FERMILAB-THESIS-2025-08

Searching for Sterile Neutrinos In the NOvA Experiment and Measurements of Hadronic Energy Response with NOvA Test Beam Data

DISSERTATION

Presented to the Faculty of the Graduate School of

the University of Cincinnati

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Physics

of the College of Arts and Sciences

by

David Francisco Dueñas Tonguino

University of Cincinnati

July, 2025

Committee Chair: Alexandre Sousa

Abstract

The NOvA (NuMI Off-Axis electron neutrino Appearance) experiment is a long-baseline neutrino oscillation experiment composed of two functionally identical detectors: a 300-ton Near Detector and a 14-kton Far Detector, separated by 809 km and placed 14 mrad off the axis of the NuMI neutrino beam created at Fermilab. This configuration enables NOvA's rich neutrino physics program, which includes measuring neutrino mixing parameters, determining the neutrino mass hierarchy, probing CP violation in the leptonic sector, studying neutrino cross sections, searching for sterile neutrinos, and more.

In this thesis I will present two independent analyses developed within the NOvA experiment. The first is a search for light sterile neutrino oscillations in the 3+1 framework using ν_μ-CC and NC samples from the NOvA Near and Far Detectors. Light sterile neutrinos are hypothetical neutral leptons that do not participate in the weak interactions but can mix with the three known active neutrinos: ν_e, ν_μ, and ν_τ. Anomalous results, such as $\nu_e/\bar{\nu}e$ appearance in a ν_μ ($\bar{\nu}_\mu$) beam observed by the MiniBooNE and LSND experiments, can be explained by the existence of sterile neutrinos. Thus, proving the existence

of this type of neutrino is essential. The latest results of these searches in the NOvA experiment, along with contributions to this effort, will be presented.

The second is the first measurement of the pion energy response using data from the NOvA Test Beam experiment. The NOvA Test Beam experiment, deployed at Fermilab, uses a scaled-down 30-ton NOvA detector to analyze tagged beamline particles. The beamline can select and identify electrons, muons, pions, kaons, and protons with momenta ranging from 0.4 to 1.4 GeV/c. Pions are an important component of the hadronic system in neutrino interactions, and understanding how the detector responds to these particles is crucial for the validation of the simulation and the reduction of detector calibration uncertainties, which remain one of the largest systematic uncertainties in NOvA analyses.

In loving memory of my beloved mother, Carmela Tonguino, and my good friend, Leonardo Cardenas. You made this world a better place.

Acknowledgements

I am deeply thankful to the many people whose support made this PhD possible.

First, I am deeply grateful to my advisor, Alexandre Sousa, for his guidance, generosity, and support over the years. Thank you for sharing your knowledge, helping me grow as a researcher, and encouraging me through the difficult times. You went above and beyond, and I will always be thankful.

I also want to thank the University of Cincinnati neutrino group for the fruitful discussions, constant support, and for generously sharing your expertise during my PhD. I could not have asked for a better group to work with.

I would like to thank my family: my mother, Carmela Tonguino, my sister, Maria Dueñas, and my niece, Luna Cardenas. Thank you for your unwavering love and support throughout my life and career. You have always been one of my greatest sources of motivation. I am also very grateful to my friend Leonardo Cardenas for helping me begin this PhD journey. To my aunt, Maria Tonguino, thank you for your incredible support during the last two years. They were the hardest part of all. I am also thankful to the Eraso

Cabrera family for being there when I needed it most and for your support.

I am especially thankful to my partner, Heather Marshall. Thank you for your patience, love, care, and for being by my side every day. I could not have done this without your support. I also thank my cat, Onyx, for being a constant and comforting companion during the most challenging moments of this journey.

Finally, a special thanks to my friends and colleagues at Fermilab, in particular the NOvA Test Beam, NOvA Operations, and NuX groups, for your encouragement, collaboration, and willingness to lend a hand whenever it was needed.

Contents

List of Figures

List of Tables

Chapter 1

Neutrinos

The history of neutrinos begins with the discovery of radioactivity. In 1896, Henri Becquerel found that certain materials spontaneously emit radiation without any external energy source. This was followed by the work of Marie and Pierre Curie, who discovered highly radioactive elements like polonium and radium. These early observations led to the identification of different types of radiation.

In 1899, Ernest Rutherford classified these emissions and identified β radiation, now known as β^- decay, where a nucleus with N neutrons and charge Z decays into another of charge $Z+1$ and $N-1$ neutrons plus an electron:

$$A(N,Z) \to A'(N-1,Z+1) + e^-. \tag{1.1}$$

In such a two-body decay, the electron should have a monochromatic energy

spectrum. In 1914, James Chadwick was the first to observe that the spectrum is actually continuous. His conclusions were controversial due to the fact that the process was violating spin statistics along with conservation of energy.

In a desperate solution to preserve these conservation principles, Wolfgang Pauli proposed in 1930 the existence of a new particle that he called neutron [25]. This neutral fermion, light and with a large penetration length, solved the problem, but Pauli thought that this new particle was undetectable.

Two years later, James Chadwick discovered the neutron we know today [26], and Enrico Fermi renamed the Pauli particle neutrino "the little neutral one". Incorporating it in his 1934 formulation of a new theory of β decay, which was also the first theory of weak interactions [27].

The next section will dive into the experiments that lead to the ν_e, ν_μ and ν_τ discovery.

1.1 The First Neutrino ($\nu_\mathbf{e}$)

More than 20 years after Fermi's theory, the neutrino now known to be an electron antineutrino ($\bar{\nu}_e$) was discovered by Fred Reines and Clyde Cowan in 1956 using a water target in which CdCl2 (cadmium chloride) has been dissolved [1]. The experiment was located near the Savannah River nuclear reactor, and measured the inverse β decay reaction. A sketch of the detector can be seen in Figure 1.1. In this process, electron antineutrinos produced in the reactor core are captured by protons in the detector as:

$$\bar{\nu_e} + p \rightarrow e^+ + n \tag{1.2}$$

The positron produced in the neutrino interaction annihilates immediately with an electron, emitting a pair of gamma rays, that are detected in two liquid scintillator tanks. The resulting neutron combines with a proton in the cadmium after ~ 10 μs of flying time. From this interaction a deuterium nucleus is created, and a photon with energy 2.2 MeV is emitted:

$$n + p \rightarrow d + \gamma \tag{1.3}$$

These two processes are reconstructed as an electron antineutrino signal, which is the primary signal used in a wide range of experiments today. Their experiment marked the first direct detection of neutrinos and confirmed their existence, more than two decades after the particle had been theoretically proposed. Frederick Reines was awarded the Nobel Prize in Physics in 1995 for this work.

1.2 The Second Neutrino (ν_μ)

In 1962 Jack Steinberger, Leon Lederman, and Melvin Schwartz discovered a new type of neutrino in interactions involving muons. As a consequence, they called it the muon neutrino (ν_μ). An experiment to detect these neutrinos was deployed at Brookhaven National Laboratory where pions created in proton-

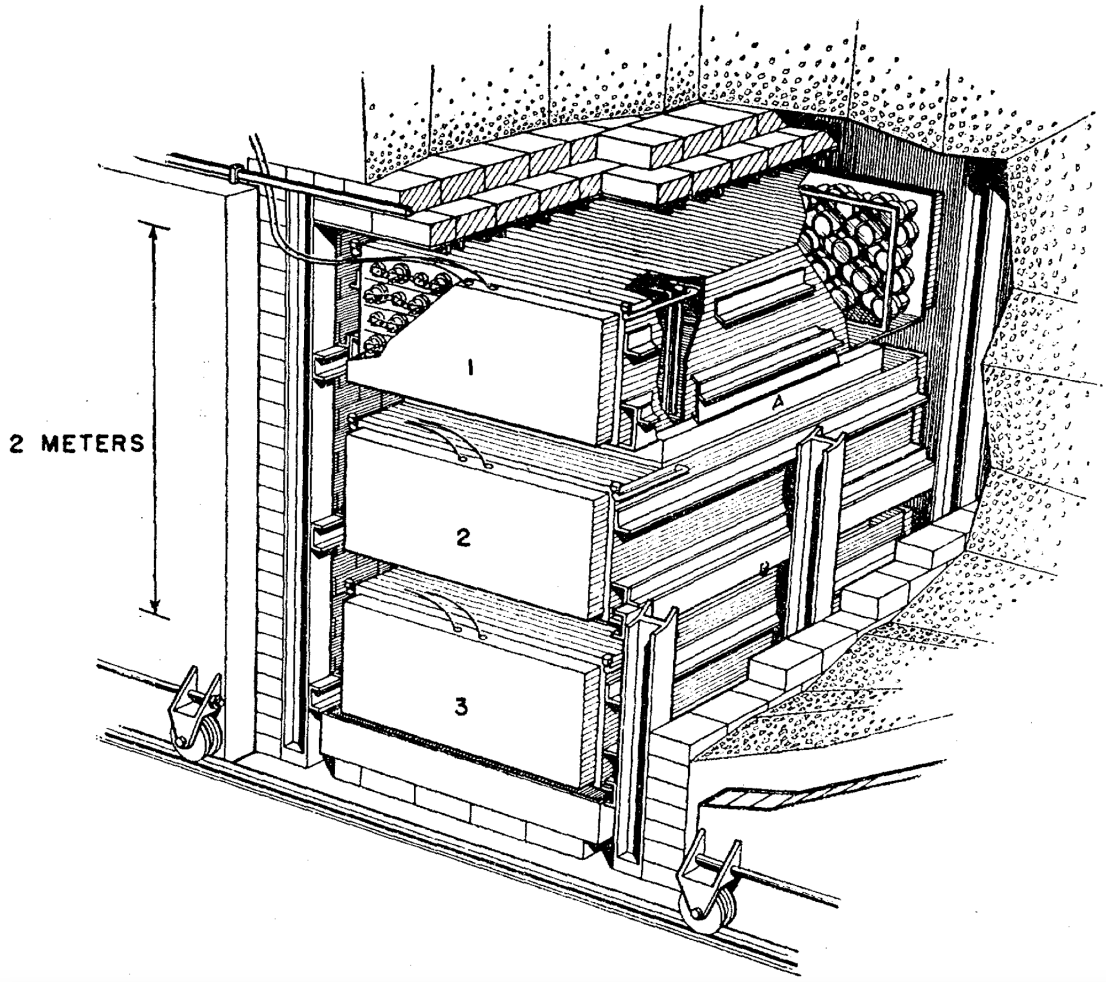

Figure 1.1: Sketch of detectors inside their lead shield. The detector tanks marked 1, 2, and 3 contained liquid scintillator solution which was viewed in each tank by 110 5-in. photomultiplier tubes. The white tanks contained the water-cadmium chloride target, and in this picture are some 28 cm deep. These were later replaced by 7.5-cm deep polystyrene tanks, and detectors 1 and 2 were lowered correspondingly [1].

target collisions decay into muons and muon neutrinos [28]:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{1.4}$$

The resulting beam was directed trough a thicker layer of shielding material to absorb all the particles except neutrinos, the muon neutrino beam then interacted in a spark chamber detector via the process:

$$\nu_\mu + n \rightarrow \mu + hadrons \tag{1.5}$$

Since a muon was one of the products of this interaction, it demonstrated that neutrinos produced by pion decay were not the same as the electron neutrinos produced in beta decay experiments. For this reason, it was clear at that point that there were at least two types of neutrinos, each associated to a different lepton.

This discovery was an important milestone, as it proved the existence of a second lepton family (ν_μ and μ). The concept of lepton families is an essential feature of the Standard Model, which organizes the fundamental particles that exist in our universe.

Their work laid down the foundations for techniques we use today to produce neutrino beams, which have served to conduct detailed studies of the weak interaction and neutrino oscillations. They were awarded the Nobel Prize in Physics in 1988.

1.3 Neutral Currents and The Third Neutrino (ν_τ)

1.3.1 Discovery of Neutral Currents

In 1973, the electroweak neutral current (NC) was discovered at CERN in the Gargamelle bubble chamber experiment. The bubble chamber was 4.8 m long and 2 m in diameter. It weighed 1 kton and held nearly 12 cubic meters of heavy-liquid freon (CF3Br). The freon used in the Gargamelle detector made it possible to detect charged particles produced by neutrino interactions, indirectly revealing the neutrinos themselves. The observations came from neutrino-nucleon and neutrino-electron scattering events that involved the exchange of a neutral Z^0 boson [29][30]:

$$\nu_\mu(\bar{\nu}_\mu) + N \rightarrow \nu_\mu(\bar{\nu}_\mu) + hadrons \tag{1.6}$$

$$\nu_\mu(\bar{\nu}_\mu) + e^- \rightarrow \nu_\mu(\bar{\nu}_\mu) + e^- \tag{1.6}$$

These neutrino interactions, in which neutrinos do not convert into its charged lepton partner in the final state, provided the first evidence for neutral weak currents as predicted by the electroweak theory. It supported the model developed by Abdus Salam, Sheldon Glashow, and Steven Weinberg [31][32][33], particularly the structure of weak interactions described by the $V - A$ framework. A NC interaction in the Gargamelle bubble chamber can be seen in

Figure 1.2.

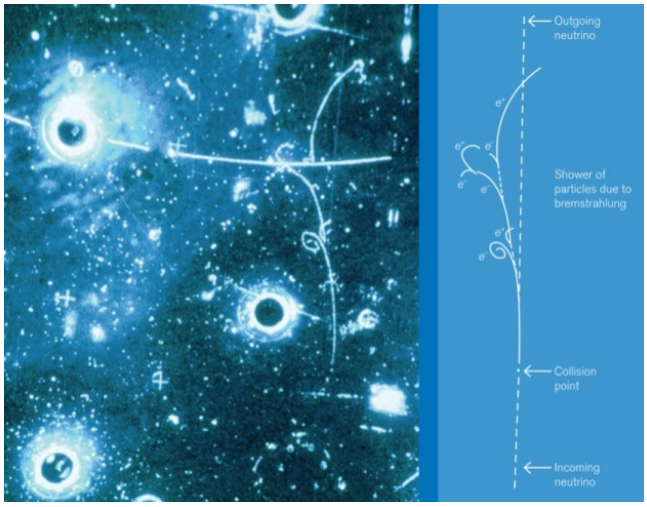

Figure 1.2: Observation of the first leptonic NC interaction. A muon antineutrino scatters off an electron, ejecting it upward, producing a distinct electromagnetic shower with electron–positron pairs [2].

1.3.2 Z Boson Decay Width and the Number of Neutrino Flavors

The discovery of neutral currents in the Gargamelle experiment confirmed that neutrinos interact via the weak neutral force. However, it did not reveal how many neutrino types exist. That came later from high-energy collider experiments at CERN, where precise measurements of the Z boson's decay width showed that only three light neutrino flavors contribute to its invisible decay channels.

In the 1990s, experiments at the Large Electron–Positron Collider (LEP) at CERN provided a precise measurement of the total decay width of the Z^0

boson [3]. This width is sensitive to all possible invisible decay channels of the Z^0, including decays into neutrinos.

The total decay width Γ_Z of the Z^0 boson was measured via the reaction:

$$e^+ + e^- \to Z^0 \to f\bar{f},$$

where f represents any fermion.

The LEP experiments scanned a range of center-of-mass energies around the Z boson mass ($\sim$91 GeV/c^2) and measured the cross sections for various visible final states, such as hadrons and charged leptons. The resulting resonance curve, described by a Breit–Wigner distribution [34], allowed precise extraction of the Z boson's mass and total width. The hadronic cross-section measurement is shown in Figure 1.3.

By subtracting the measured partial widths of the visible channels from the total width, the invisible width Γ_{inv} was determined. Assuming each neutrino flavor contributes equally to Γ_{inv}, the number of light active neutrino species (N_ν) was extracted:

$$N_\nu = \frac{\Gamma_{\text{inv}}}{\Gamma_\nu},$$

where Γ_ν is the theoretically predicted width for a single neutrino species.

The LEP experiments found:

$$N_\nu = 2.9840 \pm 0.0082,$$

in good agreement with the Standard Model prediction of three neutrino fam-

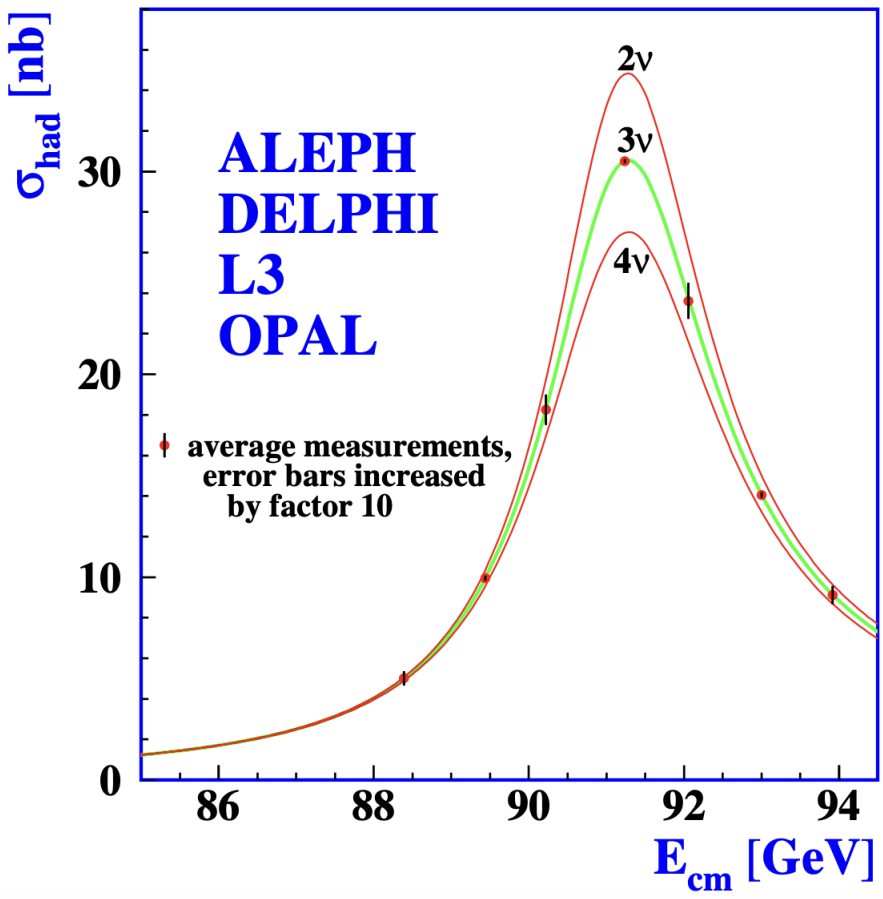

Figure 1.3: Measurements of the hadronic production cross section σ_{had} as a function of center-of-mass energy E_{cm} around the Z boson resonance at LEP. The data points represent the average measurements from the four LEP experiments (ALEPH, DELPHI, L3, and OPAL). The colored curves show theoretical predictions for two, three, and four light neutrino species. The best agreement is found with the Standard Model prediction of three neutrino flavors. Figure adapted from [3]

ilies.

This result ruled out the existence of additional active neutrinos (those that interact via the weak interactions), confirming the 3-flavor framework.

1.3.3 Discovery of ν_τ

The third neutrino flavor, the tau neutrino (ν_τ), had been implied from the discovery of the tau lepton in the 1970s [35]. Direct observation of ν_τ interactions happened much later with the DONUT (Direct Observation of the NU Tau) experiment at Fermilab, whose results were published in 2000 [36] and final results in 2008 [37]. The experiment consisted of a beam of high-energy neutrinos directed onto a nuclear emulsion detector. Nine events consistent with tau neutrino interactions were observed, displaying the production and subsequent decay of tau leptons, thus providing clear evidence for the direct detection of ν_τ.

This completed the experimental confirmation of all three neutrino flavors predicted by the Standard Model and measured at the LEP experiments.

1.4 The Solar Neutrino Problem

Nuclear fusion in the Sun produces electron neutrinos ν_e in the energy range from 0.1 to 15 MeV. The expected ν_e energy and neutrino flux along with the energy sensitivity for different experiments can be seen in Figure 1.4.

The primary process of hydrogen burning in the Sun, known as the pp cycle, occurs in three main stages [38]:

$$\begin{aligned} \mathrm{p} + \mathrm{p} &\rightarrow \mathrm{D} + e^{+} + \nu_e, \\ \mathrm{D} + \mathrm{p} &\rightarrow {}^{3}\mathrm{He} + \gamma, \\ {}^{3}\mathrm{He} + {}^{3}\mathrm{He} &\rightarrow {}^{4}\mathrm{He} + \mathrm{p} + \mathrm{p}. \end{aligned}$$

Since the binding energy of the deuteron (${}^{2}_{1}\mathrm{D}$) is only 2.2 MeV, the neutrinos from the reaction $p + p \rightarrow \mathrm{D} + e^{+} + \nu_e$ have low energies ($E_\nu < 0.5$ MeV) and are difficult to detect. As a result, experiments often focus on higher-energy solar neutrinos from rarer processes. The most energetic neutrinos originate from the β-decay of ${}^{8}\mathrm{B}$, which is produced through the following reactions:

$$\begin{aligned} {}^{4}_{2}\mathrm{He} + {}^{3}_{2}\mathrm{He} &\rightarrow {}^{7}_{4}\mathrm{Be} + \gamma, \\ {}^{7}_{4}\mathrm{Be} + \mathrm{p} &\rightarrow {}^{8}_{5}\mathrm{B} + \gamma. \end{aligned}$$

This is followed by the β-decay of ${}^{8}\mathrm{B}$:

$$ {}^{8}_{5}\mathrm{B} \rightarrow {}^{8}_{4}\mathrm{Be}^{*} + e^{+} + \nu_e, $$

which produces neutrinos with energies reaching up to 15 MeV.

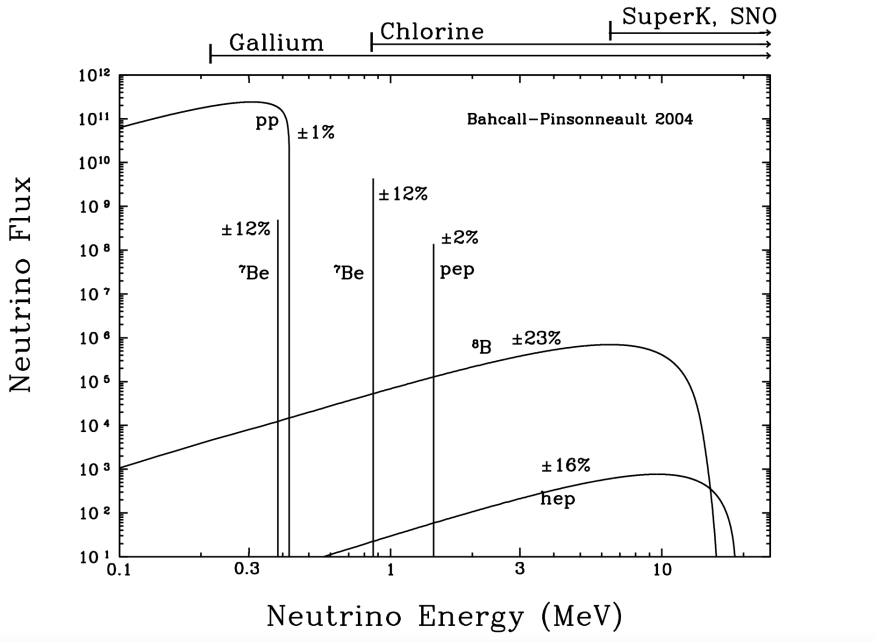

Figure 1.4: Energy spectra of neutrino fluxes from the pp chain, as predicted by the BP04 Standard Solar Model (SSM) [4]. The top axis indicates the energy ranges covered by various solar neutrino experiments. For continuous sources, the differential flux is given in units of cm^{-2} s^{-1} MeV^{-1}; for line sources, in cm^{-2} s^{-1}. The line sources include two ^{7}Be lines from the electron capture reaction $^7\mathrm{Be} + e^- \rightarrow {}^7\mathrm{Li} + \nu_e$, and the pep line from the reaction $p + e^- + p \rightarrow {}^2\mathrm{H} + \nu_e$. Percentages indicate the uncertainties in the flux values. Figure taken from [5].

Solar neutrinos can be detected by radiochemical reactions, Cherenkov detectors, and liquid scintillators. The detection of solar neutrinos revealed the so-called solar neutrino problem: the observed electron neutrino flux was only about one-third of the value predicted by the Standard Solar Model (SSM).

The first indication of this discrepancy came from Ray Davis' Homestake Solar Neutrino experiment [6]. It was designed to detect neutrinos via inverse β decay in the radiochemical reactions

$$\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^-. \tag{1.7}$$

The energy threshold for this reaction is $E_\nu^{\text{th}} = 0.814$ MeV. The detector was a large steel tank filled with 615 tons of liquid perchloroethylene (C_2Cl_4), located deep underground in the Homestake Gold Mine in Lead, South Dakota. After 25 years of operation, the experiment reported an average solar neutrino rate of $R = 2.56 \pm 0.23$ SNU. The solar neutrino unit (SNU) can be expressed as:

$$R = \sum_i \Phi_i \sigma_i \; [10^{10}\,\text{cm}^{-2}\text{s}^{-1}] \times [10^{-46}\,\text{cm}^2] = 10^{-36}\,\text{s}^{-1} \tag{1.8}$$

So, 1 SNU corresponds to 10^{-36} interactions per target atom per second. The Standard Solar Model predicted a rate of $R = 8.5 \pm 1.8$ SNU, about three times higher than the Homestake result, highlighting a major discrepancy. Results are shown in Figure 1.5.

There were three other radiochemical experiments (gallium experiments): GALLEX [39], GNO [7], and SAGE [8], all of which observed similar deficits. These experiments detected neutrinos through the reaction:

$$\nu_e + {}^{71}\text{Ga} \rightarrow {}^{71}\text{Ge} + e^-,$$

The GALLium EXperiment (GALLEX) used 101 tons of a liquid gallium chloride solution (GaCl3-HCl), containing 30.3 tons of gallium. GALLEX was located in the Laboratori Nazionali del Gran Sasso, in Italy. It operated from

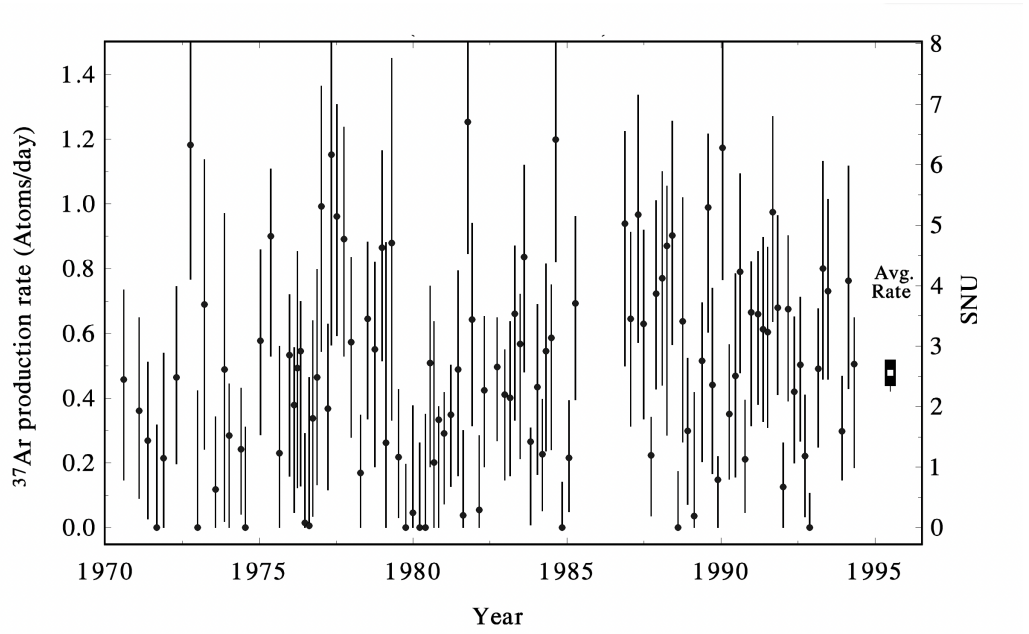

Figure 1.5: Results from 108 solar neutrino measurements recorded by the Homestake chlorine experiment between March 1970 and February 1994 [6].

1991 to 1997 and was followed by the Gallium Neutrino Observatory (GNO), which operated from 1998 to 2003 using the same detector, but with upgraded extraction equipment.

The average neutrino rate reported by GALLEX and GNO was $R = 69.3 \pm 5.5$ SNU Figure 1.6. The prediction of the SSM for this experiment is 128 SNU. This result shows a discrepancy of about half compared to the SSM prediction.

The Soviet–American Gallium Experiment (SAGE) is located 3.5 km inside Mount Andyrchi at the Baksan Neutrino Observatory in the northern Caucasus. The underground lab sits about 2000 m below the mountain's peak, with 4700 mwe of shielding that reduces the cosmic muon flux by a factor of 10^7. The experiment uses approximately 50 ton of liquid gallium, distributed across seven chemical reactors. Additional shielding suppresses neutron and

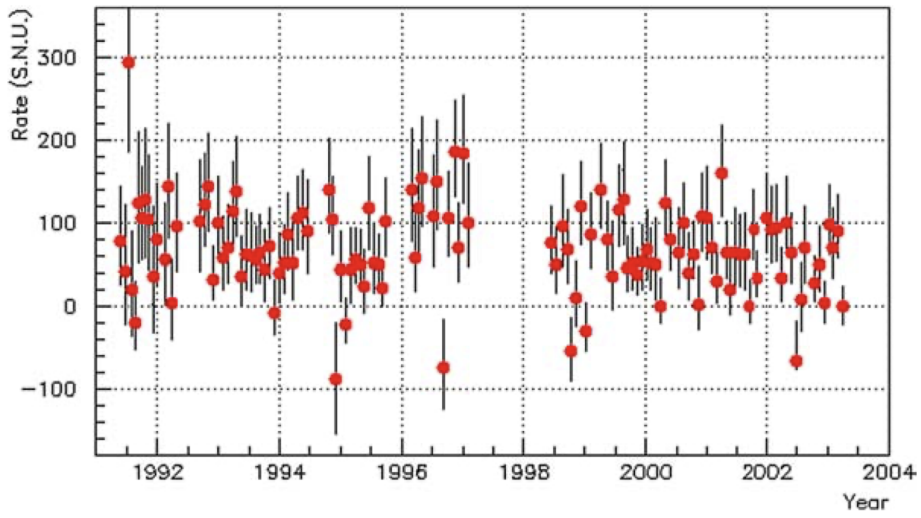

Figure 1.6: Results from individual runs of the GNO and GALLEX experiments. The plot shows the net solar neutrino production rate in SNU [7].

gamma backgrounds.

The experiment operated from 1990 to 1997. The results are shown in Figure 1.7. The average neutrino rate reported by SAGE is $65.4^{+3.1}_{-3.0}$ SNU. This result shows a discrepancy of about half compared to the SSM prediction, which is consistent with the value reported by GALLEX/GNO.

Later, Cherenkov water detectors confirmed the same reduction in the solar neutrino flux. Water Cherenkov detectors can observe neutrino interactions by detecting fast charged particles they create. When one of these particles moves faster than the speed of light in water, it gives off Cherenkov light in a cone shape around the direction of the motion of the particles. The most prominent of them are the Super-Kamiokande (SK) and the Sudbury Neutrino Observatory (SNO) experiments.

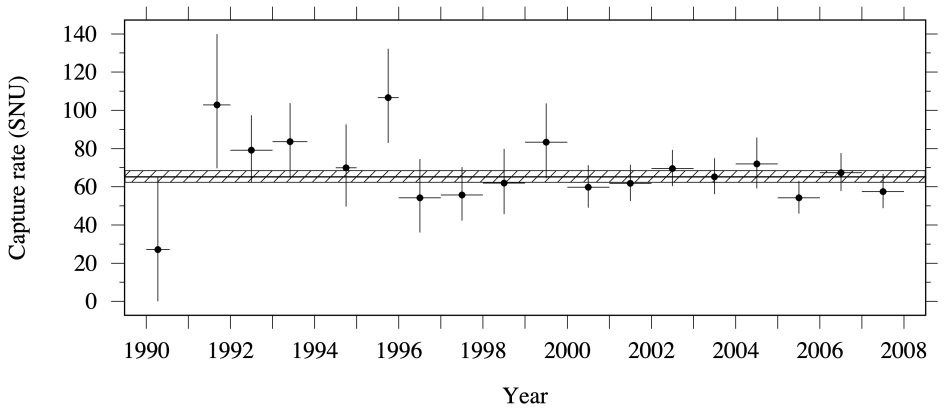

Figure 1.7: Combined SAGE results shown year by year. The shaded region represents the best-fit value and its uncertainty over the full dataset. Vertical error bars reflect statistical uncertainties [8].

The SK detector contains 50,000 tons of ultra-pure water inside a stainless steel tank measuring 39.3 meters in diameter and 41.4 meters in height. It is equipped with 11,146 20-inch photomultiplier tubes (PMTs) in the inner detector and 1,885 8-inch PMTs in the outer detector. It is situated 1000 meters underground (2700 mwe equivalent) at the Kamioka Observatory, within the Kamioka mine in Gifu Prefecture, Japan. The experiment is sensitive to ^{8}B solar neutrinos. SK determined the solar neutrino flux by observing elastic scattering primarily through the reaction:

$$\nu_e + e^- \rightarrow \nu_e + e^-, \tag{1.9}$$

with $E_\nu^{\text{th}} > 4.7$ MeV. Figure 1.8 presents the SK results of the angular distribution of solar neutrino candidates relative to the Sun's direction. A clear forward peak, caused by elastic scattering of solar neutrinos in the SK detector, can be observed. The peak's width results from both multiple scattering of

the recoil electrons and the detector's angular resolution. The shaded (dotted) region indicates background contributions.

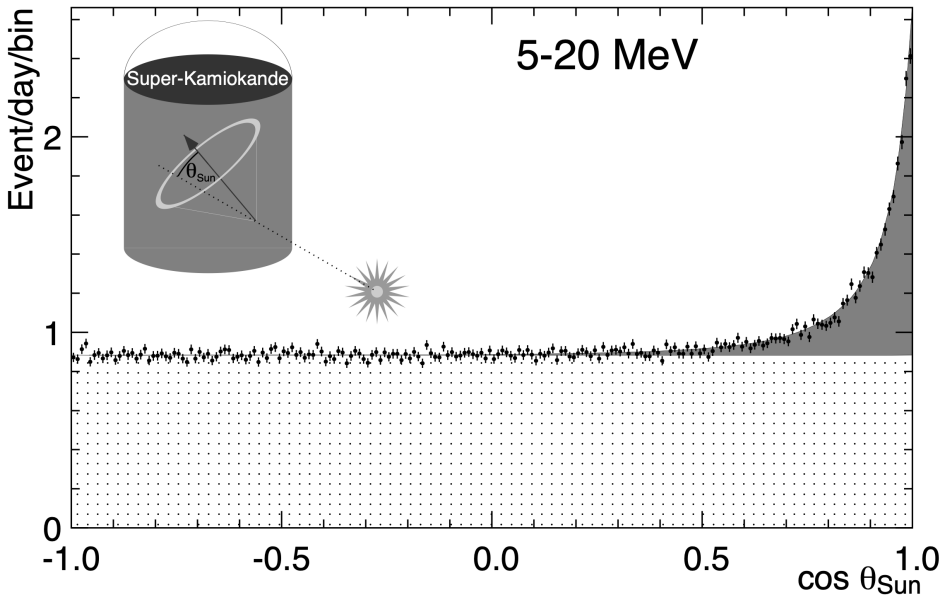

Figure 1.8: Angular distribution of solar neutrino candidates. The shaded region highlights the elastic scattering peak, while the dotted area represents background contributions [9].

The flux of solar neutrino measured by SK is:

$$\Phi_{^8\mathrm{B}}^{\mathrm{SK}} = (2.35 \pm 0.08) \times 10^6 \ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}.$$

Where the expected neutrino flux given by the SSM is $5.05 \times 10^6 \left(1^{+0.20}_{-0.16}\right)$. The observed neutrino flux by SK is smaller than half of the SSM.

Finally, the heavy water detector SNO in Ontario, Canada, provided the key result that resolved the problem [11]. Similar to SK, SNO is sensitive to ^{8}B solar neutrinos. SNO used 1,000 tons of heavy water (D_2O) contained in a

12-meter-diameter vessel, monitored by 9,600 photomultiplier tubes (PMTs). The detector is located at a depth of 2,092 meters. The cosmic ray muon background at this depth is about 65 events per day.

The main advantage of the deuterium (D_2O) is its low nuclear energy binding ($\sim$ 2.2 MeV). This provides SNO with a sensitivity to both NC and charged current (CC) neutrino events in addition to the elastic scattering (ES) events. The main reactions produced by neutrino interactions in the SNO detector are:

$$\begin{aligned} \nu_e + d &\rightarrow p + p + e^- \quad \text{(CC)}, \\ \nu_\alpha + d &\rightarrow p + n + \nu_\alpha \quad \text{(NC)}, \\ \nu_\alpha + e^- &\rightarrow \nu_\alpha + e^- \quad \text{(ES)}, \end{aligned}$$

Due to the deuteron's low binding energy, the (CC) interaction

$$\nu_e + d \rightarrow p + p + e^-$$

is kinematically allowed at solar neutrino energies. The emitted electron generates a Cherenkov ring, but because the neutrino energy is much smaller than the deuteron mass, the electron's angular distribution is nearly isotropic. Therefore, this interaction offers little directional information about the Sun and is sensitive only to the electron neutrino flux:

$$\text{CC rate} \propto \phi(\nu_e) \tag{1.10}$$

Neutrinos of all flavors can interact with deuterons via NC interactions, breaking the nucleus and producing a detectable Cherenkov signal. Since NC interactions are flavor-independent, the interaction rate depends on the total neutrino flux:

$$\text{NC rate} \propto \phi(\nu_e) + \phi(\nu_\mu) + \phi(\nu_\tau) \tag{1.11}$$

Neutrinos also scatter off electrons through ES. For electron neutrinos, both CC and NC processes contribute, while for muon and tau neutrinos only NC interactions are important. Therefore, the ES process is sensitive to all flavors, but with enhanced sensitivity to ν_e:

$$\text{ES rate} \propto \phi(\nu_e) + 0.1559\,[\phi(\nu_\mu) + \phi(\nu_\tau)]\,. \tag{1.12}$$

Here, the factor 0.1559 represents the ratio between the elastic scattering cross sections of $\nu_{\mu,\tau}$ and ν_e. The resulting fluxes measured by the SNO experiment for each interaction type are:

$$\begin{aligned}
\phi_{\text{CC}} &= 1.76^{+0.06}_{-0.05}\ (\text{stat.})^{+0.09}_{-0.09}\ (\text{syst.}) \times 10^6\ \text{cm}^{-2}\,\text{s}^{-1} \\
\phi_{\text{ES}} &= 2.39^{+0.24}_{-0.23}\ (\text{stat.})^{+0.12}_{-0.12}\ (\text{syst.}) \times 10^6\ \text{cm}^{-2}\,\text{s}^{-1} \\
\phi_{\text{NC}} &= 5.09^{+0.44}_{-0.43}\ (\text{stat.})^{+0.46}_{-0.43}\ (\text{syst.}) \times 10^6\ \text{cm}^{-2}\,\text{s}^{-1}.
\end{aligned}$$

Using Equation 1.10, Equation 1.11 and Equation 1.12 the contributions for

each flavor content is calculated:

$$\phi(\nu_e) = 1.76^{+0.05}_{-0.05}\,(\text{stat.})^{+0.09}_{-0.09}\,(\text{syst.})$$
$$\phi(\nu_{\mu\tau}) = 3.41^{+0.45}_{-0.45}\,(\text{stat.})^{+0.48}_{-0.45}\,(\text{syst.}).$$

The combined measurements are inconsistent with the hypothesis that the ^{8}B flux is composed entirely of ν_e, but are consistent with a composition of approximately 1/3 ν_e and 2/3 ν_μ and/or ν_τ. Figure 1.9 shows the constraints on the ν_e flux versus the combined ν_μ and ν_τ fluxes, determined from the measured CC, ES, and NC interaction rates.

Based on the SNO results, it is understood that the deficit of solar neutrinos is due to ν_e adiabatic flavor conversion in matter with varying density (MSW effect [40][41][42]) in the Sun. As it will be explained in section 1.7, the three active neutrino flavors are a mixture of the mass eigenstates ν_1, ν_2, and ν_3, and inversely, the mass eigenstates are a mixture of the three flavor ones (3-flavor paradigm). Since solar neutrinos are produced in the Sun's core, where the density of electrons is very high (matter effect is strong), the ν_e produced in weak interactions effectively corresponds to the matter eigenstate ν_2^m. As the neutrino propagates outward through the Sun, the electron density decreases slowly, and the neutrino remains in the same matter eigenstate. This adiabatic evolution causes ν_2^m to gradually become the vacuum mass eigenstate ν_2. The mass state ν_2 is a mixture of almost equal amounts of ν_e, ν_μ, and ν_τ, providing roughly equal detection probability for each neutrino flavor when measured on Earth.

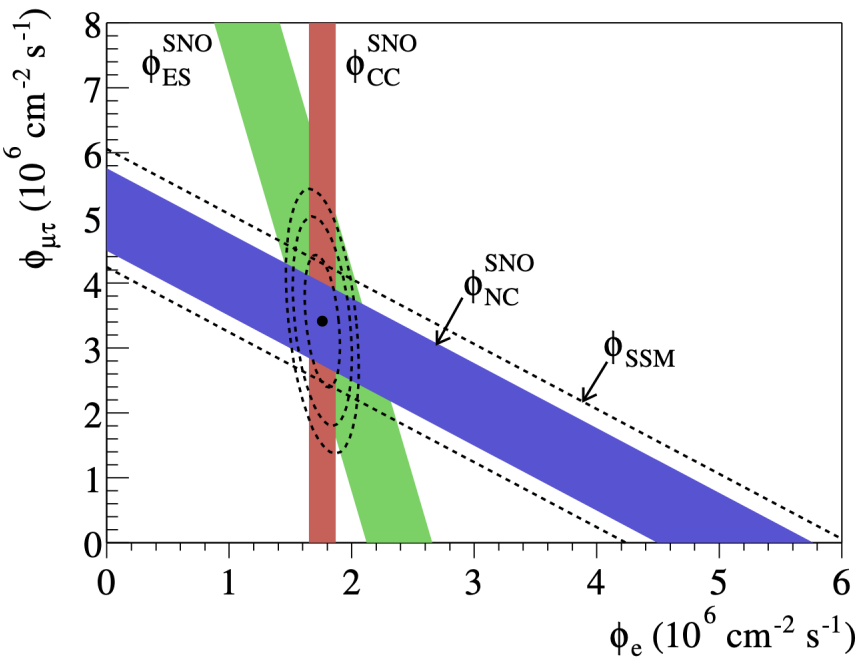

Figure 1.9: Flux of ^{8}B solar neutrinos of μ or τ flavor versus the flux of electron neutrinos, as determined from the three neutrino interaction channels measured by SNO. The diagonal bands represent the total ^{8}B flux predicted by the BP2000 SSM [10] (dashed lines) and the flux measured via the NC channel in SNO (solid band). The intersections of these bands with the axes correspond to the $\pm 1\sigma$ uncertainties. The overlap of the bands at the best-fit values for ϕ_e and $\phi_{\mu\tau}$ indicates consistency with neutrino flavor transformation.[11].

This explains why experiments like Homestake and SK (which are sensitive primarily to CC or ES processes respectively) measured a deficit in the solar neutrino flux. The breakthrough came with SNO, which was capable of detecting solar neutrinos through CC interactions (sensitive only to ν_e), NC interactions (equally sensitive to all active flavors), and ES. SNO observed that while the CC event rate was suppressed by about two-thirds, the NC event rate matched the total flux predicted by SSMs.

One may be tempted to say that this observed effect is due to neutrino oscillations, but this is not the case, since a ν_2 exiting the Sun does not have the ability to interfere (you need at least two mass states to produce interference), and therefore no oscillations occur. What SNO discovered was the adiabatic flavor transformation of solar neutrinos, which explains the measurements stated before, solving the solar neutrino problem. A detailed description of the process is found in [43].

1.5 Atmospheric Neutrinos

Another plentiful source of neutrinos is our atmosphere. High-energy heavy nuclei, such as protons or helium nuclei coming from outer space, collide with the upper atmosphere, creating particle showers that travel and interact all the way to the Earth's surface. The final particle products, such as pions, decay into muons and neutrinos. In the same manner, muons decay into electrons, muon neutrinos, and electron antineutrinos. The schematic illustration of a cosmic ray interaction and neutrino production can be seen in Figure 1.10. As a result of this decay chain, we expect roughly twice as many muon neutrinos as electron neutrinos.

$$\frac{\nu_\mu}{\nu_e} \approx 2 \tag{1.13}$$

We can calculate the ratio of the number of ν_μ over the number of ν_e, comparing theoretical predictions and experimental measurements, using the following expression:

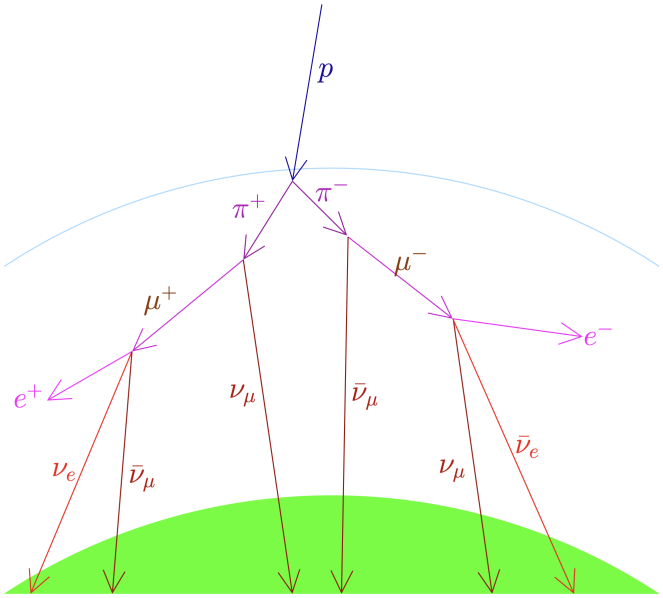

Figure 1.10: Schematic illustration of atmospheric neutrino production. A cosmic ray interacts with nuclei in the upper atmosphere, producing charged pions (π). These pions decay into muons (μ) and muon neutrinos (ν_μ). The muons subsequently decay into electrons (e), electron neutrinos (ν_e), and additional muon neutrinos (ν_μ). The resulting neutrinos can travel to detectors on the Earth. Image taken from [12].

$$R = \frac{R_{\mu e}^{\text{exp}}}{R_{\mu e}^{\text{theo}}} = \frac{\left(\frac{\nu_\mu}{\nu_e}\right)^{\text{exp}}}{\left(\frac{\nu_\mu}{\nu_e}\right)^{\text{theo}}}. \tag{1.14}$$

Here R is the ratio-of-ratios. Experiments such as SK, Soudan-2, and IMB measured the neutrino fluxes of ν_μ and ν_e as a function of the zenith angle θ_z, which ranges from 0 (downward-going neutrinos, arriving from directly from the atmosphere) to π (upward-going neutrinos, having traversed the Earth before reaching the detector). In the absence of oscillations, the ratio-of-ratios would be expected to be equal to 1 and remain constant across all zenith

angles. Any deviation from this behavior indicates an anomaly between the data and the no-oscillation prediction.

In atmospheric neutrino experiments, the fluxes of different neutrino flavors are inferred by detecting the charged leptons produced in neutrino–nucleon interactions:

$$\nu_\ell + N \to \ell^- + X, \qquad \bar{\nu}_\ell + N \to \ell^+ + X, \qquad (\ell = e, \mu, \tau). \tag{1.15}$$

The IMB (Irvine–Michigan–Brookhaven) detector was a 24 m × 18 m × 19 m water tank filled with approximately 8000 tons of water, with a fiducial mass of 3300 tons. It was located 610 m underground in the Morton Thiokol salt mine near Cleveland, Ohio (USA), under an overburden of 1570 mwe, and operated from 1982 to 1991.

The IMB experiment observed an anomaly in low-energy events (below 1.5 GeV), while no anomaly was found in upward-going muon events with energy above 0.95 GeV. The anomaly reported by IBM by the ratio-of-ratios was [44]:

$$R = 0.54 \pm 0.05 \pm 0.11. \tag{1.16}$$

Later, the SK experiment presented their measurements in 1998 for sub-GeV and multi-GeV indicated a deficit on data for upward-going neutrino fluxes [13]:

$$R = 0.63 \pm 0.03\,(\text{stat.}) \pm 0.05\,(\text{sys.}) \quad \text{sub-GeV}$$

$$R = 0.65 \pm 0.05\,(\text{stat.}) \pm 0.08\,(\text{sys.}) \quad \text{multi-GeV}$$

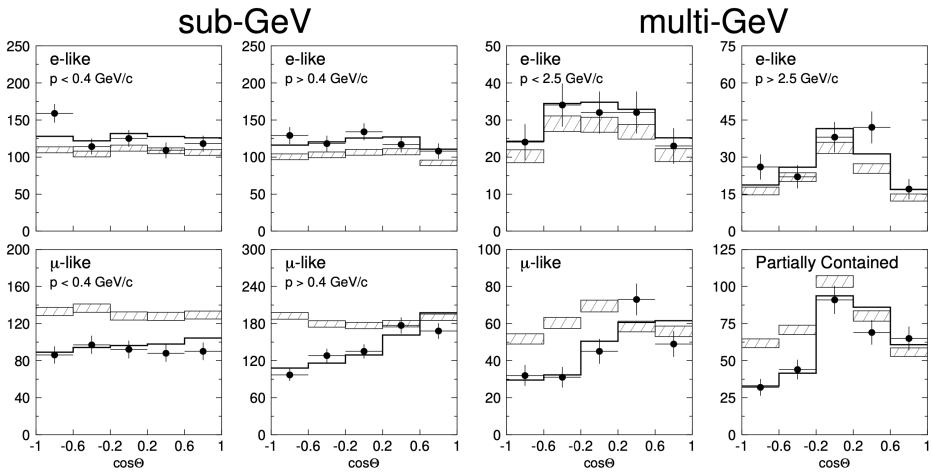

Figure 1.11: Zenith angle distributions of μ-like and e-like events from sub-GeV and multi-GeV data samples. Upward-going particles correspond to $\cos\Theta < 0$, while downward-going particles have $\cos\Theta > 0$. Sub-GeV data are divided into two momentum ranges: $p < 400$ MeV/c and $p > 400$ MeV/c. For multi-GeV e-like events, distributions are shown for $p < 2.5$ GeV/c and $p > 2.5$ GeV/c. Multi-GeV μ-like. The hatched regions represent the Monte Carlo expectation assuming no oscillations, normalized to the detector live-time and including statistical uncertainties. The bold curve indicates the best-fit prediction assuming $\nu_\mu \leftrightarrow \nu_\tau$ oscillations, with the total flux normalization treated as a free parameter [13].

SK observed fewer upward-going muon neutrinos, which came from zenith angles that required the neutrinos to travel through the Earth, than downward-going neutrinos produced in the atmosphere approximately 20 kilometers above the detector. The depletion of muon neutrino events was attributed to ν_μ oscillating into ν_τ. For this transformation to occur, neutrinos must have mass.

This observations provided the solution to the atmospheric neutrino problem.

The SOUDAN-2 experiment, provided an independent measurement, confirming the results observed by SK. SOUDAN-2 was located 710 m underground in the Soudan Mine State Park in Minnesota (USA), with an overburden of 2070 mwe. Its main detector was an iron tracking calorimeter functioning as a time projection chamber. The detector had a total mass of 963 tons and a fiducial mass of 770 tons, and was surrounded by an active shield of aluminum proportional tubes. The experiment collected data from 1989 to 2001.

SOUDAN-2 measured the ratio-of-ratios for sub-GeV events as [14];

$$R = 0.69 \pm 0.10 \pm 0.06, \tag{1.17}$$

indicating a disappearance of muon neutrinos for energies below about 1 GeV. The zenith-angle distribution of upward-going muons also showed a deficit consistent with muon neutrino disappearance as show in Figure 1.12.

