800	500	17
1800	1000	8.5

Table 4.6: Typical lateral separation of muons as a function of underground detector depth and muon energy threshold, calculated using Eq. 4.31. The separation indicates the typical transverse spread of muons from a cosmic ray shower at the detector level.

We can use Table 4.6 to determine whether any underground detector with a particular depth and area ($x \times z$) exhibits a winter or summer maximum. If the separation is much smaller than x or z, a summer maximum will be seen. If the separation is larger or comparable, a winter maximum will be observed. However, the single muon rate remains unaffected by the altitude-geometry effect. For every single muon that misses a finite-sized detector because it starts at a higher altitude in the summer, another muon that would miss the detector in the summer now hits it.

The role of summer/winter differences in the origin altitude of muons in an air shower, and its effect on shower containment in a finite detector, was also considered before by MINOS [45]. However, the conclusion in the paper was based on the distributions of observed muon separation in the detectors, instead of the true separation between muons. That reference based its conclusion on the measured muon separation distribution, assuming that it was the true distribution.

The dependence of the altitude–geometry effect as function of detector depth can also explain the seasonal behaviors observed in other experiments. For deeper detectors that require higher muon energy (TeV) thresholds, the altitude–geometry effect is minimal. The surface area of the MACRO experiment [67] is $76\,\mathrm{m}\times12\,\mathrm{m}$ and the threshold energy required for an atmospheric muon to reach the detector is $1.3\,\mathrm{TeV}$. The average separation among muons in this case, is smaller than the size of the detector. Hence, MACRO observed a summer maximum for multiple-muon events [68]. In the MINOS FD [45], the typical muon separation from an air shower is comparable to the size of the detector. So, the observed muon separation determines whether there is a summer maximum or a winter maximum. Calculations by members of the IceCube collaboration based on the altitude-geometry effect seemed to explain the MINOS and NOvA multiple-muon seasonal data [69, 70]. The altitude-geometry effect was considered responsible for the seasonal results of the DECOR experiment [71], which measured muons close to the surface of the Earth. For those muons, with an energy typically of a few GeV, other effects such as pressure changes and muon decay are important, as discussed further in Sec. 4.7.2.

4.7.2 Other seasonal explanation

The MINOS result in Ref. [45] disfavored dimuon decays or temperature effects as contributing to the explanation of multiple-muon variation. However, it considered the anti-correlation between the primary meson interaction and decay as the most plausible explanation for the observed phenomenon. According to this hypothesis, higher temperatures in summer increase the decay probability of mesons, leaving fewer mesons available to inelastically scatter and produce additional secondaries in subsequent generations of the hadronic shower. To test this independently of CORSIKA, E. Guan used a simulation of hadronic showers in the atmosphere with densities corresponding to summer and winter [72]. She followed the meson and muon content in the hadronic shower as subsequent generations of both mesons and nucleons interacted, as a function of primary cosmic ray energy. The simulation showed that a shower produces more muons when the primary meson decays rather than interacts, contrary to the hypothesis favored by MINOS [45]. In addition, this explanation would predict a winter maximum in an infinite detector, in contradiction with the CORSIKA result in Fig. 4.31.

E_μ (GeV)	N/N_0 (%) (winter)	N/N_0 (%) (summer)	Difference (%)
5	62.05	61.46	0.59
50	95.34	95.25	0.09
500	99.52	99.51	0.01

Table 4.7: Analytic calculation of the percentage of muons that reach sea level before they decay as a function of muon energy. E_μ is the muon energy, N_0 is the initial number of muons, and N is the number of muons that survive. The summer/winter altitude difference is assumed to be 2%.

Another potential explanation for the observed seasonal pattern of multiple muons is muon decay, which was not discussed in Ref. [45]. Due to relativistic effects, muon decay is more significant at lower muon energies, and thus its impact is expected to be largest for detectors located at or near the surface. An analytical estimate of the seasonal modulation due to muon decay for various energy thresholds is summarized in Tab. 4.7. Surface detectors such as the NOvA FD [47], DECOR [71, 73], and GRAPES [74], all with muon thresholds near 5 GeV, can exhibit a 0.6% excess of events during the winter. In contrast, detectors such as the NOvA ND and MINOS ND, which have muon thresholds close to 50 GeV, would only show a smaller 0.09% excess in winter, which is much smaller than the observed winter excess, and also inconsistent with Fig. 4.31. Muon decay is even less important for the MINOS FD and MACRO.

Chapter 5

Neutrino Oscillation and Interaction

Neutrinos are elementary fermions belonging to the lepton sector of the Standard Model (SM). Their existence was first proposed by Wolfgang Pauli in 1930 [75] to explain the continuous energy spectrum observed in nuclear beta decay, thereby restoring energy conservation. Enrico Fermi later incorporated the particle into his theory of weak interactions and introduced the name *neutrino* [76], meaning "little neutral one."

In the Standard Model, neutrinos are electrically neutral, spin-$\frac{1}{2}$ particles that interact exclusively via the weak interaction. They appear in three flavour states associated with the charged leptons:

$$\nu_e, \qquad \nu_\mu, \qquad \nu_\tau. \tag{5.1}$$

Due to the absence of electric charge and strong interactions, neutrinos interact very weakly with matter. Their interaction cross sections at GeV-scale energies are typically of order $10^{-38}\,\mathrm{cm}^2$, making them extremely penetrating particles. The first experimental detection of neutrinos was achieved in 1956 by Reines and Cowan [77] using reactor antineutrinos.

Originally, neutrinos were assumed to be massless within the Standard Model. However, the experimental discovery of neutrino oscillations [78, 79] demonstrated that neutrinos possess non-zero masses and that flavour states are quantum superpositions of mass eigenstates. This discovery requires physics beyond the minimal Standard Model.

5.1 Neutrino Oscillation

Neutrino oscillation [2, 80] is a direct consequence of the fact that neutrino flavour eigenstates are not identical to neutrino mass eigenstates. Neutrinos are produced and detected in flavour eigenstates through weak interactions, whereas their propagation in vacuum is governed by the free Hamiltonian whose eigenstates are mass eigenstates.

Let $|\nu_\alpha\rangle$ ($\alpha = e, \mu, \tau$) denote the flavour eigenstates and $|\nu_i\rangle$ ($i = 1, 2, 3$) denote the mass eigenstates with masses m_i. The two bases are related through a unitary transformation [81, 82]:

$$|\nu_\alpha\rangle = \sum_{i=1}^{3} U^*_{\alpha i}|\nu_i\rangle, \tag{5.2}$$

where U is the 3×3 PMNS mixing matrix satisfying

$$U^\dagger U = UU^\dagger = I. \tag{5.3}$$

Because U is unitary, the inverse relation is

$$|\nu_i\rangle = \sum_{\beta=e,\mu,\tau} U_{\beta i}|\nu_\beta\rangle. \tag{5.4}$$

Consider a neutrino produced at spacetime point

$$X_0 = (t = 0, \vec{x} = 0)$$

in a definite flavour state $|\nu_\alpha\rangle$. Using Eq. (5.2), the initial state is

$$|\nu_\alpha(0)\rangle = \sum_i U^*_{\alpha i}|\nu_i\rangle. \tag{5.5}$$

Each mass eigenstate evolves independently under the free Hamiltonian. The relativistic plane-wave solution of the Klein–Gordon (or Dirac) equation can be written as

$$|\nu_i(t,\vec{x})\rangle = e^{-i(E_i t - \vec{p}_i\cdot\vec{x})}|\nu_i\rangle, \tag{5.6}$$

where

$$E_i = \sqrt{|\vec{p}_i|^2 + m_i^2}. \tag{5.7}$$

After propagation to spacetime point $X = (t, \vec{x})$, the state becomes

$$|\nu_\alpha(X)\rangle = \sum_i U_{\alpha i}^* e^{-i(E_i t - \vec{p}_i\cdot\vec{x})}|\nu_i\rangle. \tag{5.8}$$

Substituting Eq. (5.4) into Eq. (5.8), we express the propagated state in the flavour basis:

$$|\nu_\alpha(X)\rangle = \sum_\beta \sum_i U_{\alpha i}^* U_{\beta i} e^{-i(E_i t - \vec{p}_i\cdot\vec{x})}|\nu_\beta\rangle. \tag{5.9}$$

This shows explicitly that a neutrino initially produced in flavour α becomes a coherent superposition of all flavour states during propagation.

The probability amplitude for detecting the neutrino as flavour β at position X is obtained by projection:

$$\mathcal{A}_{\alpha\to\beta}(X) = \langle\nu_\beta|\nu_\alpha(X)\rangle. \tag{5.10}$$

Using orthonormality, $\langle\nu_\beta|\nu_\gamma\rangle = \delta_{\beta\gamma}$, we obtain

$$\mathcal{A}_{\alpha\to\beta}(X) = \sum_i U_{\alpha i}^* U_{\beta i} e^{-i(E_i t - \vec{p}_i\cdot\vec{x})}. \tag{5.11}$$

The transition probability is

$$P_{\alpha\to\beta} = |\mathcal{A}_{\alpha\to\beta}|^2. \tag{5.12}$$

Substituting Eq. (5.11),

$$P_{\alpha \to \beta} = \sum_{i,j} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i[(E_i - E_j)t - (\vec{p}_i - \vec{p}_j)\cdot\vec{x}]}. \tag{5.13}$$

Neutrinos in experiments are ultra-relativistic [2], i.e.

$$E_i \gg m_i.$$

Assuming that all mass eigenstates have approximately equal momentum, $|\vec{p}_i| \approx p$, we expand the energy:

$$E_i = \sqrt{p^2 + m_i^2}. \tag{5.14}$$

Using a binomial expansion:

$$E_i = p\sqrt{1 + \frac{m_i^2}{p^2}} \approx p\left(1 + \frac{m_i^2}{2p^2}\right) = p + \frac{m_i^2}{2p}. \tag{5.15}$$

Since for ultra-relativistic particles $E \approx p$,

$$E_i \approx E + \frac{m_i^2}{2E}. \tag{5.16}$$

Therefore,

$$E_i - E_j = \frac{m_i^2 - m_j^2}{2E} = \frac{\Delta m_{ij}^2}{2E}. \tag{5.17}$$

For neutrinos traveling nearly at the speed of light,

$$t \approx L,$$

where L is the baseline (distance between production and detection). Thus the phase difference becomes

$$(E_i - E_j)t - (\vec{p}_i - \vec{p}_j)\cdot\vec{x} \approx \frac{\Delta m_{ij}^2 L}{2E}. \tag{5.18}$$

Substituting into Eq. (5.13):

$$P_{\alpha\to\beta} = \sum_{i,j} U^*_{\alpha i} U_{\beta i} U_{\alpha j} U^*_{\beta j} \exp\left(-i\frac{\Delta m^2_{ij} L}{2E}\right). \tag{5.19}$$

Separating real and imaginary parts and using trigonometric identities, the probability can be written in the standard form:

$$P_{\alpha\to\beta} = \delta_{\alpha\beta} - 4\sum_{i>j} \mathrm{Re}\left(U_{\alpha i} U^*_{\beta i} U^*_{\alpha j} U_{\beta j}\right) \sin^2\left(\frac{\Delta m^2_{ij} L}{4E}\right) \tag{5.20}$$
$$+2\sum_{i>j} \mathrm{Im}\left(U_{\alpha i} U^*_{\beta i} U^*_{\alpha j} U_{\beta j}\right) \sin\left(\frac{\Delta m^2_{ij} L}{2E}\right).$$

The final term in the oscillation probability expression contains the imaginary part of the product of PMNS matrix elements and represents CP-violating effects. For antineutrinos, the mixing matrix elements are replaced by their complex conjugates,

$$U \to U^*,$$

which reverses the sign of the CP-odd (imaginary) contribution while leaving the CP-even terms unchanged. Consequently, a difference between $P(\nu_\alpha \to \nu_\beta)$ and $P(\bar{\nu}_\alpha \to \bar{\nu}_\beta)$ signals CP violation in the lepton sector. This asymmetry vanishes if the Dirac CP phase δ_{CP} is equal to 0 or π.

Several additional properties of the oscillation probability are noteworthy:

- The probability depends on combinations of the form $U_{\alpha i} U^*_{\beta i} U^*_{\alpha j} U_{\beta j}$. If neutrinos are Majorana particles, additional Majorana phases appear as overall phase factors multiplying columns of the PMNS matrix. However, these phases cancel in the above products. Therefore, neutrino oscillation experiments are insensitive to Majorana phases. Probing these phases requires processes such as neutrinoless double beta decay.

- The oscillation probability depends explicitly on the mass-squared differences,

$$\Delta m_{ij}^2 = m_i^2 - m_j^2.$$

  For oscillations to occur, at least one neutrino mass must be non-zero, and the masses must not be degenerate. Additionally, non-zero mixing angles are required; if the mixing matrix were diagonal, flavour conversion would not occur even if the masses were distinct.

- The dependence on the ratio L/E implies a sinusoidal behavior of the transition probability. As a neutrino propagates, the probability of detecting a different flavour varies periodically with baseline and energy. This periodic variation is what is referred to as neutrino oscillation.

- When $\alpha \neq \beta$, the probability describes an *appearance* channel. When $\alpha = \beta$, it corresponds to a *survival* or *disappearance* channel. In the survival case, the imaginary (CP-violating) contribution vanishes identically due to the symmetry of the interference terms. As a result, disappearance measurements are insensitive to the complex CP-violating phase δ_{CP} and primarily constrain mixing angles and mass-squared differences.

5.1.1 Two Flavor Oscillation in Vacuum [1, 2]

The two-flavor approximation is obtained by assuming that only two neutrino states contribute to oscillations. This limit is valid when one mass-squared difference dominates the oscillation scale, and mixing with the third state can be neglected.

In this case, the mixing matrix reduces to

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}, \tag{5.21}$$

and there is only one independent mass-squared difference,

$$\Delta m^2 = m_2^2 - m_1^2. \tag{5.22}$$

Substituting this reduced mixing structure into Eq. (5.20), the summation over $i > j$ contains only one term corresponding to $(i, j) = (2, 1)$.

For the appearance channel ($\alpha \neq \beta$), the relevant product is

$$U_{\alpha 2} U_{\beta 2}^{*} U_{\alpha 1}^{*} U_{\beta 1}. \tag{5.23}$$

Using the explicit form of the mixing matrix, one finds

$$\mathrm{Re}\left(U_{\alpha 2} U_{\beta 2}^{*} U_{\alpha 1}^{*} U_{\beta 1}\right) = -\frac{1}{4}\sin^2(2\theta), \tag{5.24}$$

while the imaginary part vanishes identically since the matrix contains no complex phase. Substituting into Eq. (5.20), the appearance probability becomes

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{4E}\right). \tag{5.25}$$

Similarly, for the survival probability ($\alpha = \beta$),

$$P(\nu_\alpha \to \nu_\alpha) = 1 - \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{4E}\right). \tag{5.26}$$

Expressed in practical units, this becomes

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta)\sin^2\left(1.27\frac{\Delta m^2(\mathrm{eV}^2)\, L(\mathrm{km})}{E(\mathrm{GeV})}\right). \tag{5.27}$$

The probability therefore depends only on the ratio L/E and exhibits a sinusoidal behavior characteristic of quantum interference.

5.1.2 Three-Flavor Neutrino Oscillations in Vacuum [3]

In the three-flavor case, neutrino mass eigenstates are $(i, j = 1, 2, 3)$, hence the independent pairs are

$$(i, j) = (2, 1),\ (3, 1),\ (3, 2). \tag{5.28}$$

Therefore, Eq. (5.20) becomes

$$\begin{aligned} P_{\alpha\beta} = \delta_{\alpha\beta} - 4\Big[\mathrm{Re}(X_{\alpha\beta}^{21}) \sin^2 \Delta_{21} + \mathrm{Re}(X_{\alpha\beta}^{31}) \sin^2 \Delta_{31} + \mathrm{Re}(X_{\alpha\beta}^{32}) \sin^2 \Delta_{32}\Big] \\ + 2\Big[\mathrm{Im}(X_{\alpha\beta}^{21}) \sin 2\Delta_{21} + \mathrm{Im}(X_{\alpha\beta}^{31}) \sin 2\Delta_{31} + \mathrm{Im}(X_{\alpha\beta}^{32}) \sin 2\Delta_{32}\Big], \end{aligned} \tag{5.29}$$

where

$$X_{\alpha\beta}^{ij} = U_{\alpha i} U_{\beta i}^{*} U_{\alpha j}^{*} U_{\beta j}. \tag{5.30}$$

PMNS Matrix [3]

The lepton mixing matrix is parameterized as

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix}, \tag{5.31}$$

with

$$c_{ij} = \cos\theta_{ij}, \qquad s_{ij} = \sin\theta_{ij}.$$

Muon to Electron Appearance Probability

Setting $\alpha = \mu$ and $\beta = e$, and substituting the PMNS elements into Eq. (5.29), one obtains after algebraic simplification [83]:

$$\begin{aligned} P_{\nu_\mu \to \nu_e} &= 4s_{13}^2 s_{23}^2 \sin^2 \Delta_{31} \\ &+ 4c_{13}^2 s_{12}^2 c_{23}^2 \sin^2 \Delta_{21} \\ &+ 8s_{13}s_{12}c_{12}s_{23}c_{23}c_{13} \cos(\Delta_{31} + \delta_{CP}) \sin \Delta_{31} \sin \Delta_{21}. \end{aligned} \tag{5.32}$$

The three terms correspond to:

- Atmospheric term: $\propto \sin^2 \Delta_{31}$
- Solar term: $\propto \sin^2 \Delta_{21}$
- Interference (CP-sensitive) term

For antineutrinos, $\delta_{CP} \to -\delta_{CP}$.

Muon Neutrino Disappearance Probability

For $\alpha = \beta = \mu$, the imaginary terms vanish and we obtain:

$$\begin{aligned} P_{\nu_\mu \to \nu_\mu} = 1 &- 4|U_{\mu 1}|^2 |U_{\mu 2}|^2 \sin^2 \Delta_{21} \\ &- 4|U_{\mu 1}|^2 |U_{\mu 3}|^2 \sin^2 \Delta_{31} \\ &- 4|U_{\mu 2}|^2 |U_{\mu 3}|^2 \sin^2 \Delta_{32}. \end{aligned} \tag{5.33}$$

To leading order in Δ_{21},

$$P_{\nu_\mu \to \nu_\mu} \approx 1 - \sin^2 2\theta_{23} \sin^2 \Delta_{31}. \tag{5.34}$$

Muon to Tau Appearance Probability

For $\alpha = \mu$ and $\beta = \tau$:

$$
\begin{aligned}
P_{\nu_\mu \to \nu_\tau} = {} & 4c_{13}^4 s_{23}^2 c_{23}^2 \sin^2 \Delta_{31} \\
& - 4s_{13}^2 s_{23}^2 c_{23}^2 \sin^2 \Delta_{31} \\
& - 4c_{13}^2 s_{12}^2 s_{23}^2 c_{23}^2 \sin^2 \Delta_{21} \\
& - 8 s_{13} s_{12} c_{12} s_{23} c_{23} c_{13} \cos(\Delta_{31} + \delta_{CP}) \sin \Delta_{31} \sin \Delta_{21}.
\end{aligned} \tag{5.35}
$$

To leading order,

$$
P_{\nu_\mu \to \nu_\tau} \approx \cos^4 \theta_{13} \sin^2 2\theta_{23} \sin^2 \Delta_{31}. \tag{5.36}
$$

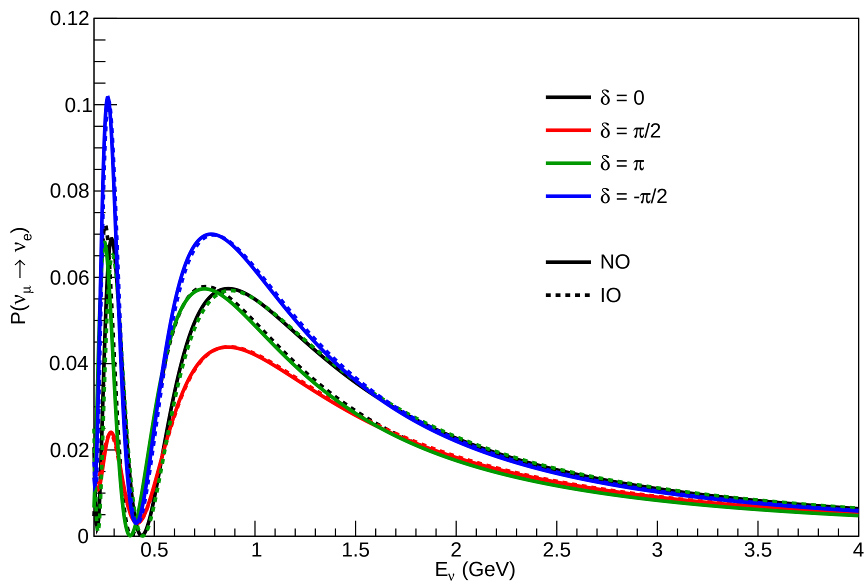

Figure 5.1: Oscillation probability as a function of neutrino energy for fixed baseline (810 km). The oscillatory behavior reflects quantum interference between mass eigenstates.

Fig. 5.1 shows the three-flavor vacuum oscillation probability as a function of neutrino energy for a fixed baseline. The dominant oscillation frequency is governed by Δm_{31}^2, while the smaller Δm_{21}^2 introduces a slow modulation. The interference term containing

δ_{CP} modifies the height and shape of the appearance peak. Fig. 5.2 shows the oscillation probability as a function of baseline at fixed energy.

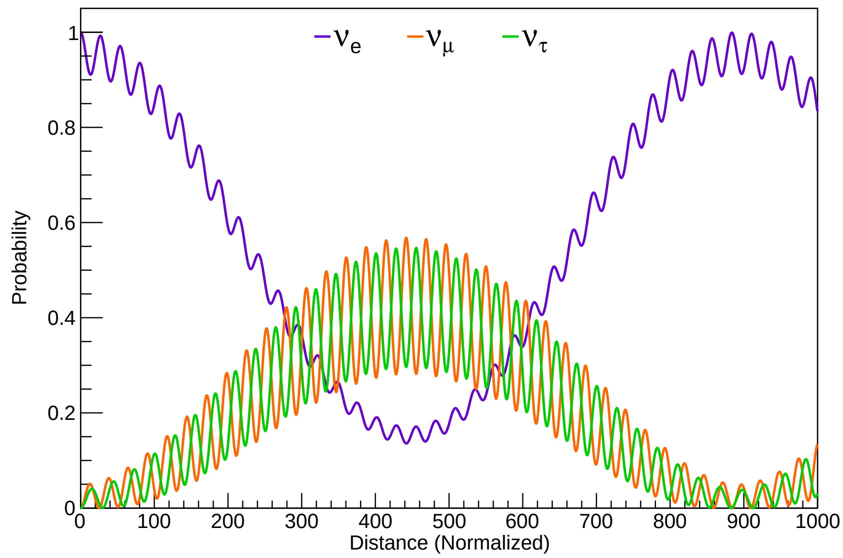

Figure 5.2: Oscillation probability as a function of baseline.

5.1.3 Oscillation in Matter

Up to this point, neutrino oscillations were derived assuming propagation in vacuum. In realistic long-baseline experiments, however, neutrinos travel through matter before reaching the detector. In the case of NOvA, neutrinos propagate through approximately 810 km of the Earth's crust.

Because ordinary matter consists primarily of electrons, protons, and neutrons, electron neutrinos experience additional coherent forward scattering with electrons via charged-current W-boson exchange. Muon and tau neutrinos do not have this interaction. This modifies the effective Hamiltonian governing neutrino propagation.

This phenomenon is known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect [84, **?**].

Matter Potential

The charged-current forward scattering generates an effective potential

$$V_e = \pm\sqrt{2}\, G_F N_e, \tag{5.37}$$

where

- G_F is the Fermi constant,
- N_e is the electron number density,
- the $+$ sign applies to neutrinos,
- the $-$ sign applies to antineutrinos.

Neutral-current interactions contribute equally to all flavors and therefore introduce only an overall phase, which does not affect oscillation probabilities.

Two-Flavor Approximation in Matter

To illustrate the impact of matter, consider first the two-flavor system (ν_e, ν_μ).

The time evolution equation becomes

$$i\frac{d}{dt}\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix} = \left[U\begin{pmatrix}\frac{m_1^2}{2E} & 0\\ 0 & \frac{m_2^2}{2E}\end{pmatrix}U^\dagger + \begin{pmatrix}V_e & 0\\ 0 & 0\end{pmatrix}\right]\begin{pmatrix}\nu_e\\ \nu_\mu\end{pmatrix}, \tag{5.38}$$

where U is the two-flavor mixing matrix with vacuum mixing angle θ.

After rotating to the flavor basis and subtracting an irrelevant trace term, the effective Hamiltonian becomes

$$H_f = \frac{\Delta m^2}{4E}\begin{pmatrix}-\cos 2\theta + A & \sin 2\theta\\ \sin 2\theta & \cos 2\theta - A\end{pmatrix}, \tag{5.39}$$

where

$$A = \frac{2EV_e}{\Delta m^2}. \tag{5.40}$$

Effective Mixing in Matter

Diagonalizing the Hamiltonian yields modified effective parameters in matter:

$$\sin 2\theta_M = \frac{\sin 2\theta}{\sqrt{(\cos 2\theta - A)^2 + \sin^2 2\theta}}, \tag{5.41}$$

$$\Delta m_M^2 = \Delta m^2 \sqrt{(\cos 2\theta - A)^2 + \sin^2 2\theta}. \tag{5.42}$$

Thus, the oscillation probability retains the vacuum form:

$$P^{\text{matter}}_{\nu_e \to \nu_\mu} = \sin^2 2\theta_M \sin^2 \left(\frac{\Delta m_M^2 L}{4E} \right). \tag{5.43}$$

In the limit $N_e \to 0$, $A \to 0$, and the vacuum result is recovered.

MSW Resonance

Resonant enhancement occurs when

$$A = \cos 2\theta, \tag{5.44}$$

for which

$$\theta_M = \frac{\pi}{4}. \tag{5.45}$$

This corresponds to maximal mixing in matter, even if the vacuum mixing angle is small. The resonance condition depends on the sign of Δm^2, making matter effects sensitive to the mass ordering.

Three-Flavor Oscillations in Matter

In the three-flavor framework, the Hamiltonian in flavor basis becomes

$$H_f = \frac{1}{2E} U \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m_{21}^2 & 0 \\ 0 & 0 & \Delta m_{31}^2 \end{pmatrix} U^\dagger + \begin{pmatrix} V_e & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{5.46}$$

The diagonalization is algebraically complex, but the qualitative result is the same: matter modifies effective mixing angles and mass splittings.

For long-baseline experiments such as NOvA, the dominant modification affects θ_{13} and Δm_{31}^2.

Appearance Probability in Matter

To leading order in $\alpha = \Delta m_{21}^2 / \Delta m_{31}^2$, the $\nu_\mu \to \nu_e$ appearance probability in matter with constant density can be approximated as:

$$P_{\nu_\mu \to \nu_e}^{\text{matter}} \approx \sin^2 \theta_{23} \sin^2 2\theta_{13}^M \sin^2 \left(\frac{\Delta m_{31,M}^2 L}{4E} \right) + \text{interference terms involving } \delta_{CP}. \tag{5.47}$$

The effective mixing angle in matter is

$$\sin^2 2\theta_{13}^M = \frac{\sin^2 2\theta_{13}}{(\cos 2\theta_{13} - A)^2 + \sin^2 2\theta_{13}}. \tag{5.48}$$

Thus, matter enhances or suppresses the appearance probability depending on:

- Neutrino vs antineutrino,
- Sign of Δm_{31}^2 (mass ordering),
- Energy.

This introduces an asymmetry between neutrinos and antineutrinos, often referred to as "fake CP violation."

Disappearance Probability in Matter

For the $\nu_\mu \to \nu_\mu$ disappearance channel, matter effects are smaller but not negligible.

To leading order,

$$P^{\text{matter}}_{\nu_\mu\to\nu_\mu} \approx 1 - \sin^2 2\theta_{23} \sin^2\left(\frac{\Delta m^2_{31,M} L}{4E}\right), \tag{5.49}$$

where the effective mass splitting $\Delta m^2_{31,M}$ includes matter corrections.

Because θ_{23} is not directly modified by the matter potential, the disappearance channel is less sensitive to matter effects than the appearance channel.

For NOvA, with $L = 810\,\text{km}$ and $E \sim 2\,\text{GeV}$, the matter parameter A is non-negligible. Matter effects:

- Enhance $\nu_\mu \to \nu_e$ probability for normal ordering,
- Suppress it for inverted ordering,
- Produce opposite behavior for antineutrinos.

This sensitivity enables NOvA to probe the neutrino mass hierarchy.

5.2 Neutrino Interactions [4, 5]

As discussed in the previous sections, neutrino oscillation experiments determine oscillation parameters by measuring interaction rates and reconstructing neutrino energies at detectors. The observed event spectrum is effectively a convolution of the incident neutrino flux $\Phi(E_\nu)$ and the interaction cross section $\sigma(E_\nu)$. Since both quantities depend strongly on neutrino energy, precise knowledge of neutrino interaction cross sections is essential for oscillation measurements.

At energies relevant to long-baseline experiments such as NOvA, the uncertainties in neutrino–nucleon cross sections remain at the level of 20–30%. These uncertainties arise from both limited experimental data and incomplete understanding of nuclear dynamics. A realistic prediction of neutrino event rates therefore requires two essential ingredients:

- A reliable description of neutrino scattering on free nucleons.
- A realistic model of the nuclear medium in which nucleons are bound.

In this section, we review the theoretical framework of neutrino interactions, beginning with weak interaction theory at the nucleon level, followed by a discussion of different interaction regimes and nuclear medium effects relevant for accelerator-based experiments.

5.2.1 Weak Interaction Framework

Neutrino interactions are mediated by the charged ($W^{\pm}$) and neutral (Z^0) weak gauge bosons. The weak interaction Lagrangian density can be written as [85]

$$\mathcal{L}_{\text{int}}^{\text{weak}} = -\frac{g}{2\sqrt{2}}\left(j_{\mu}^{CC}W^{\mu+} + \text{h.c.}\right) - \frac{g}{2\cos\theta_W} j_{\mu}^{NC} Z^{\mu}, \tag{5.50}$$

where g is the weak coupling constant and θ_W is the Weinberg angle.

The weak mixing angle satisfies

$$\sin\theta_W = \frac{e}{g}, \tag{5.51}$$

and the coupling constants are related to the Fermi constant G_F through

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8M_W^2}. \tag{5.52}$$

The first term corresponds to charged-current (CC) interactions, while the second describes neutral-current (NC) interactions.

For oscillation measurements in NOvA, charged-current interactions are of primary importance. In CC processes, the outgoing lepton carries electric charge and can be directly detected, allowing identification of the neutrino flavor. In contrast, neutral-current interactions produce a neutrino in the final state, making flavor identification impossible.

Fig. 5.3 illustrates the fundamental weak interaction vertices for CC and NC processes.

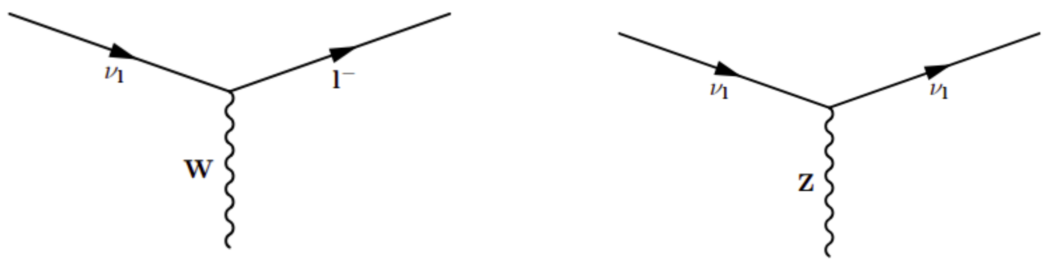

Figure 5.3: Schematic representation of weak interaction vertices. Left: charged-current interaction mediated by $W^{\pm}$. Right: neutral-current interaction mediated by Z^0.

5.2.2 Interactions on Free Nucleons

Depending on the neutrino energy, different interaction mechanisms dominate. In the few-GeV energy region relevant to NOvA, three primary charged-current processes contribute:

- Charged-Current Quasielastic (CCQE)
- Charged-Current Resonance Production (CCRES)
- Charged-Current Deep Inelastic Scattering (CCDIS)

We first describe these processes assuming scattering from a free nucleon.

Charged-Current Quasielastic (CCQE) Scattering

In CCQE interactions, the neutrino scatters from a single nucleon via $W^{\pm}$ exchange:

$$\nu_l + n \to l^- + p, \tag{5.53}$$

$$\bar{\nu}_l + p \to l^+ + n, \tag{5.54}$$

where $l = e, \mu, \tau$.

Let k and k' denote the four-momenta of the incoming neutrino and outgoing lepton, and p and p' the four-momenta of the initial and final nucleon. The four-momentum transfer is

$$q = k - k', \qquad Q^2 = -q^2. \tag{5.55}$$

The leading-order matrix element is

$$\mathcal{M} = \frac{G_F}{\sqrt{2}} \cos\theta_C \; l_\mu J^\mu, \tag{5.56}$$

where θ_C is the Cabibbo angle, l_μ is the leptonic current,

$$l_\mu = \bar{u}(k')\gamma_\mu(1-\gamma_5)u(k), \tag{5.57}$$

and J^μ is the hadronic current.

The differential cross section can be expressed in terms of the Mandelstam variables [86]:

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2 M^2 \cos^2\theta_C}{8\pi E_\nu^2}\left[A(Q^2) \mp B(Q^2)\frac{s-u}{M^2} + C(Q^2)\frac{(s-u)^2}{M^4}\right], \tag{5.58}$$

where the upper (lower) sign corresponds to neutrinos (antineutrinos).

The hadronic current is parameterized using vector and axial form factors. The axial form factor is commonly modeled using a dipole form [87]:

$$F_A(Q^2) = \frac{g_A}{\left(1 + \frac{Q^2}{M_A^2}\right)^2}, \tag{5.59}$$

where g_A is the axial coupling constant and M_A is the axial mass.

Global analyses of deuterium and electroproduction data yield

$$M_A \approx 1.03 \text{ GeV}. \tag{5.60}$$

Alternative parameterizations such as the z-expansion have been proposed to reduce model dependence.