So far, we have discussed the different experiments that were able to explain the anomalies observed in the solar and atmospheric neutrino problems. The two effects can be explained either by adiabatic neutrino flavor conversion in matter, in the case of solar neutrinos, or by neutrino oscillations, in the case of atmospheric neutrinos. The next section will provide an overview of neutrino interactions relevant to oscillation experiments such as NOvA, discuss the theory of neutrino oscillations, and address the effects that led to adiabatic neutrino flavor conversion in matter.

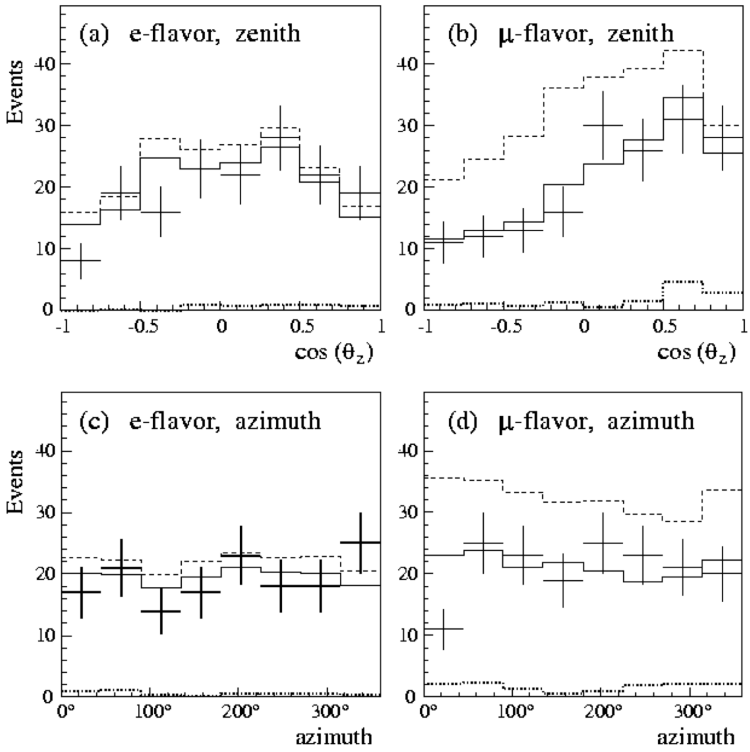

Figure 1.12: Angular distributions for e-flavor events (plots (a) and (c)) and μ-flavor events (plots (b) and (d)). Plots (a) and (b) show the cosine of the zenith angle, and plots (c) and (d) the azimuth angle. The points with error bars are the data. The dashed histograms are the sum of the predicted unoscillated neutrino distribution plus the fitted rock background contribution (neutrons and gamma rays produced by muon interactions on the surrounding rock). The solid histograms are the same but with the neutrino distribution weighted by the oscillation probability. The dotted histograms are the contribution of the rock background. Downward-going events have $\cos\theta_z = +1.0$. Note the depletion of μ-flavor events at all but the highest value of $\cos\theta_z$ [14].

1.6 Neutrino Interactions

Besides exceedingly small gravitational effects, neutrinos interact exclusively through the weak interaction, mediated by exchange of the charged $W^{\pm}$ bosons for CC interactions or the neutral Z^0 boson for NC interactions. Depending on energy and the internal structure of the target, interactions are typically classified as elastic scattering, quasi-elastic scattering, meson exchange currents (MEC), resonant production, and deep inelastic scattering (DIS).

1.6.1 Elastic Scattering

In ordinary matter, neutrinos interact via elastic scattering with electrons and nucleons. These processes are mediated by the weak interaction and proceed either through exchange of a neutral Z^0 boson (NC) or a charged $W^{\pm}$ boson (CC), depending on the neutrino flavor.

NC elastic scattering processes include Equations 1.18 – 1.21, for all $\ell = e, \mu, \tau$.

$$\nu_\ell + e^- \rightarrow \nu_\ell + e^- \,, \tag{1.18}$$

$$\bar{\nu}_\ell + e^- \rightarrow \bar{\nu}_\ell + e^- \,, \tag{1.19}$$

$$\nu_\ell + N \rightarrow \nu_\ell + N \,, \quad N = p, n \,, \tag{1.20}$$

$$\bar{\nu}_\ell + N \rightarrow \bar{\nu}_\ell + N \,, \quad N = p, n \,. \tag{1.21}$$

These interactions preserve the identities of the incoming particles and do not produce new final-state leptons. Since the final state is the same as the initial one, elastic scattering has no energy threshold. It results in a redistribution of energy and momentum between the neutrino and the target.

Elastic scattering on electrons provides a clean experimental signature due to the relatively large recoil of the electron, while scattering on nucleons is more difficult to detect because of the nucleon's larger mass. The corresponding Feynman diagrams for these NC processes are shown in Figure 1.13.

For ν_e and $\bar{\nu}_e$, the elastic scattering with electrons receives an additional contribution from the CC processes, as shown in Equations 1.22 – 1.23 and the diagrams in Figure 1.14:

$$\nu_e + e^- \rightarrow \nu_e + e^- \tag{1.22}$$

$$\bar{\nu}_e + e^- \rightarrow \bar{\nu}_e + e^- \tag{1.23}$$

These processes have the same external particles as the NC channels, but involve the CC mediators (W). The total amplitude for ν_e and $\bar{\nu}_e$ elastic scattering on electrons receives contributions from both NC and CC diagrams, leading to interference. As a result, the cross sections for ν_e and $\bar{\nu}_e$ are enhanced relative to other flavors in matter.

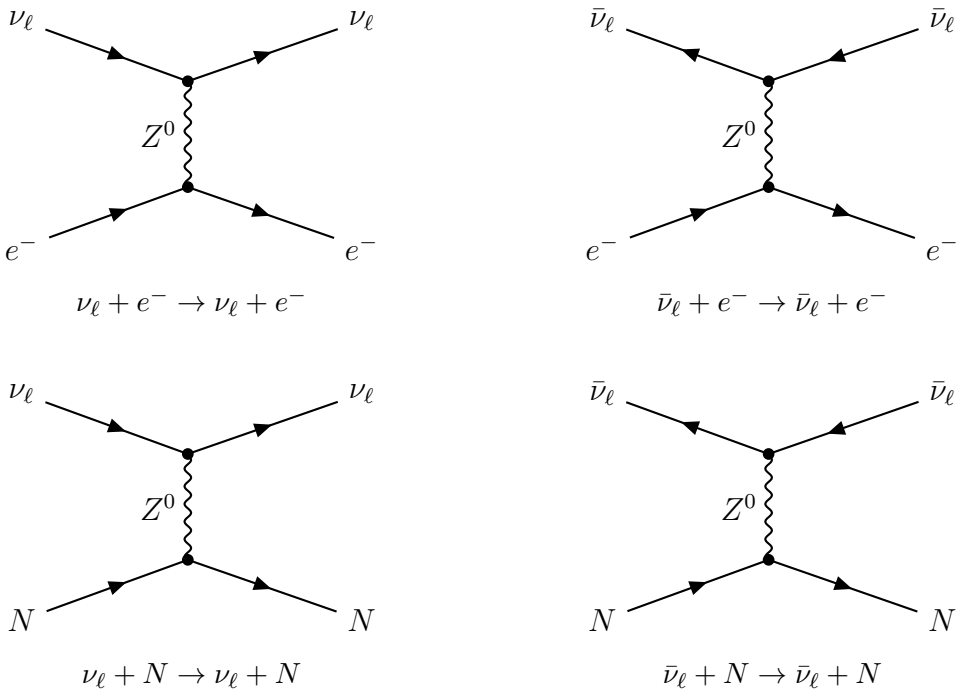

Figure 1.13: NC elastic scattering diagrams for neutrinos and antineutrinos with electrons and nucleons.

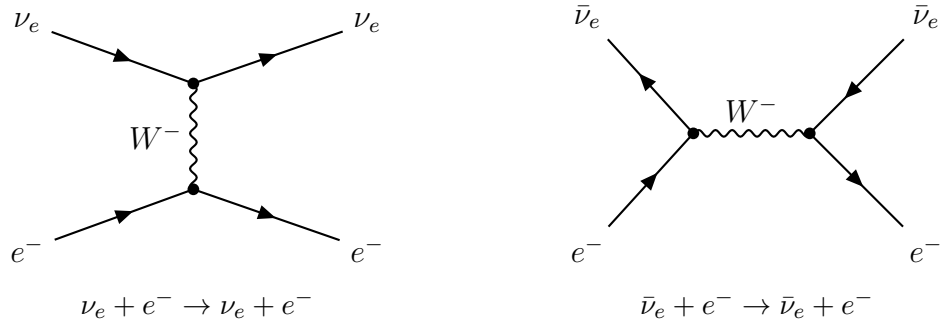

Figure 1.14: CC elastic scattering diagrams for neutrinos and antineutrinos with electrons.

1.6.2 Quasi-Elastic (QE) Interactions

Quasi-elastic (QE) interactions occur when neutrinos interact with individual nucleons within a nucleus, resulting in a charged lepton and a recoiling nucleon

in the final state. These processes are governed by the weak interactions mediated by the $W^{\pm}$ and Z^0 bosons. Unlike elastic scattering, the identity of the incoming neutrino changes due to the lepton emission, and energy is transferred to break the nucleon binding.

The typical CC QE reactions are given in Equations 1.24–1.25:

$$\nu_\ell + n \to \ell^- + p\,, \tag{1.24}$$

$$\bar{\nu}_\ell + p \to \ell^+ + n\,, \tag{1.25}$$

where $\ell = e$, μ, τ. These interactions form the primary detection mechanism in many low-energy neutrino experiments, such as reactor-based or accelerator-based detectors.

The corresponding Feynman diagrams for CC QE processes are shown in Figure 1.15. The flavor of the produced charged lepton matches that of the incoming neutrino, ensuring lepton flavor conservation.

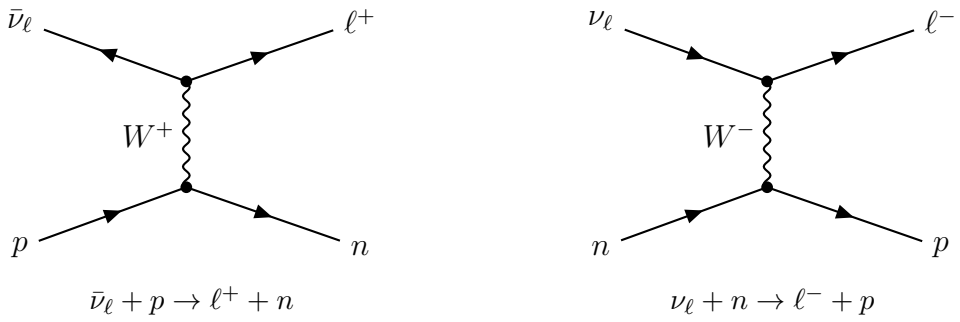

Figure 1.15: Feynman diagrams of CCQE interactions for neutrinos and antineutrinos.

1.6.3 Resonant (RES) Interactions

In resonant interactions, a neutrino excites a nucleon into a baryon resonance, typically a Δ, which decays into a nucleon and a pion. Resonant single pion production occur in CC interactions via the following channels:

$$\nu_\ell + p \xrightarrow{\Delta^{++}} \ell^- + p + \pi^+ \,, \tag{1.26}$$

$$\nu_\ell + n \xrightarrow{\Delta^{+}} \ell^- + p + \pi^0 \,, \tag{1.27}$$

$$\nu_\ell + n \xrightarrow{\Delta^{+}} \ell^- + n + \pi^+ \,. \tag{1.28}$$

The corresponding Feynman diagram for Equation 1.26 is shown in Figure 1.16.

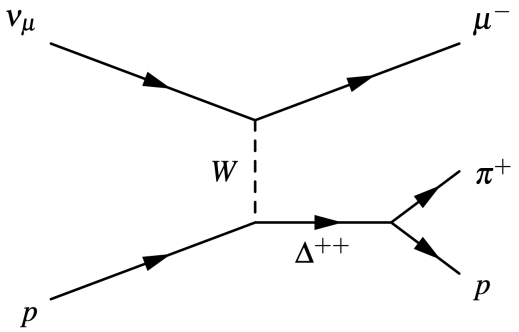

Figure 1.16: CC resonant pion production: $\nu_\mu + p \to \mu^- + \Delta^{++} \to \mu^- + p + \pi^+$.

1.6.4 Deep Inelastic Scattering

In deep inelastic scattering (DIS) interactions, the neutrino has sufficient energy to interact with individual quarks in a nucleon, as shown in Figure 1.17.

The interaction breaks apart the nucleon, producing many hadrons in the final state (X). At high neutrino energies, DIS becomes the dominant interaction mode.

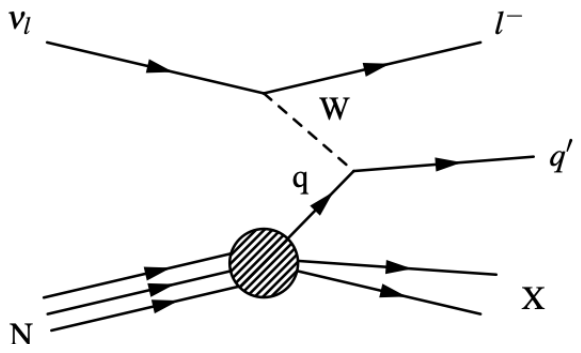

Figure 1.17: Diagram of a CC Deep Inelastic Scattering interaction.

1.6.5 Meson Exchange Current (MEC)

When a neutrino interacts with two nucleons, it can knock both out of the nucleus, leaving behind two holes. These types of events are called two-particle–two-hole (2p2h). In addition, the two ejected nucleons interact in a correlated way by exchanging a meson (typically a π), which is why the name Meson Exchange Current (MEC) is used. This interaction plays an essential role in the region between QE and RES, but it is still poorly understood and modeled.

1.6.6 Coherent Pion Production

A neutrino can interact with an entire nucleus in a coherent way, leaving the nucleus intact and producing a pion in the final state as shown in Figure 1.18. These interactions have a very forward-going pion and small four-momentum transfer to the nucleus.

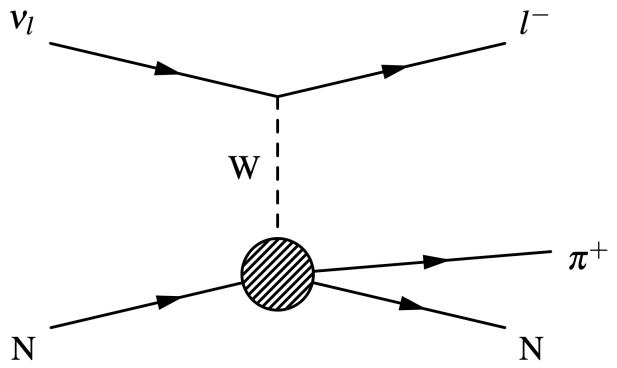

Figure 1.18: Diagram of coherent pion production: $\nu_\ell + N \rightarrow \ell^- + N + \pi^+$.

Up to this point, the most common types of neutrino interactions in experiments such as NOvA have been discussed. The next section will cover the theory of neutrino oscillations

1.7 Neutrino Oscillations

Neutrino oscillations is a quantum mechanical phenomena proposed by Pontecorvo In 1957 [45][46]. He was inspired by the known kaon oscillation process $K^0 \rightarrow \bar{K^0}$, and speculated that a similar transition, such as $\nu_e \rightarrow \bar{\nu}_e$, could

also occur. At that time, only the existence of the electron neutrino had been confirmed. After the discovery of the muon neutrino, Maki, Nakagawa, and Sakata extended the idea by suggesting that neutrinos could oscillate between different flavors [47].

For neutrino oscillations to take place, neutrinos must have finite mass, and the flavor eigenstates ν_e, ν_μ, ν_τ must differ from the mass eigenstates ν_1, ν_2, ν_3. These two sets of states are related by a unitary mixing matrix U.

Let us consider the simplest approach using plane waves. In this framework, the flavor basis can be expressed as:

$$|\nu_\alpha(t)\rangle = \sum_i U^*_{\alpha i} \, |\nu_i(t)\rangle \tag{1.29}$$

The massive neutrino states $|\nu_i\rangle$ are eigenstates of the Hamiltonian,

$$\mathcal{H} \, |\nu_i\rangle = E_i \, |\nu_i\rangle \ ,$$

with energy eigenvalues

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

The energy eigenstates evolve as (using units where $\hbar = c = 1$):

$$|\nu_i(t)\rangle = e^{-iE_i t} \, |\nu_i(0)\rangle \, , \tag{1.30}$$

and the transition probability to the flavor state β can be expressed as:

$$P(\nu_\alpha \to \nu_\beta) = \sum_{i,j} U_{\alpha i} U^*_{\beta i} U^*_{\alpha j} U_{\beta j} e^{-i \frac{\Delta m^2_{ij}}{2E} L} \tag{1.31}$$

with $\Delta m^2_{ij} = m^2_i - m^2_j$, and assuming the ultra-relativistic approximation where $L \simeq t$.

The phase of neutrino oscillations depends on the source–detector distance L and the neutrino energy E, and is given by

$$\Phi_{ij} = -\frac{\Delta m^2_{ij} L}{2E} \, .$$

This phase is governed by the mass-squared differences Δm^2_{ij}, which are physical constants. The amplitude of the oscillations is determined by the elements of the mixing matrix U, which are also fundamental constants of nature.

Neutrino oscillation experiments provide precise information on the squared-mass differences Δm^2_{ij}, but not on the absolute masses of neutrinos. However, the observation of oscillations confirms that at least two neutrinos have non-zero masses, and that their squared masses differ by more than $|\Delta m^2_{ij}|$.

1.7.1 Two-Flavor Neutrino Oscillations

In the two-flavor scenario (ν_e, ν_μ), the U matrix is parametrized using a single mixing angle θ:

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{1.32}$$

The probability for an electron neutrino to transform into a muon neutrino is given by:

$$P(\nu_e \to \nu_\mu) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2_{21} L}{4E}\right) \tag{1.33}$$

The survival probability is:

$$P(\nu_e \to \nu_e) = 1 - P(\nu_e \to \nu_\mu) \tag{1.34}$$

With physical units, the appearance probability becomes:

$$P(\nu_e \to \nu_\mu) = \sin^2(2\theta)\sin^2\left(1.27\,\frac{\Delta m^2_{21}}{\mathrm{eV}^2}\frac{L}{\mathrm{km}}\frac{\mathrm{GeV}}{E}\right) \tag{1.35}$$

Therefore, the transition probability between these two types of neutrinos can be described using two oscillation parameters: $\sin^2(2\theta)$ and Δm^2_{21}.

1.7.2 Three-Flavor Neutrino Oscillations

For the three active neutrinos ν_e, ν_μ, and ν_τ, the unitary mixing matrix known as the PMNS (Pontecorvo–Maki–Sakata–Nakagawa) matrix U contains three

mixing angles and one CP-violating phase: θ_{12}, θ_{23}, θ_{13}, and δ. This matrix takes the form:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{13}s_{23}e^{i\delta} & c_{12}c_{23} - s_{12}s_{13}s_{23}e^{i\delta} & c_{13}s_{23} \\ s_{12}s_{23} - c_{12}s_{13}c_{23}e^{i\delta} & -c_{12}s_{23} - s_{12}s_{13}c_{23}e^{i\delta} & c_{13}c_{23} \end{pmatrix} \tag{1.36}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. Atmospheric neutrino oscillations are primarily governed by θ_{23} and Δm^2_{32}, whereas solar neutrino oscillations depend on θ_{12} and Δm^2_{12}. CP violation in the lepton sector would occur if δ is neither 0 nor π.

The primary goal of the NOvA experiment is to measure ν_e appearance. In this case, the transition probability in vacuum is written as:

$$P(\nu_\mu \to \nu_e) = P_{\text{atm}} + P_{\text{sol}} + P_{\text{int}} \tag{1.37}$$

Here, P_{atm} is governed by the atmospheric oscillation parameters:

$$P_{\text{atm}} = \sin^2\theta_{23} \sin^2 2\theta_{13} \sin^2\Delta_{31} \tag{1.38}$$

and, for simplicity,

$$\Delta_{ij} = 1.27\left(\frac{\Delta m^2_{ij} L}{E}\right) \tag{1.39}$$

P_{sol} is governed by the solar oscillation parameters:

$$P_{\text{sol}} = \cos^2\theta_{13}\cos^2\theta_{23}\sin^2 2\theta_{12}\sin^2\Delta_{21} \tag{1.40}$$

The interference term P_{int} includes both solar and atmospheric contributions as well as CP-violating effects. For neutrino $(-)$ and antineutrino $(+)$ channels:

$$P_{\text{int}} = J\left[\cos\delta\cos\Delta_{32}\sin\Delta_{31}\sin\Delta_{21} \mp \sin\delta\sin\Delta_{31}\sin\Delta_{21}\right] \tag{1.41}$$

where J is the Jarlskog invariant:

$$J = \cos\theta_{13}\sin\theta_{12}\sin\theta_{13}\sin\theta_{23} \tag{1.42}$$

CP violation in neutrinos is important because it could explain why the universe has more matter than antimatter. The Big Bang should have produced both in equal amounts, but today we only see matter with the ratio of baryonic matter to radiation being $\frac{n_B}{n_\gamma} \approx 6 \times 10^{-10}$. One way to create this imbalance is through CP violation. If neutrinos and antineutrinos oscillate differently, it would be evidence of CP violation in the lepton sector.

NOvA is designed to test this. It measures how often muon neutrinos turn into electron neutrinos, and how often muon antineutrinos turn into electron antineutrinos. If there's a difference, it means CP is violated.

In addition, NOvA is also sensitive to Δm^2_{32} and θ_{23}. These oscillation param-

eters, including δ_{CP} are currently a focus of major study to determine their precise value. As we will discuss in section 1.9, the sign of Δm^2_{32} parameter is still unknown. The sign determines if $m^2_3 > m^2_2 > m^2_1$ (positive sign for Δm^2_{32}) or $m^2_2 > m^2_3 > m^2_1$ (negative sign for Δm^2_{32}), known as Normal Ordering (NO) and Inverted Ordering (IO). A graphical representation of the mass ordering, sometimes also referred to as mass hierarchy, can be seen in Figure 1.19.

Δm^2_{32} is sensitive to matter effects (see discussion in section 1.8). In atmospheric neutrino experiments and Long-Baseline experiments, such as NOvA, neutrinos travel through the earth experiencing matter effects that alter the oscillation probability and provide a handle to measure the sign of Δm^2_{32}.

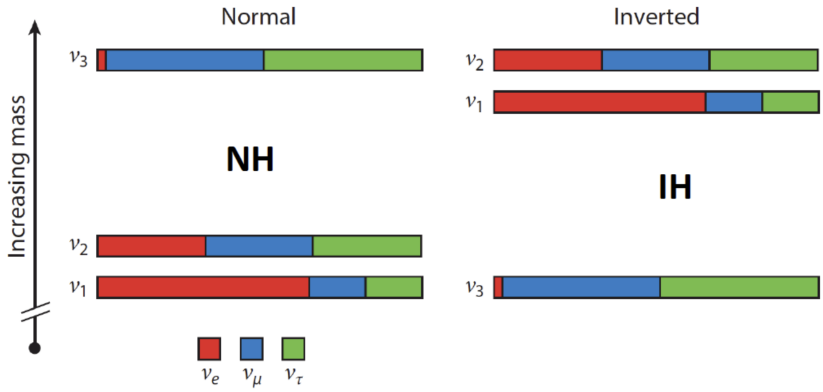

Figure 1.19: Possible neutrino mass hierarchies. Left, Normal Hierarchy and right, Inverted Hierarchy. The size each of color square represents the interaction probability for different flavor states (ν_e, ν_μ and ν_τ).

The determination of the mass hierarchy can shed light on different unsolved questions. For example, determine if neutrinos are Dirac or Majorana (neutrinos are the same as antineutrinos). This is currently explored by neutrino double beta decay experiments. In these experiments the rate is proportional

to the effective Majorana mass, $m_{\beta\beta}$ ($m_{\beta\beta} = |U_{e1}|^2 m_1 + |U_{e2}|^2 m_2 + |U_{e3}|^2 m_3$). For the IH, the lower bound for $m_{\beta\beta}$ is $m_{\beta\beta} \geq 14$ meV and for NH $m_{\beta\beta} = 0$ meV is possible. Here, the IH favors the hypothesis of Majorana neutrinos. In addition, it can also disentangle the degeneracy for the measurements of δ_{CP} and provide more information about the bounds of the total neutrino mass $\sum m_i$ from Cosmological measurements.

For the mixing angle θ_{23}, the octant has not been determined yet. θ_{23} controls the mixing between muon and tau neutrinos and determines the strength of $\nu_\mu \leftrightarrow \nu_\tau$ oscillations. It is currently the less precisely known of the three mixing angles

The three possible values of θ_{23} are:

- Maximal ($\theta_{23} = 45°$), suggesting symmetry between muon and tau flavors: ν_μ and ν_τ components are equal in ν_3
- In the lower octant ($\theta_{23} < 45°$), suggesting more tau-like component for ν_3
- Or in the upper octant ($\theta_{23} > 45°$), suggesting more muon-like component for ν_3.

The values of Δm^2_{21}, θ_{12} and θ_{13} have been determined precisely by solar and reactor neutrino experiments. The Δm^2_{21} sign is known to be positive, this implies $m_2 > m_1$, and the θ_{12} and θ_{13} are $\approx 33.4°$ and $\approx 8.57°$ respectively.

1.8 Neutrino Oscillations in Matter

While oscillations in vacuum come from differences in mass and mixing angles, the picture changes when neutrinos propagate through matter. In dense environments, such as the Sun, a supernova, or even the Earth, electron neutrinos undergo coherent forward scattering with electrons via CC interactions, while muon and tau neutrinos do not (matter does not contain muons or taus). This asymmetry introduces an additional effective potential for electron neutrinos, modifying their propagation. The resulting phenomenon is known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect [40][41][42].

The additional potential experienced by electron neutrinos is given by

$$V_{CC} = \pm\sqrt{2} G_F N_e, \tag{1.43}$$

where G_F is the Fermi constant, N_e is the local electron number density, and the positive sign is for neutrinos and negative sign for antineutrinos. This potential alters the effective Hamiltonian describing neutrino evolution in matter. For a two-flavor system (ν_e, ν_μ), the Hamiltonian in the flavor basis becomes

$$H_m = \frac{1}{2E} \begin{pmatrix} -\Delta m^2 \cos 2\theta + A & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta - A \end{pmatrix}, \tag{1.44}$$

with $A = 2\sqrt{2} G_F N_e E$. This leads to modified oscillation parameters in matter,

including a new effective mixing angle θ_m given by

$$\sin 2\theta_m = \frac{\sin 2\theta}{\sqrt{\left(\cos 2\theta - \frac{A}{\Delta m^2}\right)^2 + \sin^2 2\theta}}. \tag{1.45}$$

Note here that the mixing angle in matter can become maximal ($\theta_m = \pi/4$) even when the vacuum angle θ is small. This occurs at the resonance condition

$$\Delta m^2 \cos 2\theta = 2\sqrt{2} G_F N_e E. \tag{1.46}$$

This resonant enhancement was essential to resolve the solar neutrino problem, as discussed in section 1.4. The flavor evolution caused by the MSW effect leads to a suppressed probability of detecting solar electron neutrinos at Earth, consistent with observations from experiments such as SNO and Super-Kamiokande

As mentioned previously, neutrinos experience matter effects when they traverse the earth. In the case of the NOvA experiment, the source of neutrinos produced at Fermilab (ν_μs) can transition from ν_μ to ν_e. The first-order transition probability for $\nu_\mu \to \nu_e$, including matter effects is:

$$\begin{aligned} P_{\nu_\mu \to \nu_e} \simeq\ & 4 s_{13}^2 s_{23}^2 \frac{\sin^2\left[(1-A)\Delta\right]}{(1-A)^2} \\ & + 8 c_{12} s_{12} c_{23} s_{23} s_{13} \alpha \frac{\sin(A\Delta)\sin\left[\Delta(1-A)\right]}{A(1-A)} \cos(\Delta \pm \delta_{\mathrm{CP}}), \end{aligned} \tag{1.47}$$

where

$$\Delta \equiv \frac{\Delta m^2_{31} L}{4E}, \quad A \equiv \frac{2EV}{\Delta m^2_{31}}, \quad \alpha \equiv \frac{\Delta m^2_{21}}{\Delta m^2_{31}}. \tag{1.48}$$

Figure 1.20 illustrates the impact of matter effects on $\nu_\mu \rightarrow \nu_e$ oscillations at NOvA-relevant baseline (L) and energies (E). These effects are inverse between neutrinos and antineutrinos, and between the normal and inverted mass hierarchies, which are distinguished by the sign of Δm^2_{32}: positive for normal, negative for inverted.

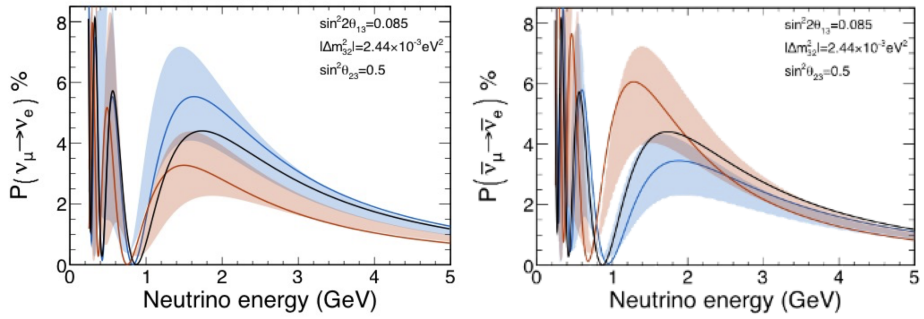

Figure 1.20: Oscillation probabilities $\nu_\mu\,(\bar{\nu}_\mu) \rightarrow \nu_e\,(\bar{\nu}_e)$ for a fixed baseline (L= 810 Km). The black curves represent the transition probability in the vacuum, while the blue curves and red curves represent the oscillation probability in matter for NH (blue) and IH (red) mass hierarchies. The shaded regions shows different variations of δ_{CP}. Left, neutrinos. Right, antineutrinos. An inverse effect for the mass hierarchy is observed for neutrinos and antineutrinos. The figure is taken from [15].

1.9 Current Status of the Neutrino Oscillation Parameters

Neutrino oscillation experiments have established that neutrinos have mass and that the flavor eigenstates $(\nu_e, \nu_\mu, \nu_\tau)$ are superpositions of three mass

eigenstates (ν_1, ν_2, ν_3). The oscillation probabilities depend on three mixing angles, two independent mass-squared differences, and one CP-violating phase.

While there are three possible mass-squared differences,

$$\Delta m^2_{21} = m^2_2 - m^2_1, \quad \Delta m^2_{31} = m^2_3 - m^2_1, \quad \Delta m^2_{32} = m^2_3 - m^2_2, \tag{1.49}$$

only two are independent because they are related by:

$$\Delta m^2_{31} = \Delta m^2_{32} + \Delta m^2_{21}. \tag{1.50}$$

Thus, the full three-flavor oscillation framework is described by:

- Two mass-squared differences: Δm^2_{21} (solar), Δm^2_{32} (atmospheric)
- Three mixing angles: $\theta_{12}, \theta_{13}, \theta_{23}$,
- One CP-violating phase: δ_{CP}.

Experiments studying neutrinos produced in the Sun are sensitive to θ_{12} and Δm^2_{21}. For example, the latest SK results [48] find:
$\Delta m^2_{21} = (6.10^{+0.95}_{-0.81}) \times 10^{-5}$ eV^2 and $\sin^2\theta_{12} = 0.306 \pm 0.013$ ($\theta_{12} \approx 33.6°$).

Antineutrinos ($\bar{\nu}_e$s) produced un reactor experiments have an extraordinary sensitivity to the θ_{13} mixing angle due to their short baselines (in the order of 1 km) and small neutrino energy production (MeV). The Daya Bay reactor experiment recently released its results [49] imposing a tight constraint on θ_{13}. The value reported for Daya Bay is $\sin^2(2\theta_{13}) = 0.0851 \pm 0.0024$ ($\theta_{13} \approx 8.5°$).

As mentioned in subsection 1.7.2, θ_{23}, Δm^2_{32} and δ_{CP} are the least understood paramters. Atmospheric and Long-Baseline Oscillation experiments such as NOvA and T2K are leading the efforts to measure these parameters. NOvA latest results [50] (more details in subsection 4.1.5) provide the measurements of $\Delta m^2_{32} = 2.424^{+0.035}_{-0.040} \times 10^{-3}$ eV^2, while $\sin^2\theta_{23} = 0.55^{+0.02}_{-0.06}$ $(\theta_{23} \approx 47.7^\circ)$ and $\delta_{\rm CP}/\pi = 0.88$. The measurements slightly favor NH and Upper Octant, and have no preference for CP violation.

On the other hand, the latests results of T2K [51], find
$\Delta m^2_{32} = 2.46^{+0.08}_{-0.09} \times 10^{-3}$ eV^2, while $\sin^2\theta_{23} = 0.537^{+0.027}_{-0.036}$ and $\delta_{\rm CP}/\pi = -0.55$. The measurements slightly favor NH and Upper Octant, and have a slight preference for CP violation.

The tension on the value of δ_{CP} between NOvA and T2K has been of great interest for the scientific community. A joint effort was conducted between the two collaborations to perform a joint analysis of the two experiments. The results of this effort yielded a strong constraint of
$\Delta m^2_{32} = 2.477 \pm 0.035 \times 10^{-3}$ eV^2 and a mild preference for the Inverted Ordering (depending of the constraints applied on θ_{13}) and it favors δ_{CP} violation, more in [50].

As we have shown, extraordinary efforts have been made to measure the oscillation parameters using different experiments involving different collaborations around the world. The future generation of neutrino experiments such as DUNE and Hyper-Kamiokande will shed light on the answers to the most important questions about neutrinos. That includes the precise measurement of δ_{CP}, the determination of the θ_{23} octant, and the neutrino mass ordering.

1.10 Anomalies in the Three-Flavor Picture

Short-baseline neutrino oscillation experiments have observed an anomalous excess of ν_e ($\bar{\nu}_e$) events in ν_μ ($\bar{\nu}_\mu$) beams at low energies. This behavior is inconsistent with standard oscillations between the three active neutrino flavors. These signals are often referred to as "anomalies" because they imply the existence of a new Δm^2 value beyond the three-flavor mixing framework.

The Liquid Scintillator Neutrino Detector (LSND) experiment, conducted at the Los Alamos Neutron Science Center from 1993 to 1998, was one of the first to report such an anomaly. In this experiment, an intense 798 MeV proton beam was directed onto a target, producing pions. Most of the neutrinos originated from pion and muon decays at rest. LSND was specifically designed to study $\bar{\nu}_\mu \to \bar{\nu}_e$ transitions. The $\bar{\nu}_e$ events were detected through inverse beta decay:

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

The LSND data covered the energy range from 20 to 60 MeV with a 30-meter baseline and showed a 3.8σ excess of $\bar{\nu}_e$-like events [16]. This result was consistent with $\Delta m^2 \gtrsim 1\,\mathrm{eV}^2$, suggesting a new mass-squared splitting that would require the existence of at least one light sterile neutrino (see section 1.11). The anomalous result can be seen in Figure 1.21.

This anomaly was further investigated by the MiniBooNE experiment at Fermilab, which began in 2002. In MiniBooNE, neutrinos are produced by 8 GeV protons from the Fermilab Booster striking a beryllium target, creating pions, which decay in flight. These pions are focused using magnetic horns whose

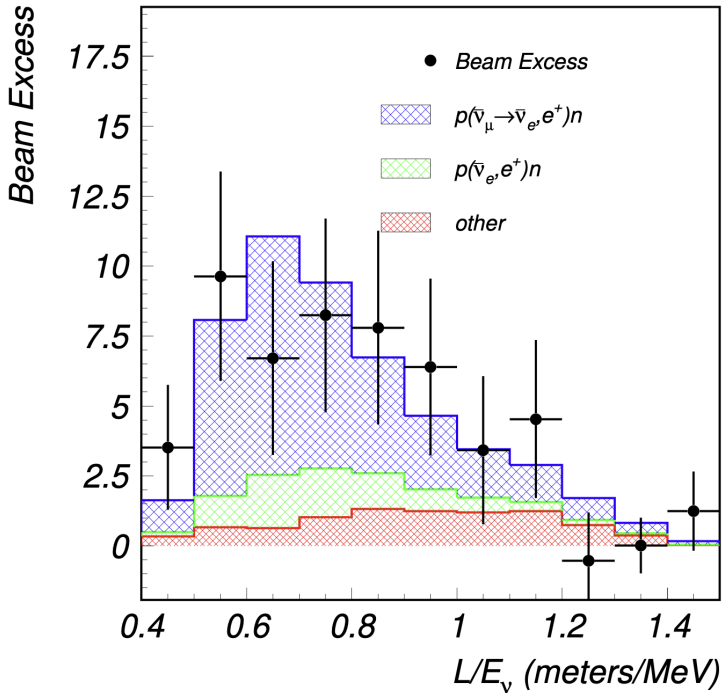

Figure 1.21: Distribution of the LSND anomalous events as a function of L/E_ν shown for the subset of selected events in the energy range of $20 < E_e < 60\,\text{MeV}$. The blue shaded region represents the best-fit two-neutrino oscillation hypothesis with parameters $\sin^2 2\theta = 0.003$ and $\Delta m^2 = 1.2\,\text{eV}^2$ [16].

polarity can be inverted to select charge. The detector is located 541 meters downstream from the target and consists of a tank filled with 818 tons of pure mineral oil, equipped with 1,520 photomultiplier tubes to detect Cherenkov light.

MiniBooNE covers an energy range from 200 MeV to 3 GeV. Initially, no excess over background was observed, resulting in a 98% exclusion of the LSND-like excess at similar L/E. However, an excess of electron-like events was observed below the 475 MeV analysis threshold. Most of this low-energy region lies outside the L/E range explored by LSND, suggesting that the MiniBooNE

low-energy excess may have a different origin.

According to MiniBooNE's 2018 results [17], this low-energy excess is best described by an effective $\Delta m^2 = 0.041\,\mathrm{eV}^2$, which differs significantly from the LSND anomaly.The energy spectrum including the anomalous event excess for MiniBooNE can be seen in Figure 1.22.

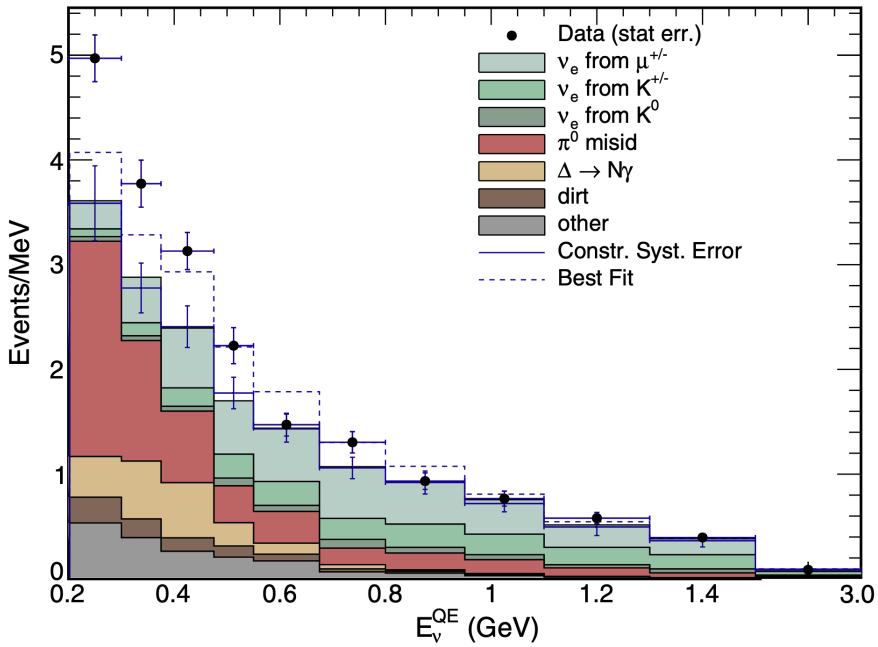

Figure 1.22: Reconstructed E_ν^{QE} distributions from MiniBooNE neutrino-mode data. Shown are ν_e CCQE candidate events (points with statistical uncertainties) and estimated background (histogram with systematic uncertainties). The dashed curve represents the best-fit result assuming a two-neutrino oscillation model [17].

1.11 Light Sterile Neutrinos

Light sterile neutrinos are hypothetical neutral leptons that mix with the ordinary active neutrinos but do not participate in standard weak interactions. Since they are invisible to current detectors, their existence can only be inferred through the effects of their mixing. A sterile neutrino with a mass around 1 eV can modify neutrino oscillation probabilities, introducing an additional mass eigenstate and a new mass-squared splitting. Searching for such effects is important for our understanding of Nature, as sterile neutrinos could open the door to physics beyond the Standard Model.

1.11.1 The 4th Flavor ν_s: 3+1 Sterile Neutrino Oscillations Model

One can extend the active neutrino flavor basis from 3 to 3+1. This model includes a light sterile neutrino, and the unitary matrix U becomes a 4×4 matrix. The U matrix can be parametrized in terms of six mixing angles and three Dirac CP-violating phases.

It is convenient to express U as the product of the following rotation matrices:

$$U = O_{34} V_{24} V_{14} O_{23} V_{13} O_{12} \tag{1.51}$$

Here, O_{ij} denotes a real rotation involving the angle θ_{ij}, and V_{ij} represents a rotation with a complex phase:

$$O_{34} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & c_{34} & s_{34} \\ 0 & 0 & -s_{34} & c_{34} \end{pmatrix}, \quad V_{24} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & c_{24} & 0 & s_{24}e^{-i\delta_{24}} \\ 0 & 0 & 1 & 0 \\ 0 & -s_{24}e^{i\delta_{24}} & 0 & c_{24} \end{pmatrix} \tag{1.52}$$

Under the approximation $\Delta_{21} \ll \Delta_{31}, \Delta_{41}$, the transition probability for $\nu_\mu \to \nu_s$ in vacuum can be written as:

$$\begin{aligned} P(\nu_\mu \to \nu_s) = {} & 4|U_{\mu 4}|^2|U_{s4}|^2 \sin^2 \Delta_{41} + 4|U_{\mu 3}|^2|U_{s3}|^2 \sin^2 \Delta_{31} \\ & + 8\,\mathrm{Re}[U^*_{\mu 4}U_{s4}U_{\mu 3}U^*_{s3}] \cos\Delta_{43} \sin\Delta_{41} \sin\Delta_{31} \\ & + 8\,\mathrm{Im}[U^*_{\mu 4}U_{s4}U_{\mu 3}U^*_{s3}] \sin\Delta_{43} \sin\Delta_{41} \sin\Delta_{31} \end{aligned} \tag{1.53}$$

where the terms are defined as:

$$\begin{aligned} \mathrm{Im}[U^*_{\mu 4}U_{s4}U_{\mu 3}U^*_{s3}] &= \frac{1}{4} c^2_{13}c^2_{24}s^2_{24} \sin(2\theta_{23}) \sin(2\theta_{34}) \sin(\delta_{24}) \\ \mathrm{Re}[U^*_{\mu 4}U_{s4}U_{\mu 3}U^*_{s3}] &= -c_{13}c_{24}c_{34}s_{23}s_{24} \left[c_{34}s_{23}s_{24} + c_{23}s_{34}\cos(\delta_{24})\right] \\ |U_{s4}|^2 &= c^2_{24}c^2_{34} \\ |U_{\mu 4}|^2 &= c^2_{23}s^2_{34}c^2_{13} + c^2_{34}s^2_{23}c^2_{24}c^2_{13} \\ &\quad + \frac{1}{2} s_{24}c^2_{13} \sin(2\theta_{23}) \sin(2\theta_{34}) \sin(\delta_{24}) \end{aligned} \tag{1.54}$$

Here, θ_{14}, θ_{24}, and θ_{34} are the new mixing angles with the fourth state, and

δ_{14}, δ_{24} are two new CP-violating phases.

While the three-flavor neutrino model explains most experimental observations, some results remain unexplained. The excesses reported by short-baseline experiments such as LSND and MiniBooNE suggest the possible existence of a light sterile neutrino, which is not part of the Standard Model.

Whether these anomalies point to new physics or unresolved experimental issues is still unknown. This highlights the need for precise measurements to test the limits of the three-flavor framework and to search for possible deviations. The NOvA experiment was built to study the three-flavor framework in detail and to look for signs of physics beyond it.

Chapter 2

NOvA Neutrino Experiment

NOvA (NuMI Off-Axis ν_e Appearance) is a long-baseline neutrino oscillation experiment designed to study the properties of neutrinos, the most elusive and abundant particle in the universe. In this experiment, neutrinos produced at Fermilab travel 810 km through the earth from Fermilab, home of the 310 kton Underground Near Detector, to the Far Detector (FD) in Ash River, Minnesota.

Over the years, NOvA has made significant progress in measuring neutrino oscillation parameters, searching for evidence of CP violation in the lepton sector, determining the neutrino mass hierarchy, and searching for sterile neutrinos. NOvA also contributes to the study of neutrino cross sections, detection of exotic cosmic phenomena, and other areas of particle physics.

The following sections of this chapter will explain how the NOvA experiment works, beginning with the creation of the Neutrinos at the Main Injector

(NuMI) beam at the Fermilab Accelerator Complex, in addition, they will cover the NOvA detector layout and construction, as well as the key components of the NOvA Data Acquisition (DAQ) system.

2.0.1 Fermilab Accelerator Complex

Neutrino production begins in the Fermilab Accelerator Complex with the Linac (Linear Accelerator) [52], which accelerates H^- ions using the RFQ Injection Line (RIL) to 400 MeV. These ions are then transferred to the Booster, a circular accelerator that removes the electrons, converting the ions into protons and increasing their energy to 8 GeV [52]. The protons are grouped into 1.6 μs long batches with a 53 MHz bunch spacing.

Next, the proton beam is directed to the Recycler Ring, a 2-mile-circumference storage ring optimized to increase beam intensity. Since the Recycler is seven times larger than the Booster, it can store up to six Booster batches. Using the slip-stacking technique, an additional six batches are merged, further enhancing beam intensity [52].

Finally, the accumulated proton batches are extracted and accelerated to 120 GeV in the Main Injector, the proton beam is extracted to different Fermilab experiments and facilities including the Test Beam facility, and high energy neutrino experiments where the NuMI beam facility is located. Figure 2.1 shows the conceptual illustration of the Fermilab Accelerator complex and its main components.

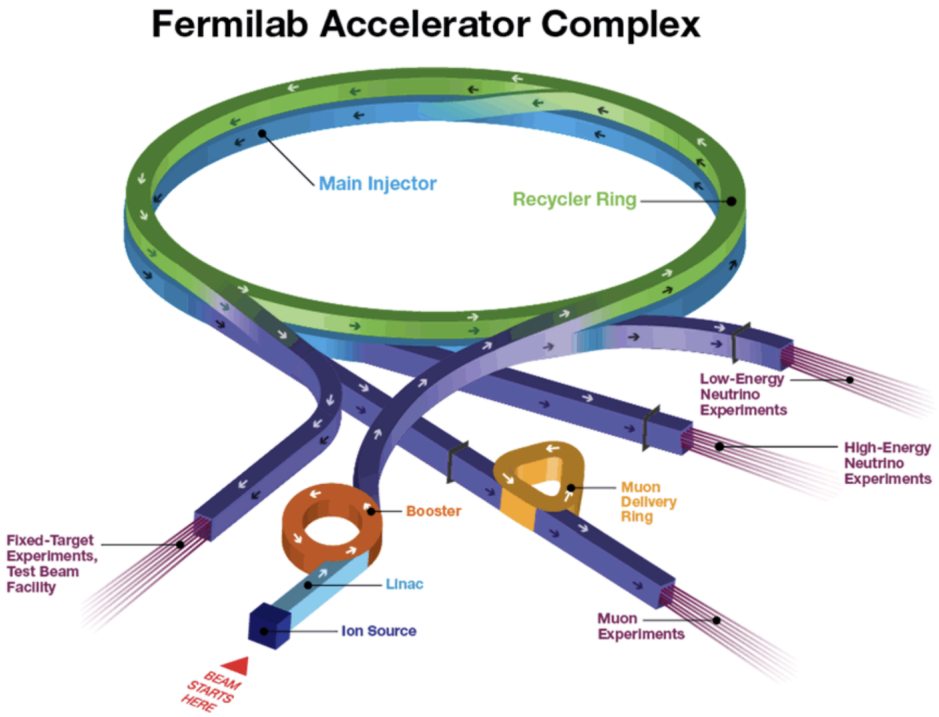

Figure 2.1: conceptual illustration of the Fermilab Accelerator Complex, it is composed of four interconnected particle accelerators and storage rings: the Linac, Booster, Recycler, and Main Injector. The complex provides high-intensity proton beams to several experiments and research programs [18].

2.0.2 NuMI Beamline

The NuMI facility impinges 120 GeV protons from the Main Injector on a graphite target, creating an enormous quantity of hadrons, primarily charged pions that are focused and charge-selected using magnetic horns. The focused hadrons eventually decay into neutrinos and muons in a 615 m decay pipe filled with helium.

The resulting muons are monitored and absorbed by the surrounding rock, while neutrinos continue traveling 1 km downstream and 14.6 mrad off-axis from the beam toward the NOvA Near Detector (Figure 2.2). Ultimately, the oscillated and non-oscilalted neutrinos are detected 809 km further away in the Far Detector in Minnesota.

Over the years, the NuMI facility and the Fermilab Accelerator Complex have undergone improvements that have significantly increased the beam power from 400 kW during initial operations to up to 900 kW. In June of 2024, the highest power beam recoded was 1 MW, making it one of the most powerful neutrino sources in the world.

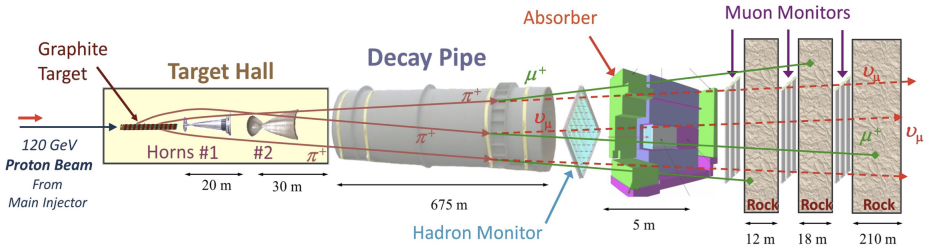

Figure 2.2: Diagram of the NuMI beam showing the target, magnetic horns, decay pipe, and absorbers [19].

2.1 The Off-Axis Concept

The NOvA detectors are strategically positioned 14.6 mrad Off-Axis from the beam. The off-axis placement results in a narrow neutrino energy flux that peaks around 2 GeV, as shown in Figure 2.3a. At this energy, the neutrino oscillation probability is near its maximum at the 810 km distance where the

NOvA Far Detector is located.

The relationship between neutrino energy and the off-axis angle arises because the neutrino energy spectrum depends explicitly on the angle between the pions and kaons created in the NuMI beam and the neutrinos in which they decay, this relationship can be expressed as Equation 2.1.

$$E_\nu = \frac{\left(1 - \frac{m_\mu^2}{m_{\pi,K}^2}\right) E_{\pi,K}}{1 + \gamma^2\theta^2} \tag{2.1}$$

Here, the energy of the neutrino is given by E_ν, θ is the angle between pions or kaons and the neutrino $\gamma = E_{\pi,K}/m_{\pi.K}$, m_π is the mass of the pion and m_K the mass of the Kaon. The neutrino energy as a function of the pion energy for several choices of θ is shown in Figure 2.3b.

The optimization of the energy-angle relationship provides a narrow beam at 1.6 GeV, which is close to the oscillation maximum at 810 km, and removes NC feed-down from higher energies of an on-axis setting. This configuration is paramount for the NOvA physics goals involving the determination and understanding of the neutrino oscillation parameters.

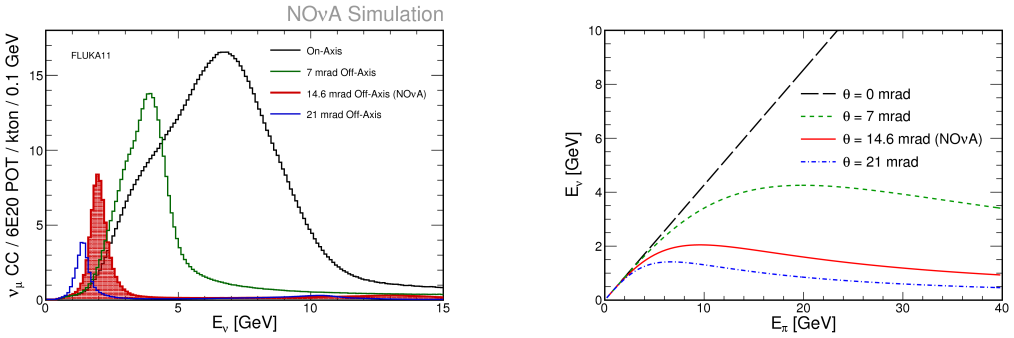

(a) Neutrino Energy Spectrum (b) Neutrino vs. Pion Energy

Figure 2.3: Neutrino flux versus neutrino energy in the FD (Left); Neutrino energy versus pion energy (Right).

2.2 NOvA Detectors

The NOvA detectors are identical tracking calorimeters composed of PVC extrusions (NOvA cells), arranged in alternating vertical and horizontal planes. Each cell contains liquid scintillator and a wavelength-shifting fiber, which connects to an Avalanche Photodiode (APD). This design allows NOvA to reconstruct neutrino interactions in three dimensions within the detector.

The experiment includes three detectors: the ND, FD, and Test Beam Detector as shown in Figure 2.4. The ND is located 105 m underground at Fermilab and consists of 214 planes, this includes a muon catcher at the back of the detector (subsection 2.2.10), the total mass of the detector is 290 tons. the FD is placed on the surface at Ash River, Minnesota, 809 km away from the ND, and comprises 892 planes with a mass of 14 kilotons. Lastly, the Test Beam detector located at the Fermilab Test Beam Facility (FTBF), containing 63 planes with a mass of 30 tons.

The ND and FD form the primary array of detectors dedicated to studying neutrino interactions and oscillations. In contrast, the Test Beam detector was specifically designed to study charged particle interactions from a well-characterized beam source. More details about the Test Beam detector are provided in Chapter 5. The following sections describe the main components and construction of the NOvA detectors.

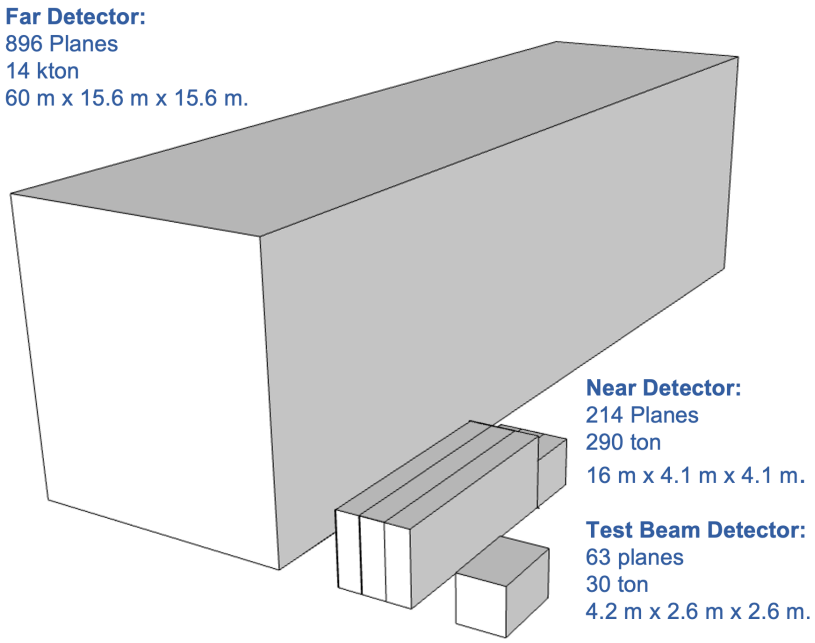

Figure 2.4: Diagram showing the size comparison of the FD, ND and Test Beam Detector. The muon catcher is also observed at the back of the ND.

2.2.1 The NOvA Cell

The NOvA cell is the fundamental building block of the NOvA detectors. Each cell is made of rigid plastic PVC and has a rectangular shape. It is

filled with liquid scintillator and contains a looped wavelength-shifting fiber, as illustrated in Figure 2.5.

The interior width (W) of the cell is 3.8 cm, positioned perpendicular to the beam, while the interior depth (D) is 5.9 cm along the beam direction. The length (L) of the cell varies depending on the detector: in the FD, the cells are the longest at 15.6 m; in the ND, they measure 3.9 m in the fully active region and 2.6 m in the muon catcher region. More details about these two regions are provided in subsection 2.2.10.

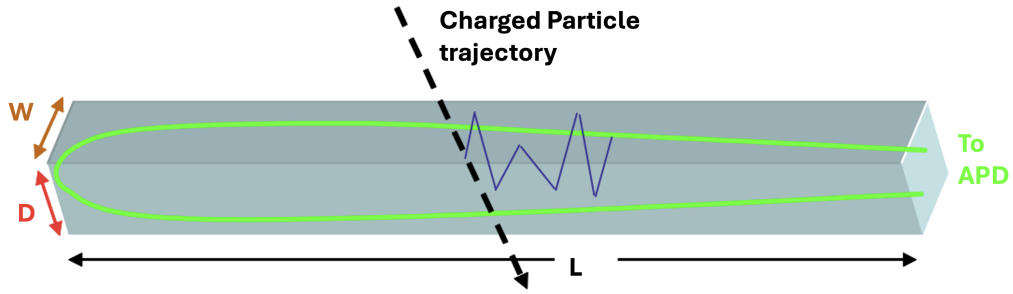

Figure 2.5: Illustration of the NOvA cell. The image shows a charge particle producing light inside of the cell that is capture by a looped wavelength-shifting fiber in green.

When a charged particle traverses the cell, it produces scintillation light in the liquid scintillator. This light reflects off the PVC walls an average of eight times before being captured by the wavelength-shifting fiber.

The light survival inside the cell depends on the reflectivity of the PVC material. To enhance reflectivity, the PVC cells are coated with more titanium dioxide (TiO_2) than commercially available PVC. This coating can achieve a

reflectivity as high as 92 %, which is near the peak of scintillation emission.

2.2.2 Extrusion Module

The cells are grouped into sets of 16 cells to form rigid PVC extrusions as seen in Figure 2.6. Two types of extrusions were designed based on the cell orientation: One with horizontal cells and another with vertical cells.

The extrusion with horizontal cells has an exterior wall thickness of 3 mm and 2 mm thick between cells. In contrast, the extrusion with vertical cells features a 4.4 mm thick exterior walls and 3 mm thick interior webs. Despite the interior differences, both extrusion types have an overall extrusion thickness of 6.6 cm.

The interior difference is necessary due to structural requirements. Horizontal extrusions rely on vertical extrusions for support, while the vertical extrusions bear the load and transfer it to the floor.

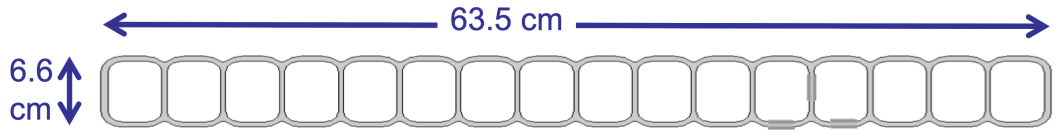

Figure 2.6: Diagram shows the top view of a PVC extrusion. Each PVC extrusion is formed by 16 cells with an exterior thickness of 6.6 cm.

The extrusion modules are created by gluing two PVC extrusions, each containing 16 cells, together side-by-side to form a 32-cell extrusion module assembly. Inside each cell, a looped wavelength-shifting fiber is placed. One end

of the extrusion assembly is sealed with a PVC plate, while the other end features a PVC manifold that routes the two fiber ends to an optical connector. The complete module assembly is shown in Figure 2.7.

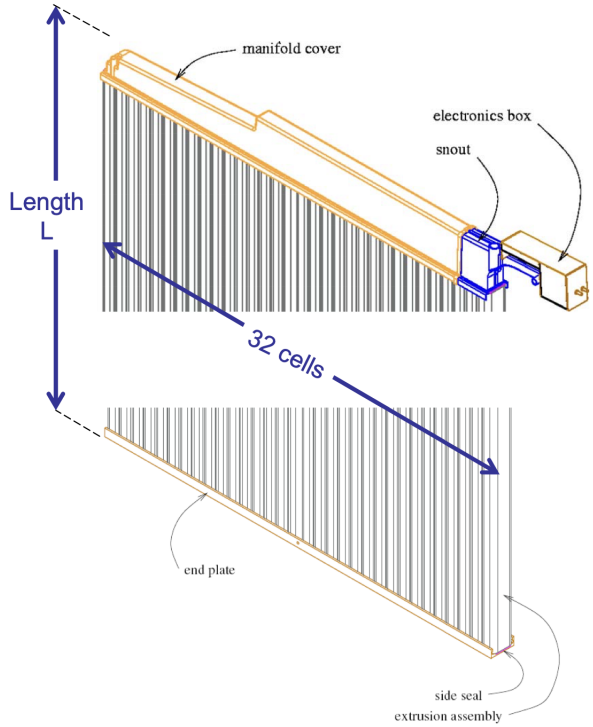

Figure 2.7: The NOvA extrusion module consists of two 16-cell PVC extrusions placed side by side and sealed at both ends to contain the liquid scintillator. One end includes a manifold that organizes the 64 fiber ends and directs them to a coupling interface connected to the electronics [20].

2.2.3 Liquid Scintillator

The liquid scintillator held inside the extrusion modules is composed primarily of mineral oil with 4.1 % pseudocumene (1, 2, 4-Trimethybenzene) as scintil-

lant, light produced by charge particles interacting with these components has a wavelength in the range of 360 - 390 nm (ultra-violet).

In addition, The liquid scintillator contains chemical additives such as PPO (2, 5-diphenyloxazole) bis-MSB (1, 4-dimethylstyryl benzene). The scintillation light produced by the mixture shifts from the UV wavelength to the range of 400 - 450 nm (blue light), which is optimal for capture by the wavelength shifter fiber [20].

2.2.4 Wavelength Shifting Fiber

Each cell extrusion contains a loop of fiber that is 0.7 mm in diameter. The fiber's core is made of polystyrene mixed with 300 parts per million of R27 dye, which acts as the wavelength shifter. Additionally, to enhance internal light reflection, two outer layers with lower refractive indices were added.

The first layer is a thin acrylic coating (PMMA, or polymethylmethacrylate), while the second is a fluor-acrylic layer. Both layers account for approximately 3 % of the fiber's diameter. The fiber captures blue light in the 400–450 nm range from the scintillator and shifts its wavelength to green light in the 490–550 nm range [20].

The absorption and emission spectra of the dye in the fiber can be seen in Figure 2.8a. Due to the overlap of the two spectra, light emitted below 500 nm is severely attenuated. In particular, wavelengths below 490 nm are completely absorbed after passing through 0.5 m of fiber. Figure 2.8b shows the resulting emission spectra for different attenuation distances. The emitted green light

is then captured by the Avalanche Photodiodes.

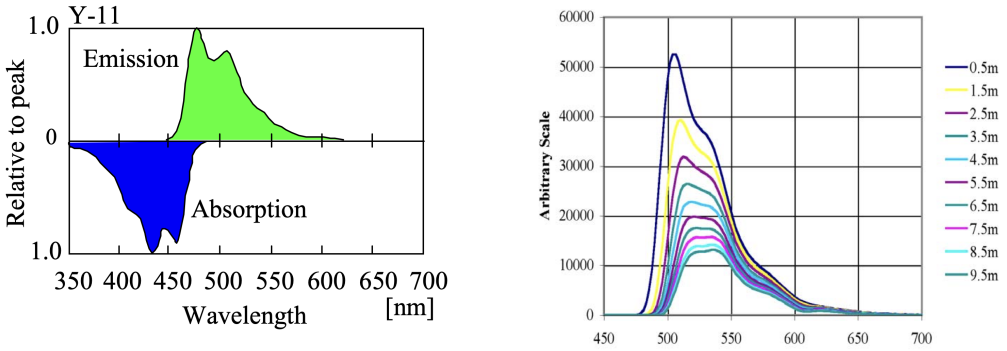

(a) (b)

Figure 2.8: Absorption and emission spectra (a). Emission spectra for different attenuation distances (b) [20].

2.2.5 Avalanche Photodiodes (APDs)

The NOvA photodetector system employs Avalanche Photodiodes (APDs) manufactured by Hamamatsu. These APDs achieve a high quantum efficiency of 85% for detecting light in the 520–550 nm wavelength range. The APDs operate with a gain of 100 under an applied voltage of approximately 375 volts. To minimize thermal noise, they are cooled to -15 °C using thermo-electric (TE) coolers operating via the Peltier effect. The heat generated by the TE coolers is managed with a water cooling system.

To prevent condensation at the cooled APD–fiber interface, a dry gas system is implemented. The dry gas circulates around the APD electronics, maintaining the local environment below the dew point and preventing condensation that could damage the electronics or degrade the signal.

When light produced by the fibers enters the APD's collection region, it is absorbed, causing the generation of electron-hole pairs. These pairs are separated by the applied electric field where the electrons are moved toward the p-n junction. At this junction, the electric field is strong, leading to a process called avalanche multiplication.

The APDs are organized into arrays of 32 pixels, each array mounted onto a carrier board substrate. This design enables two fiber ends to fit onto a single pixel. Each of the 32 pixels in the array directly corresponds to one of the 32 cells in a single PVC extrusion module [20].

2.2.6 Front End Boards (FEBs)

The Front End Boards (FEBs) process signals from the 32 APD channels. Each FEB is equipped with an Application Specific Integrated Circuit (ASIC), an Analog-to-Digital Converter (ADC), and a Field Programmable Gate Array (FPGA). These components work together to initiate the data acquisition process, ensuring that signals from the APDs are captured and processed.

The process begins with signal shaping, performed by the ASIC. The shaping defines the rise and fall times of the analog signals, preparing them for digitization. The ADC then converts the analog signals into digital signals. The digitized signals are passed to the FPGA, which extracts information such as the time and amplitude of the signals. The system is programmed to identify and forward only those signals, or "hits," with an amplitude above a predefined threshold.

In addition, the FEB is responsible for integrating the APD information into data packets. These packets contain critical details such as the amplitude, timing, and channel information for each valid signal. The FEB then transmits these data packets to the Data Acquisition (DAQ) system [20].

2.2.7 Data Acquisition (DAQ) system

The Data Acquisition (DAQ) system is responsible for collecting, processing, and storing data from the APD channels. The main goal of the system is to consolidate data into a single stream for analysis and archiving. To achieve this, the DAQ system buffers data temporarily, allowing it to determine whether the data should be recorded or discarded. Additionally, online trigger processors are employed to analyze time-stamped data and identify significant clusters of hits that indicate good events. Other functionalities include flow control, monitoring, system operations, and alarms.

The FEBs operate in an untriggered mode, continuously digitizing, time-stamping, and processing data. This data is saved for up to 20 seconds while awaiting a spill trigger message. Once a spill occurs, data from a 30-microsecond window centered around the 11-microsecond spill is recorded for further processing, saving approximately 190 GB per year. For calibration and monitoring purposes, additional off-spill data is collected at a much higher rate, storing 18 TB per year. This higher rate is created by cosmic ray muons, which occur at a frequency of approximately 180 kHz at the FD.

The DAQ system is composed of several components, including a timing and command distribution system, an array of Data Concentrator Modules

(DCMs), a Gigabit Ethernet network, and a buffer farm for data storage. Timing synchronization and command packets are handled by a Master Timer Unit and Timing Distribution Units (TDUs). Digitized data from up to 64 front-end boards is routed via CAT5 cables to the DCMs. These modules then send the data to the buffer farm nodes, where it is stored and managed using the Gigabit Ethernet network. The entire system, including the buffer farm and DCMs, is controlled by a central Run Control Computer. A Diagram of the DAQ system for the FD can be seen in Figure 2.9.

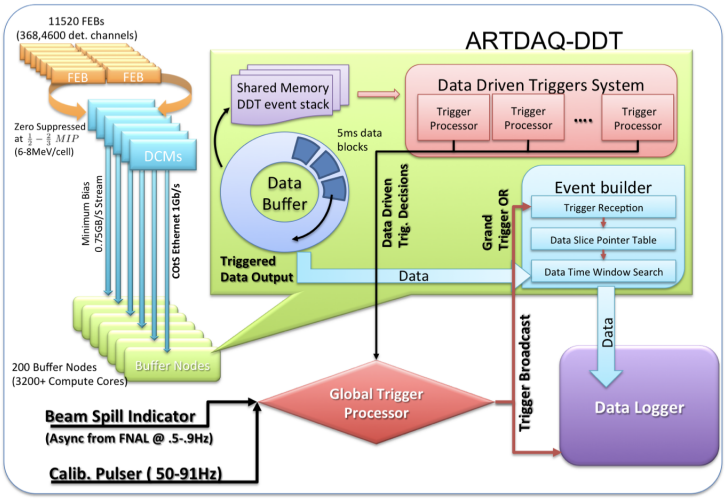

Figure 2.9: Diagram of the FD DAQ components with inclusion of a data driven triggering system [21].

2.2.8 Triggering

In NOvA, data is saved only when a trigger is received, with three types of triggers ensuring efficient data collection. A single trigger is sufficient to record data, each serving a specific purpose:

- **Beam Spill Trigger:** This trigger records interactions from beam neutrinos. A timestamp of the beam spill is sent from the accelerator and compared to the data timestamps stored in the buffer. Although a beam spill lasts only 10 microseconds, data within a 500-microsecond window around the spill is saved. This buffer window accounts for timing drift and enables the analysis of out-of-spill backgrounds. For these triggers, the neutrino time-of-flight ($\sim$ 3 milliseconds at the far detector) is also considered to ensure precise data alignment.

- **Cosmic Ray Trigger:** Operating at 10 Hz frequency, this trigger saves 500 microseconds of data for each event. Its main purpose is to capture cosmic ray interactions, which are crucial for calibration of the detector.

- **Data-Driven Triggers (DDTs):** The DDT identifies and saves events with specific properties. These triggers focus on capturing data for various rare or exotic phenomena, such as magnetic monopoles, dark matter, or supernova neutrinos.

2.2.9 Assembly of NOvA Detectors

The NOvA detectors share an identical structure, constructed from alternating layers of vertical and horizontal PVC extrusions as described in subsection 2.2.2. These layers are organized into planes enabling efficient particle tracking from neutrino interactions within the detector and 3D particle reconstruction.

The extrusion module is enclosed with an end plate and additional center and

side plates to prevent leakage. A manifold cover on top guides the fibers to the snout, which directs them to the APD pixel array. The snout also includes ports for filling and venting liquid scintillator. An electronics box attached to the snout contains the APD and FEB components.

2.2.10 Near Detector

The NOvA Near Detector (ND) is located at Fermilab, positioned 105 m underground and 1.015 km downstream from the NuMI target. It consists of two primary sections: an active region and a muon catcher. The active region is composed by 192 planes, each formed by 3 extrusion modules (or 96 cells). Each plane measures 3.9 m in height and width, and the entire full active region spans a length of 12.9 m.

The muon catcher is located at the back of the detector, its main goal is to contain muons produce during ν_μ interactions. It contains 22 planes with a 101.6-mm-thick iron plate inserted between each PVC plane. Unlike the full active region, each PVC plane in the muon catcher is formed by two extrusion modules, making them 2.6 m wide. This difference in size is due to the constraints on the available iron size plates during construction. Pictures of full active detector and the muon catcher are shown in Figure 2.10. The Near Detector (ND) of the NOvA experiment uses approximately 631 Front-End Boards (FEBs). Each FEB reads out 32 channels, giving a total of 20,192 readout channels for the detector.

Figure 2.10: Top view images of the ND. The left picture shows the front of the ND showing the full active region, the right picture shows the back of the detector showing the muon catcher.

2.2.11 Far Detector

The Far Detector (FD) is located at Ash River, Minnesota, 810 km away from the NuMI beam. It is placed on the Earth's surface with a rock overburden of 3.6 meters water equivalent (m.w.e.). Since the detector is on the surface, it is exposed to a cosmic background at a rate of 150 MHz, as discussed in subsection 2.2.8. A fraction of this background is saved and used for calibrating the detector and estimating the background in neutrino oscillation analyses.

The FD consists of 896 planes, with each plane containing 12 extrusion modules. This results in a total of 384 cells per plane, where each cell is 15.5 m long. The width and height of the detector are 15.5 m, and its total length is

60 m. Pictures of the FD in the Ash River hall can be seen in Figure 2.11. The Far Detector uses around 10,752 FEBs, with 32 channels each, resulting in 344,064 channels.

Figure 2.11: Far Detector photographs. The left picture shows the front of the FD, the right picture shows a section of the detector depicting the NOvA planes (in black), the FEBs (in gold) and the DCMs (silver).

Chapter 3

NOvA Software

Understanding and analyzing the data of our detectors play a crucial role in achieving NOvA's physics goals, requiring substantial efforts to develop and improve software for the simulation, reconstruction and identification of neutrino events.

This chapter begins with an overview of the Monte Carlo simulations used in NOvA, explaining how they model neutrino interactions and detector behavior. It then briefly covers the calibration process, which ensures the energy measurements remain reliable. Finally, it will touch on the reconstruction methods applied to both real and simulated data, which are essential for identifying and interpreting neutrino events in the detectors.

3.1 NOvA Simulation

Simulating the flux of neutrinos produced in the NuMI beamline, along with the types of neutrino interactions, particle transport within the detector, detector electronics response, and accurate modeling of backgrounds are all crucial for the NOvA oscillation analyses.

These simulations ensure reliable results when compared with experimental data, which is essential since most analyses are blind, meaning the data is only accessed after the analysis is mature. The following sections provide detailed information about the simulation chain employed in the NOvA neutrino oscillation analyses.

3.1.1 Neutrino Flux

The NOvA neutrino flux is simulated using G4NuMI [53] which is based on Geant4 [54]. The simulation begins with 120 GeV protons striking a graphite target producing a shower of particles composed mainly of kaons and pions, with the hadrons eventually decaying into neutrinos. The simulated geometry includes essential components, such as the cooling material, support structures, and magnetic horns, providing a full simulation of the NuMI beamline.

The particle interactions are modeled by a standard physics lists FTFP_BERT. The Fritiof (FTFP) model [55] describes high energy hadronic interaction above energies of 4 GeV, on the other hand, the Bertini cascade model (BERT) [56] simulates hadronic interaction in the lower energy region (<5GeV).

The complexity of simulating hadronic interactions leads to imperfections in the resulting simulated flux. NOvA applies corrections to the simulated flux using the Package to Predict Flux (PPFX) framework [57]. PPFX adjusts for mismodeling by including data from different experiments assigning correction weights to hadronic interaction in the 12-120 GeV range. To account for uncertainties, a multi-universe approach is implemented, where parameters are varied and simulations are rerun multiple times. The resulting beam flux energy spectra for neutrinos and antineutrinos are shown in Figure 3.1.

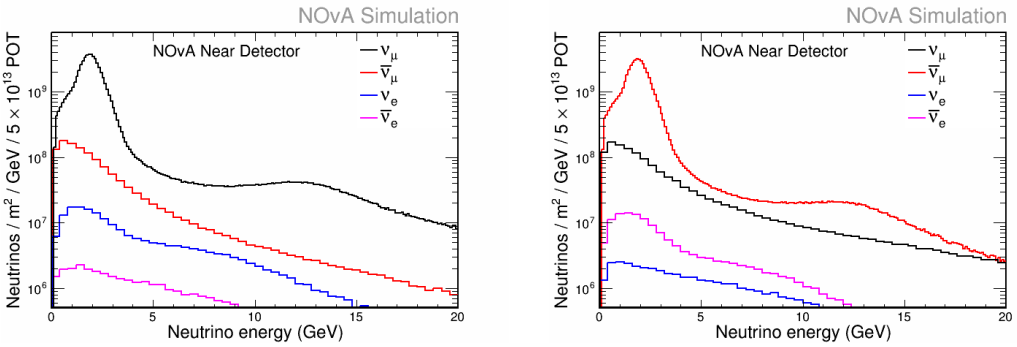

Figure 3.1: The PPFX-corrected NuMI beam flux at the Near Detector, illustrating the various components in (left) neutrino beam mode and (right) antineutrino beam mode.

3.1.2 GENIE Event Generator

Once the flux is determined, GENIE [58] simulates neutrino interactions on nuclei in the detector and its surroundings using cross-section data and different interaction types including Charge-Current Quasi-Elastic (CCQE) and Two-Particle-Two-Hole (2p2h) interactions [59] [60]. The Resonance and Coherent Pion Production is simulated using the Berger-Sehgal model [61]. Deep

Inelastic Scattering (DIS) interactions are simulated based on the Bodek-Yang model [62]. To account for how particles interact inside the nucleus before exiting, a Final State Interactions (FSI) model is also implemented using the hN semi-classical intranuclear cascade model [63].

Even after the implementation of these interactions and models, discrepancies exist between data and simulation, thus, NOvA tunes the FSI model and the 2p2h interaction contributions to improve the disagreement. The FSI model is calibrated by tuning the model parameters using different π^+ on ^{12}C scattering experiments. For 2p2h, NOvA produces weights to modify the shape and normalization of the ν_μ-CC interactions in simulation. This is done by applying weights in energy and three-momentum transfer space $(|\vec{q}|, q_0)$ (q is the four-momentum transfer) using a model based on two 2D Gaussian distributions. These tunings improve ND data and simulation agreement.

3.1.3 Detector Simulation

After GENIE simulates the final-state particles from neutrino interactions, the particles are passed to Geant4 to track them as they traverse the detector. To achieve this, Geant4 uses detailed information about the detector's geometry, materials, and atomic composition.

Geant4 simulates several processes, including the trajectories of particles, their energy deposition in detector materials, the production of scintillation light, and scattering interactions.

GEANT4 is capable of simulating scintillation light in the detectors, but op-

tical photon tracking would be computationally expensive to use in full simulations. Since the detector cells are identical, a ray-tracing simulation is run once using measured properties like scintillator emission, PVC reflectivity, and fiber attenuation.

This produces templates for photon collection and timing based on hit position. These are reused in the main simulation to estimate the light collected by the APDs. The mean number of photons collected in this step is calculated for each energy deposit and adjusted for transport losses and the APD quantum efficiency.

The electronics simulation takes these photoelectron distributions and converts them into digital signals, replicating the response of the detector's FEBs. Pulse shaping is modeled using a CR-RC circuit, and noise is added using Gaussian Markov chains tuned to pedestal data. The shaped signal is digitized, a baseline is applied, and zero suppression is performed to identify signals above threshold, simulating the actual DAQ.

For cells without energy deposits, real unclustered hits from data are overlaid instead of simulating noise, which reduces computation time. The output is formatted like raw detector data, enabling direct comparison with real events [64].

In addition, Cherenkov light is added by computing the number of UV photons produced at each GEANT4 particle step based on particle velocity and a fitted refractive index model. These photons are not tracked directly but folded into the light yield using a fixed efficiency factor and combined with scintillation light during photon transport.

The NOvA detector simulation uses different physics models, known as "physics lists". Specifically, NOvA uses the QGSP_BERT_HP_EMZ physics list, which combines four models. The QGSP model [65][66] simulates high-energy hadronic showers, such as those produced in interactions above several GeV. The BERT model (Bertini Cascade) [56] is used for intranuclear cascades in particles with low energies. The HP (High Precision) model [67] provides highly accurate cross-sections for neutron interactions, particularly for energies below 20 MeV. Lastly, the EMZ model [68] provides the best set of electromagnetic physics models selected from the low energy and standard packages.

3.1.4 Background Simulation

The ND background is primarily caused by neutrinos for which the interaction vertex is in the surrounding rock of the underground cavern, referred to as "rock events". The simulation of these events is time-consuming, therefore rock events are used multiple times in simulation to reproduce the expected background rate. At the FD, the measured cosmic data mentioned in subsection 2.2.11 is overlayed onto simulated neutrino events to produce a realistic cosmic background.

3.2 Detector Calibration

The detector calibration involves transforming the ADC output generated by hits in the detector into the physical quantity of energy deposits in GeV. This transformation is accomplished through a two-step procedure: relative cali-

bration and absolute calibration. The relative calibration ensures that ADC readings are consistent on all regions within the detector, whereas the absolute calibration converts these corrected readings into physical energy units, measured in GeV.

3.2.1 Relative Calibration

Charged particles produce light when they traverse the detector. This scintillation light is captured by the fibers and collected by the APDs producing ADC counts. Using a scaling factor, the ADC counts are converted into Photo Electrons (PE). However, as light travels through the fiber, it suffer losses due to attenuation. This means that the number of PE decreases as the distance from the readout increases.

The relative calibration process aims to account for these attenuation effects, so that the energy deposits along the cells reflect the true energy deposited by the particles. To achieve this, trough-going cosmic muons are used, avoiding the sharp energy deposits at the end of the muon track, ideal for this step in the calibration.

The calibration process is done by estimating the path length of the muons along the cell by using the tri-cell hit method that uses the adjacent hits in the two neighboring cells, as shown in Figure 3.2, to determine the particle path length inside the cell using Equation 3.1:

$$d = \frac{h}{\cos\theta}. \tag{3.1}$$

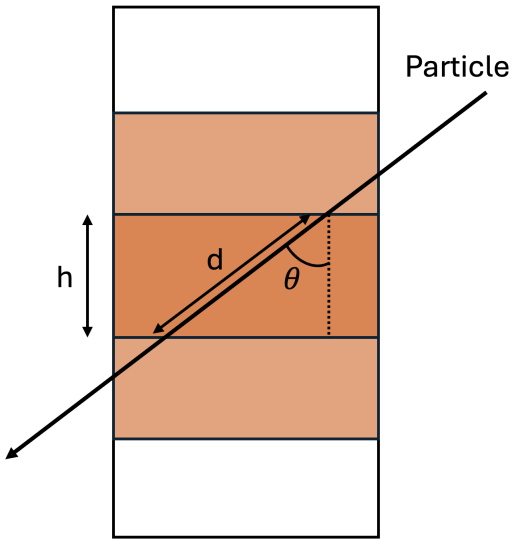

Figure 3.2: Diagram illustrating the determination of the particle path length d through a scintillator cell using the angle θ between the particle trajectory and the vertical. The vertical cell height is h, and the path length is calculated as $d = \frac{h}{\cos\theta}$. This geometry is used in the tri-cell hit method for relative calibration using through-going cosmic muons.

Using the path length from the tri-cell method, the PE and PE/cm quantities are calculated. The next step requires finding W, which is the distance along the cell. Here, W=0 represents the center, while higher values indicate closeness to the read-out electronics. The value of W is obtained by analyzing hits in adjacent detector planes. For example, to estimate the position of a hit in a vertical cell, the neighboring horizontal planes on either side are used.

Once the average signal per unit of length (PE/cm) and W are determined, thresholding and shadowing corrections are applied to PE/cm in function of W. Thresholding corrections account for light signals that fall below detection threshold. This is specially important for interactions happening at the far end of the fiber. Meanwhile, shadowing corrections factor in the distance traveled

by the muons inside the detector, since muons traversing long distances are subjected to a larger energy loss variation leading to an energy bias. The correction factor (T) for these two effects is obtained using Equation 3.2.

$$T = \frac{PE}{\lambda} \frac{E_{\text{true}}}{E_{\text{MIP}}}, \tag{3.2}$$

where PE is the number of simulated photoelectrons recorded at the readout, λ is the number of simulated photons that would be seen at the readout without statistical fluctuations, E_{true} is the true energy deposit in the cell, and E_{MIP} is the expected energy for a minimum ionizing particle based on the path length. After applying this correction, an attenuation curve is obtained for PE/cm vs W. The attenuation fit is done by Equation 3.3.

$$y = C + A\left(e^{\frac{W}{X}} + e^{\frac{-L+W}{X}}\right). \tag{3.3}$$

Here, y represents PE/cm, L is the cell length, and C, A, and X are fit parameters. The second exponential term accounts for the light traveling down the cell and looping back up to the readout. The fit is performed for values of W that are 50cm away from the end of the cell.

The remaining ranges are known as the "roll-off" regions, these fits are conducted using the LOcally WEighted Scatter plot Smoothing (LOWESS) method [ref]. An example of the fit curves for attenuation and the full fit including LOWESS in a ND cell can be observed in Figure 3.3.

The resulting fit is applied as a correction factor resulting a new unit called

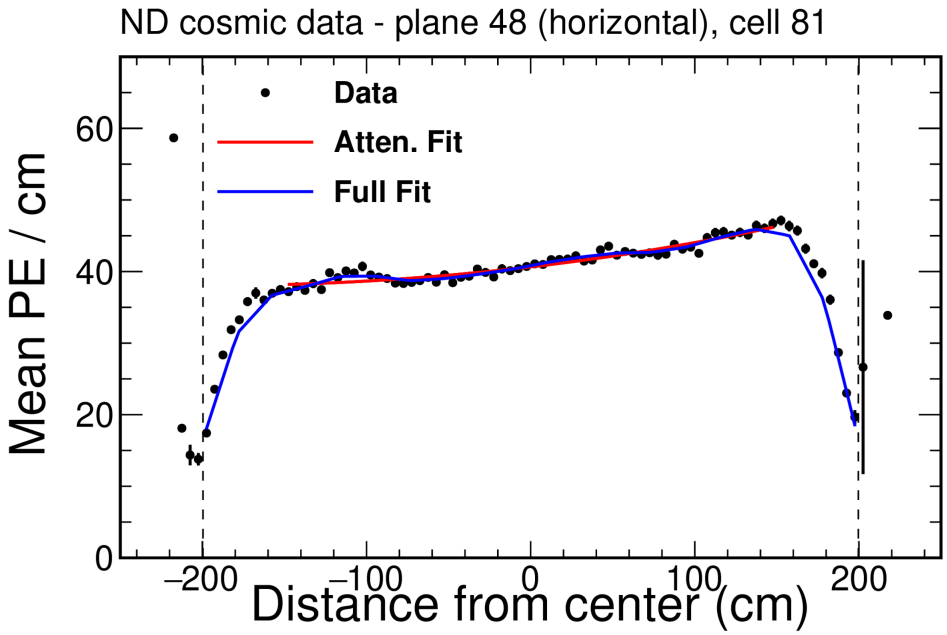

Figure 3.3: Attenuation curve showing the mean PE/cm as a function of distance from the center of the cell in ND cosmic data (plane 48, horizontal, cell 81). The red curve represents the attenuation fit described by Equation 3.3, capturing light loss along the cell. The blue curve corresponds to the full fit including attenuation and roll-off corrections, with the latter fitted using the LOWESS method. This accounts for light reflections and statistical variations near the cell ends.

PEcorr. To validate the calibration procedure, the ratio of true energy to the reconstructed energy is analyzed. An example of this can be seen in Figure 3.4 comparing the ratios before and after calibration.

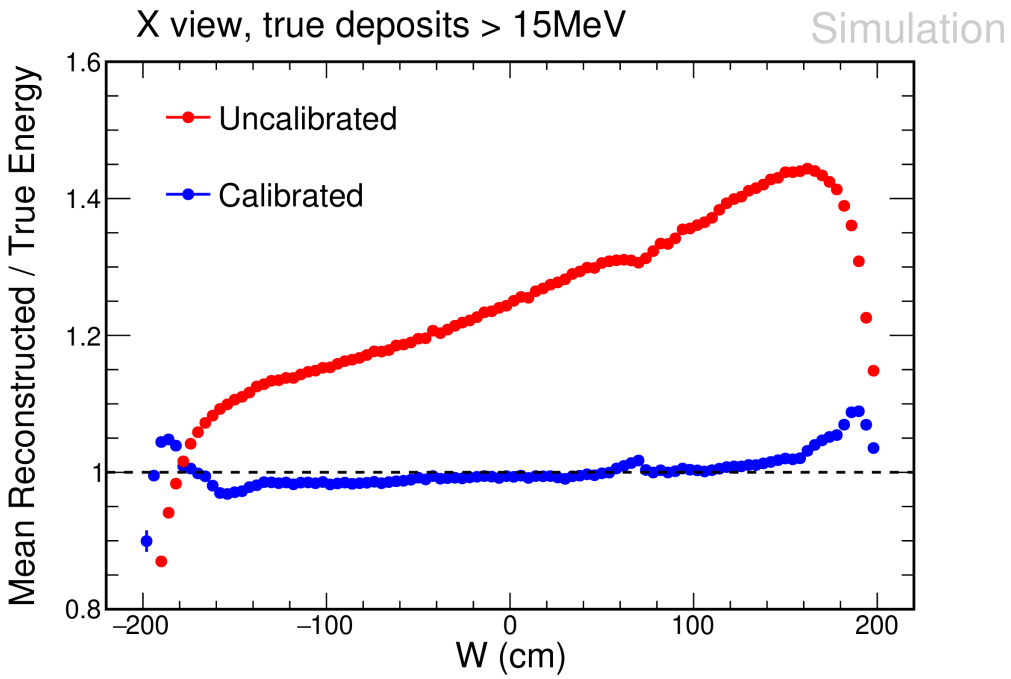

Figure 3.4: Validation of the calibration procedure by comparing the mean ratio of reconstructed to true energy as a function of W (distance from the readout electronics) for energy deposits greater than 15 MeV in simulation. The uncalibrated data (red) show a strong W-dependence, while the calibrated data (blue), using the correction factor from the attenuation fit (PEcorr), yield a uniform response closer to unity across the cell.

3.2.2 Absolute Calibration

The absolute calibration consists of transforming PECorr, obtained from the relative calibration, into GeV units. To achieve this, the "tri-cell" method is also used; however, instead of through-going muons, stopping muon samples

are used. This is motivated by how muons lose energy in the detector, as stopping muons have a well-defined minimum ionizing particle (MIP) region, where the energy deposition is assumed to be linear, an important characteristic of the Bethe-Block curve. For stopping muons traversing the NOvA detectors, the MIP region is located between 100 to 200 cm from the end of the muon track. Hits from muons in this region are selected in both data and simulation. To complete the conversion, Muon Energy Units (MEU) are defined. In data, MEU is set as the mean value of the PECorr/cm, while in simulation is defined as the mean of MeV/cm also referred to as the true MEU. The calorimetric energy scale is then obtained by taking the ratio of these two values, as shown in Equation 3.4. The full calibration procedure is performed for each detector, data collection period, and for both data and simulation.

$$\textit{Calorimetric Energy Scale} = \frac{\mathrm{MEU}_{\mathrm{Truth}}}{\mathrm{MEU}_{\mathrm{Reco}}}. \tag{3.4}$$

3.3 NOvA Reconstruction

Reconstruction is a fundamental process in which raw detector signals are transformed into meaningful physics information. When particles interact in the detector due to neutrino interactions, they leave behind traces of their passage in the form of electronic signals ("hits"). However, these hits alone are not meaningful; they must be carefully processed and interpreted to reconstruct the particle trajectories, measure their energy, and identify the type of particle traversing the detector.

Neutrino signatures in the NOvA detectors vary depending on their interaction type (see section 1.6), as shown in the event displays in Figure 3.5. The ν_μ CC interaction is relatively easy to characterize due to its unique topology, where a long muon track is accompanied by a short proton track.

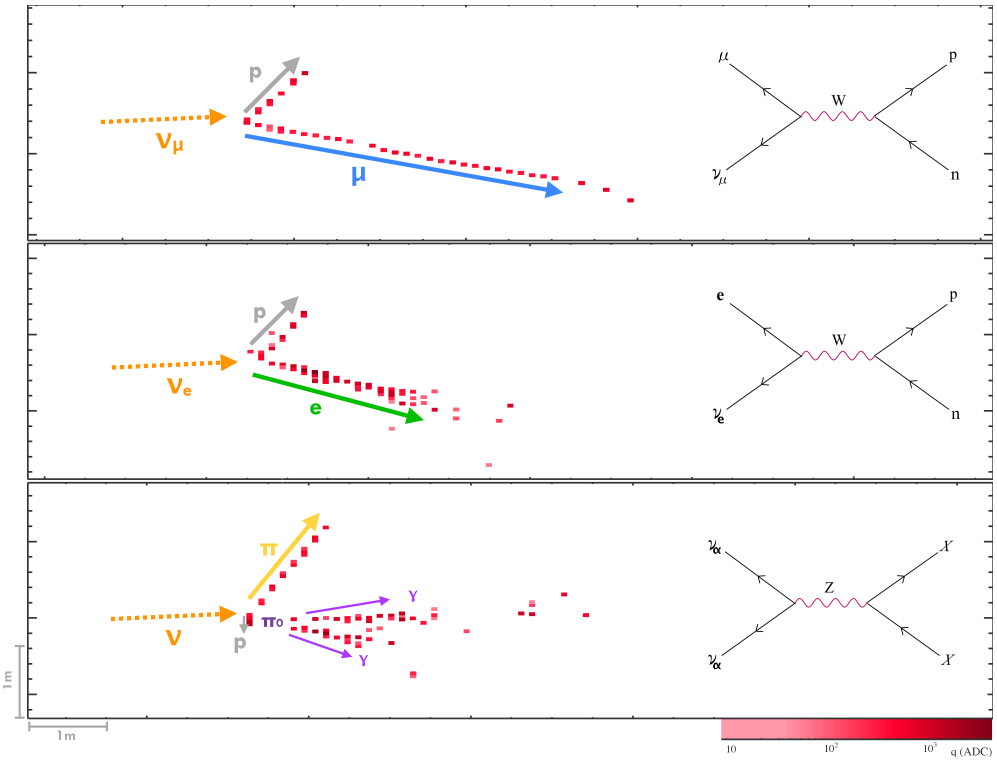

Figure 3.5: Example of neutrino topologies observed in the NOvA detector. Top: ν_μ CC interaction, middle: ν_e CC interaction, Bottom: NC interaction.

The ν_e CCQE interaction produces a proton and an electron. The electron experiences bremsstrahlung as it travels, emitting photons that quickly decay into e^+e^- pairs, which then emit more photons, creating an electromagnetic shower in the detector. This shower-like structure is a defining characteristic of ν_e CC interactions and plays a crucial role in their identification.

Finally, the NC interaction shown produces a pion and a π_0 in the final state. The π_0 decays into two photons, leaving an electromagnetic shower similar to that of a ν_e CC interaction. However, due to the low-Z composition of the detector, resulting in a 36 cm radiation length, the electromagnetic shower produced by the π_0 decay develops further away from the neutrino interaction vertex. This distinction is important for correctly identifying NC interactions, as their topology can be similar to that of ν_e CC interactions. Distinguishing these types of interactions is crucial and can be assessed by implementing efficient and accurate reconstruction tools.

The NOvA collaboration has developed several algorithms to reconstruct neutrino events in the detector. These algorithms are improved and maintained over time, and new ones are created to either replace the existing ones with something more efficient and reliable, or to focus on a specific analysis goal. The following sections will convey the steps and tools used by NOvA to reconstruct neutrino events.

3.3.1 Clustering

In NOvA, the reconstruction process begins by clustering hits that are correlated in space and time, this a technique known as slicing. The primary goal of slicing is to separate events of interest, such as neutrino interactions, from background noise. The method employed to conduct the slicing is called Time Density Slicer or "TDSlicer" [69][70]. This technique consists of three main steps:

1. **Identifying Centroids:** The algorithm starts by analyzing the density of hits in space and time. Each hit's density is determined by counting nearby hits within a specific distance and time window.

    - Each hit is assigned an isolation score based on how far it is from a denser hit.
    - Hits that are both dense and isolated are identified as *centroids*—key points that will serve as the seeds for clustering.

2. **Forming 3D Clusters:** Once centroids are identified, nearby hits are grouped together to form clusters. This is done separately in two different views:

    - One set of clusters is built in the xzt view (spatial coordinates x and z, plus time t).
    - Another set is built in the yzt view (spatial coordinates y and z, plus time t).

    Hits are added to a cluster if they are within a certain distance of a hit already in the cluster. This step ensures that related hits are grouped properly before merging.

3. **Merging into 4D Clusters:** The final step is to combine the two 3D cluster sets (xzt and yzt) into 4D clusters in $xyzt$.

    - Clusters from both views are matched and merged based on their average z-position and time values.
    - This step reconstructs the full event interaction in both space and time

To illustrate the performance of the slicer, Figure 3.6 shows a complete 500 μs window around the beam trigger for a FD event display before any slicing, where a large number of cosmic particles traverse the detector in different directions and at different times. Figure 3.7 presents the same FD event after slicing is applied, with colors representing different slices produced by the TD-Slicer algorithm. Finally, Figure 3.8 displays a slice that contains a neutrino candidate.

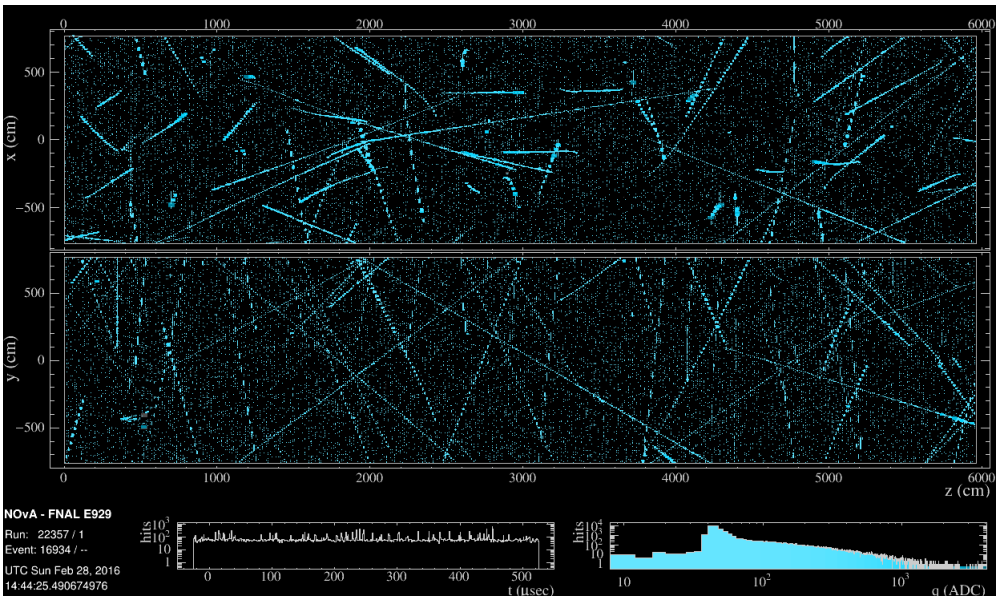

Figure 3.6: This figure shows how cosmic rays are distributed over the 550 μs time window in the Far Detector, along with the charge distribution of the raw hits.

3.3.2 Vertexing

Neutrino interactions produce leptons and hadrons that propagate in the detector, the geometrical point where these particles originate is called the neutrino

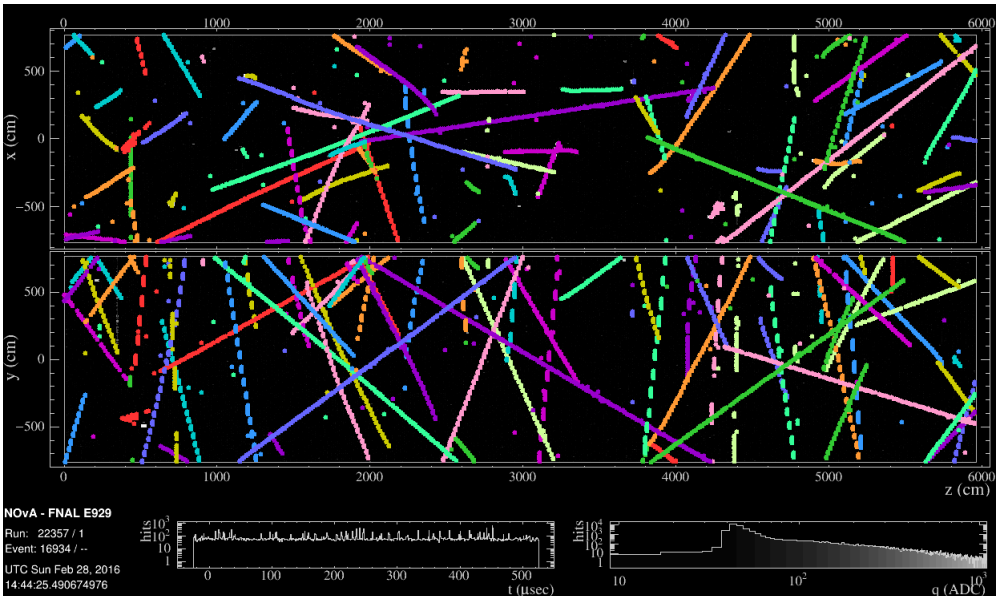

Figure 3.7: This example illustrates the distribution of cosmic rays over the 550 μs time window in the Far Detector. The reconstructed slices (Slicer4D) are shown.

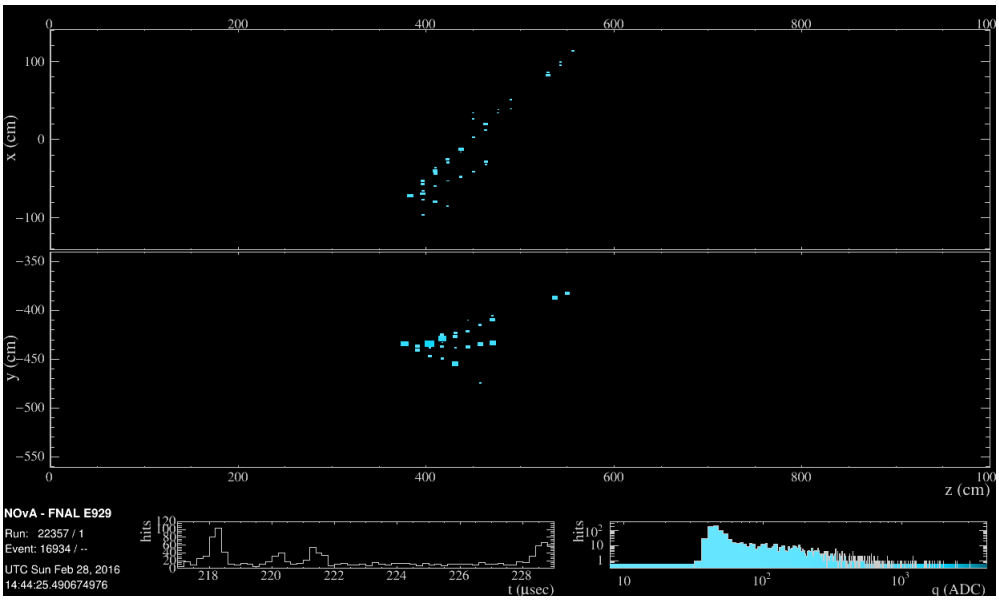

Figure 3.8: An example of a filtered slice in the Far Detector, showing a neutrino candidate.

vertex. Determining the vertex point is a crucial step in reconstruction since it allows us to accurately estimate the particle trajectories and the energy deposits.

The vertex reconstruction begins with a Hough transform with a modified voting scheme [71]. The Hough transform is performed in the x-z and y-z views separately and produces lines for each pair of hits in a slice, the transformed space uses the polar coordinates ρ and θ, for each point in the transformed space a "vote" value is assigned. The highest vote value determines the Hough line that most accurately represents the particle trajectories in the slice.

Once the Hough lines are found, they are used to seed an elastic arms algorithm [72]. This algorithm is in charge of searching for the most optimal location of the vertex by minimizing an "energy function" to find the best set of lines (arms) that better describe the hits in the slice [73]. An event display showing the Hough lines and the vertex location can be seen in Figure 3.9.

3.3.3 Prong Formation with Fuzzy k-Means

Once the global event vertex is identified using the Elastic Arms method, the next step is to assign each detected hit to a prong. In NOvA, a prong represents a particle track or shower. This is achieved using a probabilistic fuzzy-k means algorithm [74], which allows hits to have partial membership in multiple prongs and does not require a predefined number of prongs. This flexibility enables the algorithm to classify hits while treating isolated hits as noise.