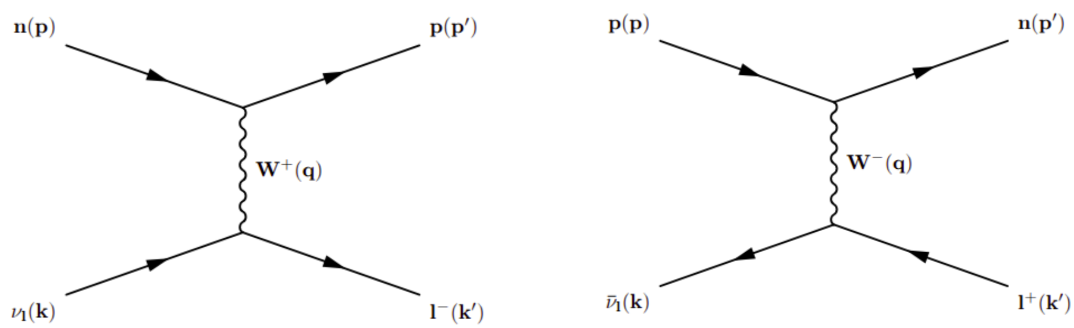

Figure 5.4: Feynman diagrams for charged-current quasielastic scattering of neutrinos (left) and antineutrinos (right).

Charged-Current Resonance Production

At higher neutrino energies, the interaction can excite the nucleon to a baryon resonance state. The dominant contribution in the few-GeV region arises from the $\Delta(1232)$ resonance [88].

The produced resonance subsequently decays, typically emitting a pion:

$$\nu_l + p \to l^- + \Delta^{++} \to l^- + p + \pi^+, \tag{5.61}$$

$$\nu_l + n \to l^- + \Delta^{+} \to l^- + p + \pi^0, \tag{5.62}$$

and analogous channels for antineutrinos.

The cross section depends on vector and axial transition form factors $C_i^V(Q^2)$ and $C_i^A(Q^2)$.

Charged-Current Deep Inelastic Scattering

At sufficiently high energies ($E_\nu \gg m_N$), neutrinos interact with individual quarks inside the nucleon:

$$\nu_l + N \to l^- + X, \tag{5.63}$$

$$\bar{\nu}_l + N \to l^+ + X. \tag{5.64}$$

The relevant kinematic variables are the Bjorken scaling variables:

$$x = \frac{Q^2}{2p \cdot q}, \qquad y = \frac{p \cdot q}{p \cdot k}. \tag{5.65}$$

The double differential cross section is [89]

$$\frac{d^2\sigma^{CC}}{dxdy} = \frac{G_F^2 s}{2\pi}\left(1 + \frac{Q^2}{M_W^2}\right)^{-2}\left[xy^2F_1 + (1-y)F_2 \pm xy\left(1 - \frac{y}{2}\right)F_3\right]. \tag{5.66}$$

5.2.3 Nuclear Medium Effects [6, 7]

In realistic detectors, nucleons are bound inside nuclei and are not free or stationary. Nuclear effects significantly modify both the interaction rate and the observed final state topology.

Important nuclear effects include:

- Fermi motion of bound nucleons
- Nuclear binding energy
- Pauli blocking

- Long-range correlations (RPA effects)
- Short-range correlations
- Meson-exchange currents (2p2h processes)
- Final State Interactions (FSI)

These effects alter reconstructed kinematics and can cause misidentification of interaction channels.

5.2.4 Final State Interactions

After the primary interaction, hadrons propagate through the residual nucleus. They may undergo rescattering, charge exchange, absorption, or secondary particle production.

Final State Interactions alter the observable final state, modifying both particle multiplicity and energy reconstruction. Accurate modeling of FSI is therefore crucial for reliable neutrino energy estimation and oscillation measurements.

5.3 Quasielastic Interaction Models

As discussed in the previous section, CCQE scattering plays a dominant role in neutrino interactions at energies of a few GeV, which are relevant for long-baseline experiments such as NOvA. In this interaction, a neutrino scatters from a nucleon inside the nucleus through charged-current weak interaction and produces a charged lepton and a nucleon in the final state.

However, the interpretation of CCQE events in nuclear targets is complicated due to nuclear medium effects such as Fermi motion, nucleon correlations, Pauli blocking, and meson-exchange currents. Therefore, accurate modeling of quasielastic scattering requires sophisticated nuclear models.

In this work, two different quasielastic interaction models are considered:

- Valencia 1p1h model with z-expansion axial form factor
- Hartree-Fock Continuum Random Phase Approximation (HF-CRPA)

These models differ in their treatment of nuclear dynamics and correlations, and therefore lead to different predictions for neutrino cross sections and reconstructed neutrino energy distributions.

5.3.1 Valencia 1p1h Model with z-Expansion Axial Form Factor

The Valencia model [7, 90] is a microscopic nuclear model developed by the Valencia group to describe neutrino interactions with nuclei in the quasielastic regime. It includes a detailed treatment of nuclear medium effects within the framework of a Local Fermi Gas (LFG) model and incorporates long-range correlations through the Random Phase Approximation (RPA).

In the quasielastic process, the neutrino interacts with a single nucleon inside the nucleus as shown in Eq 5.53 (5.54). Since only one nucleon participates in the interaction, this process is commonly referred to as the **one-particle one-hole (1p1h)** excitation.

Local Fermi Gas Description

In the Valencia approach, nucleons inside the nucleus are described using the Local Fermi Gas model. In this model, nucleons occupy momentum states up to the local Fermi momentum $p_F(r)$, which depends on the nuclear density:

$$p_F(r) = \left(3\pi^2 \rho(r)\right)^{1/3}, \tag{5.67}$$

where $\rho(r)$ represents the nuclear density distribution.

The differential cross section for quasielastic scattering can be written as

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{G_F^2 \cos^2\theta_C}{16\pi^2}\frac{|k'|}{|k|}L_{\mu\nu}W^{\mu\nu}, \tag{5.68}$$

where

- $L_{\mu\nu}$ is the leptonic tensor
- $W^{\mu\nu}$ is the hadronic tensor describing the nuclear response

The hadronic tensor contains the nuclear response functions that depend on the momentum transfer and nuclear structure.

RPA Corrections

The Valencia model incorporates long-range nucleon correlations through the RPA [7]. These correlations arise from nucleon-nucleon interactions inside the nuclear medium and modify the nuclear response functions.

The RPA correction effectively modifies the weak interaction propagator through particle-hole excitations:

$$W^{\mu\nu}_{RPA} = \frac{W^{\mu\nu}_0}{1 - V_{ph}\Pi(q)}, \tag{5.69}$$

where

- $W^{\mu\nu}_0$ is the bare nuclear response
- V_{ph} represents the particle-hole interaction
- $\Pi(q)$ is the polarization propagator

These corrections generally suppress the cross section at low momentum transfer.

z-Expansion Axial Form Factor

Traditionally, the axial form factor was parameterized using the dipole form

$$F_A(Q^2) = \frac{g_A}{(1 + Q^2/M_A^2)^2}. \tag{5.70}$$

However, this parameterization introduces model dependence. A more model-independent approach is the z-expansion method, which is based on analytic properties of form factors.

In this method, the axial form factor is expanded as [91]

$$F_A(Q^2) = \sum_{k=0}^{\infty} a_k z(Q^2)^k, \tag{5.71}$$

where the conformal variable z is defined as

$$z(Q^2) = \frac{\sqrt{t_c + Q^2} - \sqrt{t_c - t_0}}{\sqrt{t_c + Q^2} + \sqrt{t_c - t_0}}, \tag{5.72}$$

with t_c representing the threshold for particle production.

The z-expansion allows a flexible description of the axial form factor without imposing a specific functional form.

5.3.2 Hartree-Fock Continuum Random Phase Approximation (HF-CRPA)

The Hartree-Fock Continuum Random Phase Approximation (HF-CRPA) model [92, 93, 94] provides a more sophisticated description of quasielastic neutrino scattering in nuclei. Unlike the Local Fermi Gas approach, this model uses a mean-field description of nuclear structure and explicitly treats nucleon wavefunctions within the nuclear potential.

Hartree-Fock Mean Field

In the Hartree-Fock framework, nucleons move independently in an average nuclear potential generated by all other nucleons. The single-particle wavefunctions are obtained by solving the Hartree-Fock equation:

$$\hat{H}_{HF}\psi_i = E_i\psi_i, \tag{5.73}$$

where $\hat{H}_{HF}$ is the Hartree-Fock Hamiltonian and ψ_i represents the nucleon wavefunctions.

This approach provides a realistic description of nuclear shell structure and nucleon momentum distributions.

Continuum Random Phase Approximation

The CRPA framework includes long-range correlations between nucleons through particle-hole excitations in the continuum.

The nuclear response function can be written as

$$R(q,\omega) = -\frac{1}{\pi}\mathrm{Im}\left[\Pi(q,\omega)\right], \tag{5.74}$$

where $\Pi(q,\omega)$ represents the polarization propagator describing particle-hole excitations in the nuclear medium.

The CRPA equations can be written schematically as

$$\Pi_{CRPA} = \Pi_0 + \Pi_0 V \Pi_{CRPA}, \tag{5.75}$$

where

- Π_0 is the bare particle-hole propagator
- V is the residual nucleon-nucleon interaction

This formalism allows the inclusion of collective nuclear excitations that modify the neutrino scattering cross section. Compared to simpler nuclear models, HF-CRPA provides several improvements:

- Realistic nuclear shell structure
- Proper treatment of nucleon wavefunctions
- Inclusion of long-range correlations
- Improved description of low momentum transfer interactions

These features make HF-CRPA particularly suitable for describing quasielastic neutrino interactions in the few-GeV energy range.

The different treatments of nuclear dynamics in the Valencia and HF-CRPA models lead to differences in predicted cross sections and reconstructed kinematic distributions. These differences can affect neutrino energy reconstruction and, therefore, the extraction of oscillation parameters.

In the next chapter, we present a detailed analysis comparing these two models and measure the impact on the oscillation parameters.

Chapter 6

Cross-section modeling for Quasi-Elastic Interaction

Precision measurements of neutrino oscillation parameters require accurate modeling of neutrino–nucleus interactions. In long-baseline experiments such as NOvA, the reconstructed neutrino energy — the primary observable used in oscillation fits — depends directly on the assumed interaction model. Therefore, any modification to the cross-section model can propagate to the extracted PMNS parameters.

CCQE interactions play a central role in the NOvA energy regime ($E_\nu \sim$ 1–3 GeV), particularly in events with low hadronic energy. The HF-CRPA model provides an improved treatment of long-range nuclear correlations compared to the NOvA nominal model. This may modify both the normalization and kinematic distributions of QE events. This model has been incorporated in GENIE quite recently [95].

6.1 Cross-Section Comparison

The NOvA base simulation sample was generated using GENIE [96] version 3.0.6 (N18_10j_00_000), which employs the Valencia quasielastic interaction model with the standard 1p1h description.

To study the impact of improved nuclear modeling, an alternative simulation sample was generated using GENIE G21_11a_00_000, which includes the `CRPASuSAv2Hybrid-QEL` model incorporating the Hartree–Fock Continuum Random Phase Approximation (HF-CRPA) framework. This model provides a more detailed treatment of nuclear dynamics, particularly long-range nucleon correlations, which are expected to modify the quasielastic cross section especially at low momentum transfer.

Fig. 6.1 shows a comparison of the CCQE cross sections predicted by the two models for different neutrino flavors: ν_μ, ν_e, $\bar{\nu}_\mu$, and $\bar{\nu}_e$.

The HF-CRPA model predicts a systematic enhancement in the cross section relative to the NOvA base simulation. Across the relevant neutrino energy range, the cross section is increased for all flavors. This enhancement originates from the improved treatment of nuclear medium effects and long-range correlations included in the HF-CRPA framework.

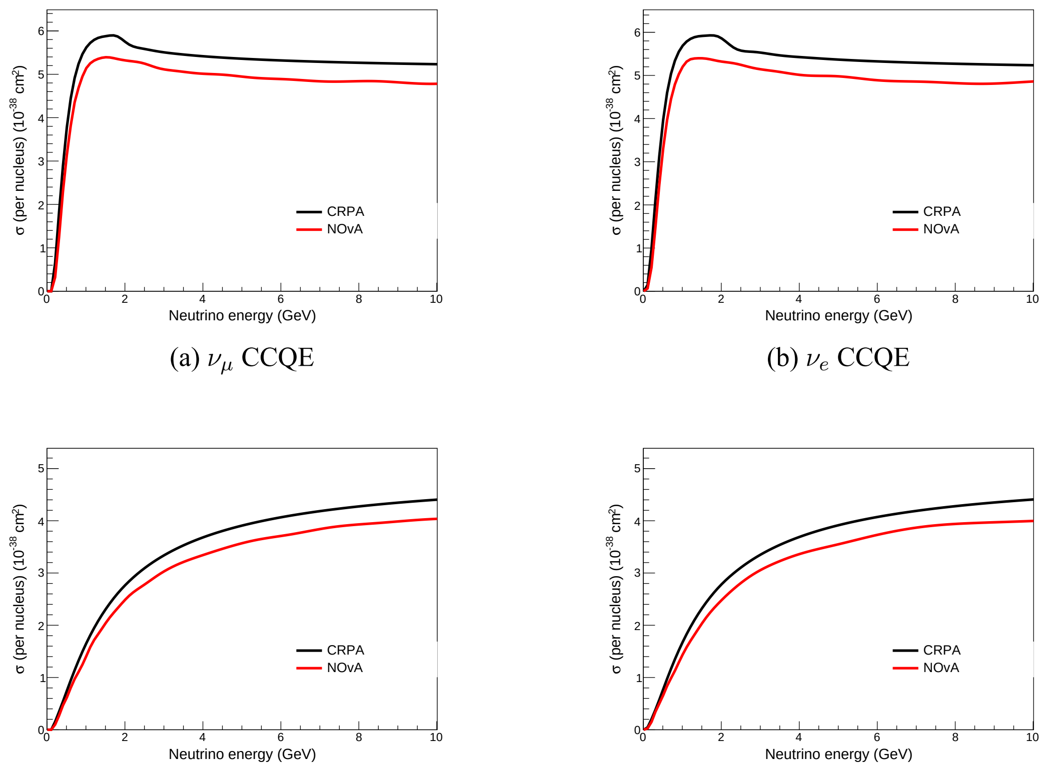

(a) ν_μ CCQE

(b) ν_e CCQE

(c) $\bar{\nu}_\mu$ CCQE

(d) $\bar{\nu}_e$ CCQE

Figure 6.1: Comparison of CCQE cross sections predicted by the NOvA base simulation (GENIE v3.0.6) and the HF-CRPA implementation (GENIE G21_11a_00_000) for different neutrino flavors.

6.2 Comparison of Truth-Level Distributions

Using the cross-section inputs described in Section 2, a total of 6.5 million events were generated for each neutrino flavor for both interaction models. The generated events were then normalized using their corresponding cross sections in order to allow a direct comparison between the predictions of the two models.

Fig. 6.2–6.5 show the comparison between the NOvA base simulation and the HF-CRPA model for different neutrino flavors. The distributions are shown as a function of several kinematic variables:

- neutrino energy (E_ν),
- angle between the outgoing lepton and the incoming neutrino,
- energy transfer (q_0),
- three-momentum transfer (q_3),
- squared four-momentum transfer (Q^2).

These variables provide important information about the underlying interaction dynamics and nuclear effects.

Across all neutrino and antineutrino flavors, a consistent trend is observed between the two models. The HF-CRPA model predicts an enhancement of events primarily in the region of low energy transfer (q_0) and low momentum transfer (q_3). This behavior is consistent with the improved treatment of long-range nuclear correlations and collective nuclear effects incorporated in the HF-CRPA framework. Similar enhancements are also reflected in the low-Q^2 region of the distributions.

The concentration of the HF-CRPA enhancement at low momentum transfer is related to the different treatment of nuclear effects in the two models. The HF-CRPA calculation incorporates long-range correlations in the nuclear response, which modify the strength and distribution of the quasielastic interaction, particularly in the low-energy- and low-momentum-transfer region. Thus, the difference between the two models is primarily a redistribution of the predicted CCQE strength across the kinematic phase space rather than a uniform increase in the interaction rate.

The HF-CRPA model predicts an enhancement of event strength primarily in the low-q0 and low-q3 region. This difference reflects the modified nuclear response arising from the treatment of long-range correlations in the HF-CRPA framework. For example, in the lowest q0 interval, the predicted event an enhanced by approximately 20-40%. In contrast, in the higher-q0 region, the ratio is almost close to 1.

6.3 Ratio-Based Modification

To incorporate the HF-CRPA nuclear effects into the NOvA base simulation, a ratio-based reweighting approach is employed. The idea is to compute the ratio of the HF-CRPA prediction to the NOvA base model in the (q_0, q_3) phase space and apply this ratio as an event weight to the NOvA simulation. The event reweighting procedure can be written formally as

$$w(q_0, q_3) = \frac{\sigma_{\text{HF-CRPA}}(q_0, q_3)}{\sigma_{\text{NOvA}}(q_0, q_3)}, \tag{6.1}$$

where $w(q_0, q_3)$ is the weight applied to each event in the NOvA base simulation, and $\sigma_{\text{HF-CRPA}}(q_0, q_3)$ and $\sigma_{\text{NOvA}}(q_0, q_3)$ represent the differential cross sections predicted by the HF-CRPA and NOvA interaction models respectively in the (q_0, q_3) phase space.