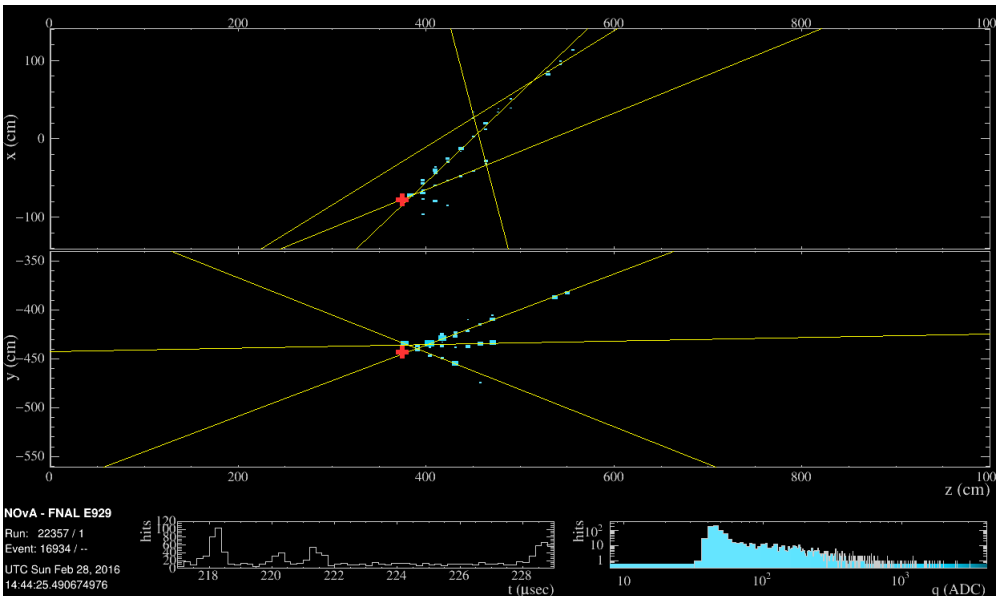

Figure 3.9: An example of a filtered slice in the Far Detector, showing a neutrino candidate. The golden lines represent the reconstructed Hough lines, while the red cross marks the reconstructed Elastic Arms vertex.

The algorithm operates by transforming hits into a one-dimensional angular space around the vertex. It identifies dense regions (prong centers) and assigns membership values based on a hit's angular distance and uncertainty. Initially, the algorithm assumes a single prong and iteratively refines its center. If additional dense regions are detected, they are added as new prongs, and the process repeats until all hits are assigned or a maximum number of prongs is reached.

Since clustering is done separately for each detector view, the final step is to match prongs across views. A temporary track is formed for each possible match, and the cumulative energy distribution along the track is compared using a Kuiper metric [75]. This ensures that the best-matching clusters are paired, refining the reconstruction of particle tracks and showers in the detec-

tor. An event display showing the prongs associated with a neutrino interaction is shown in Figure 3.9.

3.3.4 NOvA Tracking Algorithms

NOvA uses three tracking algorithms to reconstruct particle trajectories: Cosmic tracking [76], Kalman tracking [77][78], and Breakpoint Fitting (BPF) [79] [80]. Each method operates separately in each detector view before combining results into a three-dimensional (3D) track.

- *Cosmic Tracking*

  The cosmic tracker identifies cosmic muons, which typically follow straight-line paths and often appear vertical in the detector. It uses a sliding window tracking algorithm, starting with a window of n planes and fitting a straight line to the hits. Hits matching this fit are added to a two-dimensional (2D) track, and the window moves forward, incorporating new hits that align with the initial fit. This process continues until all planes in both views are examined. Finally, 2D tracks are matched to create a 3D cosmic track.

- *Kalman Tracking*

  The Kalman tracker reconstructs non-showering charged particles, such as muons from neutrino interactions. It follows three steps: track finding, track fitting, and view matching. The algorithm starts by selecting two hits within four cells and estimating the track slope. It then predicts the

location of new hits, iteratively updating the slope and direction as hits are added. Once 2D Kalman tracks are identified in each view, they are matched to form a 3D Kalman track.

- *Breakpoint Fitting (BPF) Tracking*

  BPF tracking accounts for Coulomb scattering, improving reconstruction for heavier particles like protons and pions. Using a Fuzzy-k prong as input, it calculates radiation lengths along the track and places breakpoints where the particle has traveled a multiple of the radiation length or where the scattering angle is significant. BPF then reconstructs the trajectory, repeating the process three times, assuming the particle is a muon, proton, or charged pion, to improve particle identification.

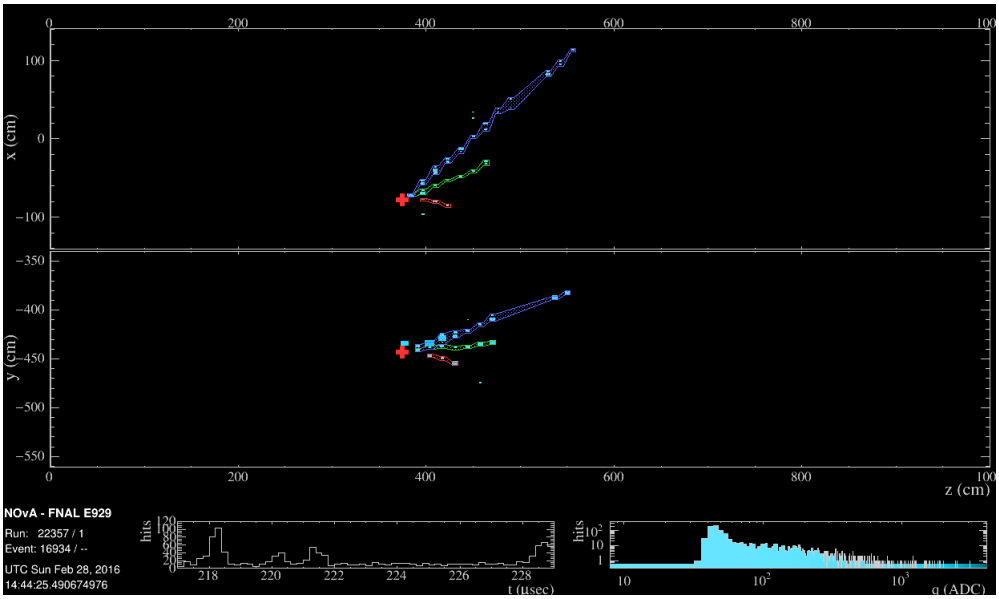

Figure 3.10: An example of a filtered slice in the Far Detector, showing a neutrino candidate. The blue, green, and red regions represent the reconstructed Fuzzy-k prongs, while the red cross marks the reconstructed Elastic Arms vertex.

3.3.5 Neutrino and Particle Event Classification

The NOvA experiment employs machine learning techniques, particularly convolutional neural networks (CNNs), to classify both neutrino events and individual particles within those events. These trained models, known as Convolutional Visual Networks (CVN) [81], analyze pixel maps images of detected hits to determine neutrino interactions and particle types.

At the event level, CVN assigns a score between 0 and 1, estimating the likelihood that an event belongs to one of several categories: ν_μ-CC, ν_e-CC, ν_τ-CC, NC, or cosmic activity. NOvA trains its CVN models using simulated neutrino interactions and cosmic ray data.

In addition to event-level classification, NOvA also employs prongCVN [82], which focuses on classifying individual reconstructed particles (prongs) within a slice. While its architecture is similar to event-level CVN, it further analyzes XY and YZ pixel maps for each particle separately.

Similar to the event CVN, prongCVN outputs scores that indicate how likely a prong corresponds to an electron, photon, muon, pion, or proton. This classification is particularly useful for electron neutrino energy estimation. By applying both event-level and particle-level classification, the use of CVNs enhances NOvA's ability to accurately identify neutrino interactions while also distinguishing individual particles within those events.

Chapter 4

NOvA Oscillation Results

This chapter presents the summary of the three-flavor oscillation and sterile neutrino searches in NOvA. These results demonstrate the efforts of the NOvA collaboration over the years in generating lading physics measurements essential to advancing knowledge in the wider community.

The first section will cover the three-flavor analysis where the main focus is to determine the atmospheric oscillation parameters Δm^2_{32} and θ_{23} using the ν_μ and $\bar{\nu}_\mu$ and disappearance channels and the CP violating phase, using the ν_e and $\bar{\nu}_e$ appearance channels.

The second part of this chapter touches on the sterile neutrino searches, where we look for anomalies in the neutrino event rate at the ND and FD for the NC and the ν_μ disappearance channels that may be consistent with a new fourth neutrino, the sterile neutrino ν_s.

The two analyses share similar strategies to look for neutrino oscillations, but differ in the way that systematic uncertainties are treated and how the fit to determine oscillation parameters is performed.

4.1 Three-Flavor Analysis

This analysis uses the largest data sample collected by NOvA during ten years of operations, including 26.6×10^{20} Protons on Target (POT) of neutrino beam exposure and 12.5×10^{20} POT of antineutrino beam.

4.1.1 Selection

The selection criterion for the analysis ensures that only events of interest are kept in the samples. These include ν_μ, $\bar{\nu}_\mu$, ν_e, and $\bar{\nu}_e$ CC events. NC neutrino interactions and cosmic events are the leading sources of contamination. The selection cuts can be divided into four main categories: quality, containment, cosmic rejection, and neutrino event selections.

- *Quality Cuts*

  Quality cuts ensure that only well-reconstructed events are included in the analysis. These cuts filter out slices with an excessive or insufficient number of hits, a lack of reconstructed prongs or tracks, or unusually high-energy slices. In addition, the selection makes sure that events with only good beam conditions get selected. This includes making sure

that the current and polarity in the magnetic horns, the beam intensity, and beam spot position and size comply with the optimal operation requirements. Runs that do not pass these beam quality checks are discarded.

For the ν_μ event reconstruction, different conditions are required. A muon track with non-zero energy must be reconstructed using the Kalman tracking algorithm, and the neutrino event must have deposited twenty hits or more in the slice and show activity in at least four planes in the detector. In addition, only events with reconstructed energy below 5 GeV are selected, since higher energy events do not contribute to the oscillation measurements.

For the ν_e event reconstruction, the quality cuts ensure that only events with an identified vertex and at least one prong are reconstructed. In addition, a specific cut removes events with too many reconstructed hits within the plane.

- *Containment Cuts*

  Containment cuts ensure that the energy in the selected slices is well contained in the detector. This is important because neutrino interactions near the edges of the detector can result in particles escaping before their energy is fully measured, biasing the event reconstruction.

- *Cosmic Rejection*

  As mentioned in previous chapters, the FD is placed on the Earth's surface with a small rock overburden of approximately 3.6 m.w.e. For

this reason, the FD is constantly bombarded by a high rate of 150 MHz. So, it is important to be able to reject as much as possible cosmic events that can potentially contaminate the neutrino samples. To achieve this, NOvA uses two different rejection algorithms.

The first algorithm uses the direction and position of tracks to identify cosmic ray events, which typically enter the detector from above and travel vertically, in contrast to neutrino beam events that enter horizontally. Events matching the cosmic ray pattern are removed from the sample. In NOvA, this is known as the "Cosmic Rejection Veto."

The second algorithm uses a CNN-based method (also called Cosmic CNN), it uses raw pixel maps of the full detector over a 16 μs time window to identify cosmic-muon like topologies in the detector. Higher scores (close to 1) are used to reject cosmic contamination effectively up to 95% while keeping neutrino candidates with about 99% efficiency.

Lastly, Boosted Decision Trees for cosmic events (Cosmic BDTs) were trained to further reduce the remaining cosmic contamination in neutrino samples. The training inputs for these BDTs included key characteristics such as the number of hits, track length, momentum, starting and ending positions, and the track's angle relative to both the beam direction and the vertical axis.

- *Rock Events Rejection*

  Rock events are produced by neutrino or cosmic rays interacting in the surrounding rock of the ND. This events can mimic ν_μ-CC events or

undergo bremsstrahlung radiation or decay-in-flight that can result in a similar signal as ν_e-CC in the detector. These events are usually filtered out by the containment cuts, since their interaction points are very close to the edge of the detector, but there may be cases where rock events go through the PVC (dead material) for some distance before interacting in the scintillator. To remove these events, a cut is implemented to reject slices where the reconstructed elastic arm vertex is less than 20 cm from each edge of the detector.

- *ν_e Selection*

  ν_e-CC interactions occur very rarely at the FD. For that reason, it is important to keep as many events as possible without sacrificing the purity of the sample. Instead of applying strict cuts that would reduce the sample size, events are classified into three groups: Core, Peripheral, and Low-Energy samples, using a combination of a CNN (or CVN mentioned in subsection 3.3.5), containment, and background rejection.

  The Core Sample includes well-contained events and passes the background rejection criteria. It is further divided by CVN scores, specifically to identify ν_e-CC events into High PID, which consists of high-purity ν_e-like events, and Low PID, which includes lower-purity events but contributes to increased statistics. This separation ensures that the High PID sample remains impactful while still benefiting from additional data from the Low PID one.

  The Peripheral Sample consists of events that fail containment, but still pass other selection requirements. These events tend to be more

background-like, so a separate Cosmic BDT is used, along with a strict CVN ≥ 0.97 cut. This allows for the inclusion of additional events while controlling background contamination.

The Low-Energy Sample (specific to the neutrino mode) includes events that do not pass the core selection but fall within a lower energy range. A dedicated BDT is applied to ensure that only high-quality, low-energy events are retained, improving sensitivity in this region without compromising purity.

- *ν_μ Selection*

  Similar to ν_e-CC events, the selection of ν_μ-CC events is based on a CVN. In addition, events must be well-contained and pass background rejection. Since many events occur in the dip region of the FD energy spectrum, where oscillation effects are most pronounced, the selection criteria are optimized to maximize the sensitivity to oscillation parameters in this range.

  This optimization is achieved through two main techniques. First, the energy bin widths are adjusted so that bins near the dip region are finer, while regions with minimal oscillations have coarser bins. Second, the ν_μ sample is divided into four quantiles based on the hadronic energy fraction, which is the portion of neutrino energy carried by the hadronic system. Events with higher hadronic energy fractions generally have poorer energy resolution and more background contamination.

4.1.2 Extrapolation in NOvA analyses

The NOvA experiment employs a near-to-far extrapolation technique to predict the expected signal at the Far Detector (FD) based on data collected at the Near Detector (ND). This method is crucial for neutrino oscillation studies, allowing for corrections based on ND data to refine FD predictions. The extrapolation process is used for both muon neutrino (ν_μ) disappearance and electron neutrino (ν_e) appearance, but the implementation slightly differs between the two.

Muon Neutrino Disappearance Extrapolation

In the muon neutrino disappearance analysis, the goal is to determine how many ν_μ neutrinos disappear as they travel from the ND to the FD. NOvA observes two survival channels: $\nu_\mu \rightarrow \nu_\mu$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$, to infer neutrino disappearance. Since ND detects many ν_μ events, this data is used to predict what should be observed at the FD.

The process begins by converting the reconstructed energy of ND ν_μ events into true energy using a 2D matrix transformation. A far/near energy ratio is then calculated to adjust for differences in detector acceptance and size. The data is modified using oscillation probabilities, reflecting how neutrinos are expected to transform. Finally, a second transformation converts the predicted true energy spectrum at the FD back into reconstructed energy, ensuring that it aligns with what the FD would actually measure. The diagram of this method can be seen in Figure 4.1.

To improve accuracy, each ν_μ quantile sample is divided into three bins of

transverse lepton momentum, p_T (12 p_T bins in total), which are extrapolated from the ND to the FD. The lepton transverse momentum is measured relative to the NuMI beam in an energy range of 0 to ∼1.5 GeV. This binning technique accounts for kinematic differences between the ND and FD due to different containment capabilities in each detector. This method improves the extrapolated prediction and reduces systematic uncertainties, with a greater impact on cross-section systematics [83][84]. The extrapolation process is repeated for each of these bins. At the end of this extrapolation process, the bins are merged back together at the FD.

Since backgrounds are relatively minor in the ν_μ disappearance analysis, the background predictions are taken directly from Monte Carlo (MC) simulations rather than being extrapolated from ND data. This simplifies the process while maintaining accuracy.

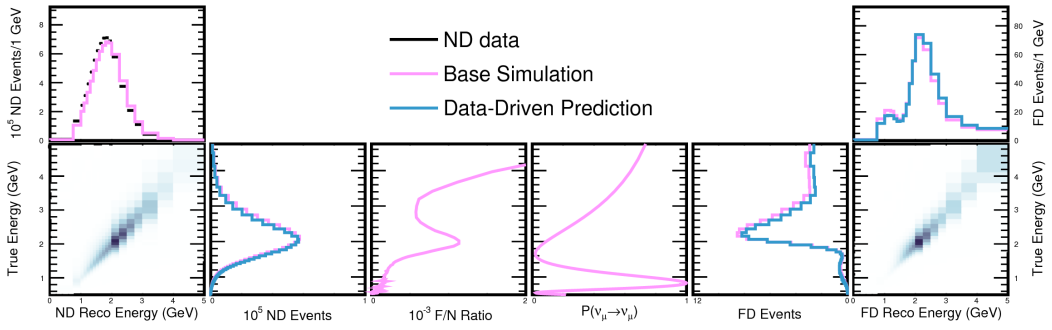

Figure 4.1: Diagram of the ν_μ extrapolation procedure, starting with ND Data/MC in the top left, through Reco→True, near→far, oscillations, true→reco, to the adjusted FD prediction in the top right

Electron Neutrino Appearance Extrapolation

The electron neutrino appearance analysis searches for ν_e events at the FD that originated as ν_μ at the ND due to oscillations. Unlike in the disappearance analysis, for which the same type of neutrino is measured at both detectors, the appearance analysis involves identifying a different flavor at the FD than was measured at the ND.

To predict the ν_e signal at the FD, NOvA applies the $\nu_\mu \rightarrow \nu_e$ oscillation probability to ND ν_μ events. The same transformation steps as in the disappearance analysis are applied: ND reconstructed energy is converted to true energy, transported to the FD using a far/near ratio, and then modified using oscillation probabilities. Finally, the true energy prediction at the FD is converted back into reconstructed energy. The diagram of this method can be seen in Figure 4.2.

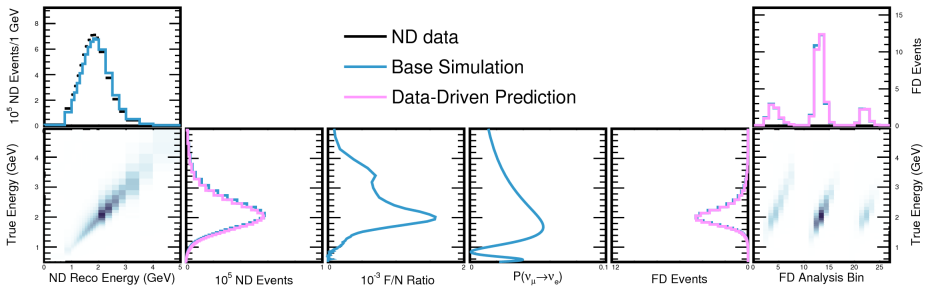

Figure 4.2: Diagram of the ν_e extrapolation procedure, starting with ND Data/MC in the top left, through Reco→True, near→far, oscillations, true→reco, to the adjusted FD prediction in the top right

Unlike in the ν_μ disappearance case, backgrounds in the ν_e appearance analysis are significant and require constraints. Instead of being extrapolated directly, the ND ν_e sample is decomposed into its true background components: intrinsic beam ν_e (from meson decays in the NuMI beam), misidentified ν_μ-CC

interactions, and NC events that mimic ν_e signals.

To constrain these backgrounds at the FD, a ratio of the ND background before and after decomposition is calculated. This ratio serves as a correction factor that is applied to the FD simulated background components. The process modifies both the normalization (total expected background events) and the shape (energy distribution of the background). The final FD background prediction is obtained by summing the individually corrected components.

In addition to these major background components, there are minor background contributions that are not constrained using the ND decomposition process. Instead of using decomposition and scaling, these minor backgrounds are taken directly from FD simulations, similar to the background treatment in the ν_μ disappearance analysis. Since their contribution is small, constraining them with ND data would have little impact, making direct simulation a more efficient approach.

Like in the ν_μ disappearance analysis, the same p_T extrapolation procedure is conducted to correct for kinematic differences between the ND and FD. Since in this analysis the ν_μ sample is not split into quantiles, the total number of p_T bins is three. This binning helps ensure that the predicted ν_e signal at the FD is as accurate as possible.

4.1.3 Estimation of Cosmic Backgrounds

Cosmic ray events significantly impact the FD, with rates around 130 kHz. As described in subsection 2.2.8, They are measured using a 500 μs data window

around the 10 μs beam spill and a 10 Hz trigger without the beam spill. Timing sidebands help estimate cosmic ray rates but have low statistics. To improve accuracy, the 10 Hz trigger determines the energy distribution shape, which is then scaled using sideband data. After extrapolation, cosmic ray events are included in the FD prediction, with uncertainties modeled using a Poisson distribution due to the low event rate.

4.1.4 Systematic Uncertainties

NOvA's neutrino oscillation analysis includes many components, each affected by systematic uncertainties. Over 100 uncertainties are considered, mainly related to neutrino interactions, beam flux, and calibration. Some are easily handled by adjusting energy spectra, while others require modifying reconstructed data to properly reflect their impact. The most complex systematic uncertainties affect fundamental aspects of simulations and calibrations, requiring additional simulated events to fully capture their effects.

- *Detector Response and Calibration*

  Calibration uncertainties are accounted for by resimulating ND and FD samples with modified systematic parameters. These uncertainties include relative and absolute energy scales, calibration shape, detector drift, Cherenkov light emission and light level.

  The uncertainty in overall energy response is determined by differences between data and simulation for proton candidates in the ND, with a discrepancy of about 5 %, which is larger than for other particles or standard

calibration metrics like Michel electrons and π_0 mass peaks. To assess this calibration systematic, nominal simulated samples are reused while adjusting the absolute calibration constants up and down by 5 % before reconstruction. Additionally, to address variations in reconstructed and true energies near the detector edges, a separate calibration shape uncertainty is included.

The detector drift systematic reflects the gradual aging of the detectors, modeled as a steady decrease in light output over time. This reduction is estimated at around 4.5 % per year, based on observed changes in the number of hits per event and chosen to match the shifts seen in the data. Because the detector response decreases over time, only positive variations of this systematic are considered. Additionally, the effect is assumed to be fully correlated between the ND and FD.

Uncertainties also exist in the Cherenkov light collection efficiency, which is determined through light level tuning. In the current simulation, this efficiency is estimated to vary by approximately 6.2 %. In addition to Cherenkov, there is a systematic uncertainty on the light level, which is tuned using a similar method and estimated at 5 %.

- *Cross Section Systematics*

  Neutrino cross-section uncertainties are one of the largest categories of systematics in NOvA's oscillation analysis. These are evaluated by reweighting the nominal simulation event rates, adjusting how often different types of interactions occur. This approach modifies the simulation to account for uncertainties in quasi-elastic scattering, multi-nucleon inter-

actions (2p2h-MEC), resonance (RES) production, deep inelastic scattering (DIS), and final-state interactions (FSI).

For quasi-elastic scattering, the primary uncertainty comes from the axial form factor, implemented in GENIE using the z-expansion model. Systematic uncertainties affect both its shape and normalization, with shifts of -20% to $+15\%$. Since GENIE does not enforce correlations between these uncertainties, NOvA applies custom weight functions to ensure proper treatment. Additionally, a central value correction based on MINERvA data is used to refine the quasi-elastic model.

NOvA applies several uncertainties related to 2p2h-MEC interactions. The MEC shape uncertainty modifies the simulation to make events appear more QE-like or RES-like, followed by a refitting of the MEC Gaussian parameters. Another uncertainty accounts for neutrino energy dependence by comparing different MEC models (SuSA and Martini) to NOvA's default Valencia model. A function describing the variation between these models defines the 1σ error.

For RES production, NOvA corrects an excess of low-Q^2 events observed in the nominal simulation. MINOS encountered a similar issue and applied a functional weight to improve agreement between simulation and data. NOvA adopts the same method, which successfully improves the simulation's accuracy in the low-Q^2 region.

- *Beam Flux Systematics*

  Beam flux uncertainties originate from hadron production corrections in

the PPFX and variations in beam transport, such as horn current and position. The PPFX applies data-driven corrections to the G4NuMI simulation, with uncertainties incorporated through event re-weighting. The main sources of uncertainty include horn current (±2 kA), horn positions in x and y (±3 mm), measured beam position on target in x and y (±1 mm), beam spot size in x and y (±0.2 mm), horn cooling water layer thickness (±1 mm), target z position (±7 mm), and beam divergence (+54 μrad). Refining the horn geometry had a negligible impact on beam flux.

To estimate hadron production uncertainties, NOvA employs Principal Component Analysis (PCA) using the PPFX framework [85]. A statistical ensemble of randomly generated universes is created, each incorporating different variations in proton-target cross sections based on external data uncertainties. A covariance matrix is constructed from the generated universes, which is then diagonalized to obtain eigenvectors and eigenvalues, representing the principal components used as systematic shifts.

- *Other Systematics Unceranties*

  NOvA incorporates a systematic uncertainty related to the cosmic content at the FD, which arises from statistical uncertainty in the cosmic ray sample used for background predictions.

  A new systematic was also introduced to evaluate neutron interactions by incorporating an alternative neutron-carbon interaction model, MENATE [86, 87]. While NOvA typically uses Geant4 for neutron simulations, this

systematic assesses the impact of replacing it with the MENATE model.

Geant4Reweight [88] is another systematic that allows modifications to Geant4 hadron interaction modeling without requiring full re-simulation. It assigns event weights based on adjusted cross-sections for different interaction types. For pions, these include absorption, quasi-elastic scattering, and charge exchange. Geant4Reweight includes 11 systematic knobs—five for π^-, five for π^+, and one for protons—while the neutron knob was excluded due to incomplete understanding of its effects. These systematics are incorporated into NOvA's three-flavor analysis using Principal Component Analysis (PCA).

In addition to these, several other systematics are considered. Normalization systematics account for multiple small effects, including uncertainties in POT, differences in detector masses, and event pile-up in the ND. Other systematics include Michel electron tagging, which affects the identification of muon decays, FD rock systematics, which account for uncertainties in interactions occurring in the surrounding rock, and matter density systematics, which address variations in Earth's density along the neutrino path.

4.1.5 Three-Flavor Oscillation Results

This section presents the most recent three-flavor oscillation results shown at Neutrino 2024. The fits were performed by comparing FD predictions to the data to determine the best-fit oscillation parameters.

The best fit spectra and data at the FD for ν_μ-CC samples are shown in Figure 4.3. Similarly, the best fit spectra and data at the FD for ν_e-CC samples are shown in Figure 4.4.

The extraction of oscillation parameters was performed using a frequentist χ^2 minimization method with a profiled Feldman-Cousins technique [89] to account for low statistics and physical boundaries and a Bayesian Markov Chain Monte Carlo technique [90]. At the time of writing this thesis, the frequentist analysis provides the Δm^2_{32} and $\sin^2\theta_{23}$ only fits.

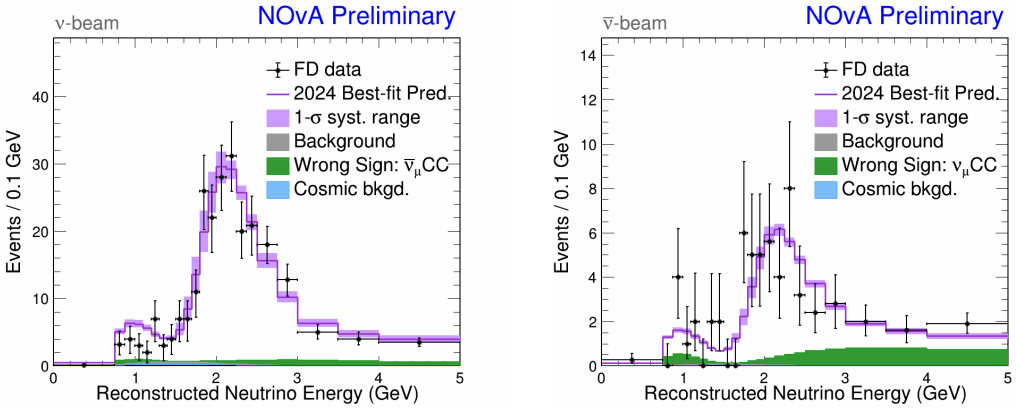

Figure 4.3: Best fit predictions at the FD for ν_μ (left) and $\bar{\nu}_\mu$ (right), the data is shown in black and the 1-σ systematic uncertainty band in magenta.

The resulting fit contours for Δm^2_{32} vs $\sin^2\theta_{23}$ and $\sin^2\theta_{23}$ vs δ_{CP} are shown in Figure 4.5 at 90% and 68% Confidence Level (CL) respectively.

The best-fit value for Δm^2_{32} is $2.424^{+0.035}_{-0.040} \times 10^{-3}$ eV^2, while $\sin^2\theta_{23}$ is $0.55^{+0.02}_{-0.06}$. The CP-violating phase δ_{CP} is constrained to two possible regions: $0.930^{+0.210}_{-0.290}\pi$ and $0.150^{+0.150}_{-0.110}\pi$.

The results indicate a 76% preference for the normal mass ordering. This

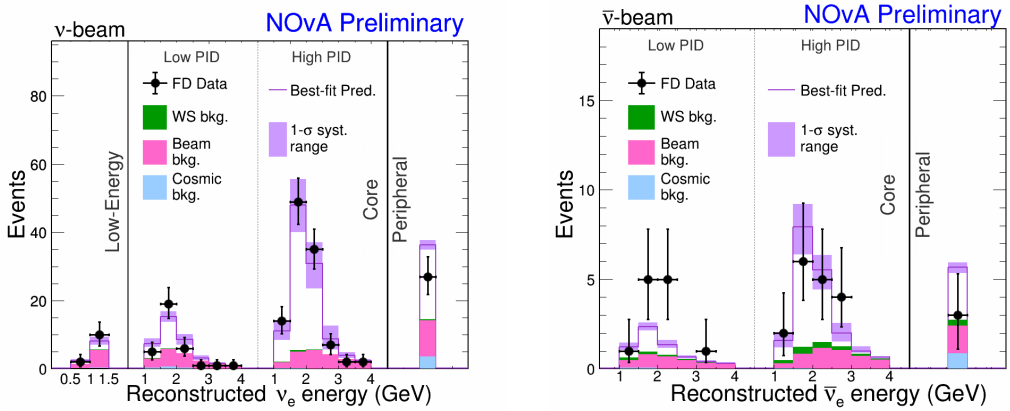

Figure 4.4: Best fit predictions at the FD for ν_e (left) and $\bar{\nu}_e$ (right), the data is shown in black and the 1-σ systematic uncertainty band in magenta.

measurement represents the most precise determination of Δm^2_{32} from a single experiment. A more detailed discussion of these results can be found in [50].

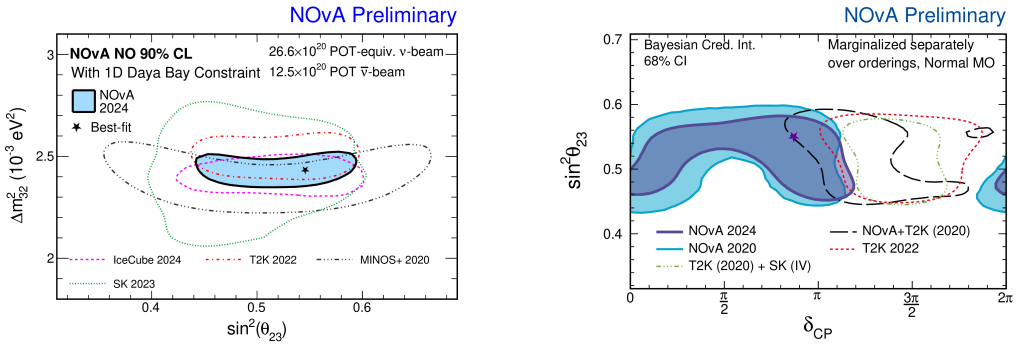

Figure 4.5: Comparison of Contours for $\sin^2\theta_{23}$ and Δm^2_{32} Across Different Experiments (Left). Comparison of for $\sin^2\theta_{23}$ and δ_{CP} Across Different Experiments (Right).

4.1.6 Impact of Systematic Uncertainties

Figure 4.6 illustrates the impact of various systematic uncertainties on two oscillation parameters at the Frequentist best-fit point. Detector calibration remains NOvA's largest source of uncertainty.

For $\sin^2\theta_{23}$, the detector calibration uncertainty is significantly larger than other systematic contributions, with the error bar pulling predominantly to the left, favoring the lower octant. However, statistical uncertainty remains the dominant source overall.

For Δm^2_{32}, the uncertainties are more symmetric, and detector calibration is less dominant compared to other systematic groups. Here, systematic uncertainties have a relatively larger effect than statistical uncertainty.

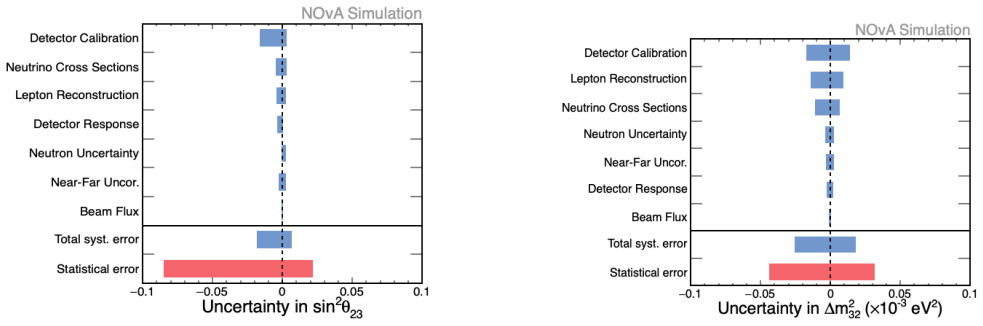

Figure 4.6: Uncertainty Estimates for Oscillation Parameters at the Frequentist Best-Fit Point for $\sin^2\theta_{23}$ and Δm^2_{32}.

One of the main goals of this thesis is to provide the tools to improve the detector calibration systematic, particularly the detector response component, which has the largest impact (subsection 4.1.6). To achieve this, a Test Beam experiment was conducted at Fermilab (see Chapter 5) to study the energy

response of single charged particles, including electrons, protons, and pions. In this work, the pion energy response is specifically studied and presented in Chapter 7. Within the three-flavor framework, the pion Test Beam data will provide valuable resources to validate the NOvA simulation and reduce the calibration systematic that directly impacts the measurements of the oscillation parameters, such as θ_{23} and Δm^2_{32}.

4.2 Sterile Neutrino Analysis

As mentioned in section 1.10, experiments such as LSND and MiniBooNE have reported an anomalous excess of $\nu_e(\bar{\nu_e})$ events in $\nu\mu(\bar{\nu_\mu})$ beams. This result deviates from the expected behavior of three-flavor neutrino oscillations. One possible explanation is mixing with the existence of a fourth neutrino, known as the sterile neutrino.

Sterile neutrinos can be probed in long-baseline neutrino experiments such as NOvA. If they exist, their presence would lead to deviations from the standard three-flavor predictions, such as a depletion of NC and ν_μ events and a possible excess of ν_e events in the detectors. This effect would be especially noticeable in NC interactions, as these are not influenced by standard three-flavor oscillations.

The UC neutrino group, as part of the NOvA collaboration, has been leading efforts in these searches. First analyses [91, 92] compare the observed event rates at the FD to predictions based on the standard three-flavor oscillation model using NC samples in the 3+1 neutrino paradigm. The new analysis

recently published in PRL [93] uses NC and ν_μ-CC samples. Unlike the three flavor analysis, and previous sterile neutrino searches at NOvA, the fit is performed using the samples at the two detectors with systematics uncertainties encoded in a covariance matrix.

This analysis uses the neutrino beam data collected up to 2020 with a total FD exposure of 13.6×10^{20} POT and the ND dataset exposure of 11×10^{20} POT. The next sections will present the analysis strategy along with my contributions and final results.

4.2.1 ν_μ-CC and NC Sample Selections

The selection criteria for the ν_μ-CC sample follows the same standards as in the three-flavor analysis. However, due to computational constraints, the sample is not subdivided into hadronic energy fraction quantiles, nor are transverse momentum extrapolated samples used.

Similarly, NC interaction events, which are characterized by hadronic activity and an outgoing neutrino that is not detected, must pass quality and containment cuts. In particular, all NC events must have a reconstructed vertex, at least one reconstructed particle, and activity in at least three contiguous planes. A CVN is also applied here where optimal scores are used to reject background events from the NC sample. Figure 4.7 shows the CVN score quantity used to reject backgrounds from the FD NC sample. A cut was set to include all events with a CVN > 0.1.

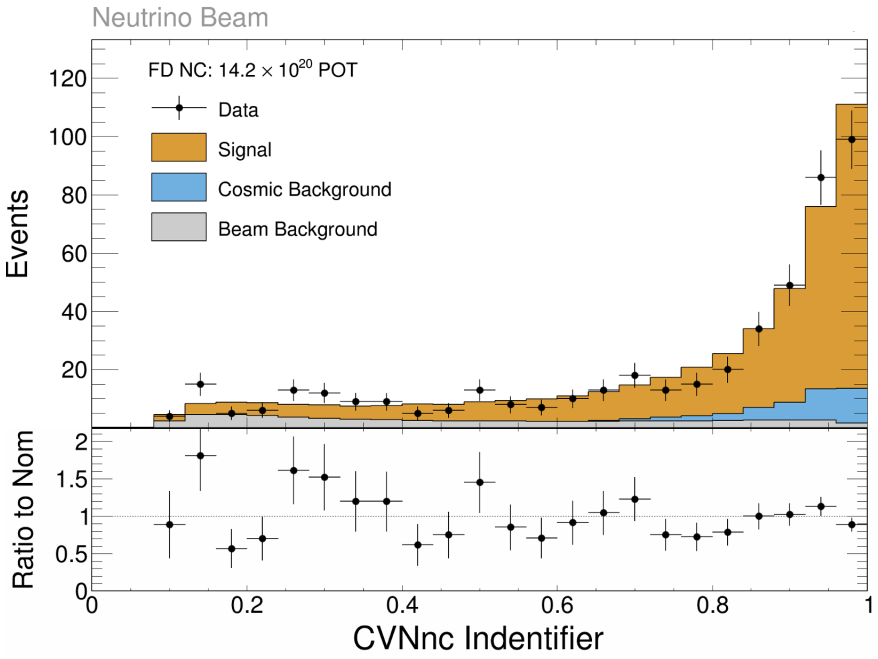

Figure 4.7: The top plot shows the CVN score used to reject background events from the FD NC sample. A cut is applied to include all events with a CVN score greater or equal than 0.1.

4.2.2 Background Rejection

As stated before, the three-flavor and sterile neutrino analyses share many techniques in selection, the same can be said about the background rejection techniques. In the case of beam backgrounds, samples of no interest can be rejected using machine learning techniques. The other main sources of background are cosmic and rock events.

The rejection of cosmic events in the ν_μ-CC and NC samples follows the same principle explained in subsection 4.1.1, where a combination of a BDT and a CNN is used to reject cosmic-like events in both samples.

Similarly, rock events are removed using topological cuts in the detector to discard events consistent with neutrinos or cosmic muons interacting in the surrounding rock at the ND.

4.2.3 Data Quality Checks

Data quality checks are an essential part of the analysis, helping ensure that the selected event sample is reliable. They allow us to verify that all events are classified correctly and that the right selection criteria have been applied. I performed these checks for NC events and found that the data and Monte Carlo (MC) simulations were in good agreement. 4.8 shows examples of data quality variables for the NC selection in the FD.

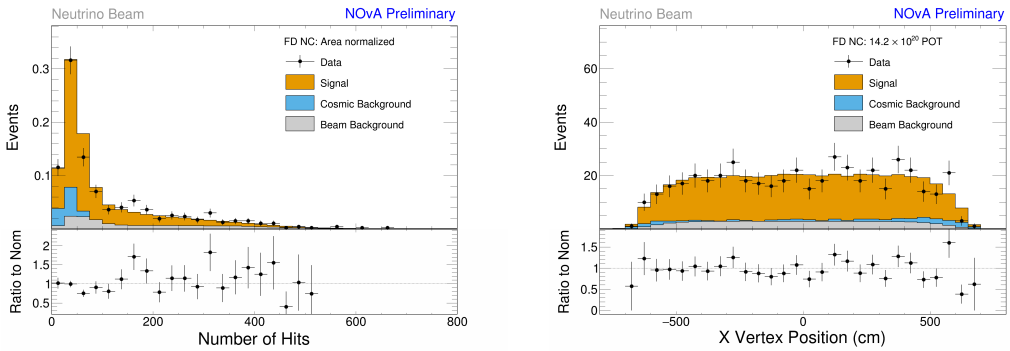

Figure 4.8: Plots comparing the predicted spectrum of neutral current events in the FD to data. Blue and gray histograms show the background component, whereas the orange histograms show the NC component. Number of Hits in the events (Left), neutrino vertex position in the x axis (Right)

4.2.4 Side Band Studies

Sidebands are typically regions of the spectrum where the signal of interest is not expected to appear or a set of events that fail a specific selection.

Studying sidebands helps identify potential issues in event selection by comparing data with Monte Carlo (MC) simulations. Ensuring a good agreement between the two is especially important in neutrino analyses, which are conducted blindly to avoid bias in important signal regions.

In this analysis, I examined the unselected portion of NC events in the FD, focusing on the energy deposited in the detector. 4.9 illustrates this distribution.

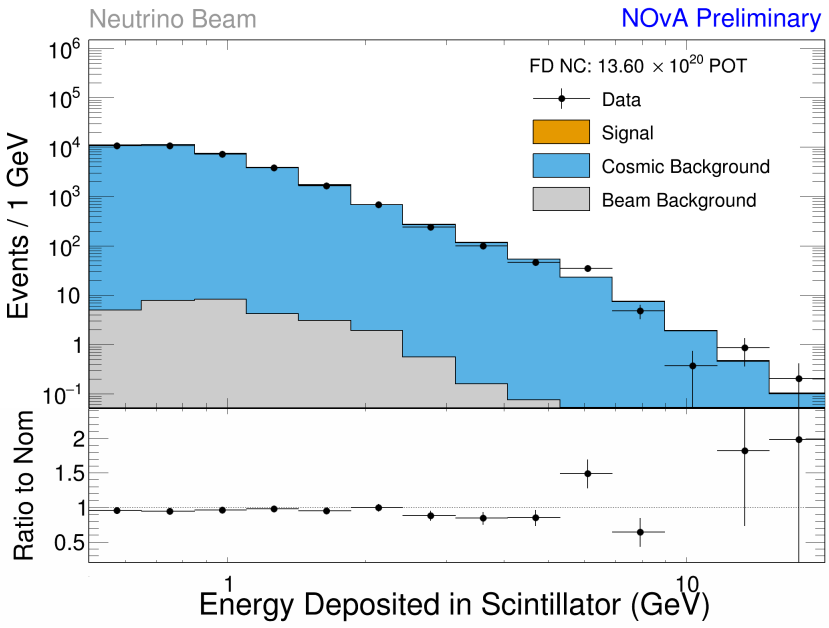

Figure 4.9: Side band plot for unselected NC events in the energy deposit spectrum, where we observe a good level of agreement between data and MC.

4.2.5 Estimation of Cosmic Backgrounds

The estimation of these events is well described in subsection 4.1.3, and the same technique is used for both ν_μ-CC and NC events. For this analysis, I was responsible for verifying that the estimated cosmic background content in the NC sample at the FD is consistent with the background estimated in the sidebands. If the two samples agree, it confirms that the estimation procedure was performed correctly. The result of this study is shown in Figure 4.10, which demonstrates good agreement between the two.

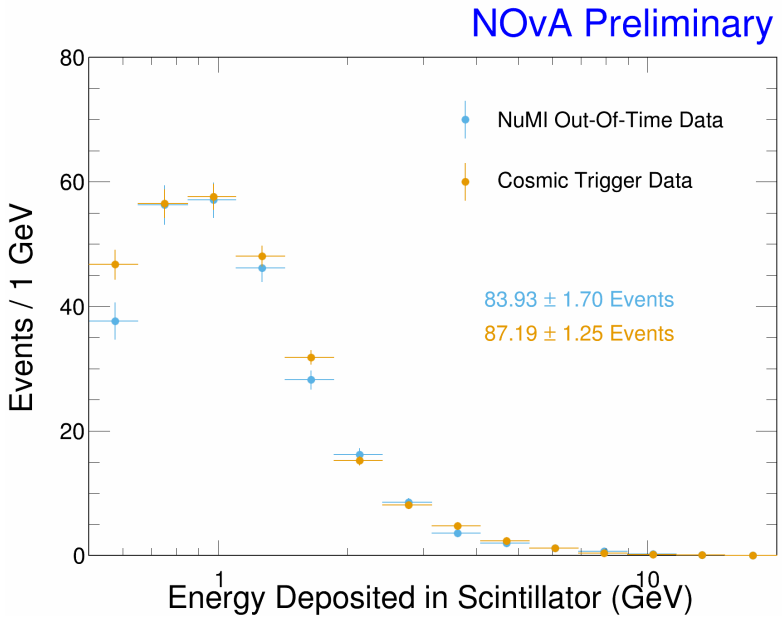

Figure 4.10: Estimated cosmic background content in the NC sample at the FD (Cosmic Trigger Data) and background estimated in the sidebands (NuMI Out-Of-Time Data), where we observe a good level of agreement between the two samples.

4.2.6 Systematic Uncertainties

The systematic uncertainties in this analysis are largely based on those from the three-flavor study, fully discussed in subsection 4.1.4, with a few differences. These differences arise primarily from the fact that no ND/FD extrapolation is applied here.

The uncertainties from beam flux and hadron production are the same as in the three-flavor analysis. However, because extrapolation is not used, we handle uncertainty bands differently. Instead of adjusting them based on changes in the far-to-near spectra in true energy, we evaluate their effects on the FD and ND spectra separately

Similarly, the three-flavor analysis accounts for calibration uncertainties by considering the ND and FD energy scale uncertainties as either fully correlated or fully anti-correlated to capture the possible range of effects. In this analysis, since each detector is fit independently, it is more appropriate to treat their energy scale uncertainties as uncorrelated.

Finally, two new sources of uncertainty were introduced: a kaon beam parent normalization uncertainty, and an uncertainty related to meson exchange current (MEC) interactions. These are discussed in more detail below.

- *Kaon Systematic Uncertainty*

  Many of the NC neutrino candidates selected for this analysis come from kaon decays. In previous studies, the beam kaon component was assigned a 30% normalization uncertainty in addition to the PPFX uncertainties

due to limited hadron production data. Instead of keeping this large uncertainty, we used samples outside the main analysis to better estimate it.

To do this, we analyzed horn-off data, which allows us to study hadron production uncertainties without the effects of the focusing horns, and a sample of uncontained high-energy ν_μ events, which provides insight into the focused kaon peak. This study examined ND horn-off data passing NC, ν_μ, and ν_e selections, as well as the uncontained ν_μ sample.

We then fit the kaon flux normalization while accounting for possible sterile neutrino oscillations. The best fit was at 1, with small variations within 10%. Since the changes were minor, we reduced the kaon flux uncertainty from 30% to 10%, which are applied in our final analysis.

- *Cross Section Systematics*

  As mentioned before, the three flavor oscillation analyses use extrapolation, where the ND is assumed to have no oscillations. This means that any differences between data and simulation are treated as mismodeling in the simulation. To correct for this, we adjust the cross-section models in the ND simulation to better match the ND data, producing updated central values and uncertainties.

  However, in a sterile neutrino search, oscillations could occur in the ND, so differences between data and simulation cannot simply be due to cross-section miss-modeling. Because of this, we do not apply the cross-section tuning and instead use the original, untuned simulation and uncertainties.

The Meson Exchange Current (2p2h-MEC) component of the simulation is not well constrained. To account for this, we introduce shape and normalization uncertainties based on differences between the Valencia, SuSA, and GENIE empirical MEC models.

4.2.7 PISCES: Analysis Framework

As discussed in section 1.11, sterile neutrino searches can be conducted using the 3+1 sterile model. This approach provides access to new δ_{CP} phases, mass-squared differences, and mixing angles in addition to the three-flavor parameters.

By using NC and ν_μ-CC samples in the sterile neutrino context, NOvA can constrain Δm^2_{41}, $\sin^2\theta_{24}$, and $\sin^2\theta_{34}$ sterile parameters. This is achieved by fitting the data from both the ND and FD to the expected predictions for three-flavor oscillations.

To accomplish this, the UC group has developed a framework known as PISCES. PISCES stands for Parameter Inference with Systematic Covariance and Exact Statistics. It uses a hybrid test statistic that handles uncertainties in two ways: Statistical uncertainties are treated using a Poisson log-likelihood approach, while systematic uncertainties are managed with a Gaussian multivariate technique. The framework is outline bellow and the full derivation of PISCES can be found in [94].

In traditional high-statistics analyses, A Gaussian multivariate technique can be used, in which statistical and systematic uncertainties and their correlation

are encoded in a covariance matrix:

$$\chi^2_{\text{cov}} = \sum_{ij} (x_i - \mu_i)\, V_{ij}^{-1}\, (x_j - \mu_j), \tag{4.1}$$

where x_i and μ_i are the data and nominal prediction in each bin i, respectively, and V_{ij} is the total covariance matrix, including both statistical and systematic uncertainties. This technique is fast since the systematic uncertainties do not need to be profiled by the fitter. However, this approach does not calculate the systematic pulls directly, therefore systematic shifts cannot be visualized. In addition, since this method uses a Gaussian approach, it is not valid when the data contain bins with very few events.

To address this, a statistically proper approach is to use the Poisson likelihood to calculate χ^2 between data and the predicted spectrum:

$$\chi^2_{\text{pull}} = 2 \sum_{i}^{N} \left(m_i - x_i + x_i \log \frac{x_i}{m_i} \right) + \sum_{j}^{K} \alpha_j^2 \tag{4.2}$$

where $m_i = \mu_i + \sum_{j=1}^{K} \alpha_j s_{ji}$ is the predicted event rate in bin i including systematic shifts. Here, α_j is the fit parameter for the j-th systematic, and s_{ji} is its corresponding reweight factor in bin i. K is the total number of systematics. Dependencies on oscillation parameters are omitted for simplicity. This method is statistically robust, but computationally expensive since the systematic uncertainties are profiled in the fit.

PISCES avoids profiling systematic uncertainties by combining the Poisson

likelihood with the Gaussian multivariate technique. Leading to:

$$\chi^2 = 2\sum_i \left[S_i - x_i + x_i \log\left(\frac{x_i}{S_i}\right)\right] + \sum_{ij}\sum_{\alpha\beta}(s_{\alpha i} - 1)F^{-1}_{\alpha i\beta j}(s_{\beta j} - 1). \quad (4.3)$$

Here, the covariance matrix encoding only systematic uncertainties is used to determine the optimal systematic pulls, $s_{\alpha i}$, for each oscillation channel α and analysis bin i. The observed data are denoted as x_i, while the predicted value, S_i, is given by $S_i = \sum_\alpha \mu_{\alpha i} s_{\alpha i}$, where $\mu_{\alpha i}$ represents the nominal prediction. The covariance matrix, $F_{\alpha i\beta j}$, accounts for systematic uncertainties and their correlations across different oscillation channels (α, β) and analysis bins (i, j).

To perform the fit, at each point in parameter space, first the χ^2 is minimized with respect to the systematic pulls, s_{γ_k}. These pulls are found by solving for the minimum χ^2 using iteratively the Levenberg-Marquardt algorithm (LMA) [95].

To determine the minimum, one solves for the roots of the gradient of the χ^2 with respect to each systematic pull. The gradient vector for each beam component γ_k and analysis bin is given by:

$$\frac{\partial\chi^2}{\partial s_{\gamma_k}} = 2\left(\mu_{\gamma_k} - \frac{\mu_{\gamma_k} x_k}{\sum_\alpha^M \mu_{\alpha k} s_{\alpha k}} + \sum_{\alpha i}^{MN}(s_{\alpha i} - 1)F^{-1}_{\alpha i\gamma_k}\right) \quad (4.4)$$

The corresponding Hessian matrix of the χ^2 with respect to a pair of systematic pulls is:

$$\frac{\partial^2 \chi^2}{\partial s_{\gamma_k} \partial s_{\delta_l}} = 2(1 + \delta_{\gamma_k \delta_l} \lambda) \left(\frac{\mu_{\gamma_k} x_k}{\left(\sum_\alpha^M \mu_{\alpha k} s_{\alpha k} \right)^2} \delta_{kl} + F^{-1}_{\gamma_k \delta_l} \right) \tag{4.5}$$

Here, λ is a stabilizing factor introduced in the LMA algorithm, which interpolates between gradient descent and Newton's method. It starts at 0.1 and is reduced iteratively by a factor of 1.5 as the method converges. When $\lambda \to 0$, pure Newton's method is recovered.

The linear system formed by the gradient vector in Equation 4.4 and Hessian in Equation 4.5 is solved via Cholesky decomposition to obtain the updated values of systematic pulls. The implementation uses LAPACK routines to handle the matrix algebra.

If a given bin has data very close to zero, the gradient may drive a pull below zero, causing instability. To ensure robustness, any bins for which $\frac{\partial^2 \chi^2}{\partial s_{\alpha i}^2}\Big|_{s_{\alpha i}=0} < 0$ are removed during minimization. Once the system converges, these bins are reintroduced and minimization is repeated, until a stable minimum is found.

Once the values of the systematic pulls have been determined, the total χ^2 is recomputed as the sum of a penalty term and a Poisson likelihood term.

First, the penalty for shifting systematic parameters from their nominal values is evaluated as:

$$\chi^2_{\text{penalty}} = \sum_{ij}^{N} \sum_{\alpha\beta}^{M} (s_{\alpha i} - 1) F^{-1}_{\alpha i \beta j} (s_{\beta j} - 1) \tag{4.6}$$

Then, the Poisson likelihood component is calculated for each systematically shifted bin:

$$\chi^2_{\text{Poisson}} = 2\sum_{i=1}^{N}\left[\left(\sum_{\alpha=1}^{M}\mu_{\alpha i}s_{\alpha i}\right) - x_i + x_i\log\left(\frac{x_i}{\sum_{\alpha=1}^{M}\mu_{\alpha i}s_{\alpha i}}\right)\right] \tag{4.7}$$

The total test statistic χ^2 is the sum of these two terms:

$$\chi^2 = \chi^2_{\text{Poisson}} + \chi^2_{\text{penalty}} \tag{4.8}$$

as described in Equation 4.3.

The process is then repeated iteratively: the Hessian matrix and gradient vector are recalculated, new systematic pulls s_{ij} are determined, and updates Equation 4.8. This process continues until the change in χ^2 between iterations is less than 10^{-10}, at which point the minimum has been located within the tolerance specified by the MINUIT fitter, which is used to perform the minimization with respect to the oscillation parameters. Convergence typically occurs within 30 iterations for standard use cases, though more iterations may be needed when χ^2 is very large.

4.2.8 PISCES: Robustness Of Fit

The fitting technique mentioned in previous section is used to search for sterile neutrinos in the 3+1 model using NC and ν_μ-CC MC selected samples, the fitting procedure is performed simultaneously in both detectors over a series of fake data experiments including Asimov fake data (simulated data where

systematic uncertainties are assumed to be at their nominal values) and fluctuated pseudo-experiments including statistical and systematics fluctuations.

I conducted extensive work to test the PISCES framework by checking the robustness of the fit. This is important since some fits can be very difficult to resolve when statistical and systematic fluctuations are large. The results of this study provided insight into the challenges we encountered when fitting different samples and oscillation parameters simultaneously. Examples of these study for the $\sin^2\theta_{24}$ vs Δm^2_{41} and $\sin^2\theta_{34}$ vs Δm^2_{41} parameter space are shown in Figure 4.11. Ultimately, these studies contributed to the development of a more sophisticated fitting framework used in the final analysis.

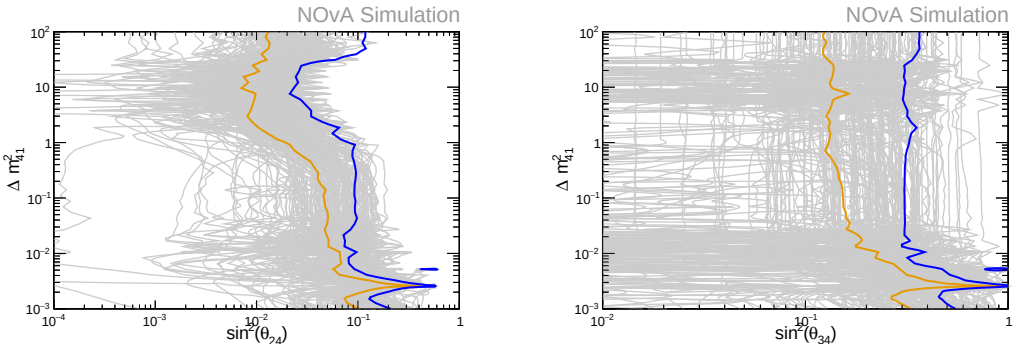

Figure 4.11: 90 % CL contours, statistical plus systematic fluctuations fit: The Median contour is represented by the orange contour, Asimov sensitivity with blue contour, and 200 individual pseudo-experiments with the gray contours. Left: $\sin^2\theta_{24}$ vs Δm^2_{41}, right: $\sin^2\theta_{34}$ vs Δm^2_{41}.

4.2.9 CMF: Analysis Framework

The Covariance Matrix Framework (CMF) is a complementary fitting technique developed by the University of Wisconsin-Madison neutrino group. This

framework is used for searches of sterile neutrinos in NOvA. Similar to PISCES, CMF encodes the systematic uncertainties in a covariance matrix, while statistical uncertainties are included in the diagonal of the covariance matrix using the combined Neyman–Pearson formalism.

4.2.10 Analysis Results

The results presented in this section make use of PISCES (Analysis 1) and CMF (Analysis 2) to search for sterile neutrinos in the NOvA experiment.

In both analyses, the oscillation parameters θ_{23} and Δm^2_{32} are varied, with a loose Gaussian constraint on Δm^2_{32} to keep the fit within a 3+1 flavor framework. This constraint is centered at $|2.51| \pm 0.15 \times 10^{-3}$ eV^2 for both mass orderings, based on a 2020 global fit to atmospheric neutrino data, and its width is set conservatively to twice the 3σ range.

The sterile parameters Δm^2_{41}, θ_{24}, θ_{34}, and δ_{24} are allowed to vary freely, while the remaining sterile parameters are fixed at zero due to constraints from solar and reactor experiments, as well as unitarity requirements..

The resulting fit spectra comparing data and predictions for the two analyses are shown in Figure 4.12

The findings show that NOvA data is consistent with three-flavor oscillations at the 90% confidence level, with limits aligning with sensitivity studies based on this model (Figure 4.13). These results include the first constraints in certain regions of phase space and exclude previously allowed parameter space

from IceCube at 90% confidence. Additionally, this work sets the most stringent limits on anomalous ν_τ appearance for $\Delta m^2_{41} \lesssim 3$ eV^2, including the strongest constraints around $\Delta m^2_{41} = 1$ eV^2.

As discussed at the end of the three-flavor analysis, the impact of systematic uncertainties can limit how precisely one can measure neutrino oscillation parameters. In the context of sterile neutrino analysis, systematic uncertainties such as detector calibration can degrade the sensitivity to sterile neutrino searches. Here, the NOvA Test Beam experiment could potentially reduce the calibration uncertainty. This is particularly essential for neutrino NC interactions, since the final state of an NC interaction is primarily composed of hadrons, including pions, which is the topic of this thesis.

4.2.11 Closing Thoughts on Oscillation Results and Uncertainties

This chapter presented the latest results on neutrino oscillations from the NOvA experiment. One key takeaway for future studies is the significant impact of statistical and systematic uncertainties on the determination of neutrino oscillation parameters, as discussed in Section 4.1.6. While statistical uncertainties can be reduced by collecting more data, systematic uncertainties are more challenging to minimize due to various factors, including a better understanding of neutrino interactions, detector response, and calibration.

As mentioned at the end of each oscillation analysis results. The NOvA collaboration deployed a Test Beam experiment to study the detector response to

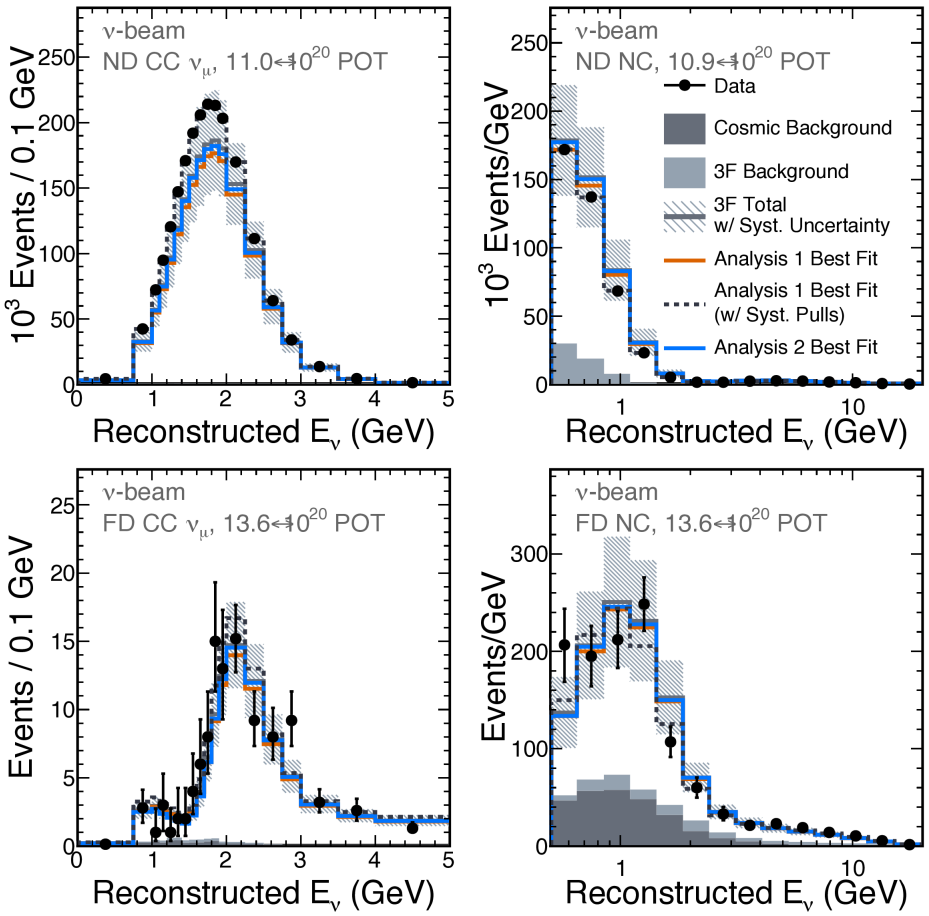

Figure 4.12: Spectra for the four samples used in the sterile neutrino analysis are presented here. The top plots display the spectra from the Near Detector, while the bottom plots correspond to the Far Detector. On the left, the charged-current muon neutrino interactions are shown, while the right-hand side illustrates neutral-current (NC) interactions. The data points are represented in black, with the best-fit results from Analysis 1 (PISCES) shown in orange and those from Analysis 2 (CMF) in blue. The dashed histogram shows the best fit for PISCES with systematic pulls applied

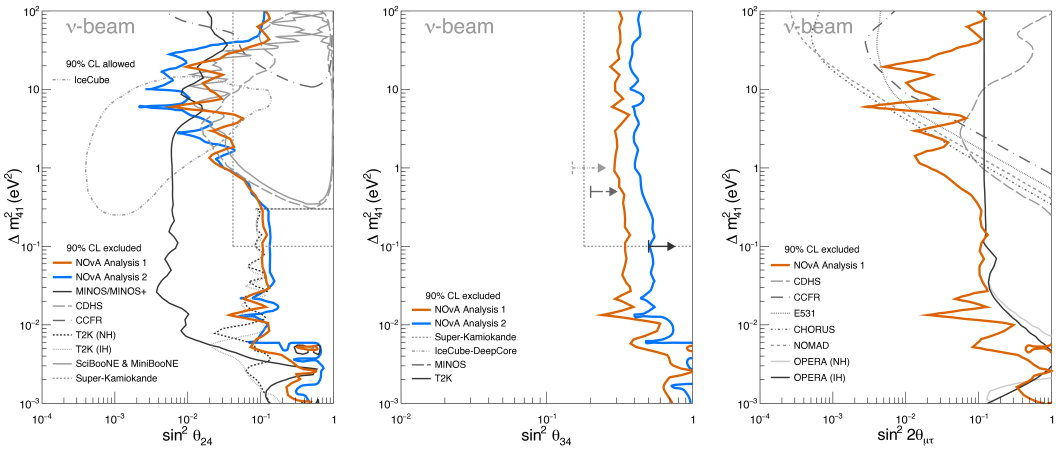

(a) (b) (c)

Figure 4.13: NOvA's Feldman-Cousins corrected 90% confidence limits are shown for (a) Δm^2_{41} vs. $\sin^2\theta_{24}$, (b) Δm^2_{41} vs. $\sin^2\theta_{34}$, and (c) Δm^2_{41} vs. $\sin^2 2\theta_{\mu\tau}$. Allowed regions and exclusion contours from other experiments are included, with areas to the right of open contours indicating excluded regions. In (a), closed contours for SciBooNE/MiniBooNE, CCFR, and CDHS also denote exclusion regions. For Super-Kamiokande, a single value of each mixing angle is reported at $\Delta m^2_{41} \geq 0.1$ eV^2. In (b), arrows indicate constraints on $\sin^2\theta_{34}$ at a specific Δm^2_{41} value, while in (c), OPERA NH/IH contours overlap for $\Delta m^2_{41} > 10^{-2}$ eV^2.

individual particles such as protons, pions, and electrons to address one of the largest systematic uncertainties, detector calibration. The following chapters will explore the NOvA Test Beam experiment and present the first results on the pion energy response.

Chapter 5

The NOvA Test Beam Experiment

The NOvA Test Beam experiment uses a scaled-down 30-ton NOvA detector to analyze tagged charged particles. A new tertiary beamline (Figure 5.1) was deployed at Fermilab to select and identify electrons, muons, pions, kaons and protons in the momentum range from 0.4 to 1.4 GeV/c . The data from the NOvA Test Beam Experiment is providing NOvA with a better understanding of some of the largest systematic uncertainties impacting NOvA's analyses, such as detector response and calibration.

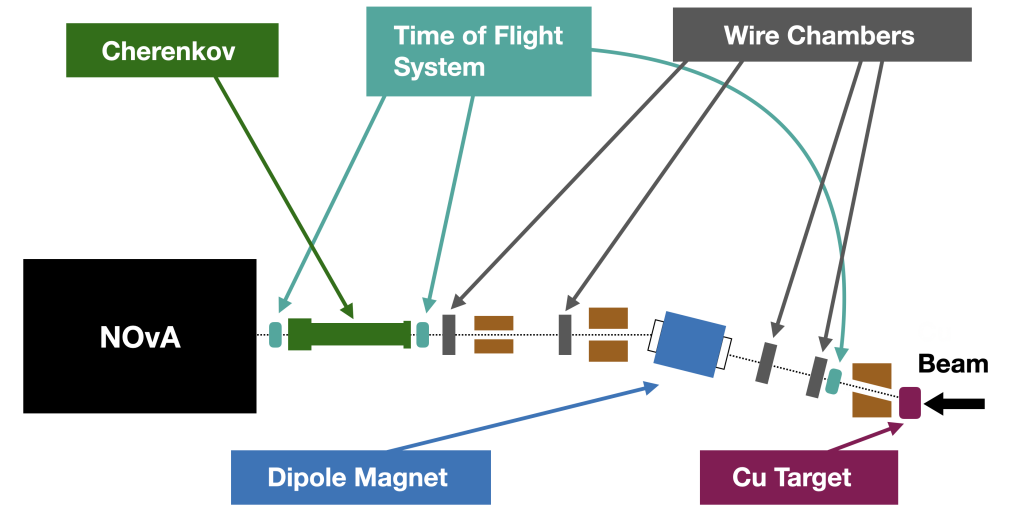

Figure 5.1: Diagram of the NOvA tertiary beamline components and NOvA detector, particle beam moves from right to left (Top). Panoramic view of the NOvA Test Beam experiment (Bottom).

5.1 The Beamline

The Fermilab Accelerator Complex delivers high-energy protons to multiple locations, including NOvA and the Fermilab Test Beam Facility (FTBF). The main Injector provides FTBF (Figure 5.2) with 120 GeV/c protons, the beam is extracted over 375000 turns of the main injector via resonant extraction and

consists of about 81 bunches, each space at 53.1 MHz, which corresponds to 18.83 ns between individual bunches. The bunches are grouped into 1.6 μs burst, and the bursts are spaced 11.2 μs apart. This pattern is repeated 93750 times over 4.2 s before the next spill begins (1 each minute). Typically, each spill delivers between 1×10^9 and 1×10^{10} protons to the primary target.

Figure 5.2: An Aerial view of the Fermilab Test Beam Facility (FTBF) is shown. FTBF has hosted numerous test beam experiments, including NOvA, MINERvA and LArIAT, among others.

protons extracted from the Main Injector interact with a primary copper target located upstream of FTBF in the MC6 enclosure (see Figure 5.3). These interactions generate a secondary beam primarily composed of protons and pions, with a momentum range spanning 1 to 80 GeV/c. The dipole and quadrupole magnets deployed in the beamline, along with shielding blocks and collimators in the enclosure, are responsible for focusing and steering the secondary beam. For the NOvA Test Beam Experiment, the secondary beam was tuned

to provide particles of 64 GeV/c for optimal tertiary particle production.

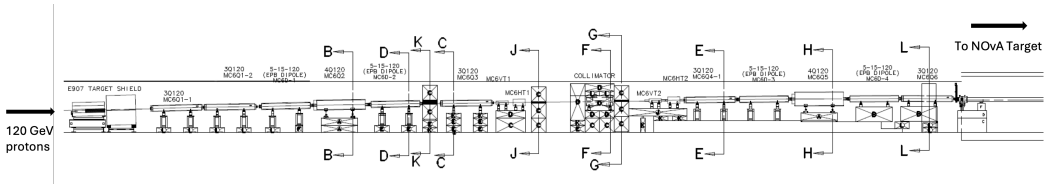

Figure 5.3: Technical drawing (side view) of the secondary beamline (MC6). The beamline instrumentation is tuned to focus 64 GeV/c particles for the NOvA Test Beam experiment. The proton beam enters from the left.

The 64 GeV/c beam interacts with the NOvA copper target in the MC7 enclosure (see Figure 5.4), producing a tertiary beam consisting of protons, kaons, pions muons and electrons. Instrumentation was deployed along the tertiary beamline to identify and select these particles, the main components of the tertiary beamline will be presented in the following sections.

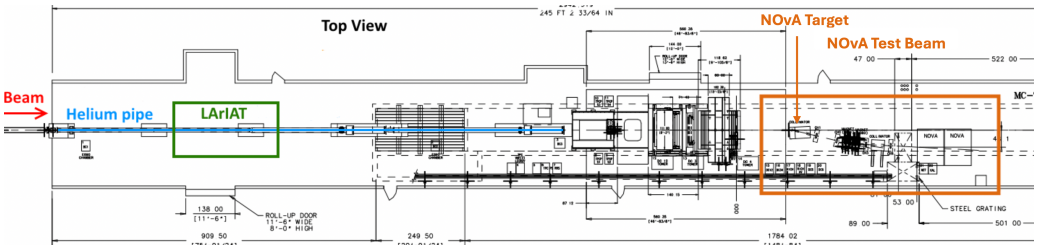

Figure 5.4: Technical drawing of the MC7 enclosure (top view) showing the NOvA Test Beam experiment and target location.

5.1.1 Target and Collimator

The first two components of the tertiary beamline are the copper target and an iron collimator. The target had a rhombohedral shape, 28 cm long, 25 cm tall, and a maximum width of 3.5 cm (Figure 5.5). Tertiary beam particles coming out of the target were absorbed and degraded by the iron collimator. The collimator assembly had an aperture that allowed only particles with a 16° trajectory to exit. The angle of collimation prevented secondary particles from being collinear with the tertiary beam.

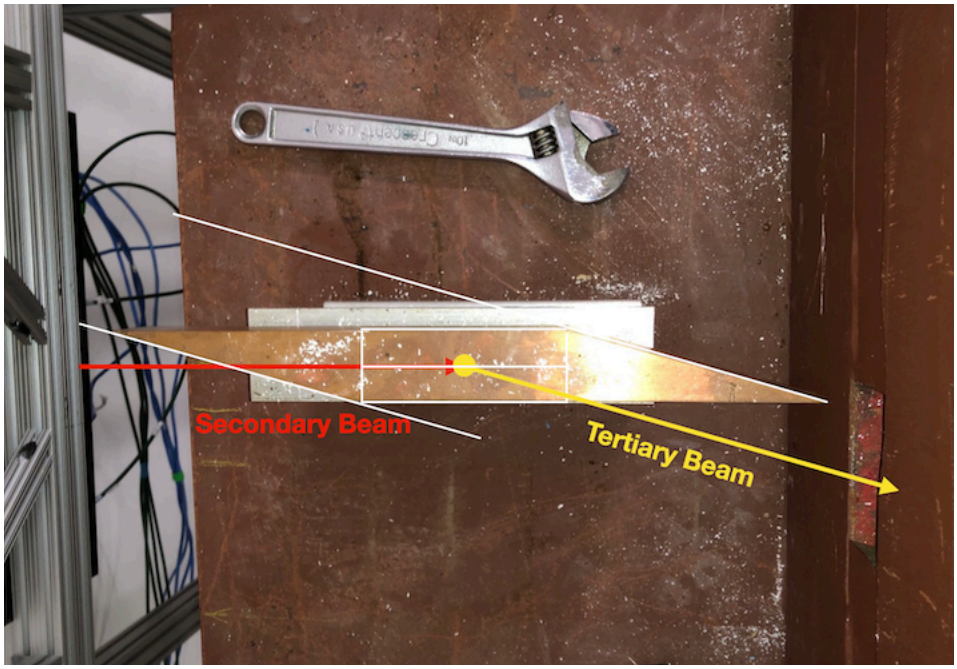

Figure 5.5: Top view of the copper target and downstream iron collimator. The collimator includes an aperture aligned at 16°, allowing only particles exiting the target at that angle to pass and preventing secondary particles from being collinear with the tertiary beam.

5.1.2 Time Of Flight (TOF) System

The system consists of three scintillator paddles: one located upstream of the beamline (US TOF) and two downstream (DS TOF1 and DS TOF2). These paddles provide precise timing information for particles traversing the beamline.

Two configurations were used during the experiment. In the first year of operation, the TOF was measured using the DS TOF1 and DS TOF2 paddles. The dimensions of these paddles were $0.6 \times 15 \times 15$ cm^3 for the US TOF and DS TOF2, and $2 \times 15 \times 15$ cm^3 for DS TOF1. In this setup, particles traveled a distance of approximately 13.16 meters between the TOFs along the beamline. It was later found that approximately 40 percent of the reconstructed beamline particle tracks were missing the DS TOF2 paddle due to scattering in the beamline material, including the thicker DS TOF1.

To address this, the DS TOF1 paddle was moved upstream to replace the US TOF and was renamed US TOF. The original US TOF paddle was relocated downstream and became the new DS TOF1. This new configuration used the US TOF (formerly DS TOF1) and DS TOF1 (formerly US TOF), with particles traveling a distance of 9.71 meters between them. The thinner DS TOF1 paddle, now in the downstream position, was placed between MWPC4 and the Cherenkov detector. This configuration proved effective in recovering the previously lost particles and was employed for the two subsequent years of operation.

5.1.3 Multi-Wire Proportional Chambers (MWPCs)

MWPCs, commonly known as wire chambers, are designed to accurately measure the position of charged particles, minimize the amount of matter in the path of the beam, and measure momentum using the dipole magnet. For the NOvA Test Beam experiment, four wire chambers were placed along the tertiary beamline: two before and two after the dipole magnet, to track particles before and after magnetic field deflection.

The wire chambers are composed of two sense planes (X, Y), with 128 gold-tungsten wires per plane, spaced 1 mm apart, where the diameter of each wire is 10 μm. These sense planes are located between three high-voltage aluminum cathode planes with a thickness of 12.5 μm. G10 spacers are used to separate each sense and cathode plane by 3.175 mm. To seal the assembly, 12.7 μm aluminum foils were used. Each chamber has an active area of 15.24 $\times$ 15.24 cm^2.

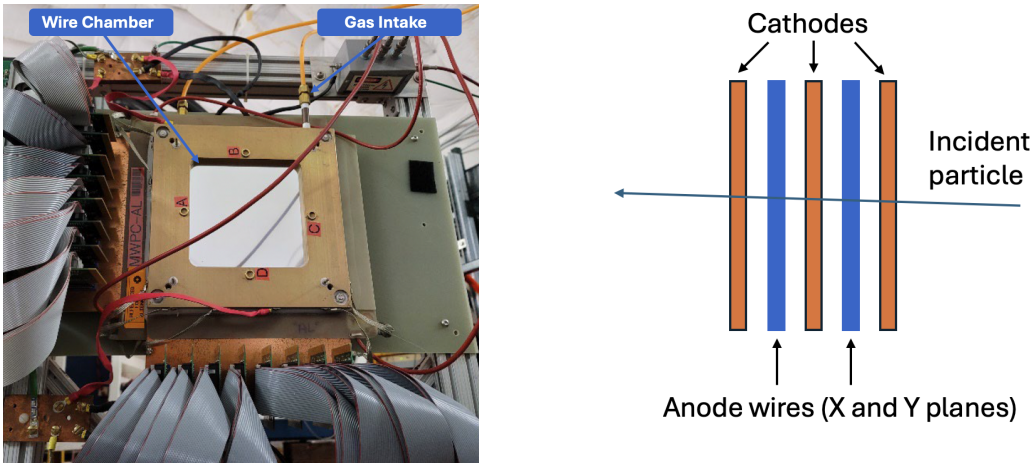

Figure 5.6: (Left) Front view of the NOvA Test Beam wire chamber. (Right) Side view representation of wire chamber with its key components (Right)

Each wire chamber was filled with 85% argon and 15% isobutane gas mixture and operated with a high voltage of ∼2.4 kV. This design provides high efficiency and precision for measuring particle location in both X and Y directions. A picture and a schematic (side view) representation of a wire chamber can be seen in Figure 5.6.

5.1.4 Dipole Magnet

The Dipole magnet is used to select the momentum and charge of the tertiary beamline particles. The active region of the magnet has an aperture of 45.1 cm (width), 8.9 cm (height) and 106.7 cm (length). It was configured to operate with currents ranging from 500 A to 1250 A, for both positive and negative polarities, a picture of the dipole magnet can be seen in Figure 5.7.

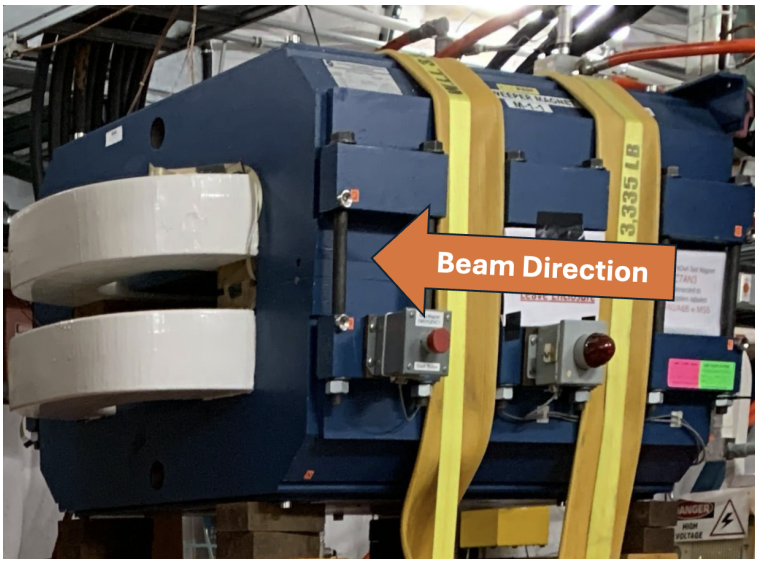

Figure 5.7: Picture of dipole magnet in the tertiary beamline for the NOVA Test Beam experiment, the arrow shows the direction of the beam.

The relationship between the average magnetic field in the active region of the magnet (B_{eff}) and the current (I) can be expressed as Equation 5.1 [22]:

$$B_{eff} = -0.0294 + \frac{|I|}{985.3} - \left(\frac{|I|}{3451.2}\right)^2. \tag{5.1}$$

The parameters in Equation 5.1 were obtained by fitting the curve to the measure data points shown in Figure 5.8.

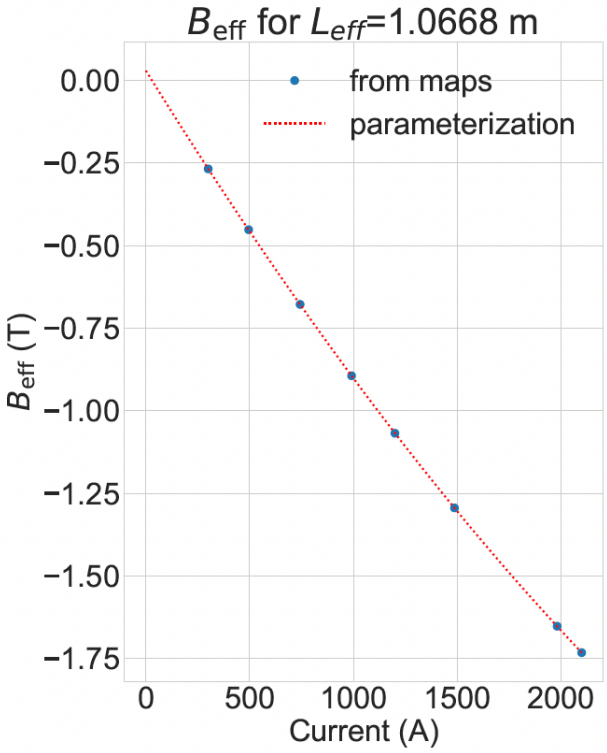

Figure 5.8: Effective Magnetic field (B_{eff}) as a function of input current (I), the red curve shows the fitted model, whereas the blue points represent the measured data [22].

5.1.5 Cherenkov Counter

In the NOvA Test Beam experiment, the Cherenkov counter plays a crucial role in distinguishing electrons from other particles by using Cherenkov light. The system consists of two stainless steel pipes connected to a L-shaped structure. A mylar mirror focuses light into photomultiplier tube (PMT) mounted at the short end of the L (Figure 5.9). To produce Cherenkov radiation, the counter is filled with CO2 at 1 atm pressure and placed downstream the tertiary beamline.

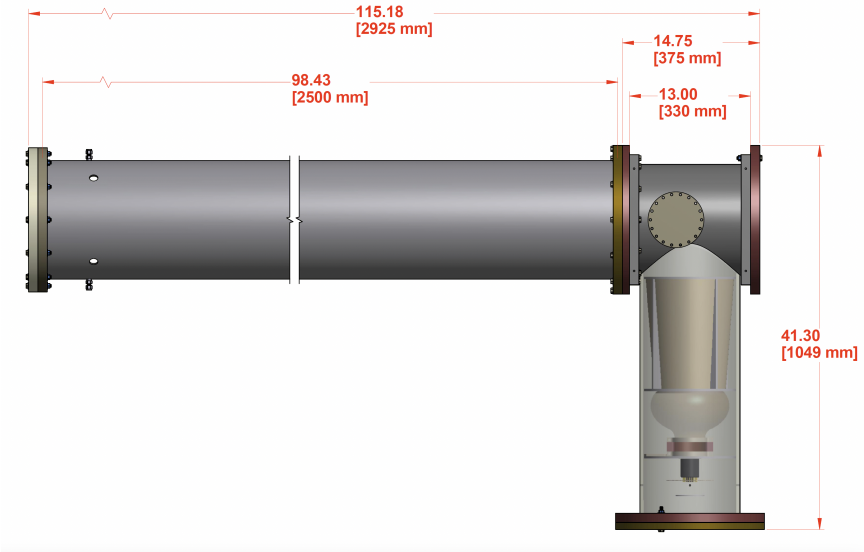

Figure 5.9: Diagram and dimensions of the Cherenkov counter

Cherenkov radiation occurs only when the charge particle's velocity exceed the speed of light in a given medium. This velocity threshold is determined by the ratio of the speed of light in the vacuum (c) to the refractive index (n) of the medium (c/n). For CO2, the refractive index depends on the gas pressure

(P), as described by Equation 5.2:

$$n = 1 + Pk \tag{5.2}$$

Here $k = 4.1 \times 10^{-4}/atm$ and P is the gas pressure. The minimum pressure ($P_{threshold}$) for CO2 to produce Cherenkov light in function of the particle momentum p is expressed by Equation 5.3:

$$P_{\text{threshold}} = \frac{\sqrt{(mc)^2 + p^2} - p}{kp} \tag{5.3}$$

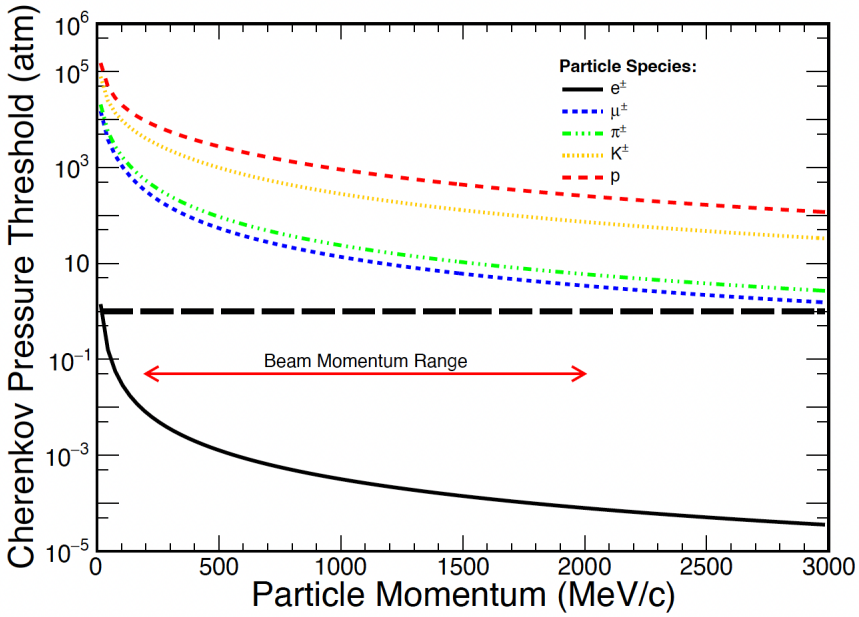

Figure 5.10: Pressure threshold curves for Cherenkov light emission as a function of particle momentum for different particle species. The dashed black line marks the operating pressure of 1 atm, chosen to reliably tag electrons within the NOvA Test Beam momentum range, shown by the red arrow.

The graphical representation of the pressure threshold in function of particle momentum in the Cherenkov counter can be seen in Figure 5.10. The relationship was critical to determine the optimal pressure for the gas inside the counter and a pressure of 1 atm was selected, this ensures reliable electron tagging by Cherenkov light in the NOvA Test Beam momentum operation range.

The components mentioned above are key to the NOvA Test beam experiment. In the next sections, operations , data taking procedures, DAQ and triggers will be presented. Then we will explore how the experiment systems work together to reliably tag particles of different species and momenta in the beamline.

5.2 Operations and Data Taking

The NOvA Test Beam experiment collected data over four different run periods, which can be classified as commissioning, optimization, and full data-taking phases. During the run, the magnet's magnetic field was set to different currents and polarities, enabling the collection of particles with different charge and momentum ranges. A summary of operations can be shown in Table 5.1.

The commissioning phase focused on verifying the experimental setup. During this stage, the team ensured that the beamline, detector, and data acquisition were working properly. Although the data quality was relatively low, this phase was essential to understand the adjustments that were needed to improve data collection and quality in future runs.

Period	Magnet Current (A)	Approximate Mean Momentum (MeV/c)	Particle Charge (q)
Commissioning (29 Mar 2019 – 7 Jul 2019)	250 to 2000 Testing	250 to 2000 Testing	+1 ±1
2020 Run: "Period 2" (5 Dec 2019 – 20 Mar 2020)	+500 +1000	500 1000	+1 +1
2021 Run: "Period 3" (12 Jan 2021 – 27 Jun 2021)	+500 ±750 +1000	500 750 1000	+1 ±1 +1
2022 Run: "Period 4" (30 Nov 2021 – 10 Jul 2022)	±500 ±750 ±1000 ±1250	500 750 1000 1250	±1 ±1 ±1 ±1

Table 5.1: Overview of the data taking operations showing the magnet current settings, the approximate particle mean momentum, and the corresponding particle charge across periods.

One of the primary challenges encountered during commissioning and period 2 was a large amount of background originating in the secondary beamline (subsection 6.2.1). This background saturated the NOvA detector's electronics and DAQ systems, limiting the ability to collect high-quality data (subsection 6.2.2). By the end of period 2, shielding blocks were placed along the beamline to mitigate the secondary background effectively.

During Period 3, extensive efforts were dedicated to improving the beam quality by tuning the angle of the beam at the primary target, selecting the configuration of the momentum collimator on the secondary beamline, and understanding sources of background that could interfere with optimal operations. Another major improvement during period 3 was the top-off of the under-filled cells in the detector (subsection 6.2.3). All these improvements were applied and set as default for the remainder of the run.

5.3 Beamline DAQ and Timing Systems

The beamline data acquisition system (DAQ) collects and records data from beamline instrumentation before it reaches the NOvA detector. It works separately from the NOvA detector's DAQ (described in subsection 5.7.1) but stays synchronized using timestamps.

The system includes a VME crate with a CAEN V1495 trigger board, a CAEN V1742 digitizer, and a CAEN V2718 controller that connects to the DAQ server via a CAEN A3818 PCIe card. The digitizer processes signals from the ToF and the Cherenkov detectors, while a custom controller handles wire chamber data. This setup ensures accurate and synchronized data collection for analysis.

Signals are processed using electronics and cables before being recorded. When a trigger occurs, an electronic signal is sent to start data readout and timing distribution. Two TDUs ensure synchronization. One triggers the NOvA detector, while the other timestamps beamline data, allowing both data streams to be matched during offline processing.

5.4 Beamline Trigger

The NOvA Test Beam detector (section 5.7) recorded data whenever it received a beam trigger, which depended on three factors: signals from dedicated trigger counters, alignment with the beam window, and availability of the readout electronics.

Each trigger counter had a PMT connected to a plastic scintillator plate (3.1 mm thick, 10.1 × 10.1 cm in size). Four counters were placed along the beamline near the wire chambers. Initially, all four needed to register a hit for a trigger, but for the final run, the fourth counter was ommiteed after tests showed it was lowering trigger efficiency.

A NIM gate generator controlled the 4.2-second beam window, using signals from the Fermilab accelerator clock to mark the start and end of each beam spill. During data collection, trigger rates ranged from 1 to 10 per spill across the last three runs.

To keep things running properly, the electronics had to meet two conditions before accepting a new trigger: the digitizer processing ToF and Cherenkov detector data couldn't be in a 'busy' state from the last trigger, and a brief 10 µs delay ensured the front-end components, especially the wire chamber controller, had enough time to reset properly.

5.5 Beamline Momentum Reconstruction

The momentum of particles traversing the beamline is estimated by reconstructing their tracks using the four MWPCs and measuring their deflection caused by the field in the dipole magnet. The first stage of track reconstruction involves identifying the hits corresponding to particles interacting with the wire chambers, followed by reconstructing the beamline track and finally estimating the momentum. The next sections will explain in detail the procedure.

5.5.1 Wire Chamber Hit Reconstruction

To achieve this task, a hit-finding algorithm is implemented, inspired by the one used by the LArIAT experiment [96]. The algorithm saves the triggered signals in a hit list array containing spatial and temporal data information. Due to ionization effects, it is possible that particles traversing the wire chambers hit multiple wires in the wire chambers.

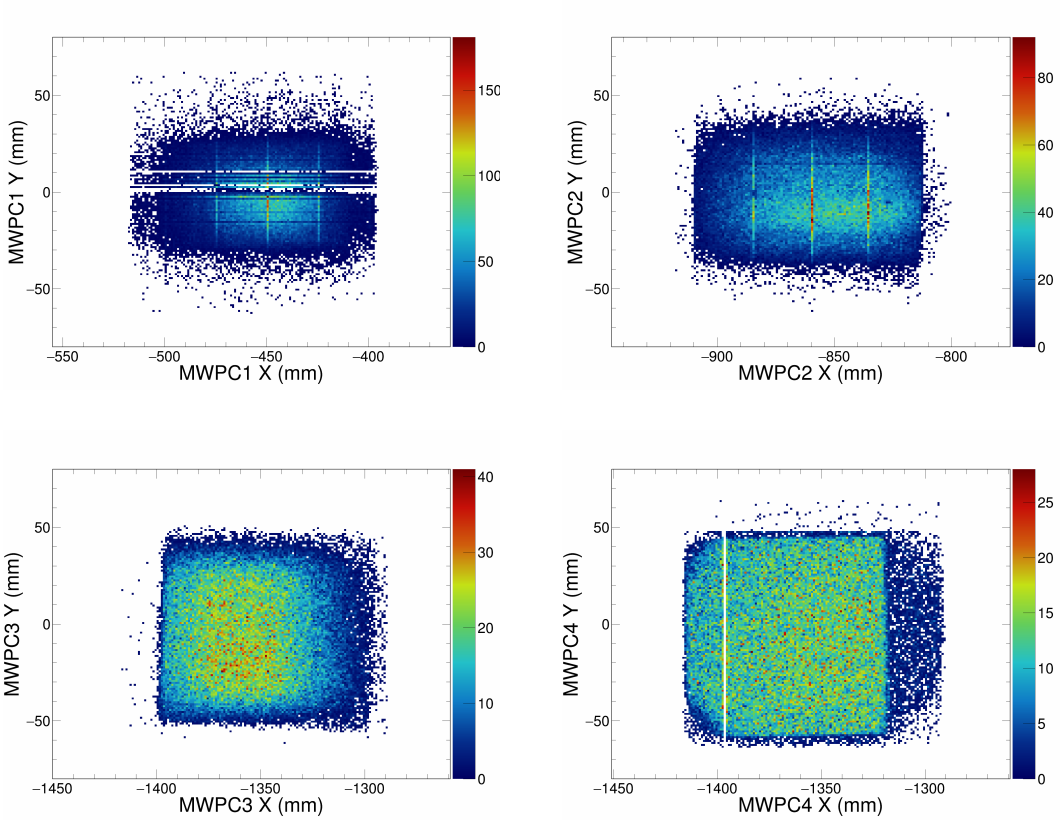

Figure 5.11: Reconstructed hits for each of the four wire chambers, each hit represents a particle traversing the wire chamber detectors in the tertiary beamline, each hit position was determined by the DBSCAN algorithm.

These hits are then passed to a DBSCAN algorithm [97] (Density-Based Spa-

tial Clustering of Applications with Noise), which groups hits into clusters based on their proximity. Since each cluster represents a particle interaction within the wire chambers, the earliest hit from a cluster from the DBSCAN determines the particle's position through the wire chamber.

An important condition for forming a valid hit is that there must be a signal on at least one vertical and one horizontal wire. For instance, if there are two vertical wire signals and one horizontal wire signal, this results in two hits. in general, the number of reconstructed hits in a wire chamber is given by the product of horizontal (nx) and vertical (ny) signals, such that $nhit = nx \times ny$ for any given wire chamber. Figure 5.11 shows hits of particles traversing the wire chambers.

5.5.2 Beamline Track Reconstruction

For every combination of reconstructed hits, a straight line in the y-z plane (non-bending view) is determined by applying a least squares fit to a straight-line, the chosen track is the one found by minimizing the residuals from the fit in Equation 5.4.

$$y_i = a + bz_i + e_i \tag{5.4}$$

Where a and b are the slope and intercept, which are determined by the fit using the hit position in each wire chamber. The residual, e_i, is the distance between the best-fit line and the hit i in the y direction. The intercept can be

found using Equation 5.5:

$$a = \frac{\sum_i y_i - b \sum_i z_i}{n} \tag{5.5}$$

Here n is the number of wire chambers and the sum is over the number of hits i. The slope is given by Equation 5.6:

$$b = \frac{n \sum_i z_i y_i - \sum_i z_i \sum_i y_i}{n \sum_i z_i^2 - \left(\sum_i z_i\right)^2} \tag{5.6}$$

the residual is evaluated for every y_i, z_i pair, and the goodness of fit e_{avg} is obtained by taking the average of the square root of the residuals.

$$e_{\text{avg}} = \sqrt{\frac{\sum_i e_i^2}{n}}. \tag{5.7}$$

The particle track is the one in which e_{avg} is the smallest.

5.5.3 Beamline Momentum Estimation

The momentum of a particle p traversing a dipole magnet of length L can be estimated using the transverse momentum p_T in Equation 5.8:

$$p = \frac{p_T}{\cos \phi} \tag{5.8}$$

Here ϕ is the angle formed by p and p_T in the y-z plane as seen in Figure 5.12

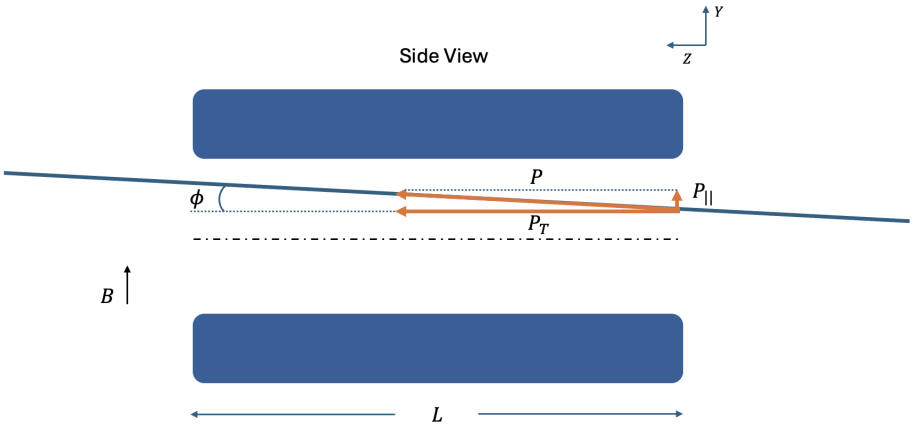

Figure 5.12: Side view representation of a particle passing through the magnet, the image shows the relationship between p_T and p.

The Lorentz force for a charged particle with a velocity v in a magnetic field B in the transverse direction is given by Equation 5.9.

$$F = qvB \tag{5.9}$$

Since the particle track bends in the magnet, the force can be expressed as the centripetal force, thus equaling the forces Equation 5.10 is obtained.

$$\frac{\gamma m v^2}{R} = qvB \tag{5.10}$$

And can be expressed as:

$$p_T = qBR \tag{5.11}$$

Equalling Equation 5.11 to Equation 5.8 and solving for p:

$$p = \frac{qRB}{\cos(\phi)} \tag{5.12}$$

The radius of the curvature of the reconstructed track R depends on the length of the magnet L, the angle of entry θ_{in}, and the outgoing angle θ_{out} of the two ends of the track as observed in Figure 5.13.

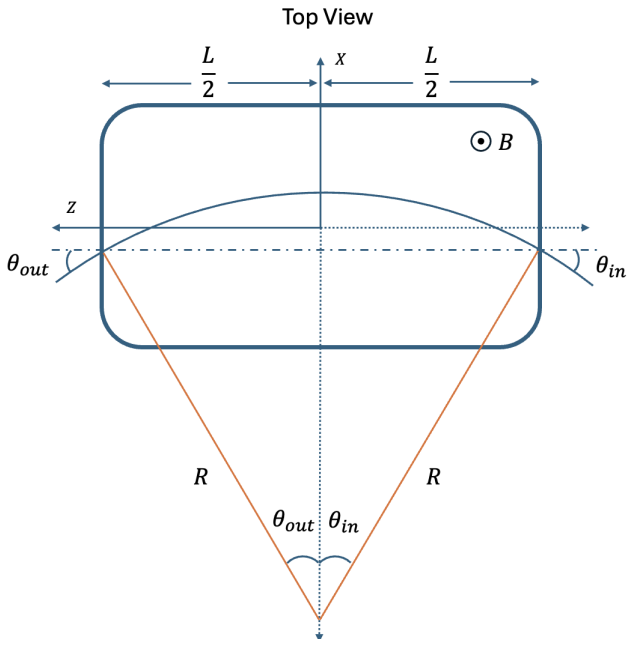

Figure 5.13: Top view representation of a particle passing through the magnet, the image shows the relationship between R, L, θ_{in} and θ_{out}.

Therefore, R is equal to:

$$R = \frac{L}{\sin\theta_{\text{in}} - \sin\theta_{\text{out}}} \tag{5.13}$$

Replacing Equation 5.13 in Equation 5.12 leads to the final expression of the reconstructed momentum of a particle traversing the beamline:

$$p = \frac{qBL}{[\sin(\theta_{in}) - \sin(\theta_{out})]\cos(\phi)} \tag{5.14}$$

Where B is equal to the effective magnetic field (B_{eff}) from Equation 5.1.

5.6 Time of Flight Estimation

The estimation starts with determining the baseline of the PMTs' digitized signal. This is achieved by calculating the mean and sigma iteratively across raw data samples. The iteration converges when the differences between the means of the samples are less than 2 ADC. The final mean and sigma are used as the baseline and noise band, respectively.

Once the baseline and noise band are determined, the Savitzky-Golay Filter [98] is used to smother the raw data signals (pulses), this helps to suppress high-frequency noise using a polynomial of order 3 and a window size of 15 points, ensuring minimal errors maintaining important features of the pulse, such as the start, peak, and width.

After filtering, the pulses are identified and reconstructed, this is achieved by establishing a threshold of 50 times the noise band, making sure that at least four points in the pulse are below this threshold ensuring a reliable signal.

Valid pulses are then analyzed, and for each of them, the start, peak, and end times are located. A Constant Fraction Discrimination (CFD) method is implemented to find the start point precisely, and a threshold of 0.7 % of the peak value of the pulse is used for the discriminator to minimize the width of the pulse. A third-order polynomial is implemented to fit the peak, this ensures a precise peak and amplitude determination.

Since each TOF has 4 PMTs, pulses that are close in time within a 5 ns window are saved into a cluster, from the collection of all the start times in the cluster, a Gaussian fit is performed to find the mean time, allowing the reduction of the statistical error by half (since the error on the mean is equal to $\sigma/\sqrt{N}$).

The final reconstruction step uses the mean times for both US TOF and DS TOF to estimate the time of flight using the Equation 5.15:

$$TOF = t_{DS} - t_{US} \tag{5.15}$$

5.6.1 Time of Flight Resolution

A Gaussian fit to the TOF distribution is performed for pions, protons, and electrons. The resolution for pions is 0.19 ns, for electrons 0.18 ns, and for

protons 0.8 ns, the TOF distributions from different momentum ranges along with the fit curves can be seen in Figure 5.14.