For each simulated NOvA event, the corresponding (q_0, q_3) bin is identified and the event weight is multiplied by $w(q_0, q_3)$. In this way, the NOvA base simulation can be modified to reproduce the kinematic distributions predicted by the HF-CRPA model while

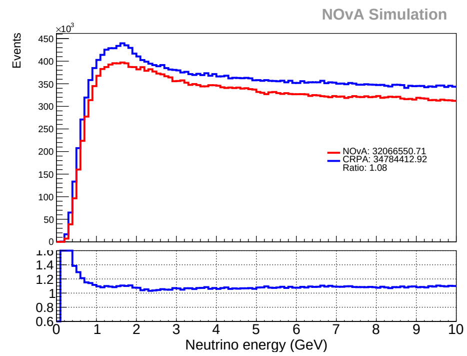

(a) Neutrino energy (E_ν)

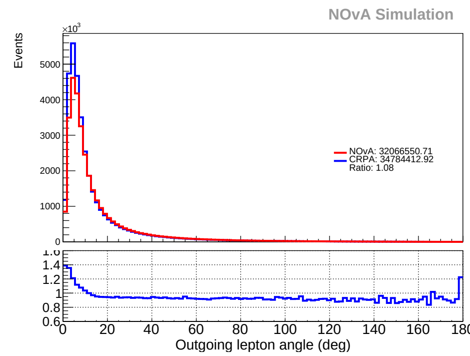

(b) Outgoing lepton angle

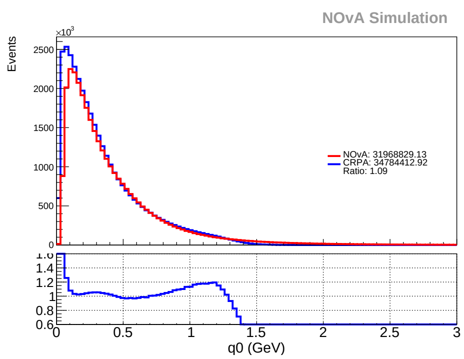

(c) Energy transfer (q_0)

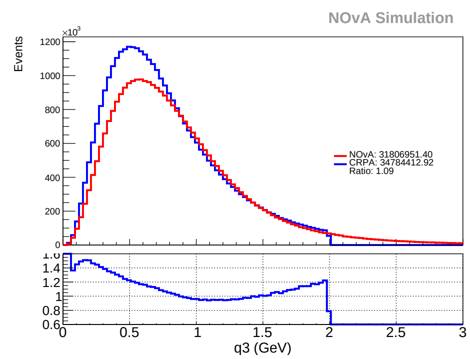

(d) Momentum transfer (q_3)

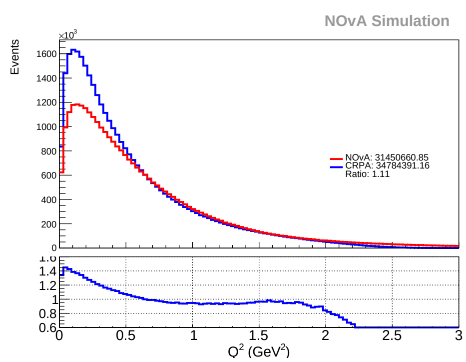

(e) Squared four-momentum transfer (Q^2)

Figure 6.2: Comparison of truth-level kinematic distributions for ν_μ charged-current quasielastic interactions predicted by the NOvA base simulation and the HF-CRPA model.

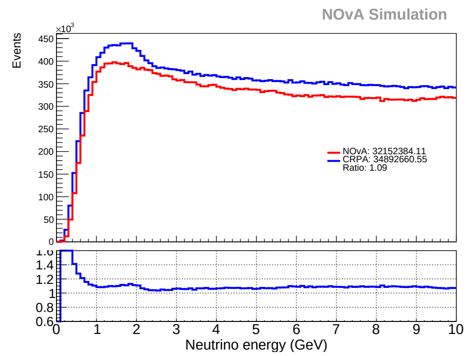

(a) Neutrino energy (E_ν)

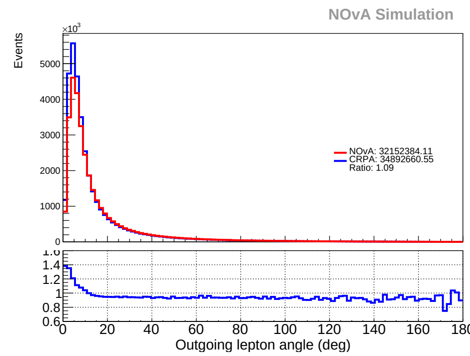

(b) Outgoing lepton angle

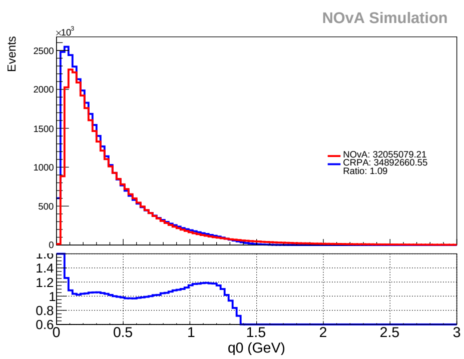

(c) Energy transfer (q_0)

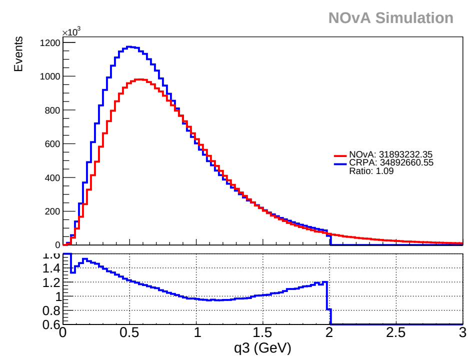

(d) Momentum transfer (q_3)

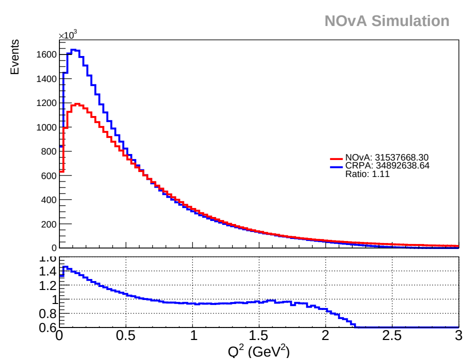

(e) Squared four-momentum transfer (Q^2)

Figure 6.3: Comparison of truth-level kinematic distributions for ν_e charged-current quasielastic interactions predicted by the NOvA base simulation and the HF-CRPA model.

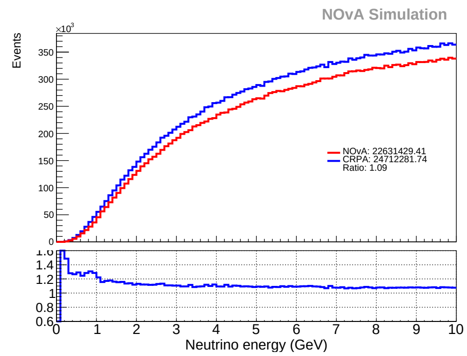

(a) Neutrino energy (E_ν)

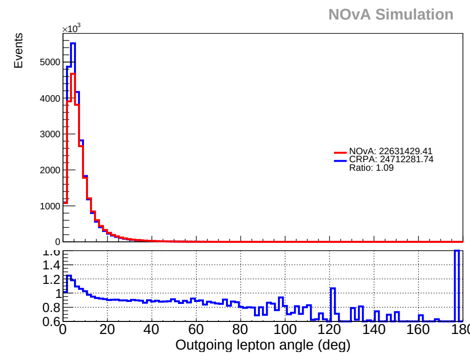

(b) Outgoing lepton angle

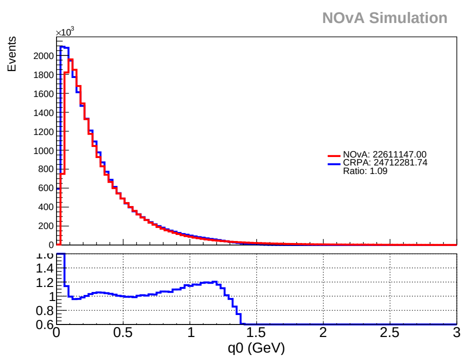

(c) Energy transfer (q_0)

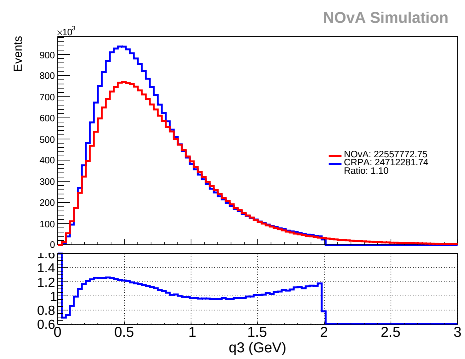

(d) Momentum transfer (q_3)

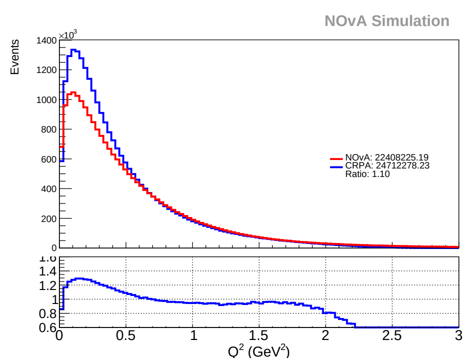

(e) Squared four-momentum transfer (Q^2)

Figure 6.4: Comparison of truth-level kinematic distributions for $\bar{\nu}_\mu$ charged-current quasielastic interactions predicted by the NOvA base simulation and the HF-CRPA model.

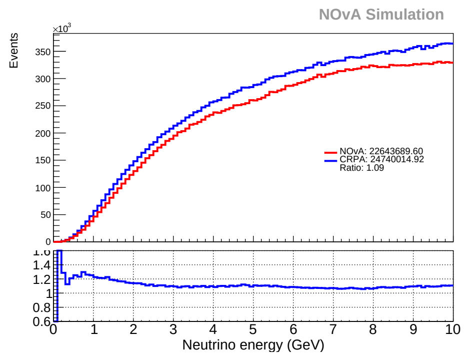

(a) Neutrino energy (E_ν)

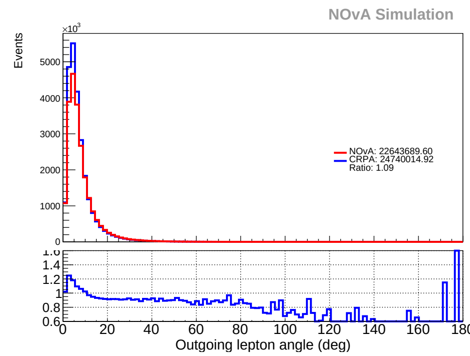

(b) Outgoing lepton angle

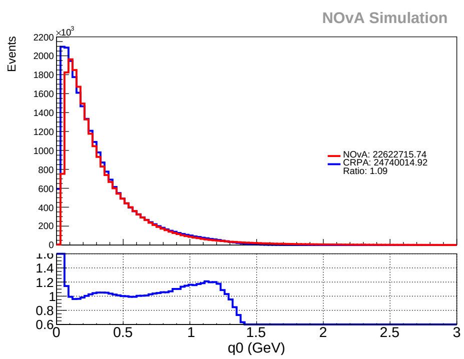

(c) Energy transfer (q_0)

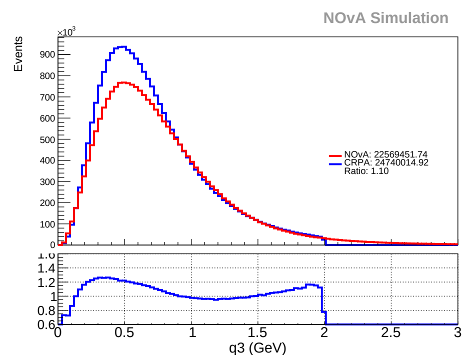

(d) Momentum transfer (q_3)

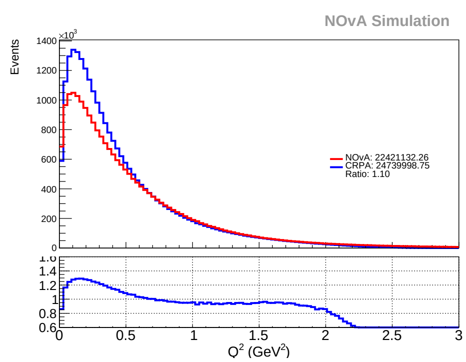

(e) Squared four-momentum transfer (Q^2)

Figure 6.5: Comparison of truth-level kinematic distributions for $\bar{\nu}_e$ charged-current quasielastic interactions predicted by the NOvA base simulation and the HF-CRPA model.

preserving the underlying event structure of the original simulation.

Fig. 6.6 shows the (q_0, q_3) distributions for the ν_μ flavor predicted by the NOvA base model and the HF-CRPA model. It can be observed that the HF-CRPA model opens up additional phase space in the kinematically allowed region, particularly at very low values of q_0. As a result, significant activity appears in the lowest q_0 bins compared to the NOvA base simulation.

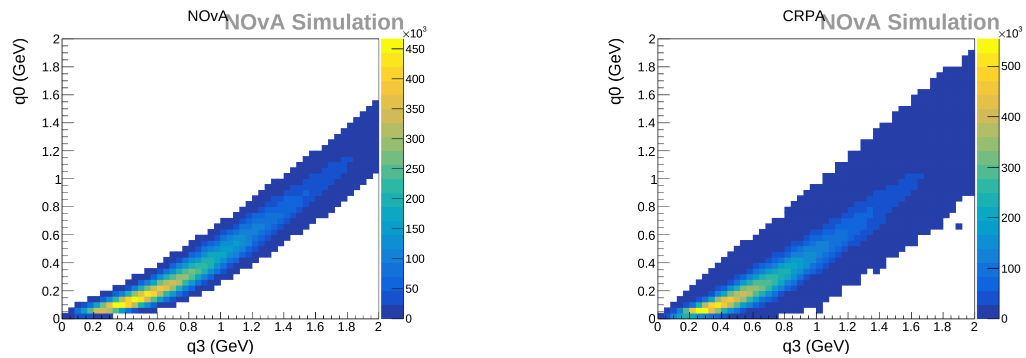

(a) NOvA base simulation

(b) HF-CRPA model

Figure 6.6: (q_0, q_3) phase-space distribution for ν_μ charged-current quasielastic interactions predicted by the NOvA base simulation and the HF-CRPA model.

The ratio of HF-CRPA to the NOvA prediction in the (q_0, q_3) phase space is shown in Fig. 6.7. Due to the enhanced population in the lowest q_0 bins for the HF-CRPA model, very large ratios appear near the boundaries of the kinematically allowed region.

When this ratio was directly applied to reweight the NOvA events, significant fluctuations appeared in the resulting distributions. In particular, sharp structures were observed in the reweighted spectra, as illustrated in Fig. 6.8a. The sudden drop around $q_3 \approx 300$ MeV arises from the rapid transition between a very large ratio and a much smaller ratio in the adjacent q_3 bin, as seen in Fig. 6.8b. This behavior is mainly caused by limited statistics in the lowest q_0 bins of the NOvA simulation.

To mitigate this issue, a modified ratio scheme was introduced. For the first four bins in

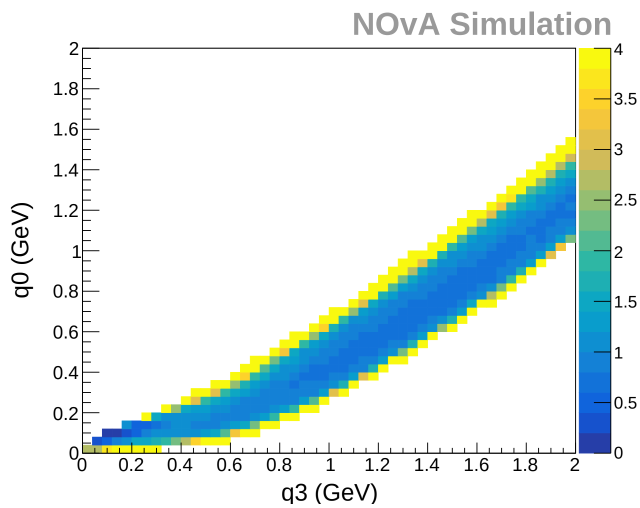

Figure 6.7: Ratio of HF-CRPA to NOvA predictions in the (q_0, q_3) phase space for ν_μ interactions.

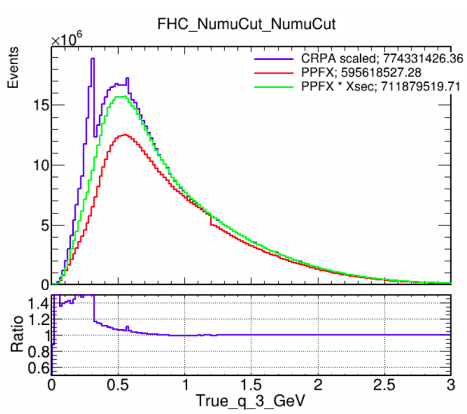

(a) CRPA-weighted distribution compared with the NOvA base model

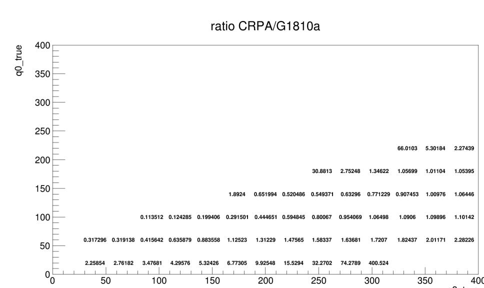

(b) HF-CRPA/NOvA ratio in (q_0, q_3) phase space

q_0, if the number of events in the NOvA simulation is very small or if the HF-CRPA/NOvA ratio exceeds 200, the events in that bin are merged with the subsequent q_0 bin for both the HF-CRPA and NOvA samples before computing the ratio. This procedure effectively smooths the ratio in the low-q_0 region and reduces large statistical fluctuations.

After applying this modification, the reweighted distributions show significantly improved agreement in both shape and normalization, as demonstrated in Fig. 6.9.