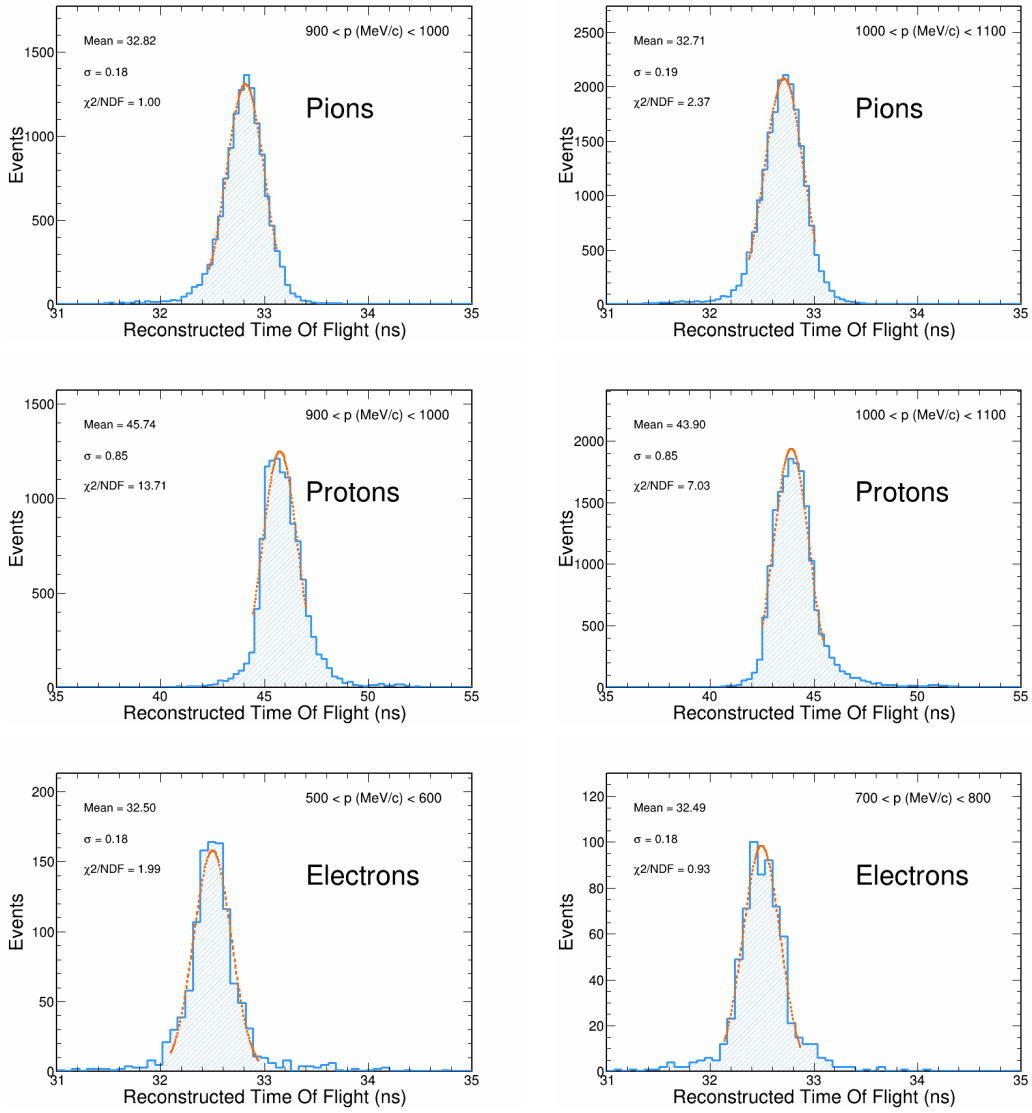

Figure 5.14: Time of flight distributions for pions, protons, and electrons. A Gaussian fit was performed on the data and the resulting curves can be seen in orange, the resolutions are given by the sigma in each plot.

5.7 NOvA Test Beam Detector

As discussed in section 2.2 the NOvA Test Beam detector is part of the set of detectors for the NOvA experiment. Specifically, this detector is used to study charged particle interactions tagged by the tertiary beam instrumentation. The detector is functionally identical to the ND and FD, but with a smaller footprint. A picture of the NOVA Test Beam detector can be seen in Figure 5.15 inside the FTBF in the MC7 enclosure.

Figure 5.15: Picture of the NOvA Test Beam detector in the MC7 enclosure.

The Test Beam detector is composed of 63 planes alternating between vertical

and horizontal cell orientations, each cell is 2.6 m long forming a 2.6 m height, 2.6 m width and 4.1 m in length detector with a total mass of 30 tons. This design was chosen to effectively contain the energy depositions of pions and electrons based on detector simulations involving charged particles in the 0.4 to 1.8 GeV/c momentum range.

The NOvA Test Beam detector was filled with a combination of two scintillator blends. One blend was repurposed from the NOvA NDOS prototype after decommissioning, while the other included a mix of scintillator stored at the University of Texas, Austin, and leftover FD scintillator from Ash River, MN. Differences in brightness between the blends were calibrated following established procedures (following section). Before filling, scintillator samples were tested for optical transmission to ensure minimal suspended impurities, meeting the NOvA-required >95% threshold.

5.7.1 Detector DAQ and Triggering

The DAQ system for the Test Beam detector is a modified version of the one used for the ND and FD (details in subsection 2.2.7), designed for continuous readout with buffered data processing.

Each APD connects to a FEB that amplifies, processes, and digitizes signals using a 32-channel ASIC before sending them to an ADC. Due to differing pileup rates, NOvA uses two FEB designs: one for the FD, which handles lower pileup, and another for the ND, which experiences higher pileup due its proximity to the NuMI beam.

To study the differences between FD and ND performance, the Test Beam detector includes both FEB types, with 118 of 126 being FD-type and the remaining 8 being ND-type.

The NOvA DAQ system can handle multiple triggers at the same time, with each one storing data from front-end buffers in its own dat file. If an event meets more than one trigger condition, it gets recorded in multiple data streams.

The main data stream for test beam data comes from the beamline trigger (section 5.4). There are also activity-based DDTs that collect data outside the beam spill, including cosmic ray interactions used for calibration (next section). Another important stream, the "beam spill" stream, saves all data recorded during a beam window, allowing analysis of the NOvA detector's status and its response to beam-induced backgrounds.

5.7.2 NOvA Test Beam Detector Calibration

The NOvA Test Beam detector uses the same calibration process as the ND and FD as described in section 3.2. The calibration ensures accurate energy measurements by correcting for light attenuation in fibers, eliminating energy deposition differences, and establishing an absolute energy scale. This enables a direct comparison of deposited energy between the Test Beam detector and the NOvA ND and FD.

The data calibration of the detector follows the same principle explained in section 3.2, this is performed in all data periods. To calibrate the simulated

ND and FD, NOvA typically uses the a Monte Carlo generator to produce a sample of cosmic muons. These muons are then processed through the same reconstruction and selection criteria as real data.

However, to improve simulation efficiency, the NOvA Test Beam instead uses a data-driven approach [99]. In this method, cosmic muons from Period 4 data are selected (best detector conditions), and their reconstructed vertex positions and four-momenta are used to seed a Geant4-based simulation of the detector. The vertex is determined by finding the intersection of the cosmic muon's track with the detector's edge, while the four-momenta are refined using an adapted break-point fitter (BPF) method.

To account for variations in detector response, the simulation divides the detector cells into 12 fiber brightness bins, based on the uncorrected average response at their centers. This correction helps compensate for differences in individual scintillators and variations in fiber brightness across the detector.

The Test Beam's calibration procedure helps validate NOvA's calibration methods, improving estimates of systematic uncertainties and providing an independent energy scale measurement. Its smaller size reduces light attenuation effects, resulting in a more uniform energy deposition. Variations in scintillators and readout electronics allow for testing detector response under different conditions. The calibration process effectively removes inconsistencies between individual cells, ensuring a nearly uniform energy response across the detector.

5.7.3 NOvA Test Beam Detector Reconstruction

The NOvA Test Beam detector reconstruction is identical to that of the ND and FD, as described in section 3.3. The key difference lies in how vertexing is performed. In the ND and FD, identifying the neutrino vertex is more challenging since neutrinos can interact in different regions of the detector. In contrast, in the Test Beam detector, the particle vertex is easier to determine because tertiary beam particles interact on the detector's front face.

To determine the particle vertex, we project the reconstructed beamline track at wire chamber four onto the detector's front face. The vertex position is then passed to the Fuzzy-k algorithm, which identifies prongs and BPF tracks for further analysis. Figure 5.16 shows an event display of a pion candidate entering the NOvA Test Beam detector, illustrating the reconstructed prong and the projected vertex.

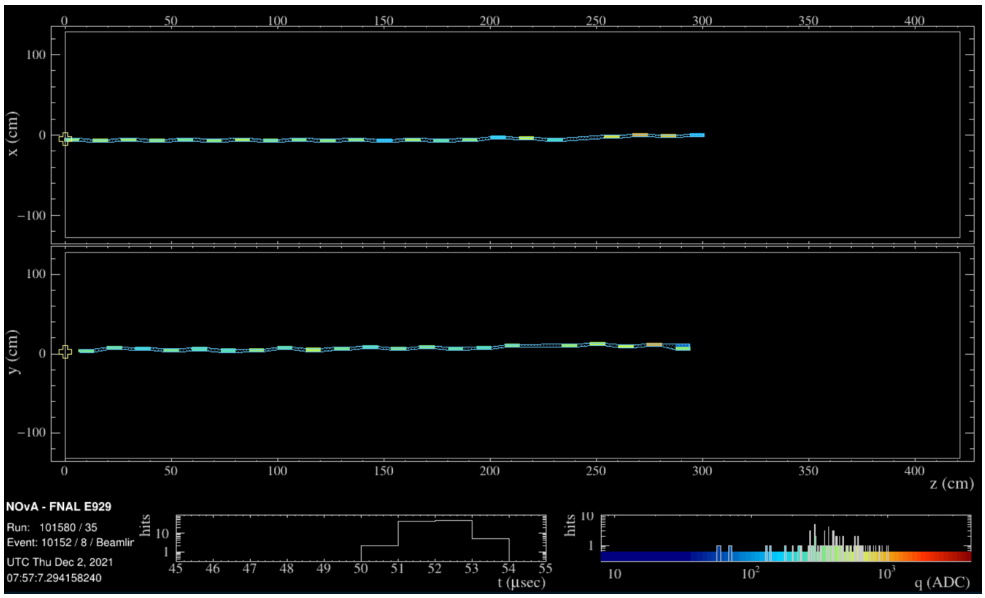

Figure 5.16: Event display of a pion candidate entering the NOvA Test Beam detector (from left to right). The reconstructed prong (blue) and vertex (yellow cross) are also shown.

Chapter 6

Pion Simulations and Data Selection in the NOvA Test Beam Experiment

This chapter outlines the simulation and selection tools used to select a clean pion sample in the NOvA Test Beam experiment. These samples form the basis for the energy response studies presented in Chapter 7.

The analysis is based on data collected during the NOvA Test Beam operations campaign described in section 5.2 and supported by simulated pion samples created by different simulation techniques which includes a detector simulation using a data-driven approach and a full beamline simulation.

The data-driven simulation provides a sample that resembles the actual data, while the beamline simulation delivers information about muon background

from pion decay and the the pion energy loss in the beamline material background estimation and energy corrections, which are essential for pion energy response studies.

The pion samples are carefully selected prioritizing high purity. To achieve this, several cuts are applied to reject other particles other than pions, keeping the background content as minimal. In addition, tools are implemented to reject pile-up, events affected by shut-offs and hardware limitations.

6.1 Test Beam Simulations

As mentioned in the introduction, simulations are an essential component of this analysis. They support the studies of particle interactions and pion energy response. In the NOvA Test Beam experiment, simulations are used as a tool to validate reconstruction, test simulation models, perform calibration in the NOvA detectors, conduct background estimation, and develop selection tools such as pile-up rejection.

For this analysis, two main simulation approaches are used for pion energy response studies: A full beamline simulation using G4Beamline, and a dedicated detector simulation using the standard NOvA simulation framework and reconstruction discussed in subsection 3.1.3 and subsection 5.7.3 respectively. The simulations are configured to address different tasks and are fully described bellow.

6.1.1 G4Beamline Simulation

G4Beamline [100] is a program based on Geant4. It is used for tracking and studying particle interactions in simulated beamlines. In the NOvA Test Beam Experiment, G4Beamline is used to simulate particle interactions traversing the tertiary beam instrumentation, study particle energy loss in the beamline material, and estimate the muon background in the pion sample.

In the simulation, each beamline detector is positioned using precise survey data. Figure 6.1 shows the simulation model of the tertiary beamline in G4Beamline. In order to perform the tasks mentioned above, the simulation is configured in two different modes, Mode 1 and Mode 2 simulations:

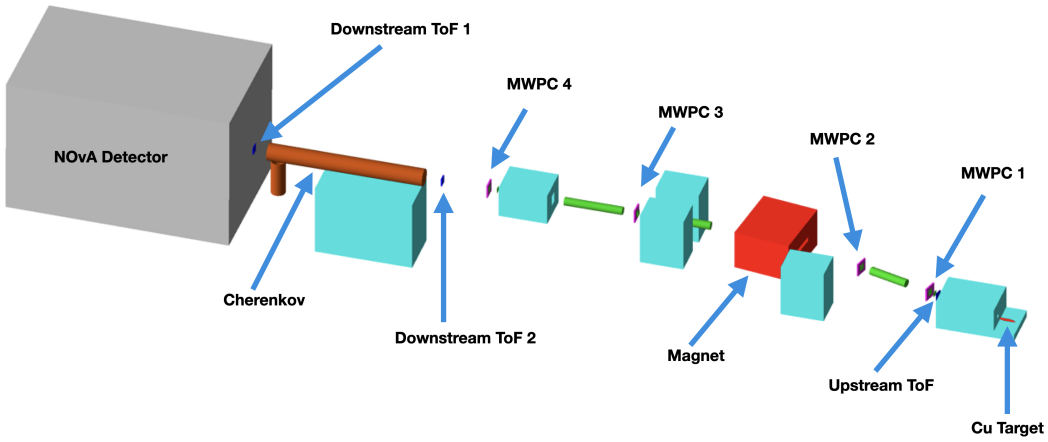

Figure 6.1: Layout of the NOvA Test Beam experiment as implemented in the G4Beamline simulation. The simulation includes all components of the tertiary beamline and the NOvA detector, including the Cu target, dipole magnet, time-of-flight (ToF) detectors, multi-wire proportional chambers (MWPCs), and the Cherenkov counter.

G4Beamline: Mode 1 Simulation

Mode 1 simulates the real conditions of the tertiary beamline, the simulation begins with 64 GeV/c protons interacting on the copper target. Particles produced by this interaction traverse the entire beamline until reaching the front face of the NOvA detector.

The true kinematic information of these particles is saved in the simulated files at each beamline detector for further reconstruction and analysis. The Mode 1 configuration allows to conduct studies about particle yield, estimation of systematic uncertainties such as particle energy loss and determination of the muon background in the pion sample (subsection 6.3.4). As an example, Figure 6.2 shows 10 protons on target simulated for Mode 1.

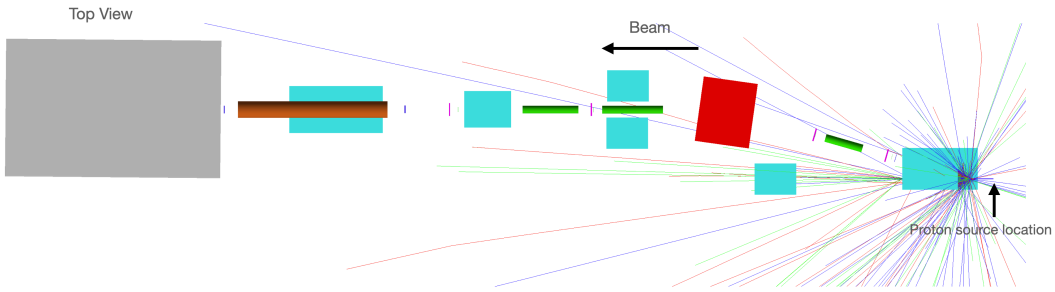

Figure 6.2: Top view of a Mode 1 G4Beamline simulation showing 10 protons with 64 GeV/c momentum interacting with the copper target. The simulation illustrates the resulting particle showers and beam propagation through the tertiary beamline toward the NOvA detector.

One of the key factors in the implementations of this simulation was to set the appropriate magnetic field inside the dipole magnet. This was achieved by transforming the magnet current distributions that were extracted from period 4 data into magnetic field distributions using Equation 5.1.

The resulting magnetic field distributions were fit to a Gaussian distribution, and the curves were used to set the magnetic field in the dipole magnet in G4Beamline by randomly sampling these distributions for each magnet polarity and current setting. An example of this implementation can be seen in Figure 6.3.

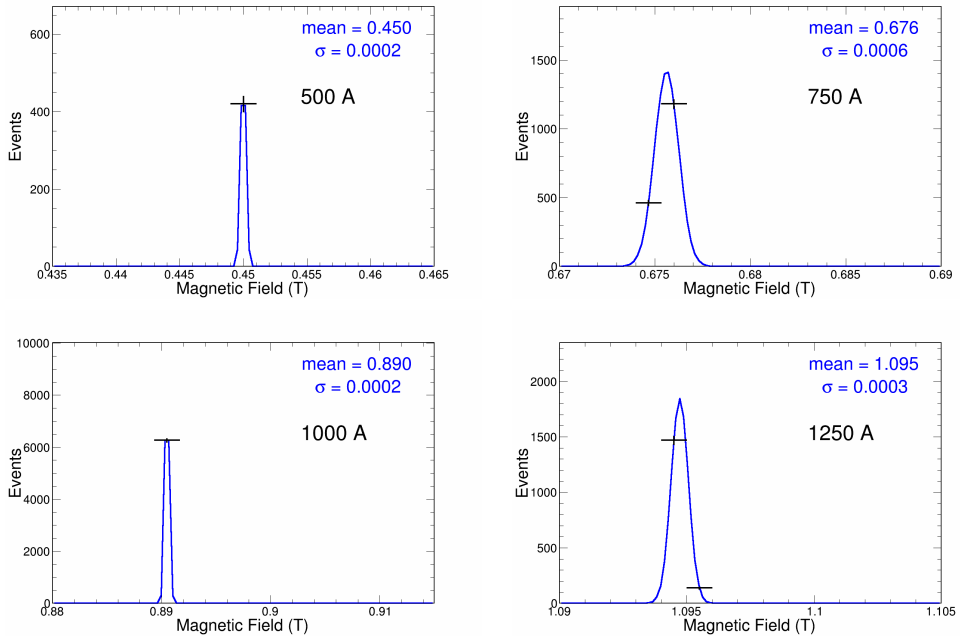

Figure 6.3: Examples of magnetic field distributions derived from magnet current data (positively charged particles in period 4) for different magnet current configurations (black), with corresponding Gaussian fits (blue). These fits were used to sample magnetic field values in the G4Beamline simulation for each magnet polarity and setting.

In this simulation campaign, several samples were created for each magnet polarity (positively or negatively charged particles) and magnetic field settings.Table 6.1 shows all the different magnet current configurations used in G4Beamline.

POT	Mean Magnet Current (A)	Mean Magnetic Field (T)
1×10^9	500	± 0.45
1×10^9	750	± 0.68
1×10^9	1000	± 0.89
1×10^9	1250	± 1.10

Table 6.1: Mode 1 simulation configuration in G4Beamline. Each configuration is run for positive (negative particles) and negative (positive particles) magnet polarities. In total, 8×10^9 protons on target were simulated, 1×10^9 for each configuration.

G4Beamline: Mode 2 Simulation

Mode 2 simulation is configured to run with a mono-energetic beam of a specific particle type such as pions. In the simulation, the particle gun is placed downstream of the target. Similar to Mode1, particles in Mode 2 traverse the entire beamline components until they reach the front face of the NOvA detector.

This simplified version of the simulation avoids the momentum spread created by particles originated by the proton interaction in the target which is a characteristic feature in Mode 1. In addition, Mode 2 allows for faster CPU time for both simulation and reconstruction.

Mode 2 is dedicated to performing studies of energy loss for particles traversing the beamline detectors as to correct the reconstructed beamline momentum for charge pions, an essential component of the energy response analyses. A detailed description of these studies along with results, is discussed in subsection 6.3.3.

Figure 6.4 shows the actual simulation of 100 π^+ for Mode 2. In this mode the magnetic field in the dipole magnet is configured similarly to Mode 1. It was further decided to add two extra samples to characterize the extremes of the particles momentum range: One at 300 A and another at 1500 A providing particles with momentum around 300 and 1500 MeV/c respectively. Table 6.2 shows the different G4Beamline configurations for Mode 2.

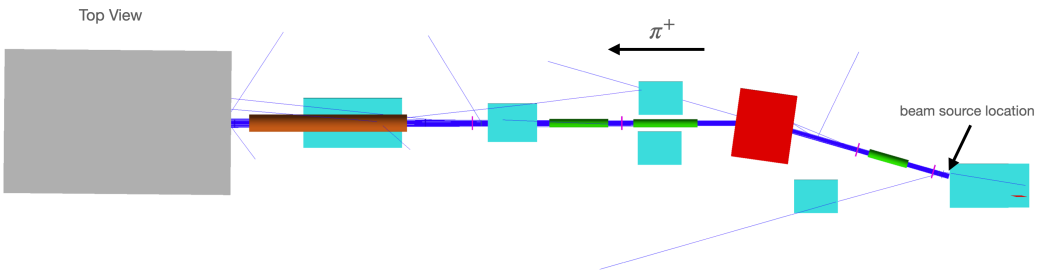

Figure 6.4: Top view of a Mode 2 G4Beamline simulation showing 100 π^+ with a momentum of approximately 1 GeV/c. In this mode, the pion source is placed downstream of the collimator to simulate tertiary beam particles without modeling the proton-target interaction.

Events	**Magnet Current (A)**	**Magnetic Field (T)**
1×10^6	300	$\pm$0.27
2.5×10^5	500	$\pm$0.46
2.5×10^5	750	$\pm$0.68
2.5×10^5	1000	$\pm$0.90
2.5×10^5	1250	$\pm$1.10
2.5×10^5	1500	$\pm$1.30

Table 6.2: Mode 2 simulation configuration in G4Beamline. Each configuration is run for positive (negative particles) and negative (positive particles) magnet polarities. In total, 1.25×10^6 pions were simulated.

6.1.2 Single Particle Simulation

Single particle simulation is a detector-level simulation performed with the NOvA simulation chain. It uses the full NOvA Test Beam detector geometry and readout simulation based on Geant4. The simulation uses Geant4 version 4.11.0 with the `QGSP_BERT_HP_EMZ` [65, 56, 67, 68] physics list. This configuration combines quark-gluon string models (QGS), binary cascade (BERT) models, high-precision neutron interactions (HP), and enhanced electromagnetic processes (EMZ).

The simulation samples are created using a data-driven approach, where real data events passing the final selection are used to extract kinematic properties (momentum, position, angles). These properties are then used as inputs (seeds) to generate simulated charged pions and muon backgrounds produced by pion decay. It also serves to study secondary muon background (subsection 6.2.1) contamination, using single particle simulation plus overlays, as described in subsection 6.1.3. This ensures that the simulation resembles the real beam data.

To reduce statistical uncertainties in the simulation, large samples are generated by reusing the seeds and smearing their position, angle, and momentum of each particle multiple times. By using this method, the sample sizes are increased by up to a factor of 160.

The simulated pion–muon samples produced by this approach serve as the primary method for performing data–simulation comparisons, providing insight into how well the current NOvA simulation models reproduce the real pion–muon data.

6.1.3 Single Particle Simulation with Overlays

The single particle simulation plus overlay sample is created by merging the single particle simulated events with randomly selected out-of-time beam spill data. The additional activity includes background originated upstream in the beamline such as high energy muons. This approach reproduces the pile-up and noise conditions observed during data-taking. This sample is used to test and develop the pile-up rejection algorithm described in subsection 6.2.4.

Each simulation sample plays a specific role in the analysis. The G4Beamline simulation is used to estimate energy loss and muon backgrounds from pion decay, while the single particle simulation and single particle simulation plus overlay samples are used to study the NOvA detector's response and develop algorithms to reject events with background and detector hardware issues. These simulations provide the foundation for the beamline momentum correction, validations, and background rejections described in the following sections.

6.2 Data Quality and Preselection

With the simulation tools in place for both the beamline and detector, the next step in the analysis is to ensure the quality and reliability of the data collected during the Test Beam campaign.

Several factors can affect data quality, including secondary beam-related backgrounds, detector shut-off, pile-up, and hardware issues such as under-filled cells in the horizontal modules. These effects can introduce unwanted back-

ground into the pion sample and bias the reconstructed energy, leading to distortions in data–simulation comparisons.

This section outlines the steps taken to identify and reject events affected by these data quality issues, ensuring that only high-quality, well-reconstructed events are used in the final pion energy response analysis.

6.2.1 Secondary Background

In the NOvA Test Beam experiment, one of the major sources of background was caused by multiple scattering of primary beam particles off materials in the secondary beamline. This scattering produced a collinear secondary beam that followed a trajectory outside the intended tertiary beam path, significantly increasing the background in the detector. The upper west side of the NOvA detector, which intersects the predicted path of secondary particles, was the most affected by this background, as shown in Figure 6.5a.

A dedicated simulation campaign using G4Beamline was conducted to investigate this background and study possible mitigation strategies. The simulation included all primary and secondary particle interactions in both the secondary and tertiary beamlines. The results showed that the most effective solution was the installation of 9 feet of concrete shielding along the secondary beamline. The impact of this shielding is shown in Figure 2.8, which compares simulated beam profiles before and after the shielding was applied.

Following the results of these studies, the concrete blocks were installed in the experimental hall during 2020, significantly reducing background contam-

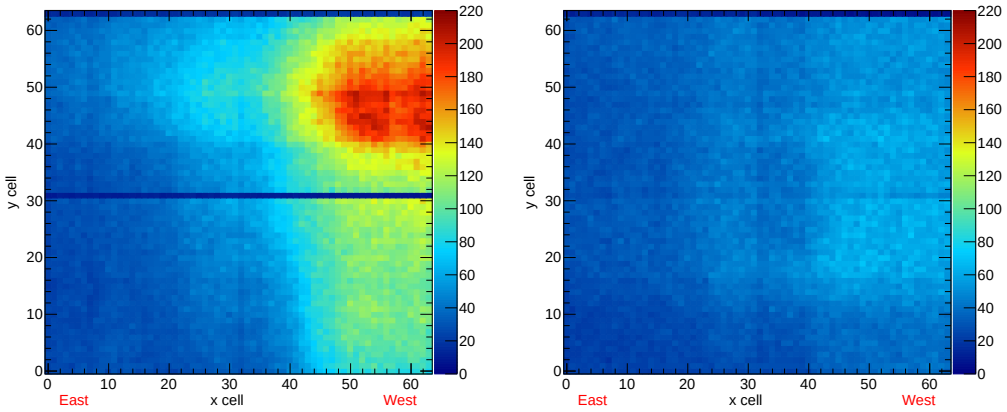

(a) 2020 Data (b) 2022 Data

Figure 6.5: Beam profiles on the front face of the NOvA detector in 2020 (left) and 2022 (right), showing the impact of additional shielding in reducing beam-induced background. In the 2020 data, a clear background hotspot is visible in the top west corner of the detector. Two low-occupancy horizontal cells, caused by under-filled detector modules, are also present in the earlier data and were corrected before final data collection (period 3 and period 4).

ination in the detector. This improvement can be seen for data taken after December 2020, as shown in Figure 6.5b, where the previously high background region on the upper west side has been substantially suppressed.

6.2.2 Electronics Shut-Off

The NOvA Test Beam detector operated in a high-background environment with a beam structure very different from what the electronics were designed for. The NOvA front-end electronics were built to handle short 10 μs spills occurring every second, with minimal activity from neutrino interactions. In contrast, the the beam delovered to the Test Beam experiment, was a continuous 4.2-seconds spill once per minute, which led to saturation issues in the

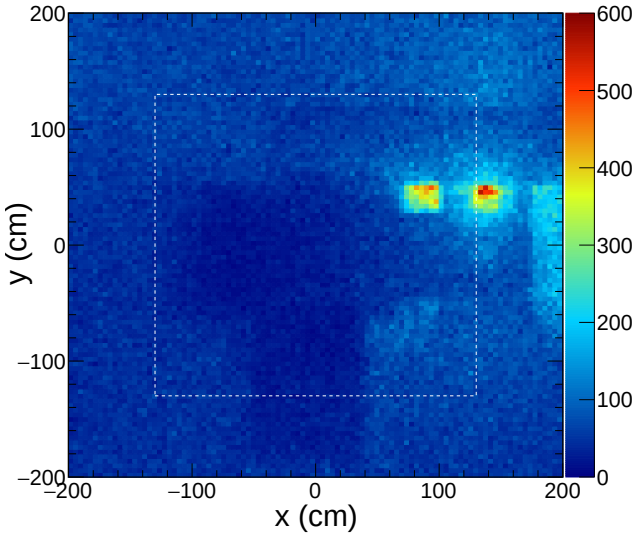

(a) Simulation before shielding

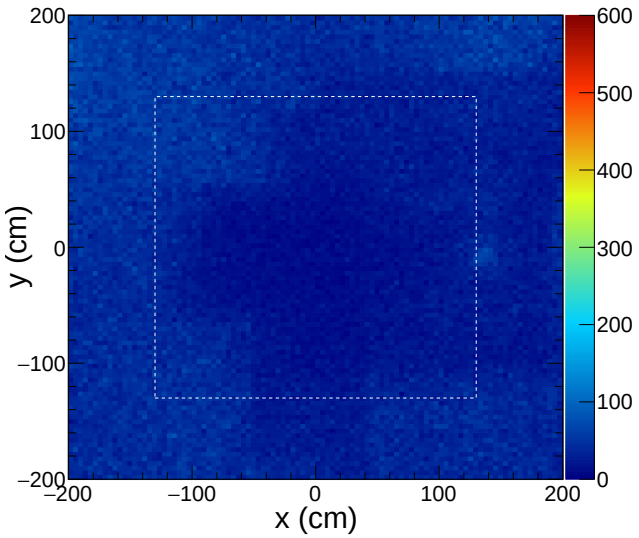

(b) Simulation after shielding

Figure 6.6: The plots show simulated beam profiles on the front face of the NOvA detector. The simulations were conducted using a full secondary and tertiary beamline model with 10^9 protons on the primary target. The top plot represents the initial beamline configuration, while the bottom plot shows the effect of added shielding. The white box indicates the approximate cross-section of the active detector volume. The darker regions are defined by the shielding within the NOvA beamline (collimators, dipole magnet, Cherenkov, etc.).

electronics, as discussed in the previous section. The most significant effect observed was the presence of dead regions in the detector, referred to as shut-offs. A clear example of an event affected by shut-offs is shown in Figure 6.7.

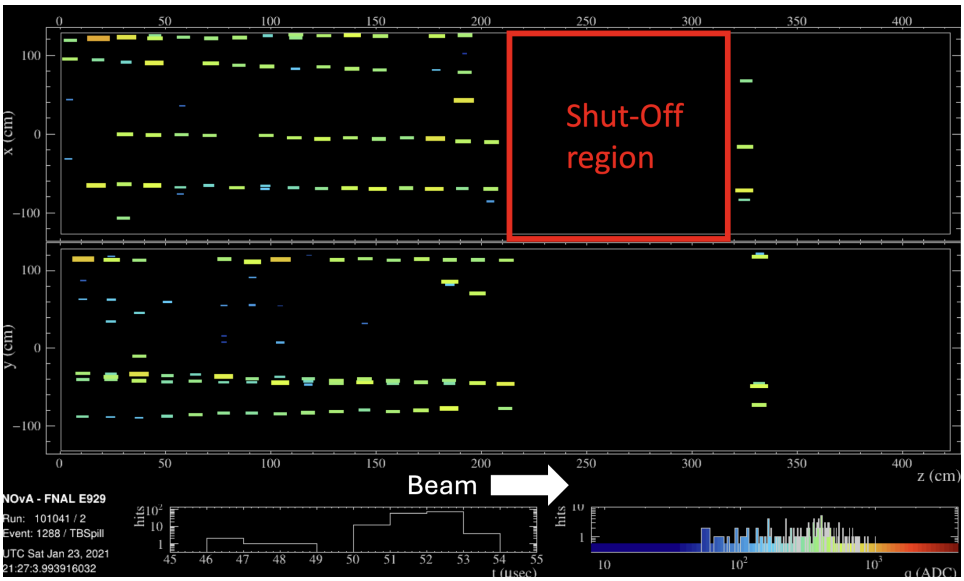

Figure 6.7: NOvA Test Beam Event Display showing the top and side view of the detector, the shut-offs can be seen as blank regions in the detector. The tertiary beam direction is indicated by the arrow.

This effect occurs when the small internal buffer in the FEB becomes full, causing the FEB to shut down until it is re-enabled by a DAQ signal which is sent ten times per second (10Hz). However, if the FEB buffer is still full, the FEB will remain off until it clears and receives a new signal.

During the run, shut-offs were mitigated using a lower beam intensity and adding concrete shielding blocks along the beamline (mentioned in previous section), which effectively reduced the background interactions with the detector.

Further studies showed that entire groups of FEBs that belong to a DCM were

also shutting down. This was known as DCM shut-offs. These shut-offs were more difficult to identify because they were not associated with any specific data quality flag associated with them, making them unpredictable.

Identifying and excluding events with FEB and DCM shut-offs is essential for the physics goals of the NOvA Test Beam Experiment, as including these events can bias the energy response studies of tertiary beamline particles in the detector.

To address these issues, a shut-off detection algorithm was implemented . The algorithm analyzes hit times in the detector, in which an event is classified as FEB shut-off when there are no hits in a 10 ms window on either side of the trigger time.

Similarly, an event is classified as DCM shut-off if no hits are recorded across all FEBs in a DCM during 80 μs within a 1 ms window on either side of the trigger time. This approach effectively tagged events with FEB and DCM shut-offs. The algorithm demonstrated an efficiency of 98-100 % in identifying these events and was implemented a preselection cut to remove shut-off events from our physics analysis.

6.2.3 Under-Filled Cells

As described in section 5.7, the NOvA test beam detector consists of 63 planes, each alternating between vertical and horizontal cell orientations. Each plane is constructed using two 32-cell extrusion modules, resulting in a total of 64 cells per plane.

During construction, the horizontal planes were intentionally tilted at an angle of 0.6° to allow proper oil settling. However, due to a manufacturing error, the fill ports were installed on the opposite side of the modules compared to those in the Near and Far Detectors.

As a result, once the oil reached the level of the fill port, air became trapped at the top of both horizontal modules. Although the very top of the detector is not critical for most analyses, the top of the lower horizontal module lies near the center of the detector, close to where the tertiary beam enters.

The lack of oil in the last cell of the lower modules caused a deficit in the expected number of hits traversing this region. The issue was resolved around the middle of period 3 by installing extensions to the fill ports, which allowed the horizontal modules to be fully filled. A small hole drilled on the opposite side allowed the trapped air to escape during filling. The results before and after the top-off of these modules can be seen in Figure 6.5.

For this analysis, runs affected by this issue were discarded to avoid potential biases in energy deposition within the detector, therefore only period 3 and period 4 events that match with a full active detector are kept in the analysis sample.

6.2.4 Pile-Up Rejection

In the NOvA Test Beam Detector, pile-up events occur when multiple particles interact in the detector that are unrelated to the tertiary beam particle but contribute significant energy and hits to the event of interest. A pile-up event

can be seen in Figure 6.8.

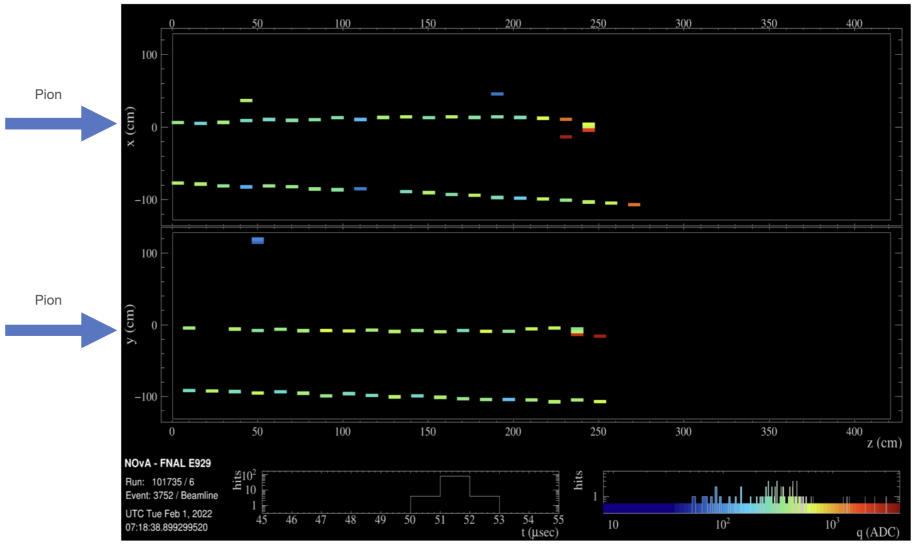

Figure 6.8: Event Display showing a pion candidate affected by pile-up. Alongside the pion track, a second particle enters the detector from below, depositing additional hits in the same slice, distorting the event reconstruction and energy estimation.

One major consequence of these events is their potential to bias energy measurements in studies of detector response. Therefore, it is essential to identify and remove such events from data samples. For this analysis, a strategy to mitigate pile-up events was implemented, which will be explained below.

The strategy focuses on the use of two simulated samples: The first one is a single particle simulated sample of pions and the second one contains the same single particles plus overlays. The second sample simulates the real conditions of tertiary particles affected by pile-up. As an example, two event displays in Figure 6.9 show the same simulated event with and without pileup.

The comparison of these two samples shed light on the effect of pile-up in pion

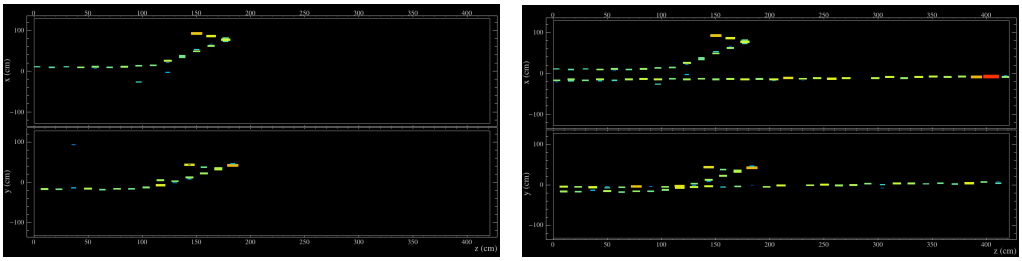

Figure 6.9: Event displays from a simulated pion used to study pile-up. The left Event Display shows a clean single-pion interaction, while the right Event Display shows the same event with pile-up overlays added, this illustrates how additional activity may alter the detector response.

events. Figure 6.10 presents a comparison of the energy deposited and the number of hits in the slice for single-particle simulated events and events with pile-up. A noticeable bias is observed in the sample with pile-up compared to the single-particle generated events.

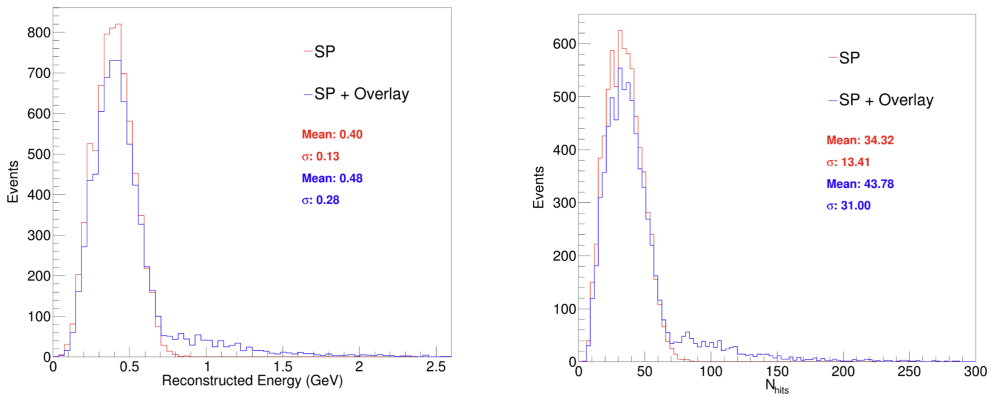

Figure 6.10: Reconstructed energy (Left) and reconstructed number of hits (Right) in the slice. The red distribution represents single-particle (SP) generated pion events whereas the blue distribution shows single-particle generated pions plus overlays (pile-up). A clear bias is observed in the sample containing pile-up compared to the one without pile-up.

One of the main characteristics of pile-up events is that unwanted particles produce hits in regions of the detector where they are not expected. For

example, secondary beam muons may enter above or below the expected entry point of a tertiary beam particle, at the same time that the tertiary particle is interacting in the detector.

The expected entry point for a tertiary particle is located around cell 28 in the Y view (side view) and cell 32 in the X view (top view). Using this as a reference, a pile-up identification algorithm was developed by defining detection regions around the particle's expected entry point in both detector views to capture pile-up-related activity.

These regions are defined by two tunable parameters: width and length. The width refers to the number of cells above and below cell 28 in Y and cell 32 in X, while the length corresponds to the number of detector planes spanned, starting from the front face of the detector (plane 0). A diagram illustrating this strategy is shown in Figure 6.12.

In addition, the last three planes with activity in a slice (Figure 6.11) are used to veto slices that contain tertiary beam particles accompanied by secondary muons exiting the detector.

Studies were conducted to identify the regions that maximize slice purity, where purity is defined as the ratio of energy deposited by the tertiary beam particle (pions) to the total energy deposited in the slice. Results from these studies showed that by filtering particles with activity in regions 12 cells or more away from the center within the first four planes, resulted in a slice purity of 99.6%. These regions were used to effectively remove pileup events. To enhance pile-up rejection, events with unusually high reconstructed energy are also removed from the sample.

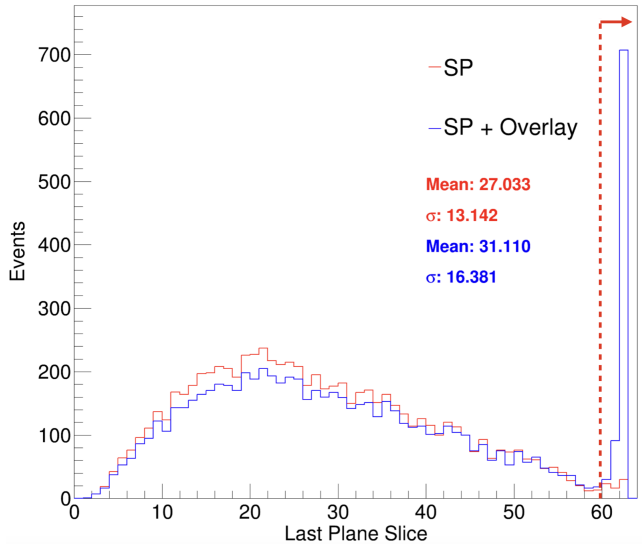

Figure 6.11: Plot showing the distribution of the last number of planes with activity in the slice for single-particle (SP) simulated pions and single-particle simulations with pile-up overlays. A distinct peak around planes 61 to 63 indicates clear signs of muons exiting the detector and contributing to pile-up. The dashed red line shows the region used to veto these events.

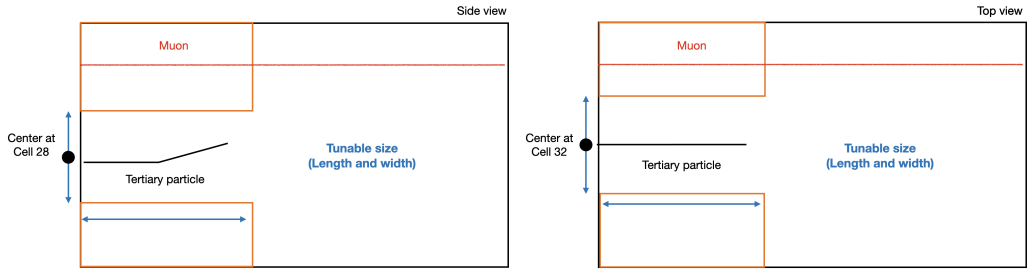

Figure 6.12: The diagram shows the two-view representation of a pile-up event in the NOvA Test Beam detector including a tertiary beam particle and a secondary muon. The orange regions are tuned and used to identify and reject pile-up events.

This pile-up rejection method forms part of the preselection and is applied before final pion event classification to ensure pile-up-free slices are used in the energy response analysis. On the other hand, events with unusual high energies were removed from the samples after full section was achieved. It was noticed that approximately 0.4% of the sample was a affected by these events, which is consistent with the 0.4% inefficiency observed in the pile-up rejection algorithm's performance.

6.3 Pion Event Selection

This section focuses on the identification of π^+ and π^- samples for the NOvA Test Beam experiment. Pion selection can be achieved by using the particle identification capabilities of the tertiary beamline instrumentation. In addition, cuts are applied to select the NOvA detector slice that contains a track that matches the beamline pion selection. These strategies will be outline in detail in the following sections.

6.3.1 Beamline Particle Identification (PID)

Beamline Particle Identification (PID) plays an essential role in the NOvA Test Beam experiment analyses. It is accomplished by measuring the particle's time-of-flight, reconstructed beamline momentum, reconstructed particle mass and Cherenkov detector information. Figure 6.13 shows time-of-flight versus reconstructed beamline momentum for particles in Period 3 and Period 4 data.

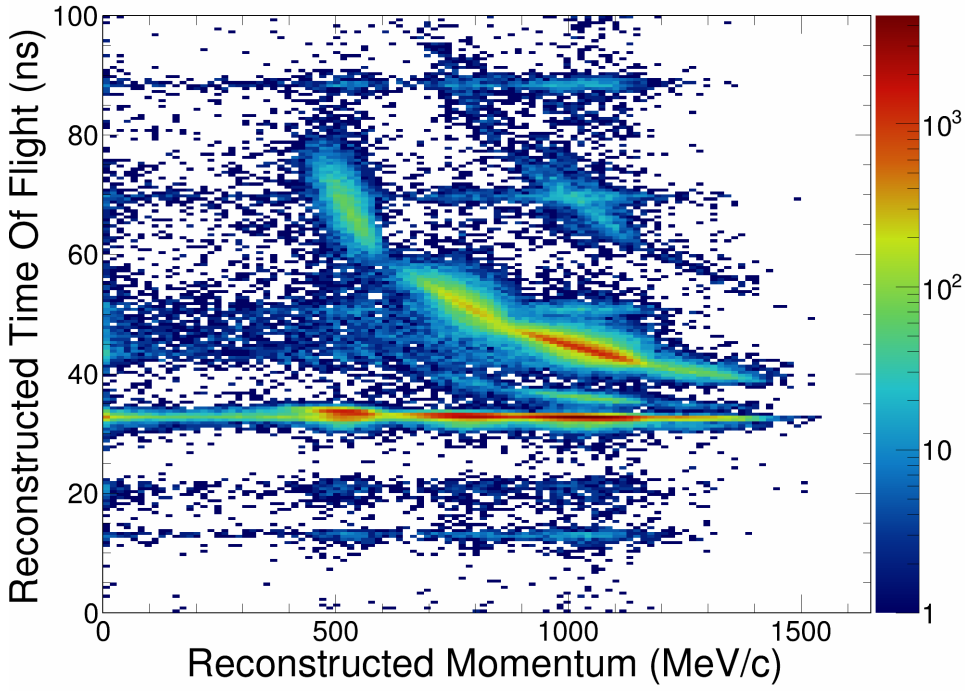

Figure 6.13: Reconstructed Momentum versus time-of-flight for all beamline events in Period 3 and Period 4 data with no selection cuts applied. The Z axis represents the accumulation of events in parameter space. The reconstructed time-of-flight versus momentum data reveals several structures. Horizontal structures spaced by 18.83 ns correspond to signal pairs from light particles in different Main Injector bunches and do not represent physical particles. Repeating sloped structures around 1000 MeV are identified as protons from different bunches. Additionally, diffuse structures at approximately 21 ns and 42 ns likely result from a light particle and a proton originating from separate bunches.

The first PID selection ensures that only pions, muons, electrons, kaons, and protons pass the selection, removing non-physical particles observed in Figure 6.13, as they are the only particles of interest for NOvA Test Beam Experiment. This is conducted by applying selection cuts in the data sample according to the expected time-of-flight and momentum in the beamline for these particles.

Equation 6.1 is used to calculate the time-of-flight as a function of momentum for each tertiary beam particle using a path length of 9.71 m. Here t is the time of flight, m is the particle mass, p the particle momentum, L the distance and c the speed of light. Figure 6.14 shows the expected time-of-flight vs momentum curves for different particle types in Period 3 and Period 4.

$$t = \frac{L}{c}\sqrt{1 + \frac{m^2}{p^2}} \tag{6.1}$$

The curves demonstrate the separation of fast particles (electrons, muons, and pions) from slow particles (protons and kaons). Using these calculations, a loose two-dimensional cut is applied around the regions of parameter space where particles are expected to be in the 2D plot. These cuts effectively remove most of the unwanted non- physical particles. The resulting sample after applying these loose PID selection cuts can be seen in Figure 6.15.

With the loose PID cuts in place, the next step is to further refine the PID selection process. Electrons are identified by the reconstructed Cherenkov Counter signal, which enables electron identification and allows their separa-

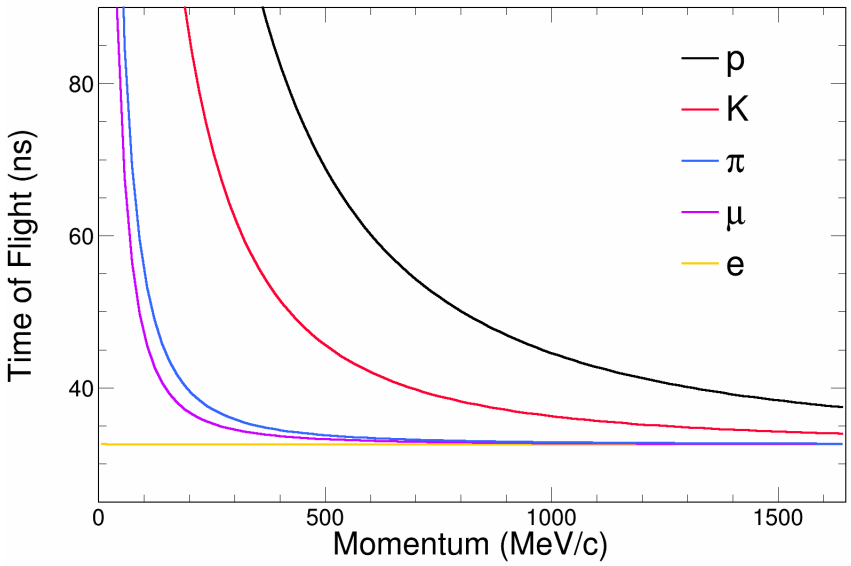

Figure 6.14: Time-of-Flight vs Momentum curves for different particle types. Here, the distance traveled by the particles is 9.71 m (Period 3 and Period 4 particle path length).

tion from pions and muons. In addition, the reconstructed mass of the particles is also considered to enhance the beamline selection. In order to achieve this, the reconstructed particle mass (m) is calculated using Equation 6.2:

$$m = p\sqrt{\frac{c^2t^2}{L^2} - 1}. \tag{6.2}$$

Here again, p is the reconstructed beamline momentum, t is the measured time-of-flight, c is the speed of light and L the particle path length (9.71m). The reconstructed mass for the tertiary beam particles passing the loose PID cuts can be seen in Figure 6.16.

The mass distribution peaks are closely around the expected true mass of the

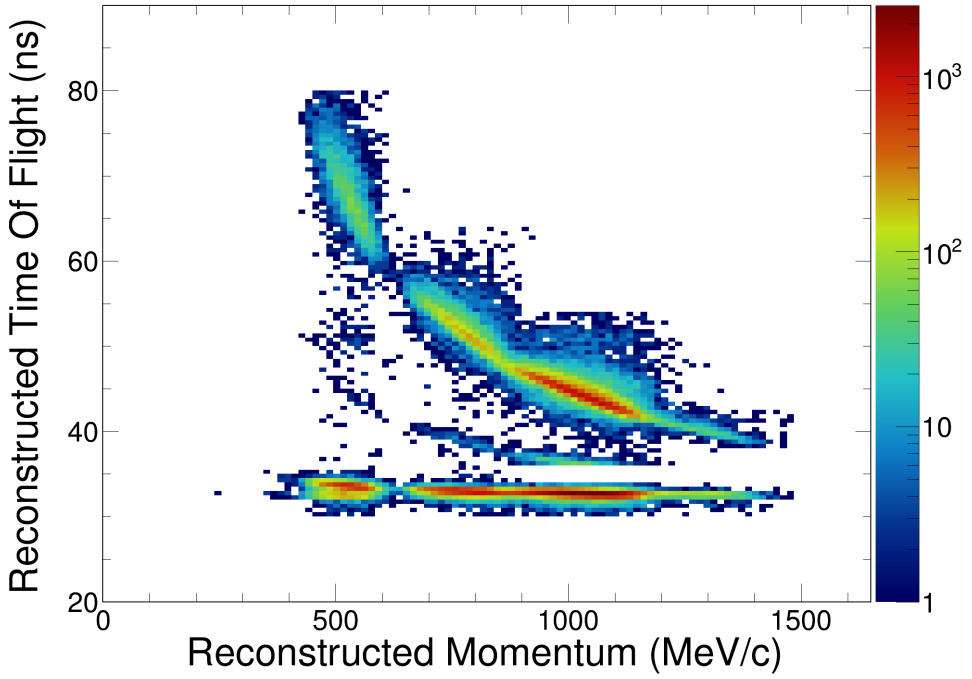

Figure 6.15: Reconstructed Momentum vs time-of-flight for beamline events passing the loose PID cuts, the Z axis represents the accumulation of events in parameter space, the resulting sample shows the characteristic shape expected by the calculated curves.