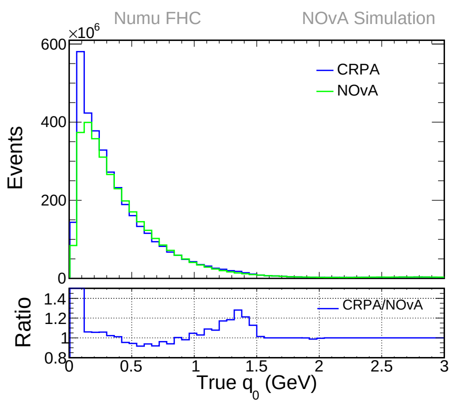

(a) True q_0 distribution

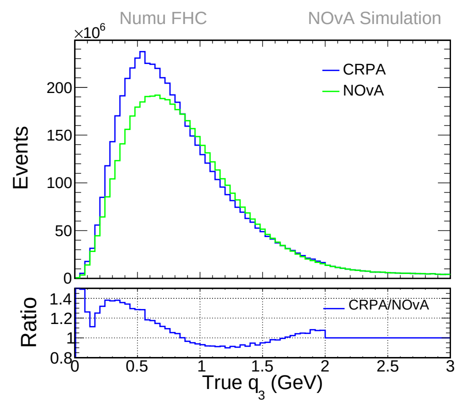

(b) True q_3 distribution

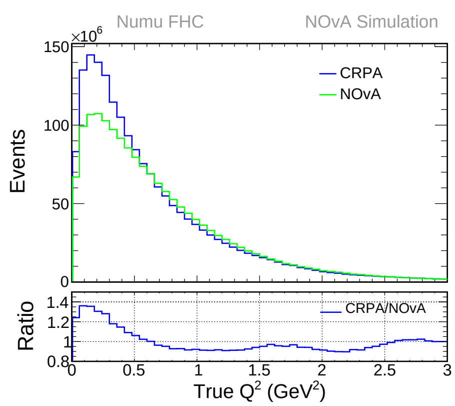

(c) True Q^2 distribution

Figure 6.9: Truth-level kinematic distributions comparing the NOvA base simulation and the HF-CRPA reweighted events after applying the modified ratio scheme.

6.4 Event Selection

The identification of neutrino interactions in the NOvA detectors requires efficient rejection of background events while maintaining high signal efficiency. Backgrounds in the ν_μ and ν_e selections arise from two main sources: cosmic-ray induced events and beam-related neutrino interactions. Cosmic backgrounds are particularly relevant for the far detector because it operates at the surface. Since neutrinos do not deposit energy in the detector prior to interacting, cosmic-ray events often have reconstructed vertices close to the detector boundaries. This property is used to suppress cosmics through fiducial-volume and containment requirements. In addition, events must be consistent with the direction of the NuMI beam in order to be considered beam-related interactions.

Fig. 6.10 illustrates the overall event selection strategy for both ν_μ disappearance and ν_e appearance channels used in the oscillation analysis [97].

6.4.1 ν_μ Disappearance Selection

The ν_μ charged-current (CC) sample provides the primary constraint for the $\nu_\mu \rightarrow \nu_\mu$ disappearance measurement. This sample is relatively pure because the outgoing muon typically produces a long, well-defined track in the detector.

The selection begins with a set of basic quality requirements ensuring sufficient detector activity, including the presence of a reconstructed track and a minimum number of hits distributed across several detector planes. Events clearly identified as cosmic rays, such as tracks entering or exiting through detector walls, are rejected using dedicated veto criteria. Additional containment requirements ensure that the interaction vertex and reconstructed prongs are sufficiently far from detector boundaries, particularly the top surface where cosmic contamination is largest.

Following these geometric selections, several particle-identification (PID) algorithms

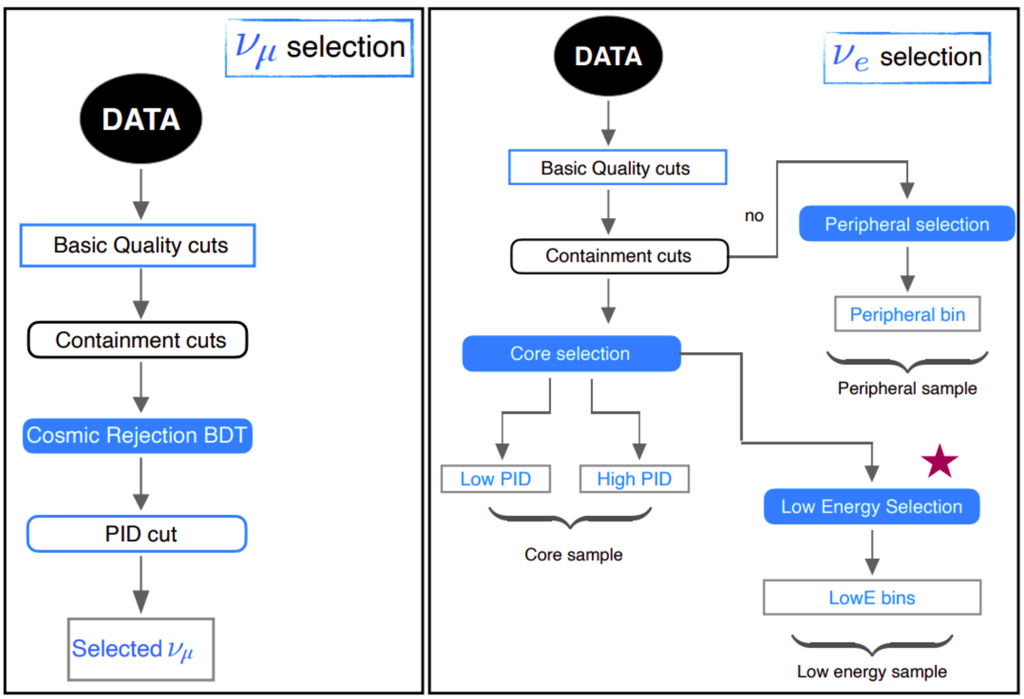

Figure 6.10: Schematic overview of the NOvA event selection procedure for ν_μ and ν_e samples.

are applied. A convolutional neural network classifier (EventCVN) provides a probability that the event corresponds to a ν_μ CC interaction. Additional multivariate classifiers, including a boosted decision tree (BDT) and the RemID algorithm, further suppress cosmic-ray and neutral-current backgrounds using track and event-level variables. The final PID thresholds are optimized to maximize the signal sensitivity of the ν_μ sample.

6.4.2 ν_e Appearance Selection

The ν_e appearance selection targets charged-current ν_e interactions produced through the oscillation $\nu_\mu \to \nu_e$. Compared to the ν_μ sample, this selection is more challenging due to the smaller signal rate and the presence of several backgrounds. Beam-related backgrounds include neutral-current interactions, misidentified ν_μ CC events, and intrinsic ν_e present in the neutrino beam.

The primary signal signature is an electromagnetic shower produced by the outgoing electron. Machine-learning techniques are therefore used to distinguish these showers from hadronic activity or photon conversions originating from neutral-current interactions.

The ν_e selection is divided into three complementary samples:

- **Core sample:** This sample forms the main contribution to the oscillation sensitivity. Events passing basic quality and containment requirements are evaluated using the EventCVN ν_e classifier and an additional BDT trained to separate signal from background. The selected events are further divided into low- and high-PID categories based on the classifier output.

- **Peripheral sample:** Events that fail the containment requirement but otherwise satisfy quality criteria are included in the Peripheral sample. Although these events have higher background contamination, dedicated classifiers are used to recover additional signal acceptance.

- **Low-energy (LowE) sample:** An additional sample targets events at lower reconstructed energies that may fail the main PID requirements. A separate BDT trained on event topology and classifier outputs is used to identify these candidates. This category increases the sensitivity to oscillation parameters by recovering signal events near the lower-energy region of the neutrino spectrum.

The combination of these three samples maximizes the signal efficiency for ν_e appearance while maintaining adequate background suppression.

6.5 Energy Reconstruction

Precise determination of neutrino oscillation parameters requires measuring the event rate as a function of neutrino energy at a fixed baseline. Therefore, an accurate reconstruction of the neutrino energy is a crucial component of the analysis. The calibration procedures used to convert detector signals into deposited energy were described in the previous chapter. After event reconstruction and particle identification, dedicated energy estimation methods are applied separately for ν_e and ν_μ charged-current interactions. These methods exploit the topology of the final-state particles to achieve better performance than a simple calorimetric energy sum.

6.5.1 ν_e Energy Reconstruction

For ν_e charged-current interactions, the visible energy in the detector arises primarily from two components: the electromagnetic shower produced by the outgoing electron and the hadronic activity originating from the recoil system. The reconstructed neutrino energy is therefore expressed as a function of these two contributions.

The ν_e energy estimator is parameterized as

$$E_\nu = p_0 \left(p_1 E_{\text{EM}} + p_2 E_{\text{had}} + p_3 E_{\text{EM}}^2 + p_4 E_{\text{had}}^2 \right), \tag{6.2}$$

where E_{EM} and E_{had} denote the energies of the electromagnetic and hadronic components respectively. These quantities are obtained by summing the calibrated visible energy of the reconstructed prongs identified as electromagnetic or hadronic using the prong-level convolutional neural network classifier (ProngCVN).

The coefficients p_i are determined from simulation by fitting the parameterization to the relationship between reconstructed quantities and the true neutrino energy. To avoid bias toward the region near the beam energy peak, the fitting procedure employs weights that flatten the true energy spectrum. The calibration is performed separately for FHC and RHC running conditions. Using this method, the reconstructed ν_e energy achieves an overall resolution of approximately 10.8% for FHC data and 8.5% for RHC data.

6.5.2 ν_μ Energy Reconstruction

In ν_μ charged-current interactions, the outgoing muon typically forms a long and well-defined track in the detector. The muon energy can therefore be estimated from the track length, which provides a more precise measurement than calorimetric energy deposition alone. The uncertainty in the reconstructed track length is small compared to the overall detector energy scale uncertainty, enabling a relatively accurate determination of the muon energy.

The hadronic component of the interaction is reconstructed by summing the visible energy of detector hits that are not associated with the muon track. The neutrino energy is then estimated as

$$E_\nu = E_\mu + E_{\text{had}}, \tag{6.3}$$

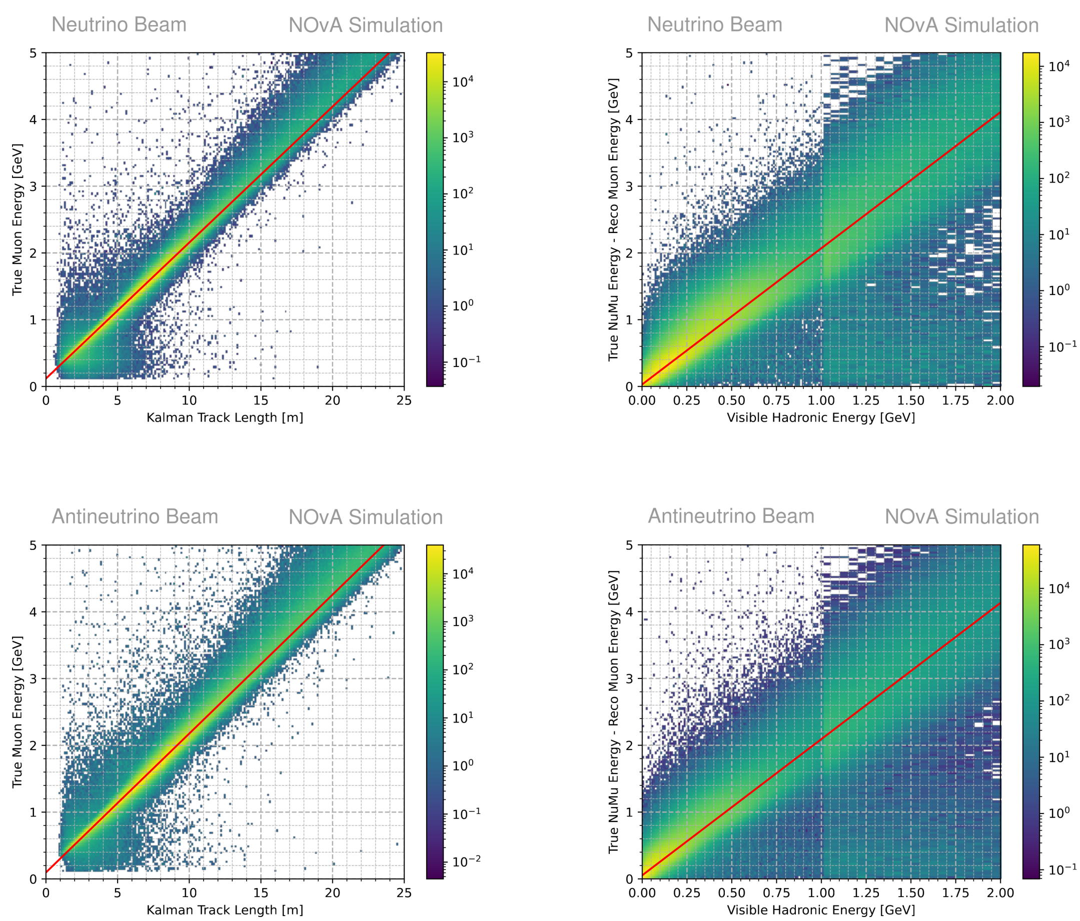

Figure 6.11: Relationship between reconstructed observables and true particle energies obtained from simulation. The red lines show the linear fits used to calibrate the reconstructed quantities to the corresponding true energies for FHC (top) and RHC (bottom) beam configurations. The left panels correspond to the muon component, where the reconstructed variable is the Kalman track length, while the right panels show the hadronic component based on the visible hadronic energy. The narrower spread observed in the muon distributions indicates the better energy resolution achievable for the muon component compared to the hadronic energy. Taken from [26].

where E_μ represents the muon energy derived from the track length and E_{had} corresponds to the visible hadronic energy.

Both the muon and hadronic energy estimators are further calibrated using simulated events. Two-dimensional histograms are constructed relating reconstructed quantities, such as Kalman track length or visible hadronic energy, to the corresponding true particle energies. For each bin of the reconstructed variable, the distribution of true energy is examined and its most probable value is determined after removing small tails of the distribution. A Gaussian fit is then applied to obtain the mean true energy associated with that bin.

These mean values are subsequently used as points in a linear calibration fit that maps the reconstructed variable to the true energy. The procedure is illustrated in Fig. 6.11, which shows the resulting relationships for both the muon and hadronic components in FHC and RHC modes. The narrower spread of the muon distributions demonstrates the superior energy resolution achieved for the muon component compared to the hadronic energy.

6.6 Effect on ND Reconstructed Energy

The HF-CRPA model predicts an enhancement of approximately $\sim 10\%$ in the CCQE cross section compared to the NOvA nominal interaction model. However, CCQE interactions represent only a fraction of the total neutrino interaction rate in the NOvA energy range, where quasielastic, resonant, and deep inelastic scattering processes all contribute. Consequently, when all interaction channels are considered, the overall change in the event rate due to the CRPA model is reduced to approximately 3–4% in the true energy distributions.

The ν_μ and ν_e event selections used in this analysis were described in Section **??**. After applying these selections to events in the ND, the CRPA reweighting described in the previous section is applied to study its impact on reconstructed neutrino energy.

Fig. 6.12 shows the reconstructed neutrino energy spectra for the ν_μ sample at the ND for both FHC and RHC running modes. The comparison between the NOvA nominal simulation and the CRPA-reweighted events indicates that the effect of the CRPA model on reconstructed energy is smaller than that observed at the truth level. This reduction occurs because the reconstructed energy combines contributions from both the muon and hadronic components, which partially smear the differences introduced by the interaction model.

The corresponding ratios of CRPA to the NOvA nominal prediction are also shown. The largest deviations appear in the lowest energy bins, where the CCQE contribution is most significant. However, these variations remain within the 1σ systematic uncertainty band of the nominal NOvA prediction.

A similar comparison for the ν_e sample is shown in Fig. 6.13. The overall impact of the CRPA reweighting is again small, with only minor variations observed in the lowest reconstructed energy bins. As in the ν_μ case, these differences remain well within the systematic uncertainty envelope of the nominal prediction.

6.7 Extrapolated prediction at the FD

In long-baseline neutrino oscillation experiments such as NOvA, the ND plays a crucial role in constraining systematic uncertainties associated with the neutrino flux, interaction cross sections, and detector response. Because the ND and FD are designed to be functionally identical, measurements at the ND can be used to improve the prediction of the event spectra at the FD through a procedure known as *near-to-far extrapolation*.

The goal of the extrapolation procedure is to use the high statistics ND data to correct the simulated prediction and propagate this information to the FD. This approach allows many systematic effects that are correlated between the two detectors to be significantly reduced.

The extrapolation procedure begins by decomposing the ND Monte Carlo (MC) predic-

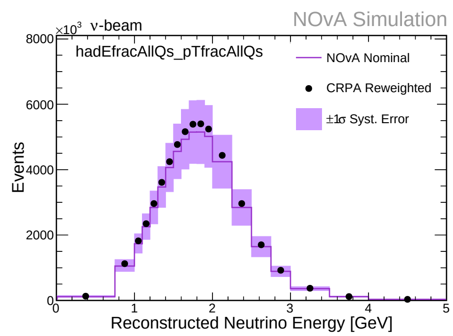

(a) ν_μ reconstructed energy in FHC running.

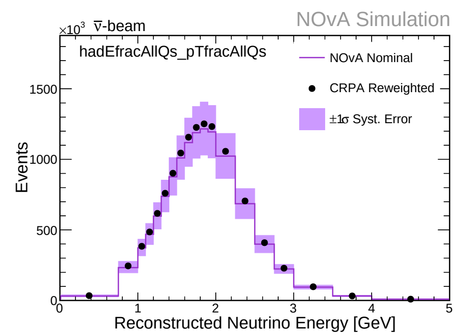

(b) ν_μ reconstructed energy in RHC running.

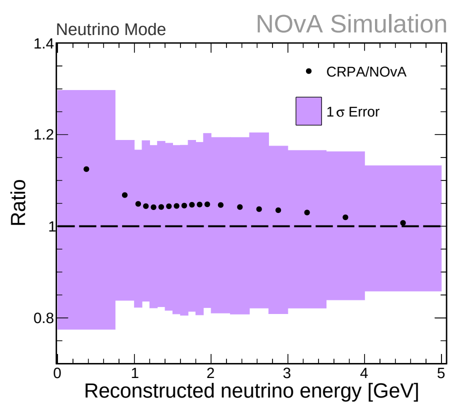

(c) CRPA/NOvA ratio for FHC running.

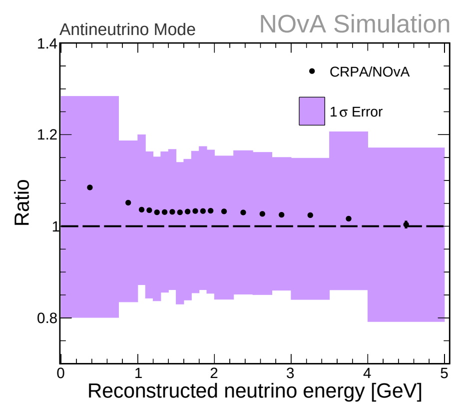

(d) CRPA/NOvA ratio for RHC running.

Figure 6.12: ND reconstructed neutrino energy distributions for the ν_μ sample comparing the NOvA nominal simulation with the CRPA-reweighted prediction. The lower panels show the ratio of CRPA to the nominal model.

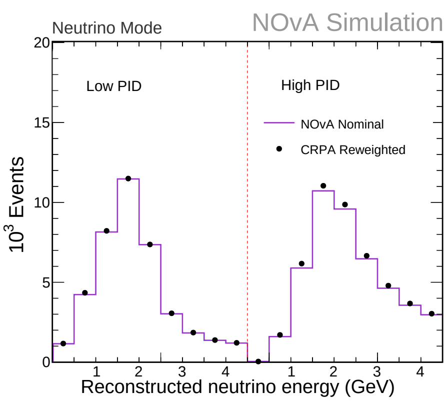

(a) ν_e reconstructed energy in FHC running.

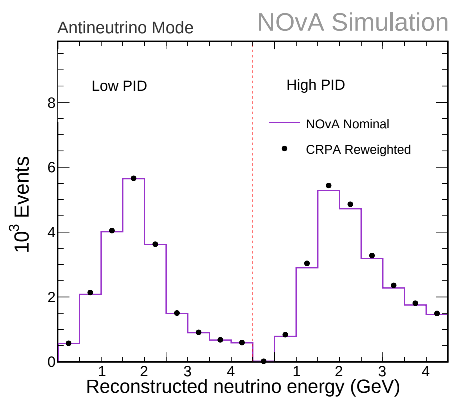

(b) ν_e reconstructed energy in RHC running.

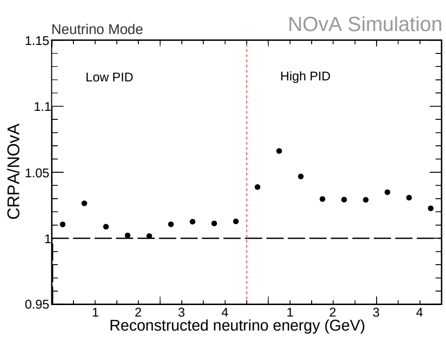

(c) CRPA/NOvA ratio for FHC running.

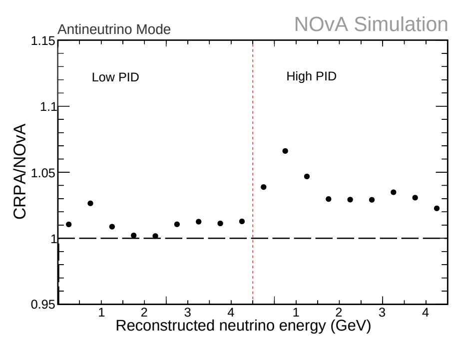

(d) CRPA/NOvA ratio for RHC running.

Figure 6.13: ND reconstructed neutrino energy distributions for the ν_e sample comparing the NOvA nominal simulation with the CRPA-reweighted prediction.

tion into its individual interaction components, such as charged-current signal and various background processes. The reconstructed energy spectrum predicted by the simulation is then compared with the ND data. Any observed difference between data and simulation is attributed to the individual components according to their fractional contributions.

For each interaction component, a migration matrix relating true neutrino energy to reconstructed energy is constructed using simulated events. Separate matrices are produced for the ND and FD detectors. These matrices encode the detector response and allow the true energy distribution inferred from ND measurements to be translated into the expected reconstructed energy spectrum at the FD in the absence of oscillations.

The corrected ND information is then propagated to the FD using the far-to-near ratio derived from simulation. Finally, neutrino oscillation probabilities are applied as weights to the true energy distributions in order to produce the predicted FD reconstructed energy spectrum.

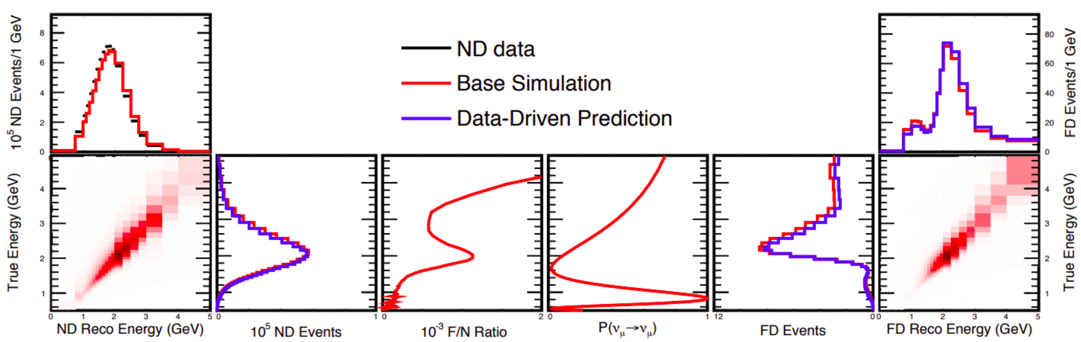

Figure 6.14: Schematic illustration of the near-to-far extrapolation procedure used in the NOvA oscillation analysis. The measured energy spectrum at the ND is used to correct the simulation through a decomposition into interaction components. The corrected distributions are then propagated to the Far Detector using detector response matrices and the simulated far-to-near ratio. Finally, oscillation probabilities are applied to obtain the predicted FD reconstructed energy spectrum.

This extrapolation strategy substantially reduces systematic uncertainties that are correlated between the ND and FD, including those related to the neutrino flux and interaction cross sections. As a result, the oscillation analysis relies primarily on differences between the ND and FD spectra that arise from neutrino oscillations rather than from common systematic effects.

6.8 Effect on FD reconstructed neutrino energy

After applying the CRPA reweighting to the ND Monte Carlo sample, the modified prediction was propagated to the FD using the near-to-far extrapolation procedure described in the previous section. In this procedure, the ND data are used to constrain the prediction while the CRPA weights modify the interaction model in the simulation.

The resulting FD reconstructed neutrino energy spectra obtained with the CRPA reweighted MC are compared with the NOvA nominal prediction. The comparison for the ν_μ sample is shown in 6.15a for the FHC beam configuration and in 6.15b for the RHC configuration. The corresponding ratios of the CRPA prediction to the NOvA nominal prediction are shown in 6.15c and 6.15d. These comparisons indicate that the largest differences appear in the low reconstructed energy region, where quasi-elastic interactions contribute most significantly. However, the CRPA reweighted prediction remains well within the 1σ systematic uncertainty band of the NOvA nominal model across the entire energy range.

A similar comparison was performed for the ν_e sample, as shown in 6.16. The behavior observed for ν_e events is consistent with that seen for ν_μ events, where the CRPA reweighted prediction lies within the systematic uncertainty band of the nominal simulation.

To further investigate the dependence of the CRPA effect on the event topology, the FD sample was divided into four quartiles based on the fraction of hadronic energy in the event. The first quartile corresponds to events with the smallest hadronic energy fraction, while the fourth quartile contains events with the largest hadronic energy fraction. Because the NOvA detector has a relatively poorer resolution for hadronic energy compared to muon

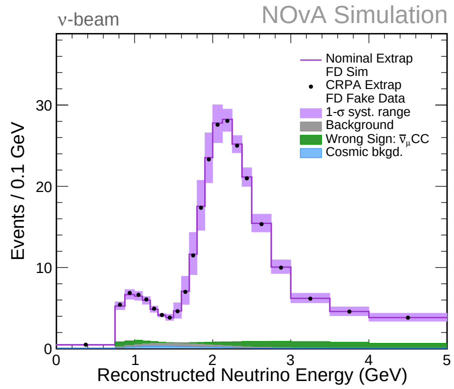

(a) Comparison of NOvA nominal and CRPA reweighted MC for the FHC beam configuration with the 1σ systematic uncertainty band.

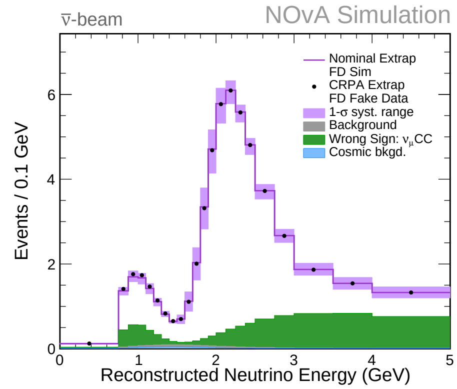

(b) Comparison of NOvA nominal and CRPA reweighted MC for the RHC beam configuration with the 1σ systematic uncertainty band.

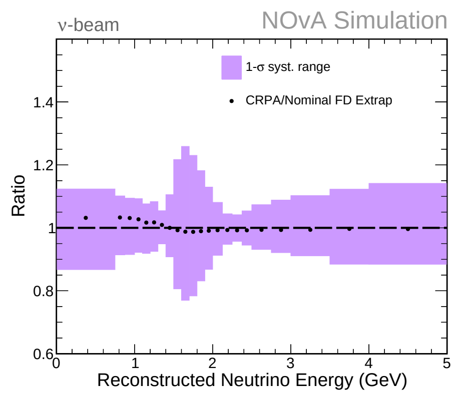

(c) Ratio of CRPA to NOvA nominal prediction for the FHC beam configuration.

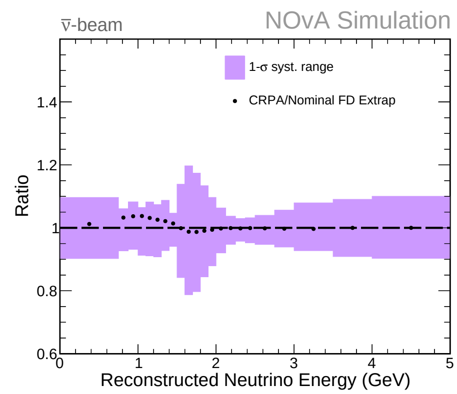

(d) Ratio of CRPA to NOvA nominal prediction for the RHC beam configuration.

Figure 6.15: FD reconstructed ν_μ energy spectra comparing the NOvA nominal prediction with the CRPA reweighted prediction.

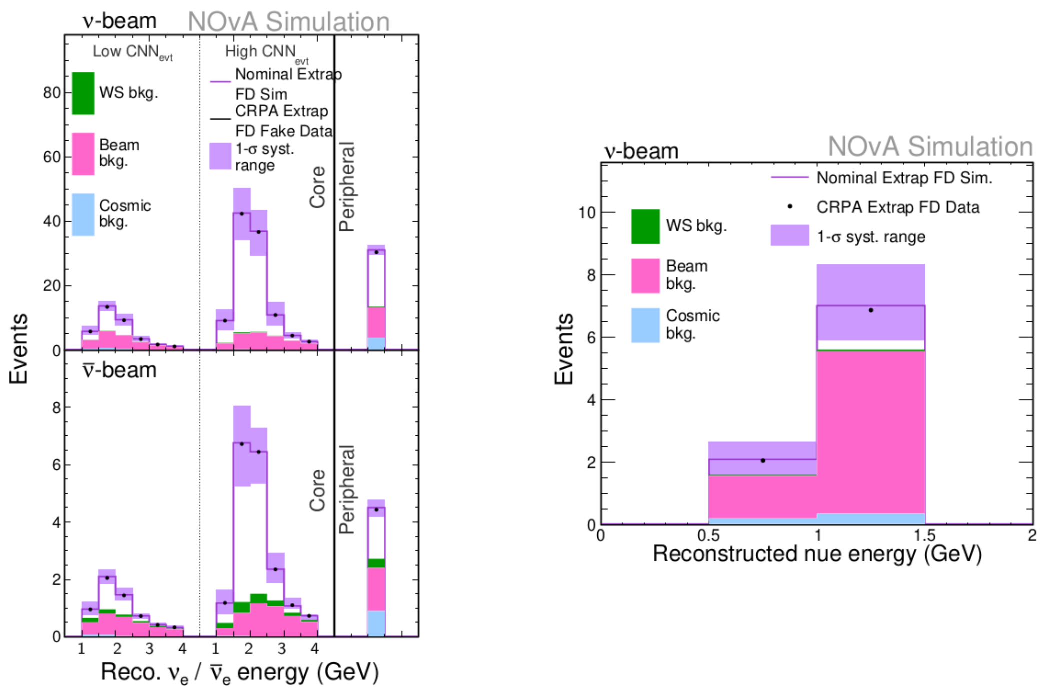

Figure 6.16: FD reconstructed ν_e energy spectra for neutrino and antineutrino beam configurations. The events are divided into core and peripheral samples. The core sample is further categorized into low-purity and high-purity events based on the CVN score. The CRPA reweighted prediction remains consistent with the NOvA nominal prediction within the estimated systematic uncertainties.

energy, events in the first quartile typically have the best reconstructed energy resolution, whereas events in the fourth quartile have the worst resolution.

Since quasi-elastic interactions generally produce minimal hadronic activity, the CRPA model primarily affects events in the first quartile. Consequently, the largest deviations between the CRPA reweighted prediction and the NOvA nominal prediction are observed in this quartile, while the fourth quartile shows the smallest effect. This behavior can be seen in Fig. 6.17 for FHC and 6.18 for RHC.

6.9 Oscillation parameter fitting methodology

The determination of neutrino oscillation parameters in the NOvA experiment is performed using a statistical comparison between the observed event spectra and the predicted spectra obtained from simulation. In this work, the oscillation parameters are extracted using a frequentist fitting framework based on a binned likelihood method.

The predicted event rate in each analysis bin depends on the oscillation parameters as well as various systematic uncertainties associated with neutrino flux, interaction cross sections, and detector response. The best-fit values of the oscillation parameters are obtained by minimizing the binned Poisson log-likelihood function

$$-2\ln L(\vec{\theta},\vec{\phi}) = -2\sum_{i=1}^{N}\left[E_i(\vec{\theta},\vec{\phi}) - O_i + O_i \ln\left(\frac{O_i}{E_i(\vec{\theta},\vec{\phi})}\right)\right] + \sum_{j=1}^{M}\frac{\phi_j^2}{\sigma_j^2}, \tag{6.4}$$

where $E_i(\vec{\theta},\vec{\phi})$ represents the predicted number of events in analysis bin i, which depends on the oscillation parameters $\vec{\theta}$ and the set of systematic parameters $\vec{\phi}$. The observed number of events in bin i is denoted by O_i. The second term corresponds to a penalty contribution that constrains the systematic parameters around their nominal values. Here ϕ_j denotes the shift of systematic parameter j from its central value, and σ_j represents its associated 1σ

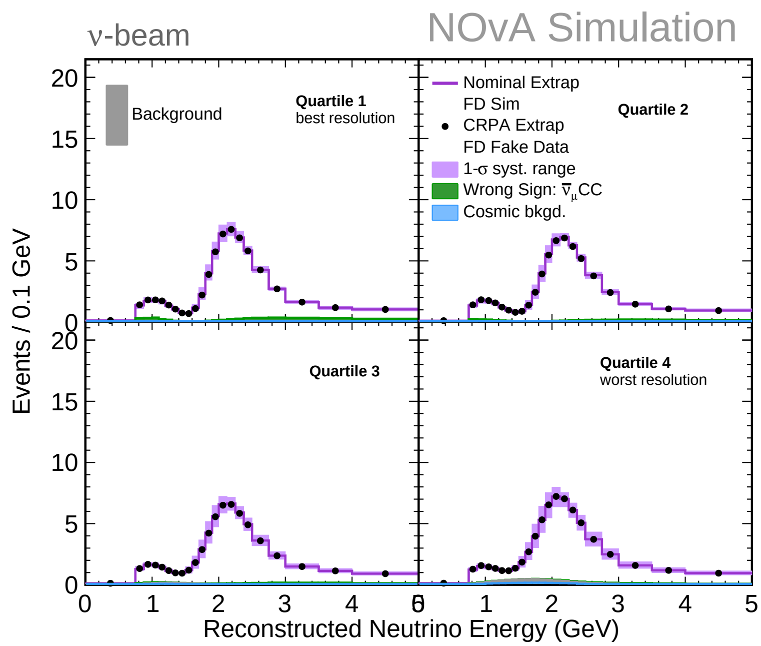

(a) FHC prediction separated into hadronic-energy quartiles.