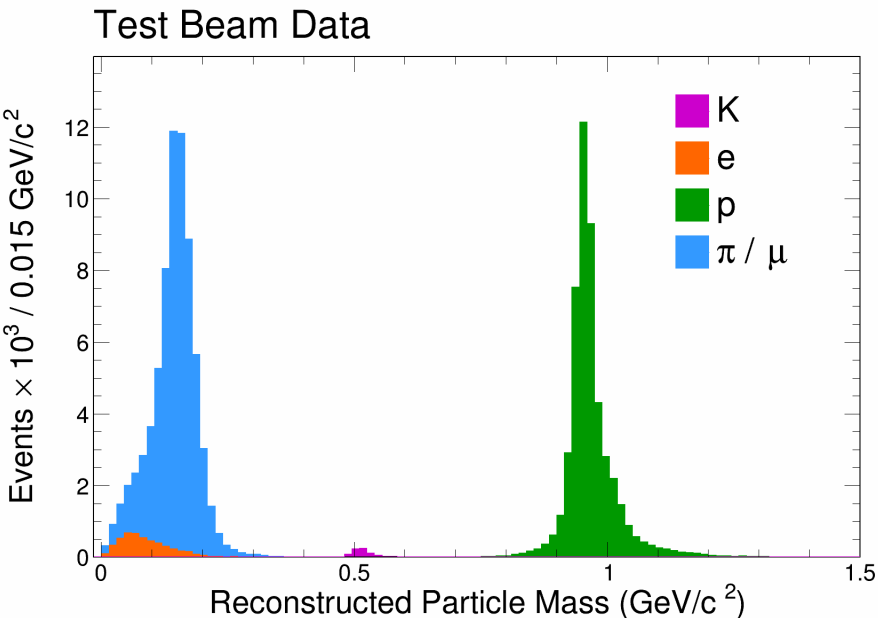

Figure 6.16: Stacked Histogram: Reconstructed particle mass for beamline events. Three distributions are observed: Left distribution (electrons, muons, and pions); middle distribution (kaons); and right distribution (Protons).

particles with exception of electrons, where the distribution is largely smeared and the location of the peak is approximately two orders of magnitude larger than the electron mass. This is a direct consequence of the time-of-flight resolution discussed in subsection 5.6.1. The large 0.18 ns resolution for electrons is a limiting factor and directly affects the mass resolution and location of the peak. Fortunately, as mentioned previously, electrons are tagged by the Cherenkov counter in the beamline, which allows us to effectively identify this sample.

The proton, kaon and pion-muon mass distributions look approximately symmetrical and they also show a clear separation between them. Therefore, we impose selection cuts in these distributions requiring that the mass of the protons must be within ±20% of the true value proton mass, while pion and muon

events must have masses below 250 MeV/c^2.

After applying this selection, the improved reconstructed time of flight versus momentum distribution along with their theoretical curves are shown in Figure 6.17. The one-dimensional representation of these variables for the reconstructed beamline momentum and time-of-flight can be seen in Figure 6.19 and Figure 6.18 respectively.

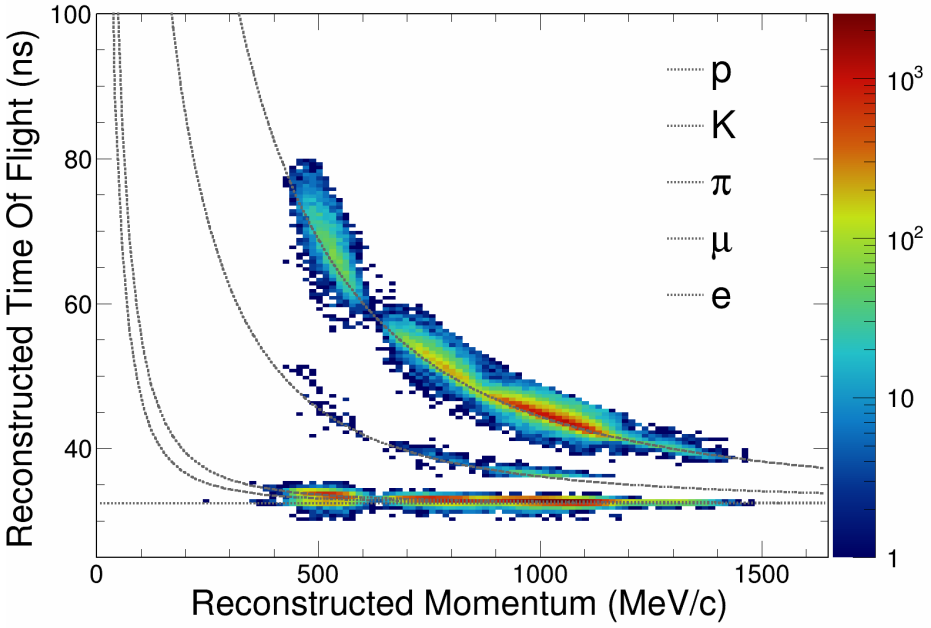

Figure 6.17: Reconstructed momentum versus time-of-flight for all beamline events in period 3 and period 4 data after particle mass selection. Particle labels are shown in the plot along with the theoretical curves. The Z-axis represents the accumulation of events in parameter space.

Since the final goal is to isolate pion events from everything else, a selection cut from 30 to 40 ns in time-of-flight is implemented. Unfortunately, it is not possible to separate pions from muons by using the beamline information alone. The mass and the time-of-flight of these particles overlap in these recon-

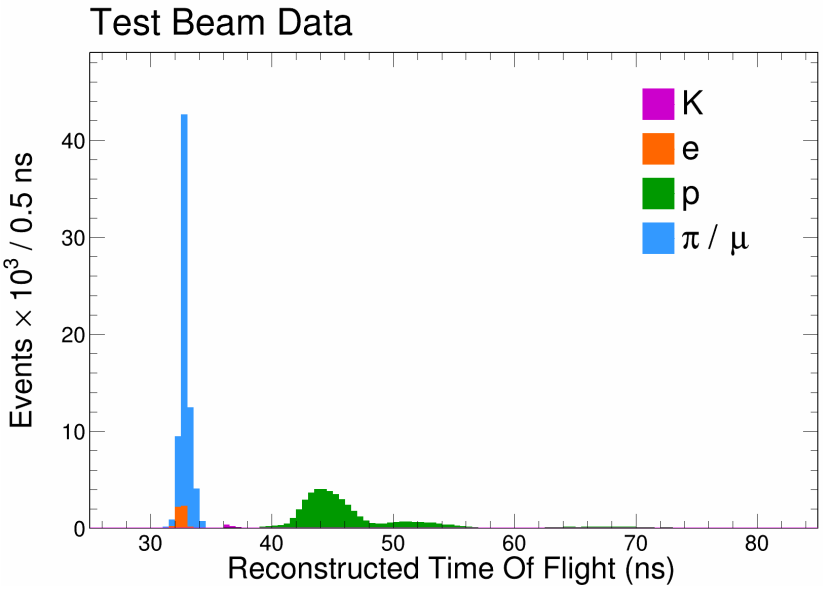

Figure 6.18: Time-of-flight (staked histogram) for all beamline events in period 3 and period 4 data after beamline selection. Particle ID labels are shown in the plot.

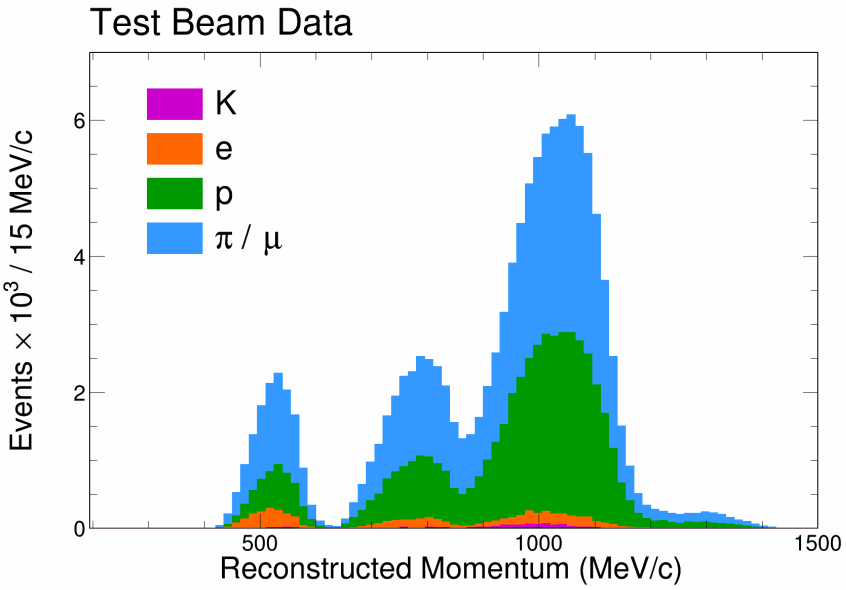

Figure 6.19: Reconstructed Momentum (stacked histogram) for all beamline events in period 3 and period 4 data after beamline selection. Particle ID labels are shown in the plot.

structed quantities. Therefore, tertiary beam muons are a source irreducible background in the pion sample. This muon component is relatively small and is estimated using simulation, as described in subsection 6.3.4.

6.3.2 NOvA Detector Selection

Following the identification of pion candidates using the beamline instrumentation, the next step on the pion selection is to ensure that events passing the beamline cuts match the events observed in the NOvA detector. In addition, the reconstructed pion tracks (BPF tracks) must be within the fiducial volume of the NOvA detector. In order to achieve this, the following selection cuts are applied to the pion candidate tracks in the detector:

- For every beamline trigger, the NOvA detector reads 150 μs of data. Tertiary particles are expected to enter the NOvA detector at around 50 μs during this trigger window. Thus, a selection cut on the BPF tracks is applied, requiring the minimum hit time of the track to be within 45 to 55 μs, this ensures full containment of the pion event in time window that matches to the expected beamline event.

- The next selection cut guaranties that the track start position is located within the expected region of the tertiary beam spot ($|x| < 20$ cm, $-25 < y < 15$ cm). In addition, only tracks with their vertex located in the first three planes are accepted ($z < 18$ cm).

- In order to avoid misreconstructed tracks in the detector, a minimum number of hits is required in the track in each view. Since the BPF

track algorithm requires at least 10 hits in both views to reconstruct a track, we require at least four hits per view.

- BPF tracks in the detector should not extend beyond 380 cm, ensuring proper vetoing of exiting particles.

The total number of events passing both beamline and detector selections are shown in Table 6.3

Magnet Current	**Selected Events**	
	$(\pi^+ + \mu^+)$	$(\pi^- + \mu^-)$
500 A	1412	571
750 A	2397	1493
1000 A	2401	979
1250 A	556	394
Total	**6766**	**3437**

Table 6.3: Selected pion-muon events at different magnet current setting.

Before studying the NOvA detector response for charged pions, two important aspects need to be considered: the energy loss of pions in the beamline material; and the muon background component from pion decay.

The energy loss is estimated using simulation and is used to correct the reconstructed beamline momentum at the detector front face. On the other hand, the muon background cannot be separated from the pion sample in data, and is instead estimated using simulation to quantify its contribution to the pion sample.

6.3.3 Beamline Momentum Correction

The NOvA detector response depends on the initial momentum of the tertiary particles. While the beamline instrumentation allows us to reconstruct this momentum accurately, it does not account for the energy lost by particles as they pass through materials downstream in the beamline, such as the two time-of-flight detectors, the air in the hall, and the Cherenkov counter.

This section presents the results of studies conducted to quantify and correct the energy loss experienced by pions as they traverse the beamline material prior to interacting in the NOvA Test Beam Detector.

Pion Energy Loss Estimation

The estimation of energy loss by pions traversing the downstream part of the NOvA beamline instrumentation is performed using Mode 1 simulation (subsection 6.1.1). The mono-energetic beam at different momentum settings provides well-defined distributions of the true momentum at each beamline detector.

Changes in momentum observed at each beamline detector can be attributed to energy loss as pions traverse the beamline. Since we aim to correct the reconstructed beamline momentum, determined by curvature in the dipole magnet, the energy loss of interest is the one occurring between Wire Chamber 4 and the front face of the NOvA detector.

The momentum fraction between Wire Chamber 4 and the detector allows us

to quantify the magnitude of the energy loss. This fraction is less than or equal to one, where a value of one indicates no energy loss, and lower values correspond to increasing energy loss. Figure 6.20 shows the momentum fraction distributions for various momentum settings between these two locations.

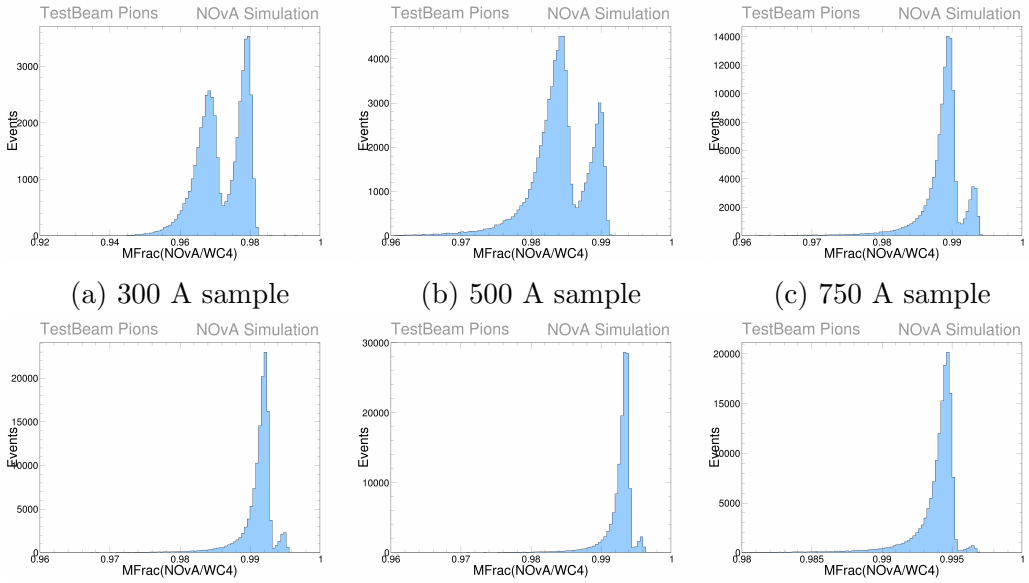

(a) 300 A sample (b) 500 A sample (c) 750 A sample

(d) 1000 A sample (e) 1250 A sample (f) 1500 A sample

Figure 6.20: Momentum fraction between WC4 and the NOvA detector for different magnet current settings. Each distribution shows a double-peak structure, corresponding to particles that hit or missed the downstream TOF 2 detector. It is also observed that lower energy (lower current) particles have a largest contribution for particles that missed DS TOF 2 (right peak of the distribution)

The distributions for all momentum settings exhibit a two-peak structure. The smaller momentum fraction peak in each distribution corresponds to particles that traverse all downstream detectors. In contrast, the larger momentum fraction peak is attributed to pions that miss the most downstream time-of-flight detector (DS TOF 2). A diagram illustrating this classification is shown in Figure 6.21.

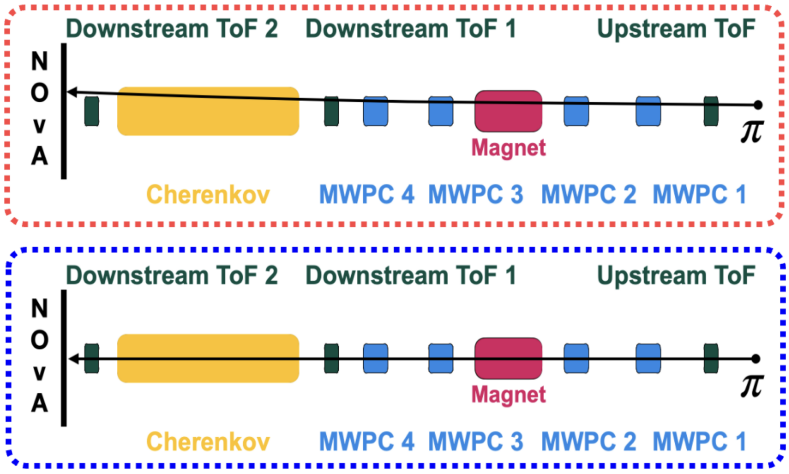

Figure 6.21: Diagram showing two possible pion trajectories through the beamline. In the top case, the particle misses the downstream TOF 2 detector, whereas in the bottom case, the particle hits it. This helps visualize how particles can be classified as missing or hitting DS TOF 2 based on their trajectories a direct consequence of pion scattering.

Another characteristic of these distributions is that the number of pions missing DS TOF 2 decreases as the momentum increases. This is a direct consequence of the relationship between a particle's initial momentum and the angular deviation it experiences while traversing the detectors. Higher-energy pions undergo smaller deviations from their original trajectory, making them more likely to pass through all downstream detectors.

Therefore, in order to correctly assess the pion energy loss in the beamline material, the original sample must be split into two: one consisting of pions that missed DS TOF 2; and one consisting of pions that hit DS TOF 2. This separation is relatively straightforward to perform in simulation, as we have access to the true hit information at each beamline detector, allowing for accurate sample classification.

Once the sample is separated, fits are performed for each distribution to determine the most probable value of momentum loss fraction for each momentum setting. Then, these data points are interpolated using a polynomial of order five represented in Equation 6.3.

$$p_{frac}(p_{true}) = a_0 + a_1\, p_{true} + a_2\, p_{true}^2 + a_3\, p_{true}^3 + a_4\, p_{true}^4 + a_5\, p_{true}^5. \quad (6.3)$$

Where p_{true} is the mean true momentum for each distribution, the value of the coefficients given by the interpolation for π^+ is shown in Table 6.4 and for π^- in Table 6.5.

Coefficient	**π^+ (Missed DS TOF 2)**	**π^+ (Hit DS TOF 2)**
$\boldsymbol{a_0}$	0.923054	0.883241
$\boldsymbol{a_1}$	3.23235×10^{-4}	4.86744×10^{-4}
$\boldsymbol{a_2}$	-6.19026×10^{-7}	-9.25784×10^{-7}
$\boldsymbol{a_3}$	6.05726×10^{-10}	9.0057×10^{-10}
$\boldsymbol{a_4}$	-2.94843×10^{-13}	-4.362×10^{-13}
$\boldsymbol{a_5}$	5.65877×10^{-17}	8.33803×10^{-17}

Table 6.4: Polynomial coefficients obtained by interpolating the momentum fraction for π^+ events that hit or missed the downstream TOF 2 detector. These coefficients are used to correct the reconstructed momentum at the front face of the NOvA detector.

The interpolated curves for π^+ and π^- are shown in Figure 6.27. The results of the interpolation show no significant difference between π^+ and π^-, as both lose energy primarily through ionization in the beamline material, a process that affects positive and negative pions equally.

Coefficient	π^- (Missed DS TOF 2)	π^- (Hit DS TOF 2)
a_0	0.916345	0.88199
a_1	3.72933×10^{-4}	5.05108×10^{-4}
a_2	-7.46056×10^{-7}	-9.83394×10^{-7}
a_3	7.54347×10^{-10}	9.74753×10^{-10}
a_4	-3.76514×10^{-13}	-4.79145×10^{-13}
a_5	7.36892×10^{-17}	9.26496×10^{-17}

Table 6.5: Polynomial coefficients obtained by interpolating the momentum fraction for π^- events that hit or missed the downstream TOF 2 detector. These coefficients are used to correct the reconstructed momentum at the front face of the NOvA detector.

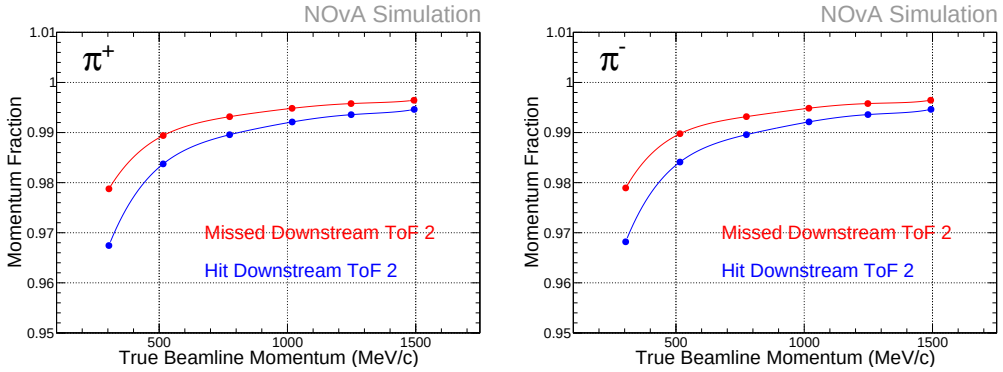

Figure 6.22: Momentum fraction as a function of true beamline momentum for simulated π^+ (left) and π^- (right). Red and blue curves are the results of the interpolation and correspond to particles that missed and hit the DS ToF 2, respectively. Higher momentum fractions for missed events reflect reduced energy loss. Both charge types show the same behavior, as energy loss is dominated by ionization. Note that the y-axis range starts at 95%.

Pion Energy Loss Correction

The momentum fraction functions are used to correct the reconstructed beamline momentum to account for energy loss by pions in the beamline material. This correction is applied by multiplying the reconstructed momentum for each event, p_{reco}, by the corresponding momentum fraction, as defined in Equation 6.3:

$$p_c(p_{reco}) = p_{reco} \cdot p_{frac}(p_{reco}) \tag{6.4}$$

To observe the effect of this correction, we use the reconstructed beamline momentum in simulation, specifically using G4Beamline in the Mode 1 configuration. This setup closely replicates the conditions of the Test Beam experiment and provides a realistic momentum range comparable to what is observed in data. Figure 6.23 shows the corrected p_c and uncorrected p_{reco} momentum distributions for both π^+ and π^-, across magnet current settings of 500 A, 750 A, 1000 A, and 1250 A.

The corrected distributions show a shift in the mean momentum of approximately 7 MeV/c. While this correction is straightforward to apply in simulation, where we can separate events based on whether pions hit or missed DS TOF 2, it is not trivial to apply to real data because we do not have access to the true information that would allow us to distinguish between these two cases.

Although this separation is not possible in real data, we can use the recon-

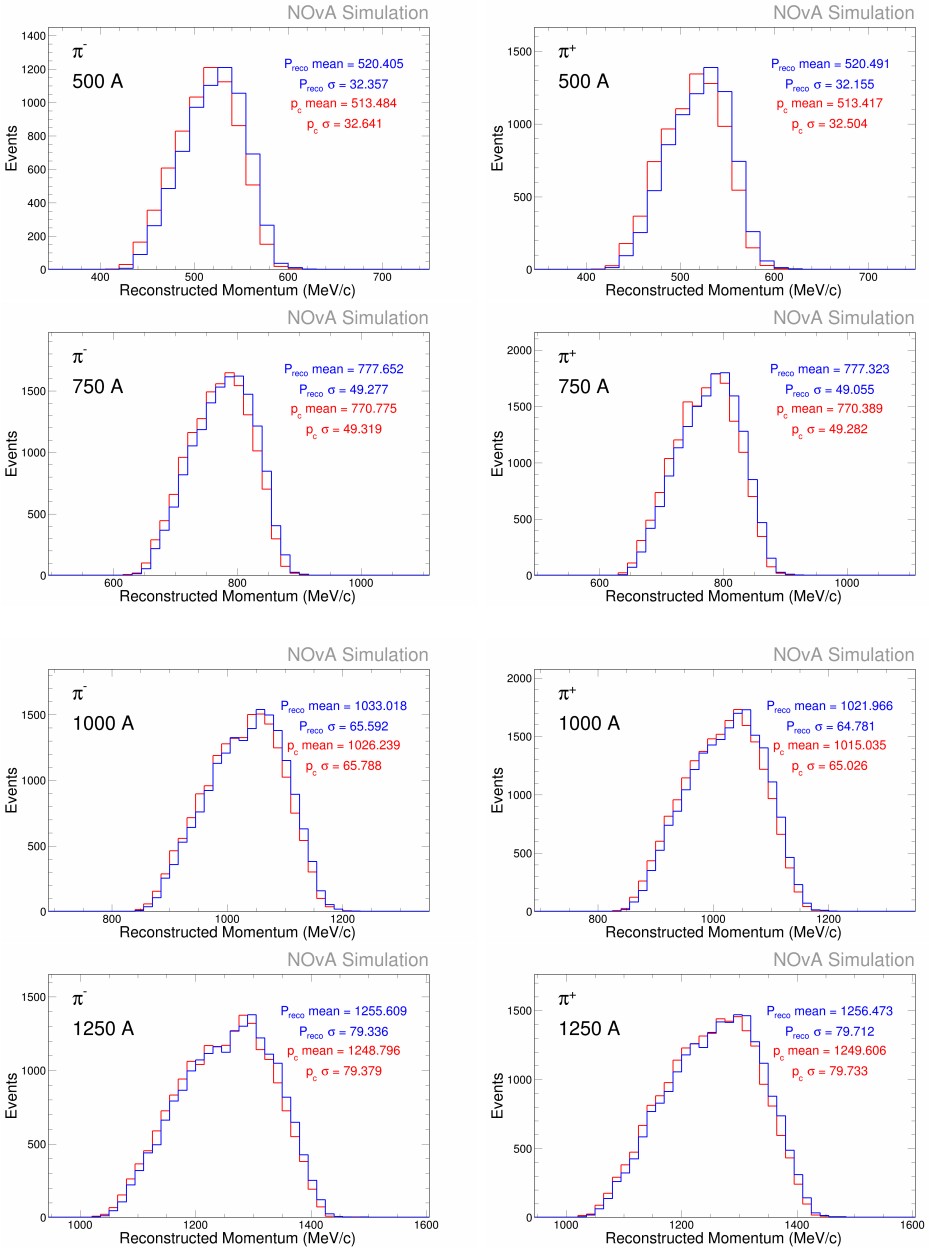

Figure 6.23: Reconstructed beamline momentum distributions for simulated π^- (left column) and π^+ (right column) at four different magnet current settings: 500 A, 750 A, 1000 A, and 1250 A. The blue histograms represent the uncorrected reconstructed beamline momentum (p_{reco}), while the red histograms show the corrected momentum (p_c) after applying the momentum correction. A small shift of approximately 7 MeV/c in the mean momentum is observed after correction.

structed projected beamline track positions on the front face of the NOvA detector to classify the samples. To achieve this, we use the projected positions of the tracks along with the known location and size of the DS TOF 2 detector. This information defines two regions that allow us to categorize the events, as illustrated in Figure 6.24. Pions with projected beamline track positions inside the DS TOF 2 region are classified as having hit DS TOF 2, while those with projections outside this region are classified as having missed DS TOF 2.

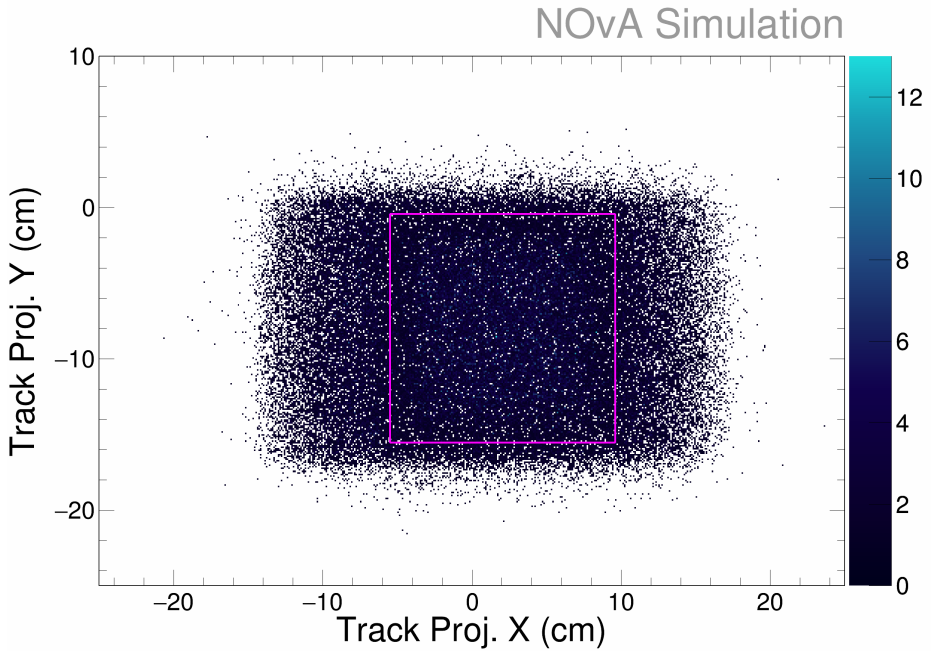

Figure 6.24: Projected beamline track positions on the front face of the NOvA detector in G4Beamline, Mode 1 simulation. The magenta box represents the geometric region corresponding to the location and size of the DS TOF 2. Tracks with projections inside this region are classified as having hit DS TOF 2, while those outside are considered to have missed it.

This method can classify the two samples with some limitations. One of the main issues with this approach is that it relies on the reconstructed track

projections, which are propagated from Wire Chamber 4 to the front face of the NOvA detector. This procedure assumes that pions do not interact with other elements in the beamline, such as DS TOF 1, the air in the hall, or the Cherenkov detector. In reality, such interactions can alter the pion trajectory, introducing a bias in the projected position and affecting the accuracy of the classification at the NOvA detector front face. This misclassification of events can be visualized in Figure 6.25.

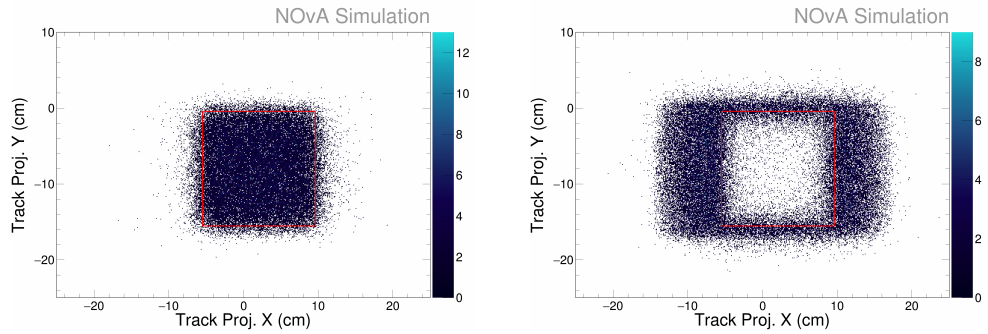

Figure 6.25: Projected beamline track positions on the front face of the NOvA detector for events where the pions hit DS TOF 2 (left) and missed DS TOF 2 (right). The red box represents the geometric region corresponding to the physical location of the DS TOF 2 detector, used to classify events based on reconstructed projections. The spread and overlap at the edges illustrate the limitations of the method, as interactions upstream can bias the projected track positions and lead to misclassification.

This classification, referred to as the nominal classification, correctly identifies 89.9% of pions that hit DS TOF 2 and 82.5% of those that missed it. To improve this imbalance, we carried out a series of studies where the size of the classification region was adjusted. The region was expanded outward and reduced inward in steps of 1 millimeter in both the X and Y directions, with ten steps in each direction. The goal was to see how these changes affected classification accuracy. The percentage of correctly classified events for each configuration is shown in Table 6.6.

The results of this study showed that reducing the classification region by 3 mm in each direction yields 86.4% and 86.6% correctly classified events for pions that hit and missed DS TOF 2, respectively. This region provides a more balanced classification between the two samples and is used to separate the two samples.

The beamline momentum correction is applied to real data using Equation 6.4 for pions that hit and missed DS TOF 2, based on the method described above. While this procedure is effective, it is not perfect, as the classification method cannot separate the two samples with complete accuracy. To evaluate the impact of this limitation, dedicated studies were carried out to understand how misclassification affects the corrected beamline momentum. The resulting variation is treated as a source of systematic uncertainty, which is estimated in subsection 7.3.1.

6.3.4 Muon Background in the Pion Sample

In addition to correcting for energy loss in the beamline, it is also important to account for background contamination in the selected pion sample

As mentioned in section 6.3, one of the main sources of background in the selected pion sample is tertiary beam muons due to pion decay. This section will cover the theory of pion decay in the NOvA Test Beam experiment and the use of simulation to quantify the muon background in the pion sample as a function of momentum.

Size Region (X and Y)	Hit DS TOF 2 Classification (%)	Missed DS TOF 2 Classification (%)
+10 mm	97.1	64.4
+9 mm	96.6	66.4
+8 mm	96.2	68.3
+7 mm	95.6	70.2
+6 mm	95.0	72.2
+5 mm	94.4	74.1
+4 mm	93.6	75.9
+3 mm	92.8	77.6
+2 mm	91.9	79.4
+1 mm	90.9	81.0
Nominal	89.9	82.5
-1 mm	88.8	84.0
-2 mm	87.7	85.3
-3 mm	86.4	86.6
-4 mm	85.1	87.7
-5 mm	83.7	88.8
-6 mm	82.2	89.8
-7 mm	80.6	90.7
-8 mm	79.0	91.5
-9 mm	77.4	92.2
-10 mm	75.7	92.9

Table 6.6: Classification for pions that hit and missed DS TOF 2 as a function of the classification region size. The region was expanded and reduced in steps of 1 millimeter in both X and Y directions relative to the nominal region.

Pion Decay Modeling

Tertiary beam pions produced at the secondary target, can decay before reaching the NOvA detector. The mean lifetime (τ) of a charged pion is 2.6033×10^{-8} s, and their primary decay modes are $\pi^+ \to \mu^+ + \nu_\mu$ and $\pi^- \to \mu^- + \bar{\nu}_\mu$, with a branching ratio of 0.999877. The mean decay length for charged pions is given by:

$$L_{\text{decay}} = \frac{p}{m} c\tau \tag{6.5}$$

where p is the pion momentum, m is its mass, c is the speed of light, and τ is the mean lifetime of the pion. This expression shows that higher-momentum pions are more likely to reach the detector before decaying. The momentum of tertiary beam pions ranges from 400 MeV/c to 1400 MeV/c, which yields a mean decay length of approximately 22.4 m to 78.3 m.

The probability that a charged pion decays at a particular position x along its path follows an exponential distribution. Starting from the decay law in time, and applying relativistic time dilation and a change of variable to distance, the decay probability density function can be expressed as:

$$\frac{dP}{dx} = \frac{1}{L_{\text{decay}}} e^{-x/L_{\text{decay}}} \tag{6.6}$$

Here, L_{decay} is the mean decay length defined in Equation 6.5. This function describes the probability per unit length that a pion decays at position x from

the point of its production. The decay probability is highest near the origin and falls off exponentially, implying that most pion decays occur closer to the target.

Tertiary beam pions travel approximately 16.7 meters from the target to the front face of the NOvA detector. The probability per unit length for pions decaying along this distance, evaluated at six different momentum settings, is shown in Figure 6.26.

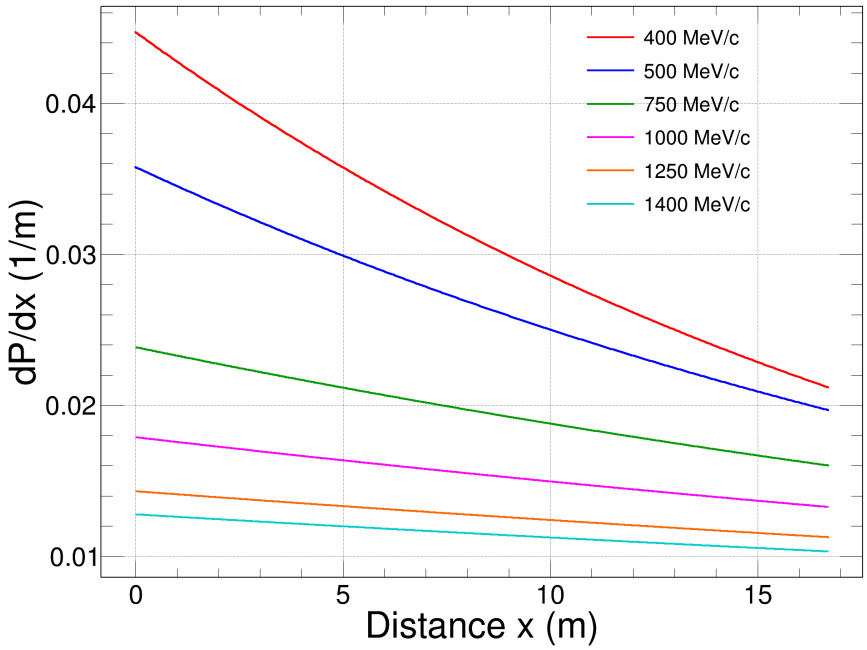

Figure 6.26: Decay probability density $\frac{dP}{dx}$ as a function of distance from the target, for charged pions with momenta from 400 MeV/c to 1400 MeV/c. The curves follow an exponential distribution with decay length $L_{\text{decay}} = \frac{p}{m} c\tau$, showing that lower-momentum pions are more likely to decay near the target, while higher-momentum pions decay more gradually along the beamline. The plot covers the 16.7 m distance to the NOvA Test Beam Detector.

The analysis of the decay probability density of charged pions shows that the

likelihood of decay is highest near the point of production and decreases exponentially with distance. This effect is more pronounced for lower momentum pions, which have shorter mean decay lengths and are therefore more likely to decay within the first few meters of the beamline. In contrast, higher momentum pions are more likely to reach the detector without decaying.

Given that the NOvA Test Beam detector is located 16.7 meters from the secondary target, a larger fraction of lower momentum pions (for example, those in the range of 400 to 600 MeV/c) are expected to decay before reaching the detector. These decay muons form an irreducible background in the selected pion sample. To properly account for their contribution, we estimate the muon background using simulation.

Muon Background Estimation

As described in section 6.3 the beamline instrumentation alone is not capable of distinguishing between pions and muons. As a result, the muon content in the selected sample is estimated using G4Beamline simulation in Mode 1.

This simulation models pions produced at the secondary target and tracks their possible decays along the beamline. The resulting pion-muon sample is then passed through the same selection criteria described in section 6.3 to determine the fraction of muons that pass the cuts.

The results show that the selected muon fraction is highest at lower momentum (lower current settings), with muon content ranging from approximately 7% at 500 A to 4% at 750 A. For higher settings, such as 1000 A and 1250 A, the

muon contribution becomes negligible. These results are presented for positive and negative muons in Figure 6.27, Table 6.7, and Table 6.8.

The muon background is also estimated as a function of the reconstructed beamline momentum as shown in Figure 6.28. This result is important because it provides the expected muon content in the pion sample as a function of momentum. These muon fractions are used to generate pion–muon simulation samples with the appropriate background content, as described in subsection 6.1.2.

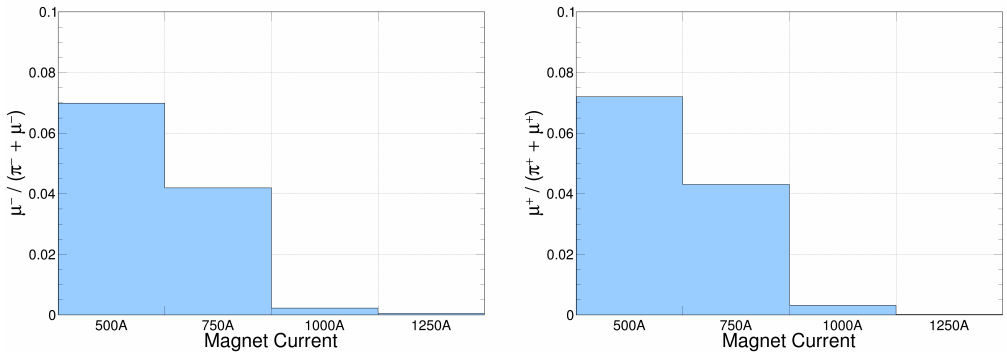

Figure 6.27: Fraction of muons relative to the total number of same-sign charged particles as a function of magnet current. The left panel shows the μ^- to $(\pi^- + \mu^-)$ ratio, and the right panel shows the μ^+ to $(\pi^+ + \mu^+)$ ratio. In both cases, the muon contamination decreases significantly with increasing magnet current, indicating effective pion selection at higher magnetic settings.

This chapter described the steps taken to select a clean pion sample in the NOvA Test Beam detector, including background estimation and beamline momentum corrections. With this sample in place, the next step is to study how the detector responds to charged pions.

Magnet Current (A)	$\boldsymbol{\mu^+/(\pi^+ + \mu^+)}$ (%)	**Stat. Error** ± (%)	**Data Events** $(\pi^+ + \mu^+)$
500	7.20	0.27	1412
750	4.30	0.15	2397
1000	0.32	0.04	2401
1250	0.02	0.01	556

Table 6.7: Percentage of positively charged muons (μ^+) relative to the total number of positively charged particles (μ^+ and π^+) as a function of magnet current. The results show a clear decrease in muon contamination with increasing magnet current. The last column shows the number of selected data events (see Table 6.3) that is impacted by the muon background.

Magnet Current (A)	$\boldsymbol{\mu^-/(\pi^- + \mu^-)}$ (%)	**Stat. Error** ± (%)	**Data Events** $(\pi^- + \mu^-)$
500	6.99	0.28	571
750	4.19	0.15	1493
1000	0.23	0.03	979
1250	0.05	0.02	394

Table 6.8: Percentage of negatively charged muons (μ^-) relative to the total number of negatively charged particles (μ^- and π^-) as a function of magnet current. The results show a clear decrease in muon contamination with increasing magnet current. The last column shows the number of selected data events (see Table 6.3) that is impacted by the muon background.

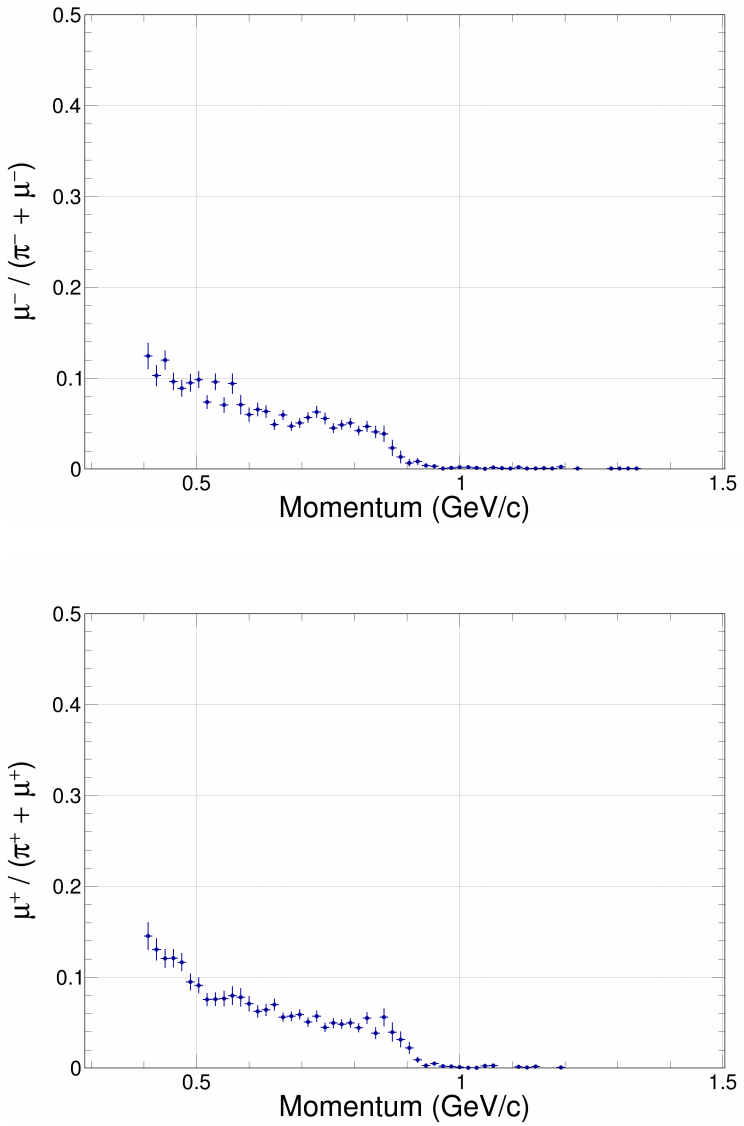

Figure 6.28: Fraction of muons from pion decay in the selected sample as a function of momentum, based on G4Beamline simulation in mode 1. The top panel shows the ratio of μ^- to $(\pi^- + \mu^-)$, and the bottom panel shows the ratio of μ^+ to $(\pi^+ + \mu^+)$. The muon content is highest at low momentum and becomes negligible at higher momenta.

Chapter 7

Pion Energy Response and Resolution in the NOvA Test Beam Detector

This chapter presents the first results of measuring the pion energy response and resolution in the NOvA Test Beam Detector. The primary objective is to evaluate the performance and ability of the NOvA detectors to accurately reconstruct charged pion events and assess their energy response and resolution using a well-characterized charged particle beam.

This work is part of a broader effort to improve the modeling of the hadronic response in the NOvA detectors. The results presented here contribute to the understanding of the energy calibration uncertainty, which remains one of the largest impacting the NOvA oscillation analyses described in Chapter 4.

By carefully analyzing the energy response for positively and negatively charged pions in the NOvA Test Beam experiment, this study can provide a foundation for future improvements in the simulation and calibration strategies used in NOvA oscillation analysis, while providing guidance in quantifying mismodeling and response-related systematic uncertainties.

7.1 Pion Energy Reconstruction

The energy response and resolution in the detector depends on the particle energy measured by the NOvA beamline instrumentation (incident energy), the particle type (charged pions in this study), and the measurement of the energy deposited in the detector.

The incident pion energy E is calculated using Equation 7.1:

$$E = \sqrt{p^2 + m^2} \tag{7.1}$$

where p is the corrected beamline momentum and m is the pion mass. This provides the total pion energy at the front face of the detector, and it is used as the reference for the pion energy response and resolution studies.

In the NOvA Test Beam experiment, the reconstructed energy measured by the detector is calculated by summing all the calibrated energy deposits in the active region of the detector (scintillator). For each selected BPF pion track, the energy in the slice is obtained and used to measure the pion energy

response. This method is applied to both data and simulation.

The Single Particle Simulation (subsection 6.1.2) approach is used to generate the pion samples, where the beamline kinematic information from pion data is used to seed the initial conditions of the simulation. This approach directly uses the corrected beamline momentum as input, providing an accurate comparison between data and MC.

In this analysis, no additional corrections are applied to recover energy lost in passive materials like the plastic cell walls, except in section 7.4, where the energy response between NOvA and MINERvA Test Beam results are compared. In the NOvA analyses, such as three-flavor or sterile neutrino searches, the energy loss corrections are typically applied at a later stage using algorithms to reconstruct the neutrino energy [101][102]. The algorithms known as "neutrino energy estimators" properly match the true and reconstructed neutrino energy.

In the context of the NOvA Test Beam analyses, in particular, the energy response analyses, energy estimators are not used, since one of the purposes of the Test Beam experiment is to validate the NOvA simulation framework, including the energy response to single particles.

The reconstructed energy in the Test Beam detector includes all visible deposits in the active cells, including contributions from secondaries produced in pion interactions. This reflects the type of visible energy typically seen in hadronic interactions for neutrino events.

7.2 Pion Energy Response

The detector response is studied by comparing the reconstructed energy in the detector to the incident pion energy at the front face as a function of the incident energy. This comparison can be performed separately for positively and negatively charged pions using both data and simulation. The simulation samples include pions and muons (background) from pion decay and are generated using the same reconstruction and detector selection as in data.

To evaluate the energy response across the full range of pion beam energies, the samples are separated into several energy regions. These will be used to study the level of agreement between data and simulation. Later in this section, different sources of systematic uncertainty are evaluated to quantify their impact on the response results.

7.2.1 Reconstructed Energy Distributions

The selected samples in data and simulation are divided into eleven different incident energy regions ranging from 0.45 to 1.4 GeV. For each energy region, the reconstructed energy in the slice (TotESlice) and the incident pion energy distributions are shown separately for both negatively and positively charged pions in Figures 7.1 to 7.4 and Figures 7.5 to 7.8, respectively. Additionally, a table showing the number of data events for each energy region is shown in Table 7.1.

One of the primary features observed is the increase of the mean reconstructed

energy as the incident pion energy increases. Additionally, it is observed that the agreement between data and simulation is reasonable across all energy regions, although the most noticeable differences appear in the π^+ samples, where the mean reconstructed energy in the detector in data is consistently higher than in simulation.

The excellent agreement between data and simulation shown for the incident energy spectra is a consequence of the single particle simulation generation. Since the initial kinematic conditions of the simulated particles are seeded with the data samples in the simulation, the agreement is expected by construction.

The muon background component is also shown for each energy region. The muon distributions tend to have higher and narrower energy peaks compared to pions. It is also worth noting that the muon sample fraction decreases with increasing incident energy. From the 0.95 GeV region onward, the selected muon content is negligible. In the next subsection, the energy response will be evaluated in two cases, where the muon background is or is not subtracted from data and simulation.

7.2.2 Pion Energy Response Results

The mean energy response is calculated by dividing the reconstructed energy in the slice by the incident energy for each event. For the resulting energy response distributions, the mean and the error in the mean for each spectrum are extracted. These data points are plotted as a function of the incident energy.

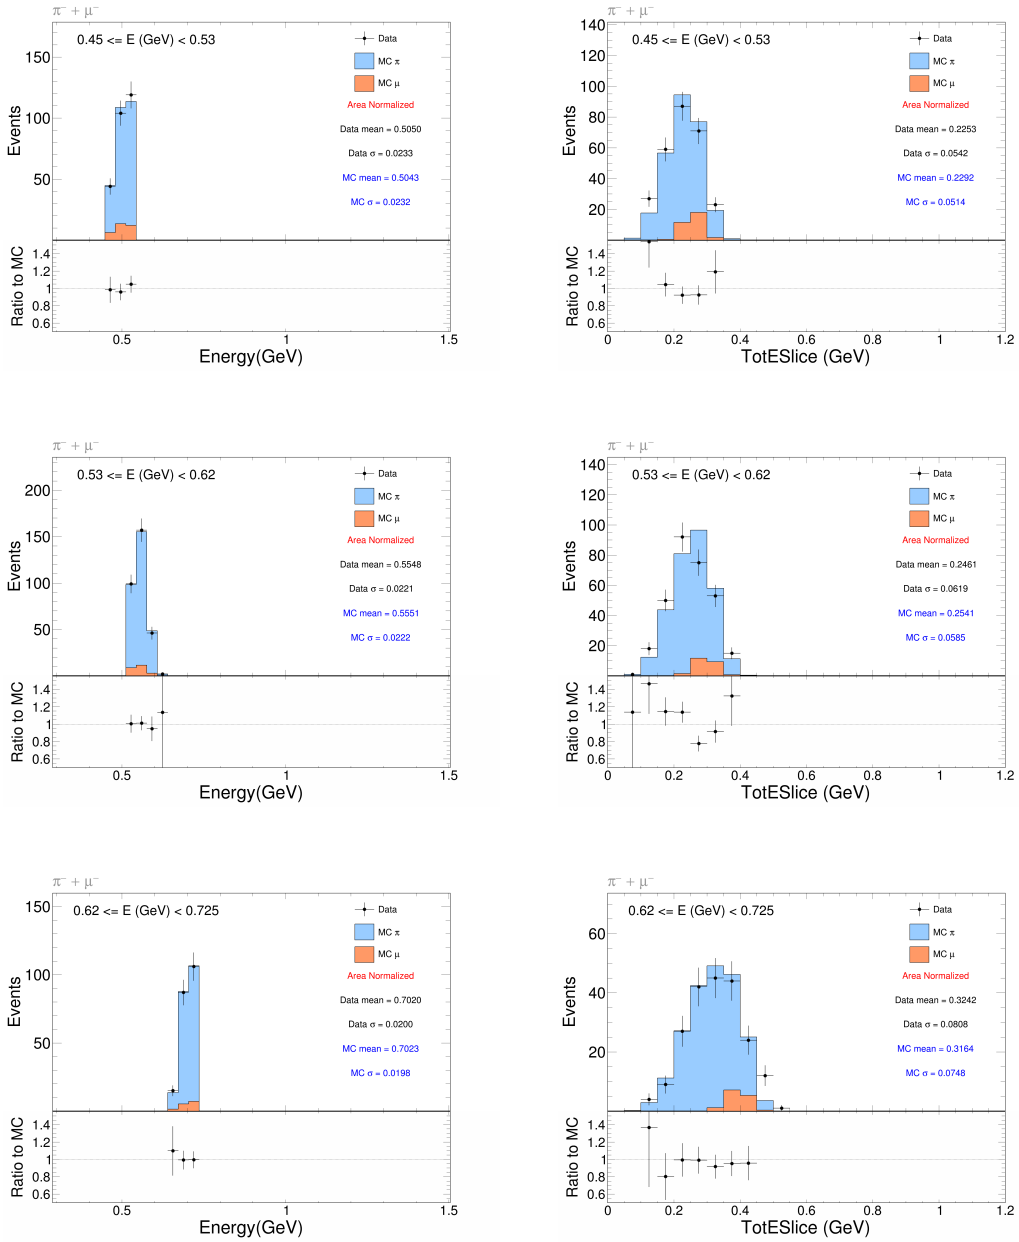

Figure 7.1: Incident energy (left) and reconstructed energy (right) distributions for negatively charged pions in data and simulation, shown for incident energy ranges from 0.45 to 0.725 GeV.