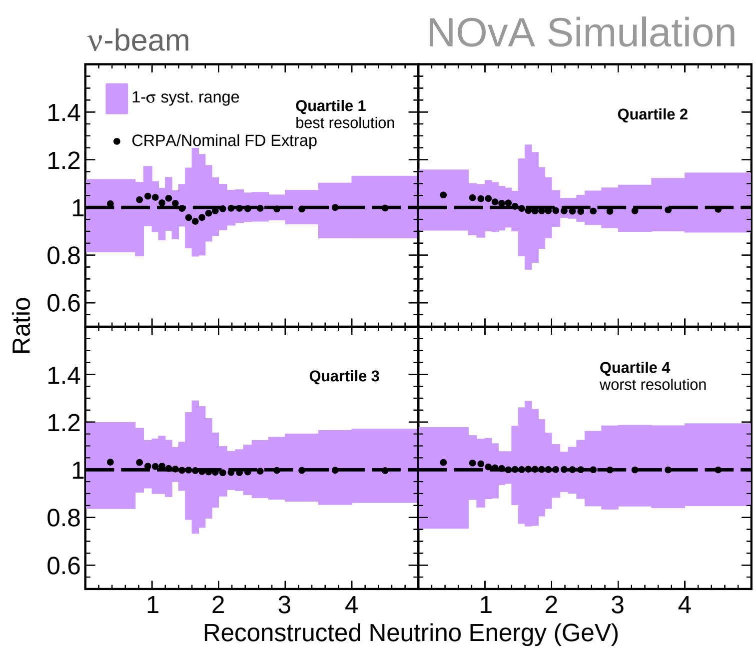

(b) Ratio of CRPA to NOvA nominal prediction for FHC in each quartile.

Figure 6.17: FD reconstructed ν_μ energy spectra for FHC divided into four quartiles based on the hadronic energy fraction in the event. The first quartile corresponds to events with the smallest hadronic energy fraction, while the fourth quartile contains events with the largest hadronic energy fraction.

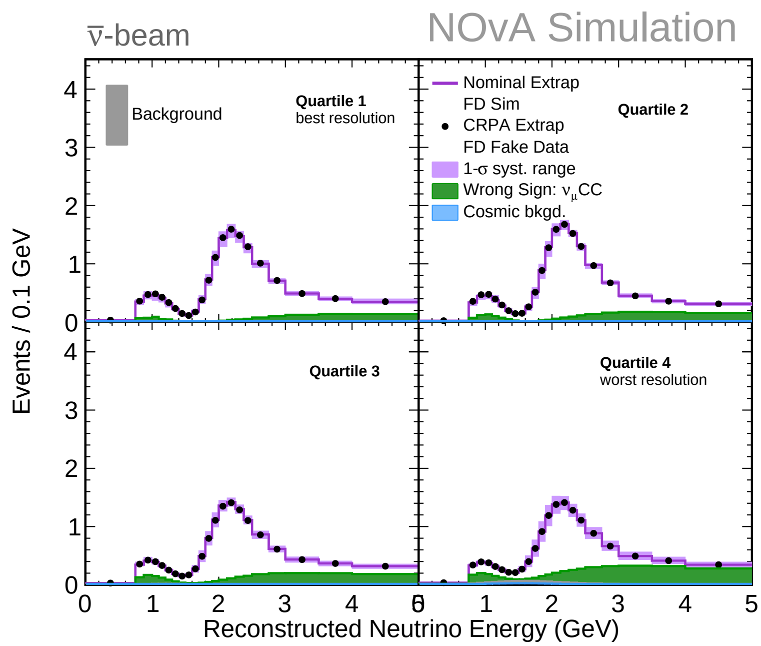

(a) RHC prediction separated into hadronic-energy quartiles.

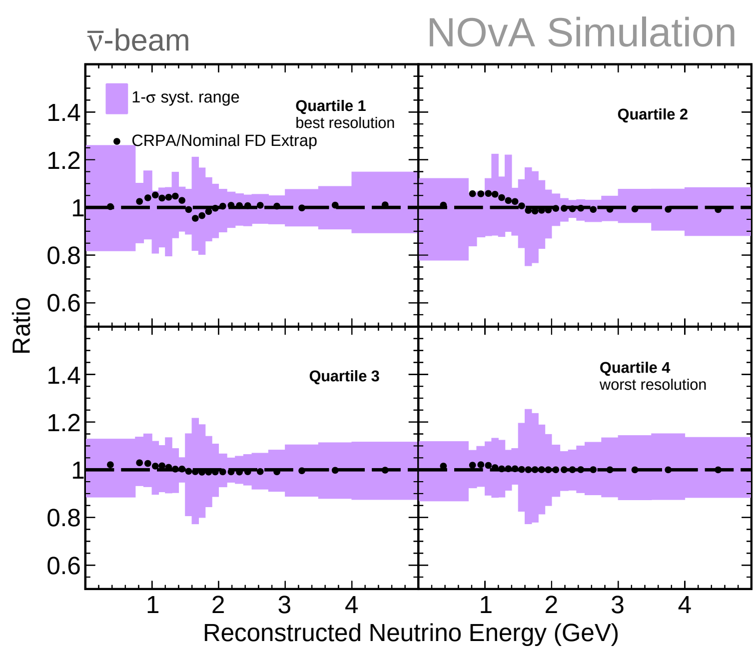

(b) Ratio of CRPA to NOvA nominal prediction for RHC in each quartile.

Figure 6.18: FD reconstructed ν_μ energy spectra for RHC divided into four quartiles based on the hadronic energy fraction in the event. The first quartile corresponds to events with the smallest hadronic energy fraction, while the fourth quartile contains events with the largest hadronic energy fraction.

uncertainty. This penalty term prevents systematic parameters from varying unphysically during the minimization.

The prediction of the FD event spectra incorporates the full three-flavor neutrino oscillation probability in matter. For the NOvA experiment, the baseline and average Earth density used in the oscillation probability calculation are fixed to

$$L = 810 \text{ km}, \qquad \rho = 2.74 \text{ g/cm}^3. \tag{6.5}$$

The oscillation parameters associated with solar neutrino mixing are not strongly constrained by NOvA data and are therefore fixed to the globally measured values reported by the Particle Data Group (PDG),

$$\Delta m^2_{21} = 7.53 \times 10^{-5} \text{ eV}^2, \qquad \sin^2\theta_{12} = 0.307. \tag{6.6}$$

The remaining oscillation parameters, Δm^2_{32}, $\sin^2\theta_{23}$, δ_{CP}, and $\sin^2 2\theta_{13}$, are allowed to vary during the fit along with the systematic uncertainties.

The sensitivity of NOvA to $\sin^2 2\theta_{13}$ is limited because its effects can be partially degenerate with other oscillation parameters. However, this parameter has been measured with high precision by reactor neutrino experiments. In particular, the Daya Bay experiment provides the most precise measurement,

$$\sin^2 2\theta_{13} = 0.0851 \pm 0.0024. \tag{6.7}$$

This measurement can therefore be incorporated as an external constraint in the oscillation fit, improving the overall precision of the parameter estimation.

6.9.1 Frequentist fitting approach

Within the frequentist framework, the quantity

$$\chi^2(\vec{\theta}, \vec{\phi}) = -2 \ln L(\vec{\theta}, \vec{\phi}) \tag{6.8}$$

is treated as the test statistic used to determine the best estimate of the oscillation parameters. The parameter set that minimizes χ^2 defines the best-fit point, denoted as χ^2_{best}.

Confidence intervals for the oscillation parameters are constructed using the test statistic

$$\Delta\chi^2(\vec{\theta}) = \chi^2(\vec{\theta}) - \chi^2_{\text{best}}. \tag{6.9}$$

For a given point in the oscillation parameter space, $\Delta\chi^2$ is computed while minimizing the likelihood with respect to all remaining oscillation and systematic parameters. This procedure, commonly referred to as *profiling*, allows the construction of confidence regions in the parameter space.

Under the assumption of Gaussian statistics, specific values of $\Delta\chi^2$ correspond to standard confidence levels. For example, in the case of two degrees of freedom the 1σ, 2σ, and 3σ confidence regions correspond to

$$\Delta\chi^2 < 2.30, \qquad \Delta\chi^2 < 6.18, \qquad \Delta\chi^2 < 11.83. \tag{6.10}$$

However, several aspects of neutrino oscillation experiments deviate from the Gaussian approximation. In particular, some analysis bins contain relatively small event counts and are therefore governed by Poisson statistics. In addition, the oscillation parameter space contains physical boundaries, such as the allowed range $0 \leq \sin^2\theta_{23} \leq 1$, as well as the cyclic nature of the δ_{CP} parameter. These features can lead to deviations from the ideal Gaussian behavior of the test statistic.

Despite these complications, the frequentist approach provides a robust and widely used framework for determining oscillation parameters and constructing confidence regions based on the comparison between observed data and predicted event spectra.

6.10 Impact on oscillation parameters

To evaluate the impact of the CRPA interaction model on the oscillation measurement, the FD predictions obtained with CRPA reweighting were fitted to the observed FD data using the frequentist fitting framework described in the previous section.

The fit was performed around the best-fit oscillation parameters obtained in the NOvA ANA2024 analysis, which are

$$\begin{aligned} \sin^2\theta_{13} &= 0.0850525, \qquad \delta_{CP} = 0.874617, \\ \Delta m^2_{32} &= 2.4333\times10^{-3}\ \text{eV}^2, \qquad \sin^2\theta_{23} = 0.545742. \end{aligned} \tag{6.11}$$

Using these parameters as the reference point, fits were performed for both the NOvA nominal interaction model and the CRPA reweighted prediction. The resulting best-fit values of the oscillation parameters are summarized in Tab. 6.1, 6.2, and 6.3 for Δm^2_{32}, $\sin^2\theta_{23}$ and δ_{CP} respectively for both the normal ordering (NO) and inverted ordering (IO).

	NOvA nominal	CRPA reweighted
NO	$2.433^{+0.037}_{-0.051}$	$2.436^{+0.037}_{-0.049}$
IO	$-2.474^{+0.022}_{-0.022}$	$-2.477^{+0.024}_{-0.025}$

Table 6.1: Best-fit value of Δm^2_{32} obtained using the NOvA nominal interaction model and the CRPA reweighted prediction along with their 1-σ systematics error.

	NOvA nominal	CRPA reweighted
NO	$0.546^{+0.024}_{-0.024}$	$0.544^{+0.025}_{-0.085}$
IO	$0.544^{+0.015}_{-0.023}$	$0.541^{+0.017}_{-0.030}$

Table 6.2: Best-fit value of $\sin^2\theta_{23}$ obtained using the NOvA nominal interaction model and the CRPA reweighted prediction along with their 1-σ systematics error.

To further illustrate the effect of the CRPA model on the oscillation measurement, the $\Delta\chi^2$ profiles for the relevant oscillation parameters were evaluated. The profiles for Δm^2_{32},

	NOvA nominal	CRPA reweighted
NO	$0.869^{+0.209}_{-0.340}$	$0.860^{+0.347}_{-0.514}$
IO	$1.488^{+0.140}_{-0.148}$	$1.494^{+0.154}_{-0.169}$

Table 6.3: Best-fit value of δ_{CP} obtained using the NOvA nominal interaction model and the CRPA reweighted prediction along with their 1-σ systematics error.

$\sin^2\theta_{23}$, and δ_{CP} are shown in Fig. 6.19, 6.20, and 6.21, respectively. Each figure compares the behavior obtained with the NOvA nominal interaction model to that obtained with the CRPA reweighted prediction.

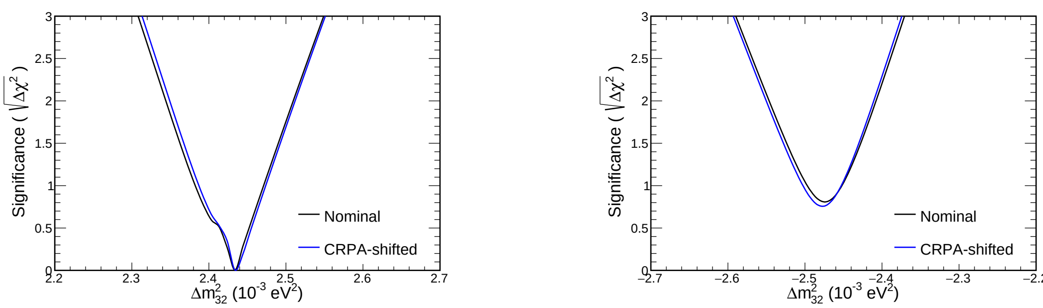

(a) Normal hierarchy

(b) Inverted hierarchy

Figure 6.19: $\Delta\chi^2$ profiles for Δm^2_{32} obtained with the NOvA nominal interaction model and the CRPA reweighted prediction for both normal and inverted mass hierarchies.

Overall, the shifts observed in the best-fit oscillation parameters due to the CRPA reweighting are relatively small compared to the current experimental uncertainties. The $\Delta\chi^2$ profiles also indicate that the inclusion of the CRPA model does not significantly alter the overall shape of the likelihood surfaces for the oscillation parameters. This suggests that, within the current level of experimental precision, the CRPA modification produces only a modest impact on the extracted oscillation parameters.

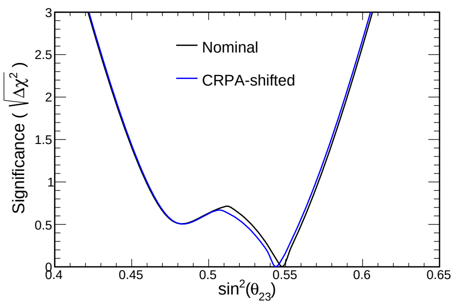

(a) Normal hierarchy

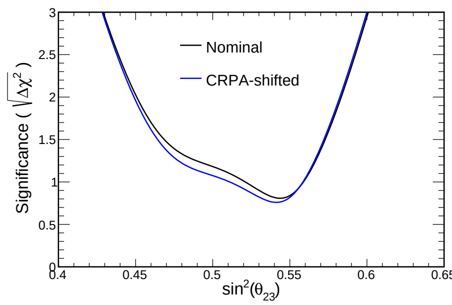

(b) Inverted hierarchy

Figure 6.20: $\Delta\chi^2$ profiles for $\sin^2\theta_{23}$ for the NOvA nominal model and the CRPA reweighted prediction.

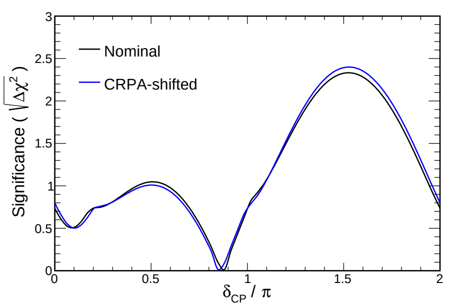

(a) Normal hierarchy

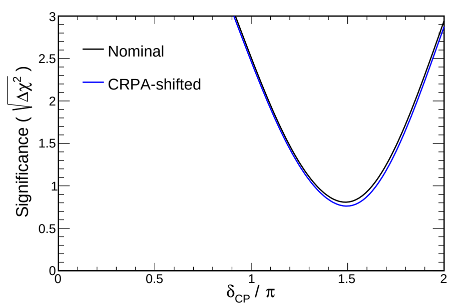

(b) Inverted hierarchy

Figure 6.21: $\Delta\chi^2$ profiles for the CP-violating phase δ_{CP} comparing the NOvA nominal interaction model with the CRPA reweighted prediction.

Chapter 7

Summary

This thesis presents two complementary studies performed using the NOvA ND that address important questions in cosmic-ray physics and neutrino interaction modeling. Both topics are relevant for improving our understanding of particle interactions in complex environments such as the Earth's atmosphere and atomic nuclei, and they contribute to reducing systematic uncertainties in precision neutrino measurements.

The first part of this thesis focuses on the study of cosmic-ray induced muons and their seasonal variation. Cosmic-ray interactions in the Earth's atmosphere produce secondary mesons, primarily pions and kaons, which decay into muons and neutrinos. The rate of atmospheric muons observed in underground or surface detectors is known to vary with atmospheric temperature due to the competition between meson decay and interaction processes. Measuring this variation provides valuable information about high-energy hadronic interactions in the atmosphere and the propagation of secondary particles.

Using data collected with the NOvA ND over four years of operation, the seasonal modulation of the cosmic muon rate for multiple muon events was measured with high precision. This work investigated multiple-muon events originating from high-energy cosmic-ray air showers. The analysis was extended with simulations using the CORSIKA framework in order to understand the underlying atmospheric effects responsible for the observed variations.

The results demonstrate that the seasonal behavior of cosmic muons is strongly influenced by changes in the atmospheric temperature profile, which affects the production and decay probabilities of secondary mesons. Furthermore, the study provides an explanation for the long-standing observation that single-muon and multiple-muon events exhibit different seasonal phases. The simulations indicate that this difference arises from the geometric altitude distribution of meson production and decay in the atmosphere. These results improve our understanding of atmospheric cosmic-ray interactions and provide important inputs for future cosmic-ray and neutrino experiments.

The second part of this thesis investigates the impact of neutrino interaction modeling on oscillation parameter measurements in the NOvA experiment. In the energy range relevant for long-baseline experiments, charged-current quasi-elastic interactions play a central role in neutrino event reconstruction and energy estimation. However, theoretical uncertainties in the description of neutrino–nucleus interactions can introduce biases in the reconstructed neutrino energy and consequently affect the determination of oscillation parameters.

To study these effects, the HF-CRPA model was implemented through event reweighting of simulated neutrino interaction samples. This model provides a more detailed description of nuclear dynamics by incorporating mean-field effects and long-range nuclear correlations. The modified interaction model was propagated through the full NOvA oscillation analysis chain using the Near-to-Far extrapolation technique that constrains the FD prediction using ND data.

Fits to the FD fakedata were performed in order to evaluate the effect of the CRPA model on the extracted oscillation parameters. The results show for NO:

- Δm^2_{32}: NOvA nominal: 2.43363, bias due to CRPA: 0.00298 ($\sim$0.1%)
- $\sin^2(\theta_{23})$: NOvA nominal: 0.546967, bias due to CRPA: 0.00232 ($\sim$0.4%)
- δ_{CP}: NOvA nominal: 0.869513, bias due to CRPA: 0.0095 ($\sim$1%)

Similar effect has been observed for IO. So, the inclusion of the HF-CRPA interaction model introduces small shifts in the best-fit values of Δm^2_{32}, $\sin^2\theta_{23}$, and δ_{CP}. These shifts are found to be smaller than the current experimental uncertainties, indicating that the impact of the CRPA modification on the NOvA oscillation measurement is modest within the present statistical precision.

Although the effect observed in this study is small, experiments with high neutrino flux in the low-energy region ($\sim$ a few hundred MeV) are expected to show significant improvements in the reconstruction of Q^2 and neutrino energy. Experiments such as DUNE and Hyper-Kamiokande are therefore likely to achieve higher sensitivity with this model.

In summary, the work presented in this thesis advances our understanding in two areas of particle physics. The study of cosmic muon seasonal variation provides new insights into atmospheric cosmic-ray interactions and resolves important features of the observed modulation patterns. The investigation of neutrino interaction modeling quantifies the influence of nuclear physics uncertainties on oscillation parameter measurements in the NOvA experiment. Together, these studies contribute to improving the precision and reliability of measurements in both cosmic-ray and neutrino physics.

References

[1] V. D. Barger, K. Whisnant, S. Pakvasa, and R. J. N. Phillips. Matter effects on three-neutrino oscillations. *Phys. Rev. D*, 22(11):2718–2726, 1980.

[2] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.

[3] S. Navas et al. Review of particle physics. *Physical Review D*, 110:030001, 2024.

[4] Joseph A. Formaggio and Geoffrey P. Zeller. From ev to eev: Neutrino cross sections across energy scales. *Rev. Mod. Phys.*, 84(3):1307–1341, 2012.

[5] Teppei Katori and Marco Martini. Neutrino-nucleus cross sections for oscillation experiments. *J. Phys. G*, 45(1):013001, 2018.

[6] Omar Benhar, Daniele Day, and Ingo Sick. Inclusive quasielastic electron-nucleus scattering. *Rev. Mod. Phys.*, 80(1):189–224, 2008.

[7] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. Inclusive charged-current neutrino-nucleus reactions. *Phys. Rev. C*, 83(4):045501, 2011.

[8] Wikimedia Commons contributors. Standard model of elementary particles. Accessed: 2026-03-16.

[9] D. Barak, B. Harrison, and A. Watts. *Accelerator division—operations department*. Concepts Rookie Book, 2020.

[10] P. Adamson et al. The numi neutrino beam. *Nucl. Instrum. Meth. A*, 806:279–306, 2016.

[11] M. D. Tutto and G. Kafka. *Blessing Package - Beam Simulation Plots*. NOvA Internal Document DocDB 13524, 2016.

[12] A. Booth. *Electron neutrino appearance at the NOvA experiment*. PhD thesis, Indiana University, 2021.

[13] Stefano Castro Tognini. *Observation of multiple-muon seasonal variations in the NOvA Near Detector*. PhD thesis, Goias U., 2018. DOI: 10.2172/1468447.

[14] G. S Davies. Images of the far detector (blessed package), 2017. NOvA Internal Document, DocDB 20792.

[15] P. Lukens B. Miller and L. A Corwin. Photographs of progress of fd construction (blessed package), 2018. NOvA Internal Document, DocDB 8284.

[16] C. Patrignani et al. Review of particle physics. *Chinese Physics C*, 40(10):100001, 2016.

[17] Konrad Bernlöhr. `https://www.mpi-hd.mpg.de/hfm/CosmicRay/Showers.html`. Accessed: 2026-02.

[18] M. Ambrosio et al. (MACRO Collaboration). Seasonal variations in the underground muon intensity as seen by macro. *Astropart. Phys.*, 7:109–124, 1997.

[19] A. Bouchta. Seasonal variation of the muon flux seen by AMANDA. In *Proceedings of the 26th ICRC, Salt Lake City, Utah, USA* (1999), volume 2, pages 108–111, 5.

[20] M. Selvi *et al.* (LVD Collaboration). In *Proceedings of the 31st ICRC, Lodz, Poland* (2009).

[21] P. Desiati *et al.* Seasonal Variations of High Energy Cosmic Ray Muons Observed by the IceCube Observatory as a Probe of Kaon/Pion Ratio. In *Proceedings of the 32nd ICRC, Beijing, China* (2011), volume 1, pages 78–81.

[22] G. Bellini *et al.* Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth. *J. Cosmol. Astropart. Phys.*, 05:015, 2012.

[23] T. Abrahao *et al.* (Double Chooz Collaboration). *J. Cosmol. Astropart. Phys.*, 2017(02):017, 2017.

[24] P. Adamson et al. (MINOS Collaboration). Observation of muon intensity variations by season with the minos far detector. *Phys. Rev. D*, 81:012001, 2010.

[25] R. P. Kokoulin *et al.* Atmospheric effects in the intensity of muon bundles and geometrical mechanism of their formation. *J. Phys. Conf. Ser.*, 675(3):032034, jan 2016.

[26] B. Jargowsky. Blessed package for prod5.1 energy estimation for 3f ana2024, 2024. NOvA Internal Document, DocDB 61952.

[27] D. S. Ayres et al. (NOvA Collaboration). The nova technical design report, 2007. FERMILAB-DESIGN-2007-01.

[28] E. D. Niner. *Observation of electron neutrino appearance in the NuMI beam with the NOvA experiment*. PhD thesis, Indiana University, 2015.

[29] S. M. Lein. *Muon neutrino contained disappearance in NOvA*. PhD thesis, University of Minnesota, 2015.

[30] K. Kephart. Near detector under ground (ndug) parameters, 2014. NOvA Internal Document, DocDB 9860.

[31] Michael David Baird. *An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter*. PhD thesis, Indiana U., 2015.

[32] Thomas K. Gaisser, Ralph Engel, and Elisa Resconi. *Cosmic Rays and Particle Physics*. Cambridge University Press, 2 edition, 2016.

[33] Malcolm S. Longair. *High Energy Astrophysics*, volume 2. Cambridge University Press, 2011.

[34] Victor F. Hess. Über beobachtungen der durchdringenden strahlung bei sieben freiballonfahrten. *Phys. Z.*, 13:1084–1091, 1912.

[35] K.-H. Kampert and A. A. Watson. Extensive air showers and ultra-high energy cosmic rays: A historical review. *Eur. Phys. J. H*, 37:359–412, 2012.

[36] Victor F. Hess and Carl D. Anderson. The nobel prize in physics 1936. `https://www.nobelprize.org/prizes/physics/1936/summary/`, 1936. Nobel Prize Foundation.

[37] T. Stanev T. K. Gaisser and S. Tilav. Cosmic ray energy spectrum from measurements of air showers. *Front. Phys.*, 8:748–758, 2013.

[38] E. Fermi. On the origin of the cosmic radiation. *Phys. Rev.*, 75:1169–1174, 1949.

[39] R. Blandford and D. Eichler. Particle acceleration at astrophysical shocks: A theory of cosmic ray origin. *Phys. Rep.*, 154:1–75, 1987.

[40] A. M. Hillas. The origin of ultra-high-energy cosmic rays. *Annu. Rev. Astron. Astrophys.*, 22:425–444, 1984.

[41] A. M. Hillas. *Cosmic Rays: Recent Progress and Some Current Questions*. Springer, 2005.

[42] R. U. Abbasi et al. First observation of the greisen–zatsepin–kuzmin suppression. *Phys. Rev. Lett.*, 100:101101, 2008.

[43] Pierre Auger Collaboration. Measurement of the energy spectrum of cosmic rays above 10^{18} ev. *Phys. Lett. B*, 685:239–246, 2010.

[44] Paul H. Barrett, Lowell M. Bollinger, Giuseppe Cocconi, Yehuda Eisenberg, and Kenneth Greisen. Interpretation of cosmic-ray measurements far underground. *Rev. Mod. Phys.*, 24:133–178, Jul 1952.

[45] P. Adamson *et al.* Observation of seasonal variation of atmospheric multiple-muon events in the minos near and far detectors. *Phys. Rev. D*, 91:112006, Jun 2015.

[46] M. A. Acero et al. (NOvA Collaboration). Observation of seasonal variation of atmospheric multiple-muon events in the nova near detector. *Phys. Rev. D*, 99:122004, 2019.

[47] M. A. Acero *et al.* Seasonal variation of multiple-muon cosmic ray air showers observed in the nova detector on the surface. *Phys. Rev. D*, 104:012014, Jul 2021.

[48] D. Hagmann, C. Lange and D Wright. Cosmic-ray shower generator (cry) for monte carlo transport codes. 2007, DOI: 10.1109/NSSMIC.2007.4437209.

[49] D. Heck *et al.* CORSIKA: A Monte Carlo code to simulate extensive air showers. Report No. FORSCHUNGSZENTRUM KARLSRUHE FZKA-6019 (1998).

[50] NCEI Geomagnetic Modeling Team and NOAA National Centers for Environmental Information. Noaa geomagnetic field calculator, 2017. Accessed: 2022.

[51] S. Agostinelli et al. Geant4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003.

[52] J Allison *et al.* Geant4 developments and applications. *IEEE Trans. Nucl. Sci.*, 53(1):270–278, 2006.

[53] J Allison *et al.* Recent developments in geant4. *Nucl. Instrum. Meth.*, 835:186–225, 2016.

[54] Maurizio Spurio. *Particles and Astrophysics: A Multi-Messenger Approach*. Springer, Cham, Switzerland, 2018. Section 4.2, p. 90.

[55] NOAA Air Resources Laboratory. Global data assimilation system (gdas1) archive, 2017. Accessed: 2017–2018.

[56] Paul V. C. Hough. Methods and means for recognizing complex patterns, December 1962. U.S. Patent 3,069,654.

[57] Dana H. Ballard. Generalizing the hough transform to detect arbitrary shapes. *Pattern Recognition*, 13(2):111–122, 1981.

[58] Leandro A. F. Fernandes and Manuel M. Oliveira. Real-time line detection through an improved hough transform voting scheme. *Pattern Recognition*, 41(1):299–314, 2008.

[59] Mark Baird. Global vertex reconstruction beginning with a modified hough transform. Technical Report DocDB-8241, NOvA Collaboration, 2012. Internal report.

[60] C. Forbes, M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. Wiley, Hoboken, NJ, 4th edition, 2011. Section on Erlang distribution.

[61] D. Dee *et al.* The era-interim reanalysis: Configuration and performance of the data assimilation system. *Q. J. R. Meteorol. Soc.*, 137:553–597, 04 2011.

[62] D. P. Dee et al. The era-interim reanalysis: configuration and performance of the data assimilation system. *Q. J. R. Meteorol. Soc.*, 137(656):553–597, 2011. DOI: 10.1002/qj.828.

[63] Hersbach. et al. The era5 global reanalysis. *Q. J. R. Meteorol. Soc.*, 146(730):1999–2049, 2020. DOI: 10.1002/qj.3803.

[64] European Centre for Medium-Range Weather Forecasts (ECMWF). Grib overview, 2024. Accessed: 2026-02-12.

[65] European Centre for Medium-Range Weather Forecasts (ECMWF). eccodes, 2024. Accessed: 2026-02-12.

[66] E. Grashorn *et al.* The atmospheric charged kaon/pion ratio using seasonal variation methods. *Astropart. Phys.*, 33(3):140–145, 2010.

[67] M. Ambrosio *et al.* The macro detector at gran sasso. *Nucl. Instrum. Methods Phys. Res.*, 486(3):663–707, 2002.

[68] F. Ronga. Seasonal variations of the rate of multiple-muons in the gran sasso underground laboratory. *EPJ Web of Conferences*, 136:05004, 2017.

[69] Thomas K. Gaisser and Stef Verpoest. Profiles of energetic muons in the atmosphere. *Astropart. Phys.*, 133:102630, 2021.

[70] Thomas Gaisser and Stef Verpoest. Seasonal variation of atmospheric muons. *PoS*, ICRC2021:1202, 2021.

[71] E. A. Yurina *et al.* Variations in the intensity of cosmic ray muon bundles according to DECOR data 2012-2017. *J. Phys. Conf. Ser.*, 1189(1):012010, 2019.

[72] Ellen Guan and Maury Goodman. Monte carlo simulations of factors influencing seasonal variation of multiple muon events, 2023.

[73] N. V. Tolkacheva *et al.* Atmospheric effects in the intensity of cosmic ray muon bundles. *Bulletin of the Russian Academy of Sciences: Physics*, 75(3):377–380, 2011.

[74] K.P. Arunbabu *et al.* Dependence of the muon intensity on the atmospheric temperature measured by the grapes-3 experiment. *Astropart. Phys.*, 94:22–28, 2017.

[75] W. Pauli. Dear radioactive ladies and gentlemen. *Phys. Today*, 31N9:27, 1978.

[76] Enrico Fermi. An attempt of a theory of beta radiation. *Zeitschrift für Physik*, 88:161–177, 1934.

[77] Frederick Reines and Clyde L. Cowan. Detection of the free neutrino: A confirmation. *Science*, 124(3212):103–104, 1956.

[78] Y. Fukuda et al. (Super-Kamiokande Collaboration). Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[79] Q. R. Ahmad et al. (SNO Collaboration). Direct evidence for neutrino flavor transformation from neutral-current interactions. *Phys. Rev. Lett.*, 89:011301, 2002.

[80] Samoil M. Bilenky. *Introduction to the Physics of Massive and Mixed Neutrinos*. Springer, 2010.

[81] Bruno Pontecorvo. Mesonium and antimesonium. *Soviet Physics JETP*, 6:429, 1957.

[82] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.

[83] A. Cervera, A. Donini, M. B. Gavela, J. J. Gomez-Cadenas, P. Hernandez, O. Mena, and S. Rigolin. Golden measurements at a neutrino factory. *Nucl. Phys. B*, 579(1-2):17–55, 2000.

[84] Lincoln Wolfenstein. Neutrino oscillations in matter. *Phys. Rev. D*, 17:2369–2374, 1978.

[85] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995.

[86] C. H. Llewellyn Smith. Neutrino reactions at accelerator energies. *Phys. Rep.*, 3(5):261–379, 1972.

[87] V. Bernard, L. Elouadrhiri, and Ulf-G. Meißner. Axial structure of the nucleon: Topical review. *J. Phys. G*, 28(1):R1–R35, 2002.

[88] Dieter Rein and Lalit M. Sehgal. Neutrino excitation of baryon resonances and single pion production. *Annals of Physics*, 133(1):79–153, 1981.

[89] Chris Quigg. Gauge theories of the strong, weak, and electromagnetic interactions. *Rev. Mod. Phys.*, 59(4):1067–1088, 1987.

[90] J. Nieves, J. E. Amaro, and M. Valverde. Inclusive quasielastic charged-current neutrino-nucleus reactions. *Phys. Rev. C*, 70(5):055503, 2004.

[91] Richard J. Hill and Gil Paz. Model independent extraction of the proton charge radius from electron scattering. *Phys. Rev. D*, 82(11):113005, 2010.

[92] M. Martini, M. Ericson, G. Chanfray, and J. Marteau. A unified approach for nucleus-nucleus reactions: Quasielastic and multinucleon excitations. *Phys. Rev. C*, 80(6):065501, 2009.

[93] N. Jachowicz, K. Heyde, J. Ryckebusch, and S. Rombouts. Continuum random phase approximation approach to charged-current neutrino-nucleus scattering. *Phys. Rev. C*, 85(2):024606, 2012.

[94] T. Dieminger, S. Dolan, D. Sgalaberna, A. Nikolakopoulos, T. Dealtry, S. Bolognesi, L. Pickering, and A. Rubbia. Uncertainties on the $\overset{(-)}{\nu}_e/\overset{(-)}{\nu}_\mu$ and $\nu_e/\overline{\nu}_e$ cross-section ratio from the modeling of nuclear effects at 0.2 to 1.2 gev neutrino energies and their impact on neutrino oscillation experiments. *Phys. Rev. D*, 108:L031301, Aug 2023.

[95] S. Dolan, A. Nikolakopoulos, O. Page, S. Gardiner, N. Jachowicz, and V. Pandey. Implementation of the continuum random phase approximation model in the genie generator and an analysis of nuclear effects in low-energy transfer neutrino interactions. *Phys. Rev. D*, 106:073001, Oct 2022.

[96] C. Andreopoulos et al. The genie neutrino monte carlo generator. *Nucl. Instrum. Meth. A*, 614(1):87–104, 2010.

[97] R. J Nichol E. Catano-Mur and Z. Vallari. Ana2024 3f executive summary, 2024. NOvA Internal Document, DocDB 61655.