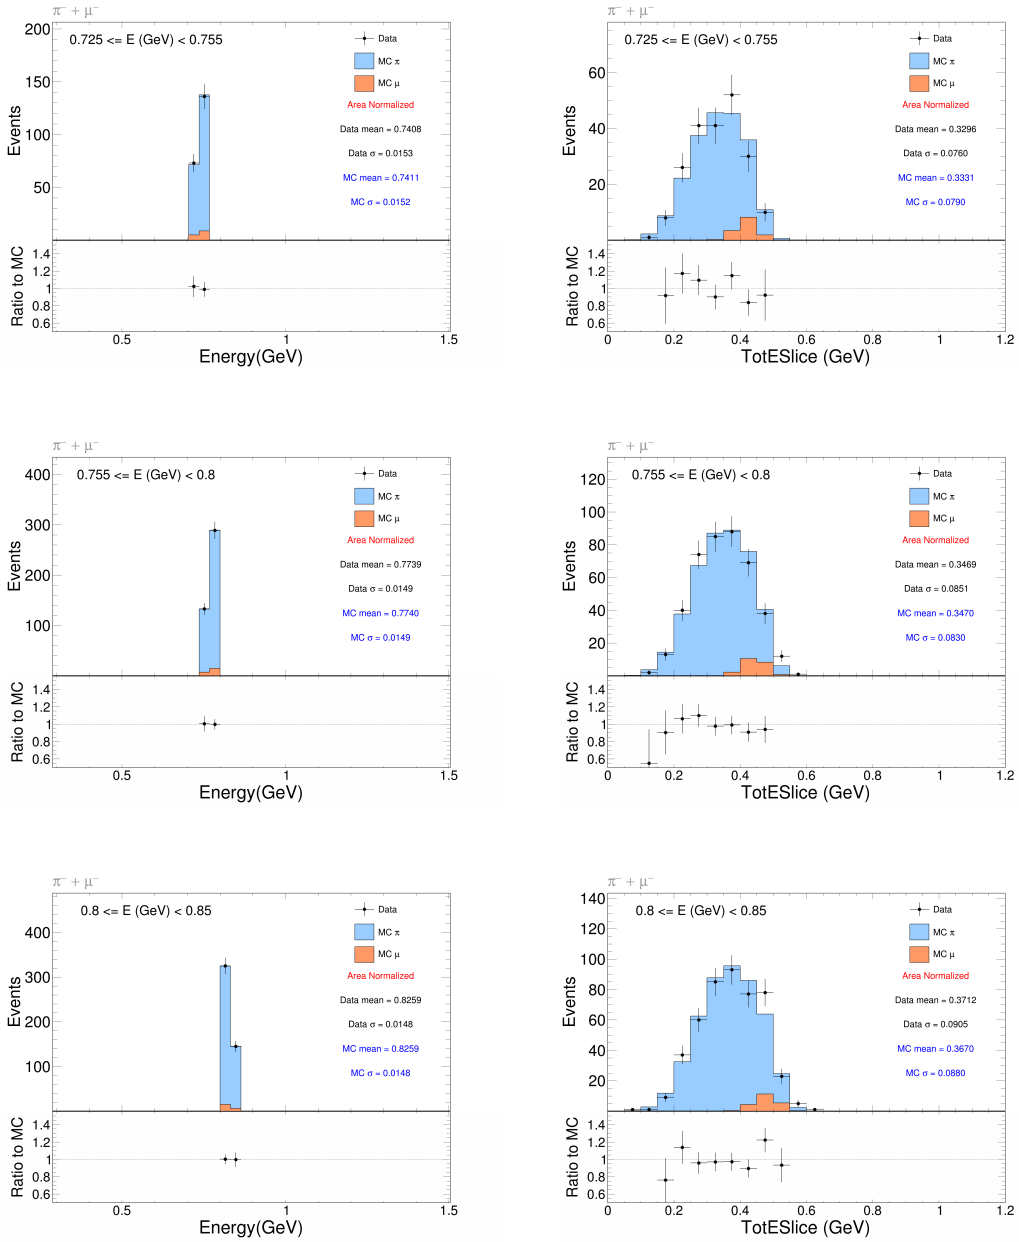

Figure 7.2: Incident energy (left) and reconstructed energy (right) distributions for negatively charged pions in data and simulation, shown for incident energy ranges from 0.725 to 0.85 GeV.

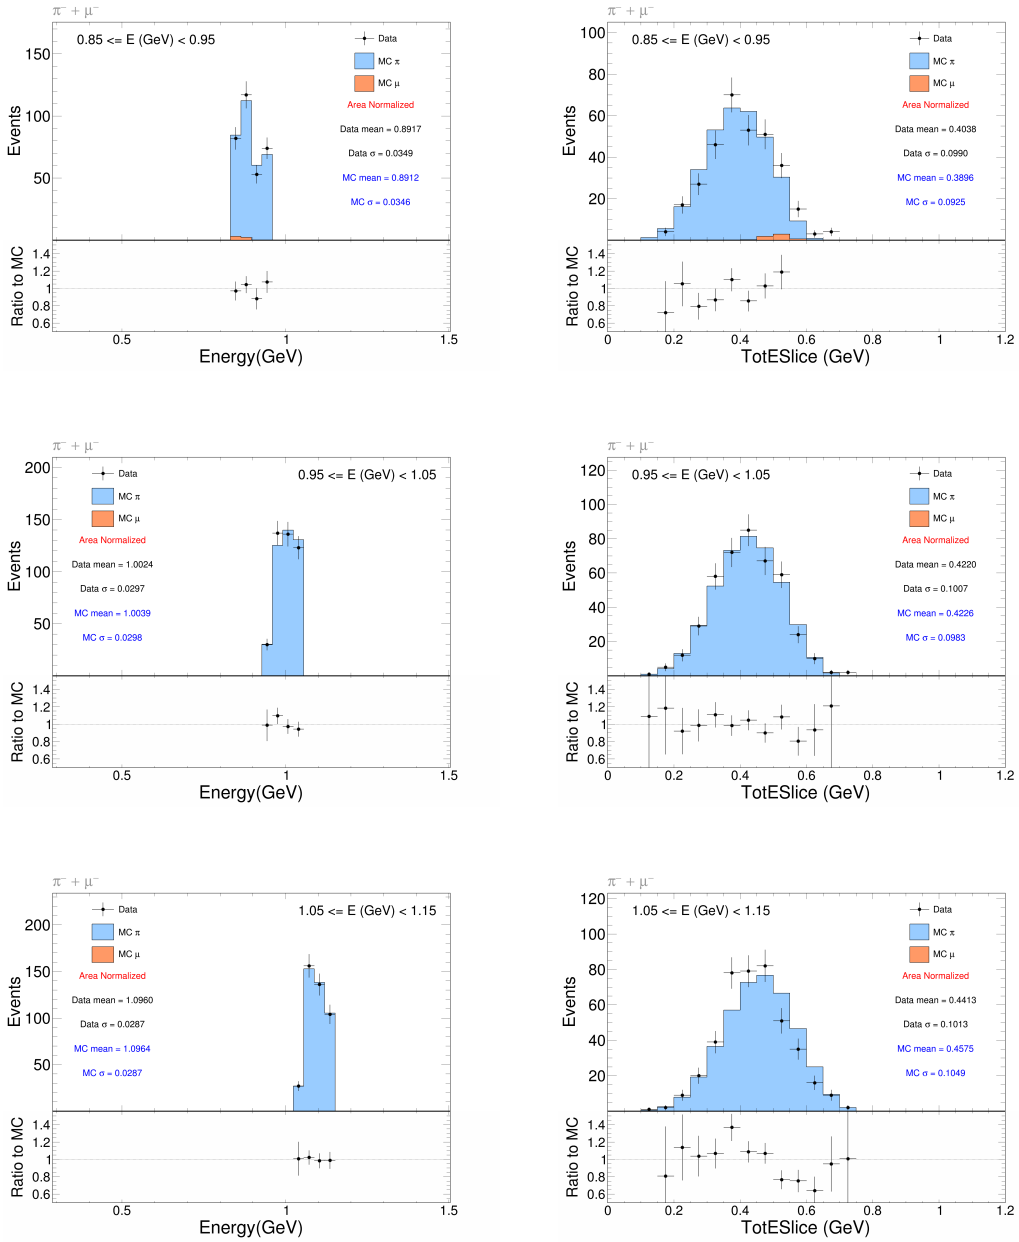

Figure 7.3: Incident energy (left) and reconstructed energy (right) distributions for negatively charged pions in data and simulation, shown for incident energy ranges from 0.85 to 1.15 GeV.

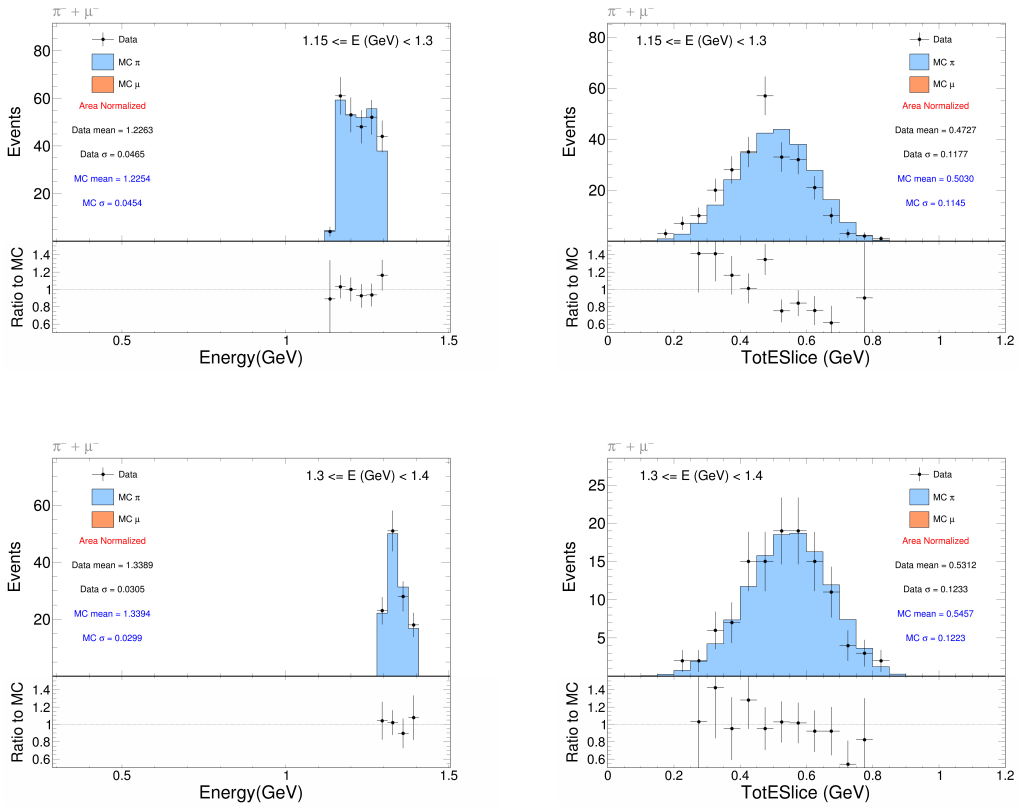

Figure 7.4: Incident energy (left) and reconstructed energy (right) distributions for negatively charged pions in data and simulation, shown for incident energy ranges from 1.15 to 1.4 GeV.

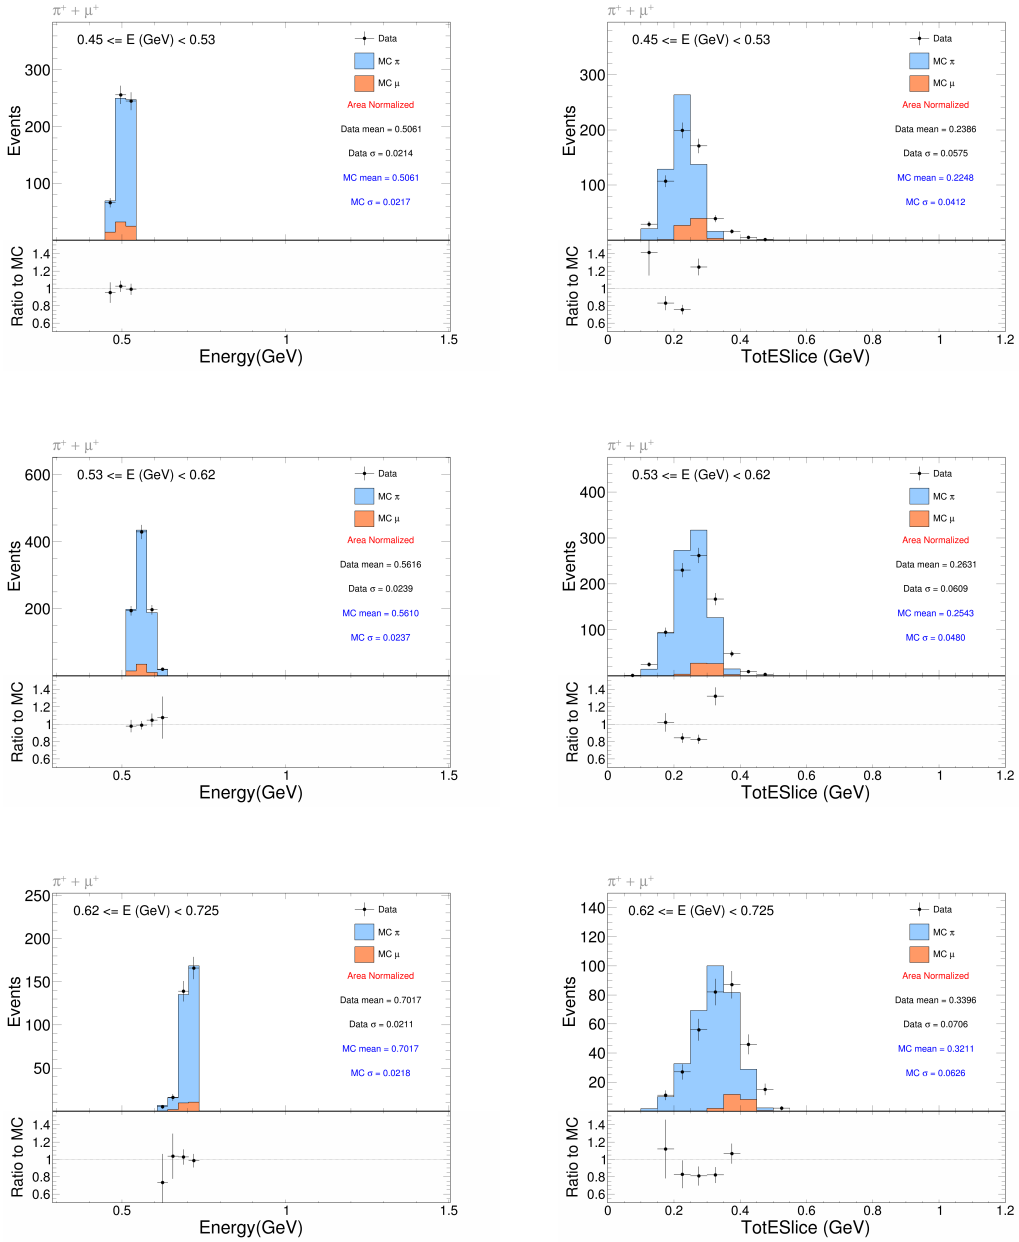

Figure 7.5: Incident energy (left) and reconstructed energy (right) distributions for positively charged pions in data and simulation, shown for incident energy ranges from 0.45 to 0.725 GeV.

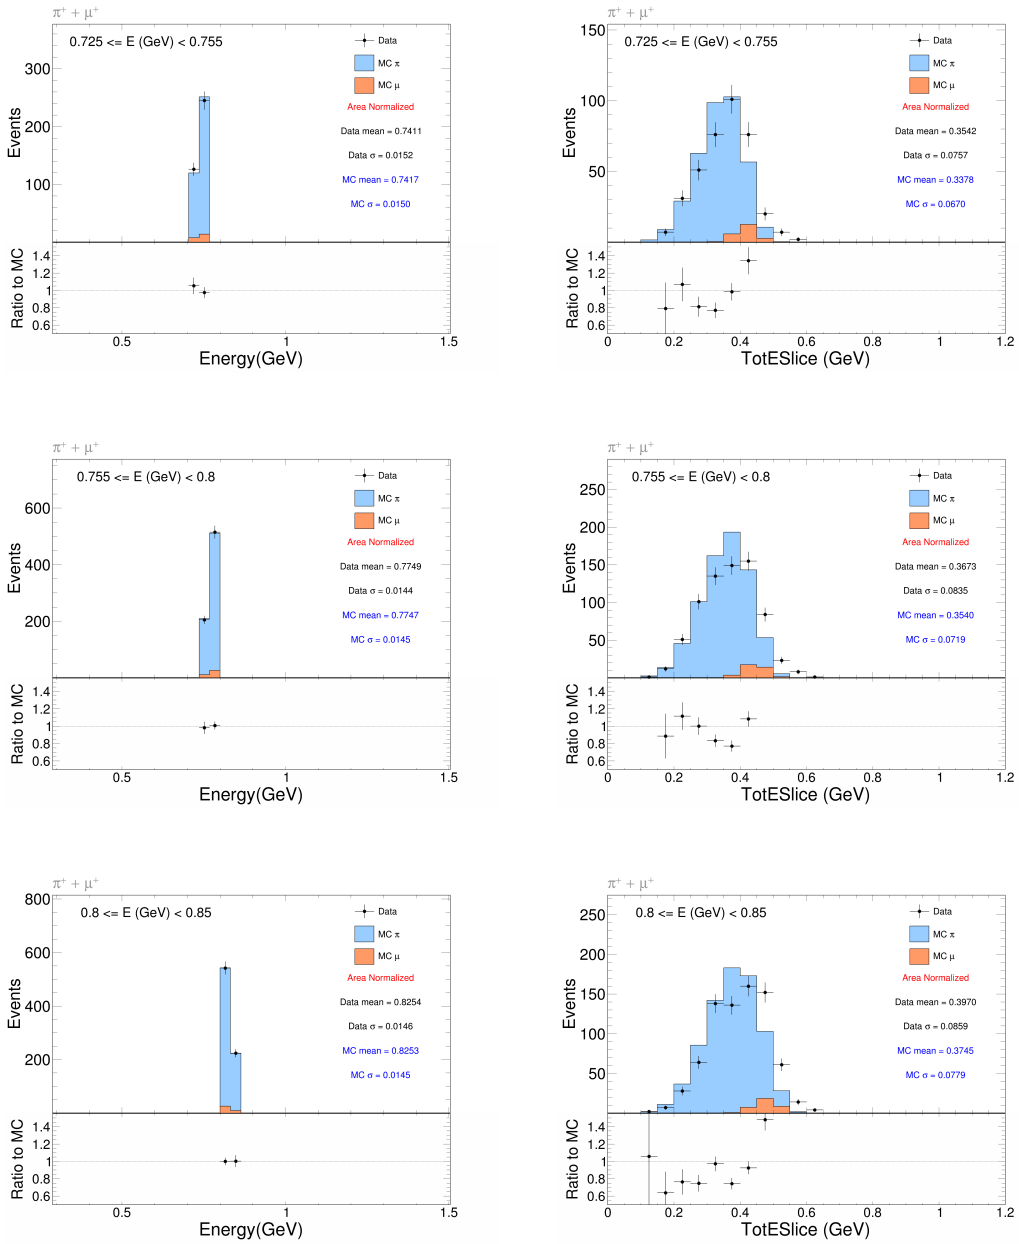

Figure 7.6: Incident energy (left) and reconstructed energy (right) distributions for positively charged pions in data and simulation, shown for incident energy ranges from 0.725 to 0.85 GeV.

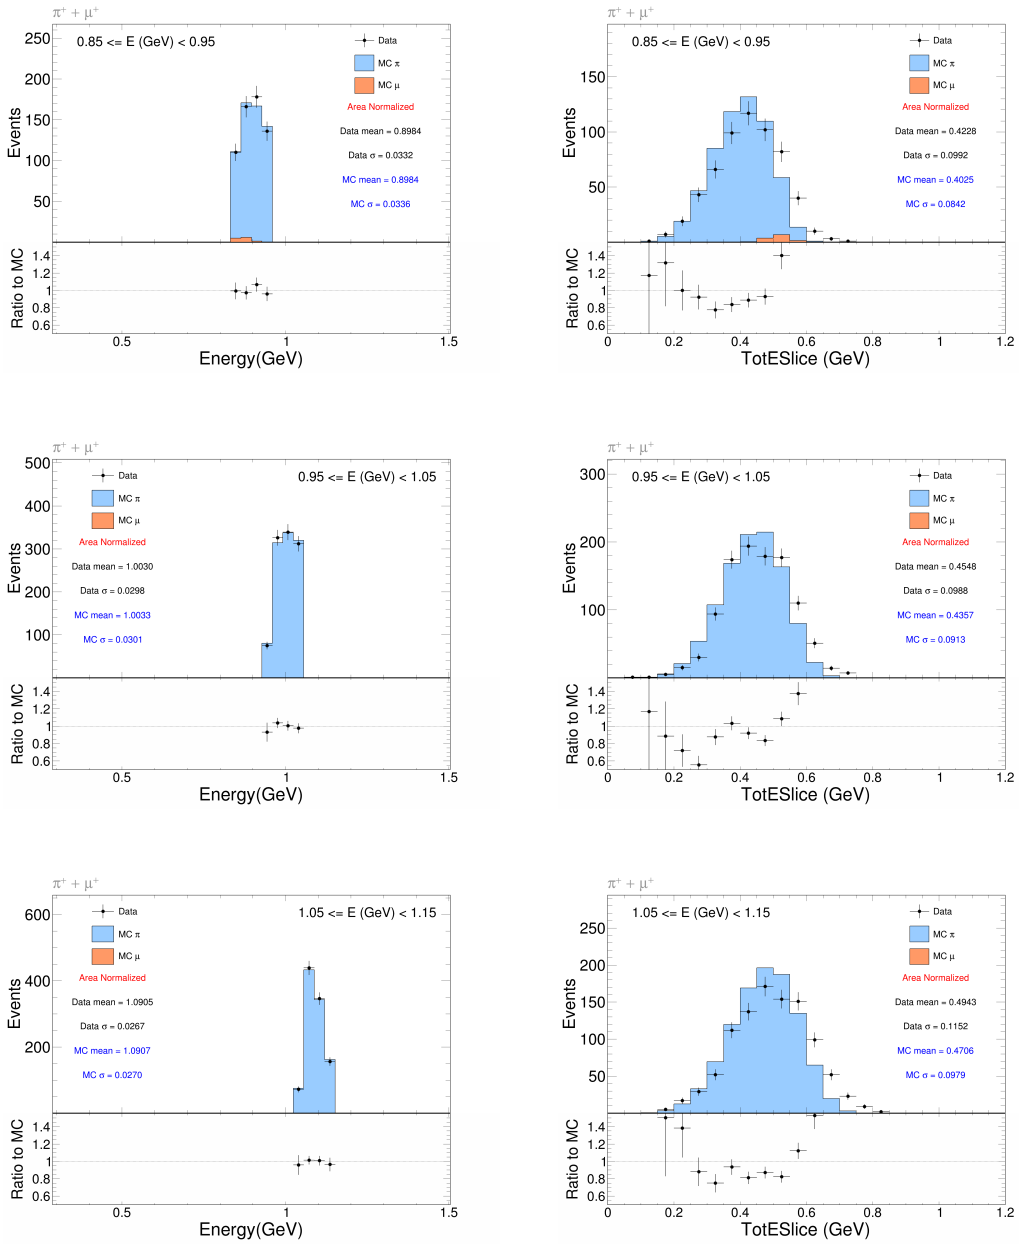

Figure 7.7: Incident energy (left) and reconstructed energy (right) distributions for positively charged pions in data and simulation, shown for incident energy ranges from 0.85 to 1.15 GeV.

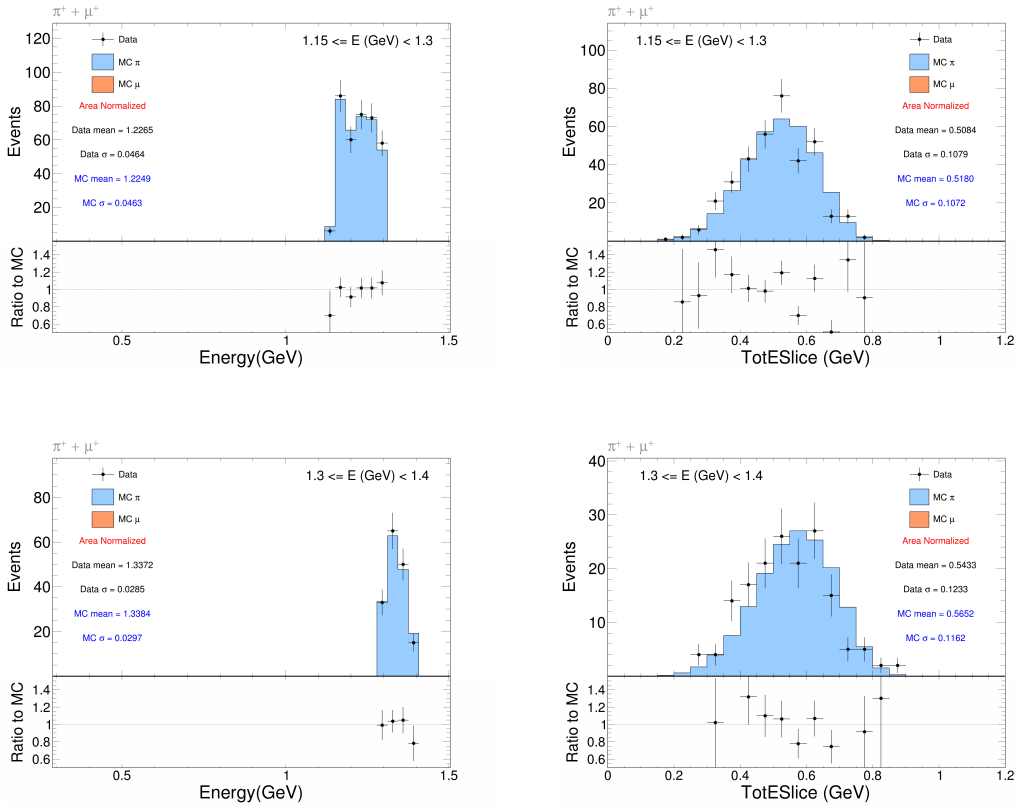

Figure 7.8: Incident energy (left) and reconstructed energy (right) distributions for positively charged pions in data and simulation, shown for incident energy ranges from 1.15 to 1.4 GeV.

Energy Region (GeV)	**Number of Events**	
	$(\pi^+ + \mu^+)$	$(\pi^- + \mu^-)$
$0.45 \leq \mathrm{E} < 0.53$	567	267
$0.53 \leq \mathrm{E} < 0.62$	840	304
$0.62 \leq \mathrm{E} < 0.725$	326	208
$0.725 \leq \mathrm{E} < 0.755$	371	209
$0.755 \leq \mathrm{E} < 0.8$	720	422
$0.8 \leq \mathrm{E} < 0.85$	766	470
$0.85 \leq \mathrm{E} < 0.95$	590	326
$0.95 \leq \mathrm{E} < 1.05$	1052	426
$1.05 \leq \mathrm{E} < 1.15$	1013	423
$1.15 \leq \mathrm{E} < 1.3$	358	262
$1.3 \leq \mathrm{E} < 1.4$	163	120
Total	**6766**	**3437**

Table 7.1: pion-muon events by energy region.

Figure 7.9 shows the energy response results for pions and muons in the sample. Each data point represents the energy response for each incident energy region discussed in the previous subsection. The error bars in the data are statistical and are calculated as $\sigma/\sqrt{N}$ for each sample. Here, the statistical errors in the simulation are negligible. The total systematic uncertainty is also included and is shown around the mean response in the simulation. Section 7.3 fully discusses the different systematic uncertainties included in the analysis.

For negative pions and muons, the energy response tends to be flat up to approximately 0.8 GeV in incident energy. From there, it decreases by up to $\sim$ 4% in energy response for the largest incident energy. The data and simulation show good agreement for incident energies $\leq$ 1 GeV. A larger disagreement is observed for incident energies within 1.25 and 1.3 GeV, where the maximum difference is around 2.5% in energy response compared to the nominal simulation.

Similarly, positive pions and muons tend to have a flat energy response at lower energies, with a decrease at higher energies. In contrast to the negatively charged particles, data and simulation for these samples show larger differences. Here, the data consistently have a higher energy response compared to simulation, with a slight improvement at incident energies beyond 1.2 GeV.

Since the main focus is to study the pion energy response, the estimated muon background is subtracted from each energy response distribution. The results of the energy response after background subtraction can be seen in Figure 7.24. By isolating the pion sample, the mean energy response is affected. This is especially noticeable in the lower incident energy regions, where most of the

muon background is located.

The observed difference is a reduction in the mean energy response when comparing the pion-plus-muon to the pion-only sample. The response reduction in the pion-only samples reaches a maximum of $\sim 1\%$ at lower energies and becomes completely suppressed at incident energies ≥ 1 GeV. These differences can be observed in Figure 7.11. In this figure, it can be seen that the differences in data between the pion-only and pion-plus-muon samples are well within the statistical error.

Another essential point to note is that for the positively charged particles, subtracting the muon background does not affect the observed differences. The data for the pion-only sample show a mean energy response that is significantly higher in data compared to simulation across most of the incident energy range.

In addition, these differences are not captured by the systematic error band, which may indicate a possible mismodeling of pion interactions in Geant4. Particularly cross section modeling. The next section covers the discussion of systematic uncertainties, where alternative models are also implemented in the simulation and compared to the nominal model used in the NOvA experiment.

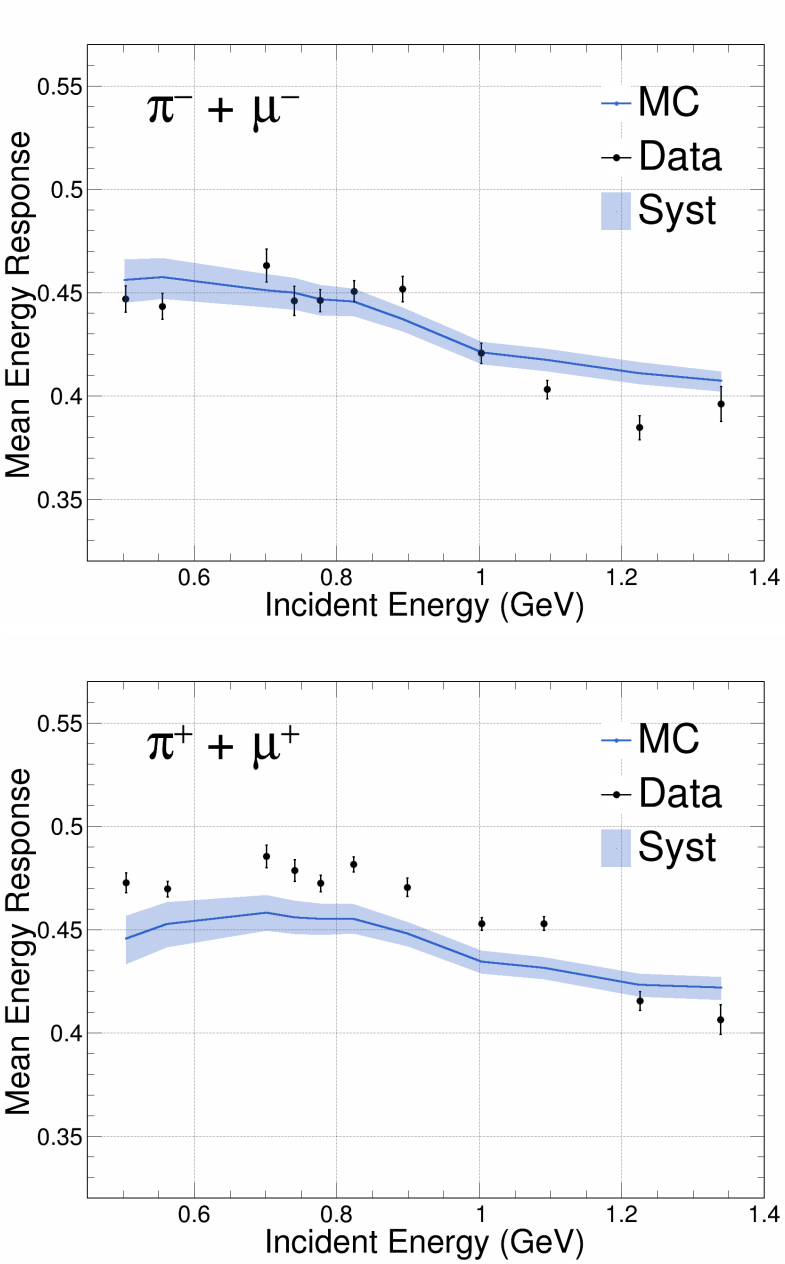

Figure 7.9: Mean energy response as a function of incident energy for negatively charged (top) and positively charged (bottom) pion and muon samples. The data points are shown in black and compared to simulated predictions in blue, with the shaded bands indicating systematic uncertainties. For $\pi^- + \mu^-$, the data generally follow the simulation trend but show some disagreement at higher energies. In contrast, the $\pi^+ + \mu^+$ response differs more noticeably from the simulation across most of the energy range.

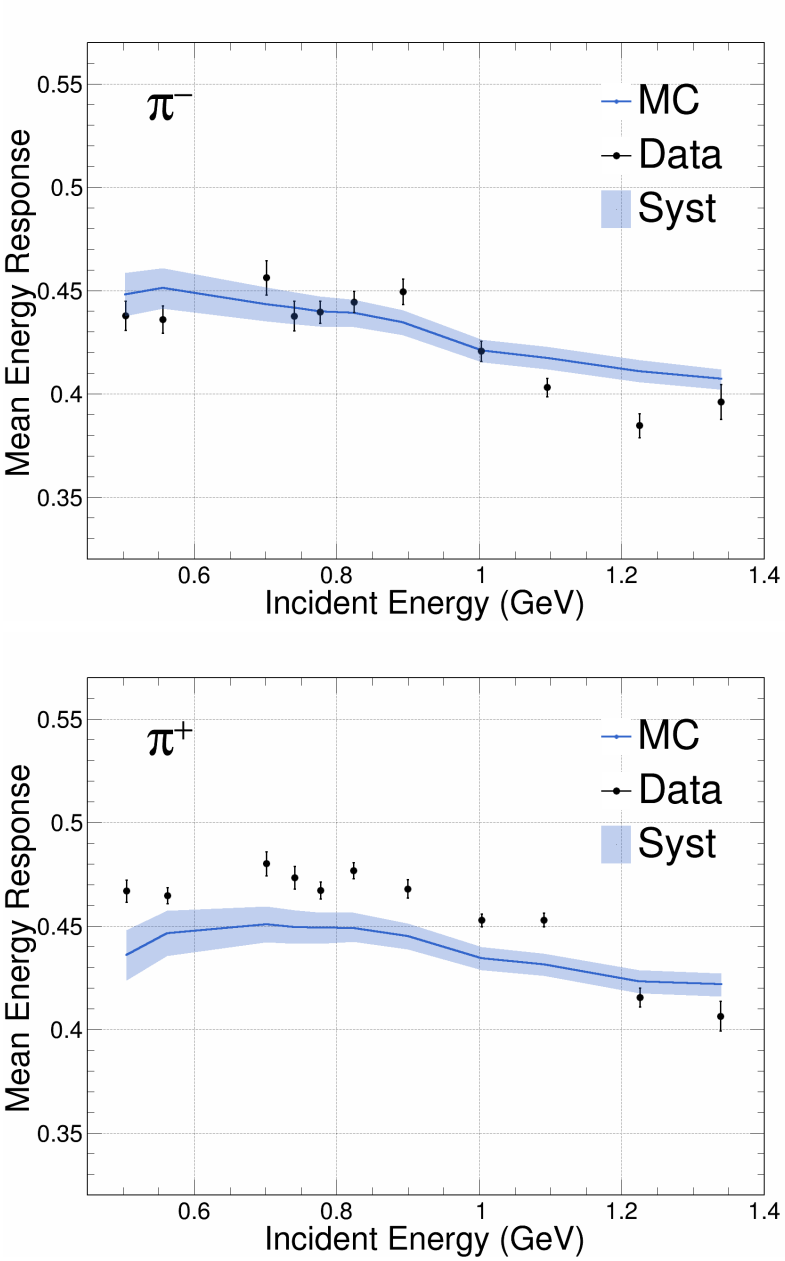

Figure 7.10: Mean energy response as a function of incident energy for negatively charged (top) and positively charged (bottom) pion samples. The data points are shown in black and compared to simulated predictions in blue, with the shaded bands representing systematic uncertainties. For π^-, the data closely follow the simulation with some deviation at higher energies. In contrast, the π^+ sample exhibits a more consistent overestimation of the response in data relative to simulation across much of the energy range.

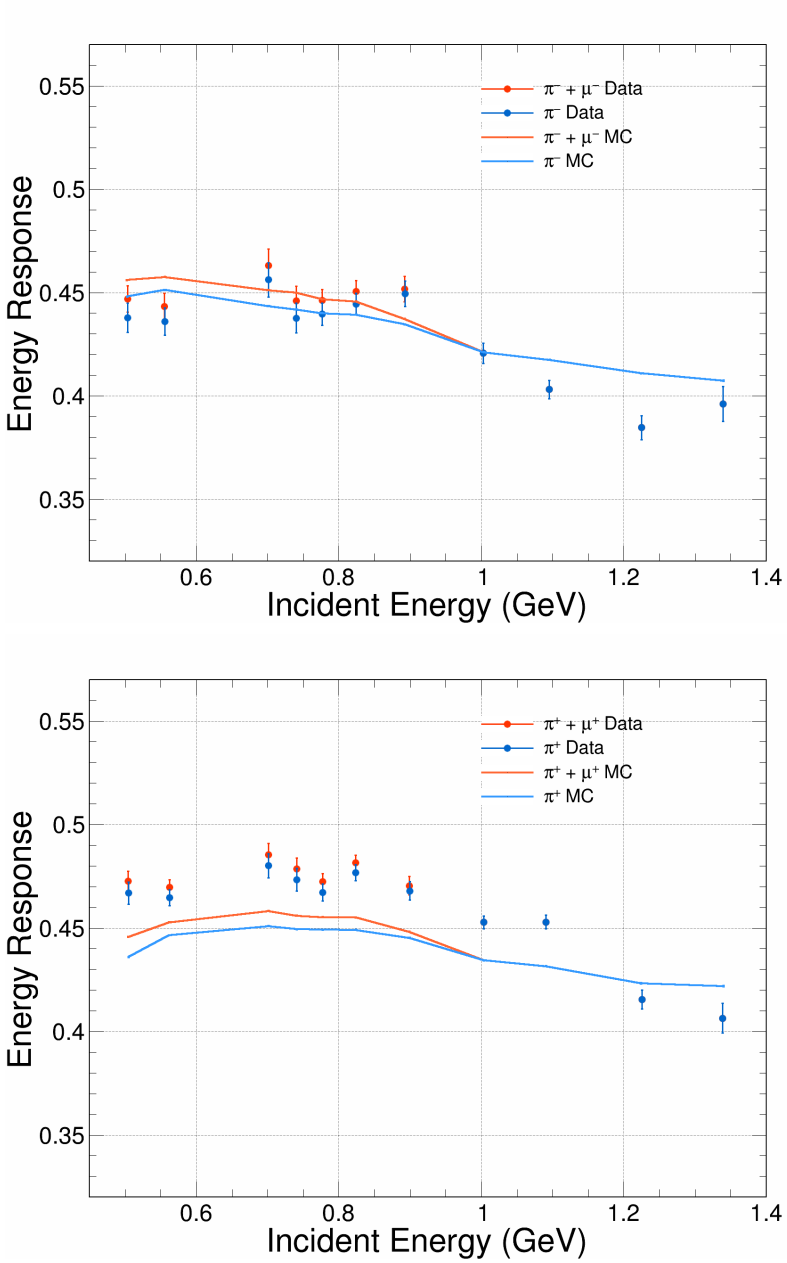

Figure 7.11: Mean energy response for pion-only and pion+muon samples as a function of incident energy, for negatively charged (top) and positively charged (bottom) particles. Subtracting the muon background reduces the response in both cases. While the π^- response remains consistent across samples, the π^+ sample still shows a significant data–simulation discrepancy after subtraction.

7.3 Systematic Uncertainties

Systematic uncertainties play an essential role in the interpretation of the pion energy response results discussed in subsection 7.2.2. This section presents the methods used to study these uncertainties and determine their effect on the pion energy response. In this analysis, systematic uncertainties can be divided into two main categories: Beamline Systematics and Detector Systematics. These systematic groups will be discussed in the following subsections.

7.3.1 Beamline Systematics

Beamline systematics assess the level of uncertainty in the reconstructed beamline momentum and the background content in the pion sample. This includes the energy loss in the beamline material uncertainty, the dipole magnet uncertainty, wire chamber uncertainty, and muon background estimation.

Pion Energy Loss Uncertainty

The procedure to estimate the energy loss for pions in the downstream beamline material is fully described in subsection 6.3.3. This method provides a correction for these energy losses for pions that either hit or missed DS TOF 2. These corrections are subsequently applied to the beamline data to obtain the corrected beamline momentum. Since it is not possible to accurately separate these two samples in real data, a systematic uncertainty due to this misclassification must be estimated.

To determine the uncertainty, two sets of samples for the corrected beamline momentum are compared. For the first set of samples, the momentum correction due to energy loss is applied using the truth information in the simulation. The truth information provides the exact location of particles entering the front face of the detector. This makes the distinction between particles that hit and missed DS TOF 2 straightforward. Thus, the energy loss correction on these samples is precisely applied. These samples are named "Perfect corrected samples."

For the second set of samples, the correction is applied without using true information. This correction relies on the geometry of the DS TOF 2 detector as reference and the projected beamline tracks on the front face of the detector to determine if a pion hit or missed DS TOF 2. In subsection 6.3.3, it was determined that the most optimal size of the selection cut to separate between particles that hit and missed DS TOF 2 was -3 mm (meaning a DS TOF 2 detector reduced 3 mm in the x and y directions). This second set of samples simulates the momentum correction applied to data.

The comparisons of the perfect momentum corrections and the corrected momentum for different magnet current settings can be seen in Figure 7.12. One can note a good agreement between these samples, where the variations on the mean are < 1 MeV. To account for the overall uncertainty, the ratio in the figures is used. For this calculation, the mean around the 1 σ range of the distribution is obtained ($[Mean - \sigma/2, Mean + \sigma/2]$). This region was chosen to avoid the large statistical errors on the tails of each distribution.

The means for each distribution are shown in Figure 7.13. Here, a maximum of 0.4% disagreement is observed. This value is used to shift the momentum

up and down in the single particle simulation to account for the correction in the beam momentum. This approach provides two individual pion samples to study the effect of this uncertainty in the pion energy response.

The results presented in this study successfully validate the method of momentum correction applied in data using the projected beamline tracks on the front face of the detector. In addition, the method assigns a small uncertainty of 0.4% on the corrected beamline momentum.

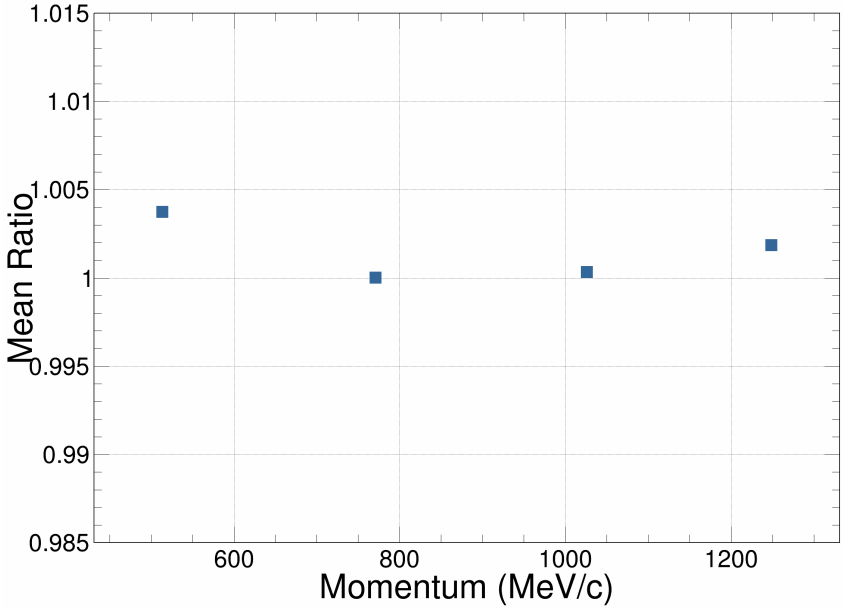

Figure 7.13: Mean ratio of reconstructed momentum between the "Projected corrected samples" and the "Perfect corrected samples" as a function of beam momentum. The ratios are computed over the 1σ range of each distribution to minimize the impact of statistical fluctuations in the tails. A maximum deviation of 0.4% is observed across all settings. This value is taken as the systematic uncertainty on the corrected beamline momentum and is used to shift the momentum up and down in single particle simulations to assess its impact on the pion energy response.

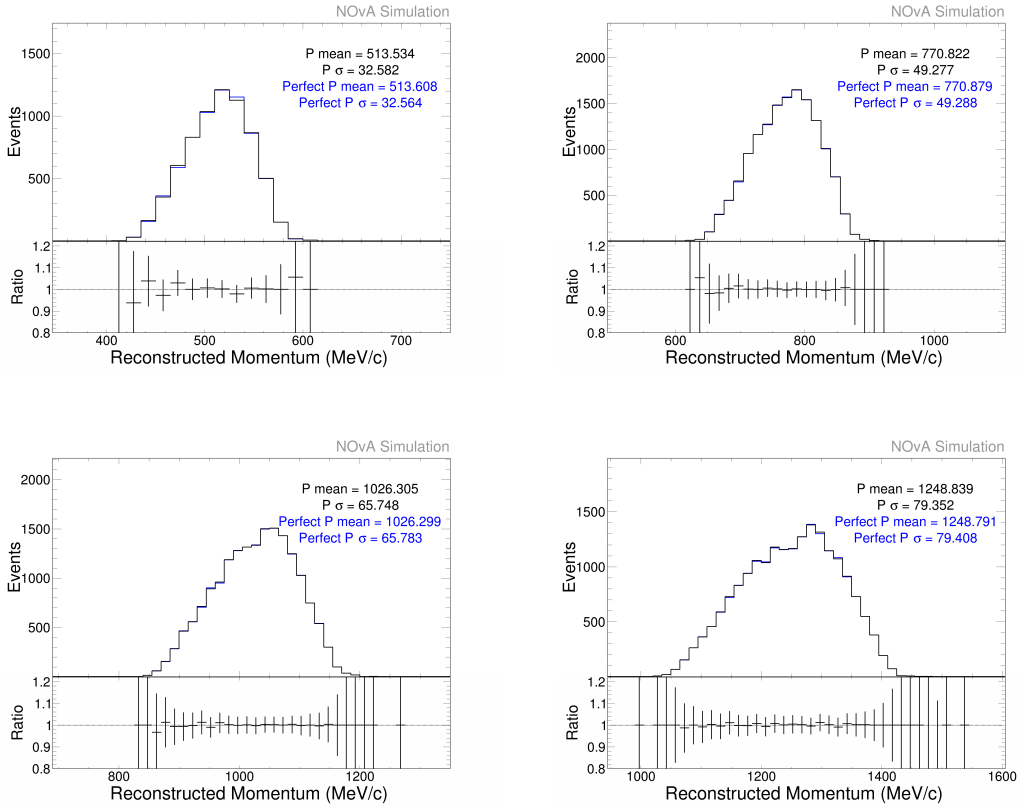

Figure 7.12: Comparison of reconstructed momentum distributions for "Perfect corrected samples" (blue) and "Projected corrected samples" (black) for four beam settings: 500 A (top left), 750 A (top right), 1000 A (bottom left), and 1250 A (bottom right). The perfect correction uses true simulation information to classify whether pions hit or missed DS TOF 2, while the projected correction relies on detector geometry and projected track information, emulating the procedure used in data. The lower panels show the ratio between the two methods. The reconstructed momentum means agree to within less than 1 MeV, validating the projected correction approach

Dipole magnet uncertainty

The relationship between applied current and the resulting magnetic field in the dipole magnet can be affected by magnetic hysteresis, where the field depends not only on the current but also on how the current was varied over time. This effect can lead to small deviations in the effective magnetic field, which in turn impact the beam momentum calibration.

To estimate the uncertainty due magnetic hysteresis. First, the integrated magnetic field along the magnet axis was measured as a function of the magnet current using a rotating coil placed in the magnet's air gap (see Figure 7.14).

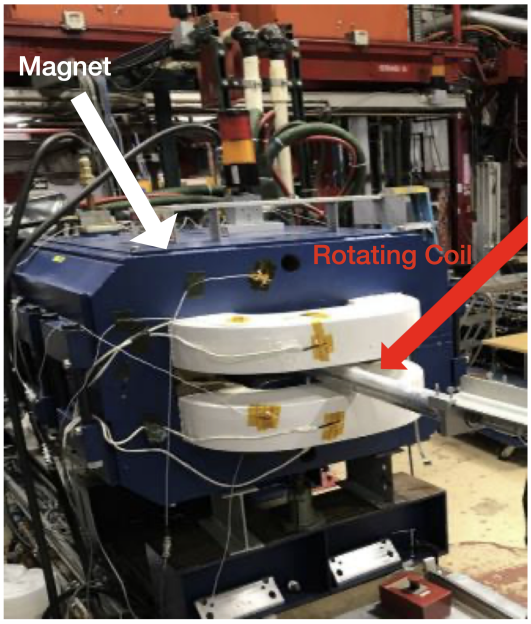

Figure 7.14: Test Beam dipole Magnet and rotating coil. In this picture the rotating coil is placed within the magnet's air gap along the z axis.

The coil rotated at constant angular velocity with its axis aligned along the beam direction. As the coil rotated, the changing magnetic flux through its

area induced a voltage (Faraday's law), which provided with measurement of the integrated magnetic field along the axis. This was performed in steps of 100 A from 0 to 3000 A and then back to 0 A [103].

Figure 7.15 shows the integrated magnetic field curves in function of the current. This results were obtained by ramping up the magnet's current up to 3000 A (lower part of the curves) and back down to 0 A (upper part of the curves) in steps of 100 A. The effect in the curve's behavior is noticeable and can be seen in the zoom in plots. The magnitude of this effect is shown in Figure 7.16. The curve is calculated by taking the ratio of the upper to lower curves minus one, representing the difference divided by the lower value.

We can use this resulting curve to estimate the uncertainty. Since NOvA's dipole magnet operated with currents from 500 A to 1250 A, a 1% systematic uncertainty is assigned to account for possible hysteresis-related deviations not captured by the current reading alone. To study the effect on the energy response, single-particle simulated samples were created, shifting the incident momentum up and down by 1%.

Wire Chamber Uncertainty

A systematic uncertainty on the reconstructed beam momentum is assigned to account for possible misalignments or resolution effects in the wire chambers. This estimate is based on a dedicated study in which hit positions were shifted by $\pm$ 1 mm in each chamber. This shift size is comparable to the wire spacing in the wire chambers.

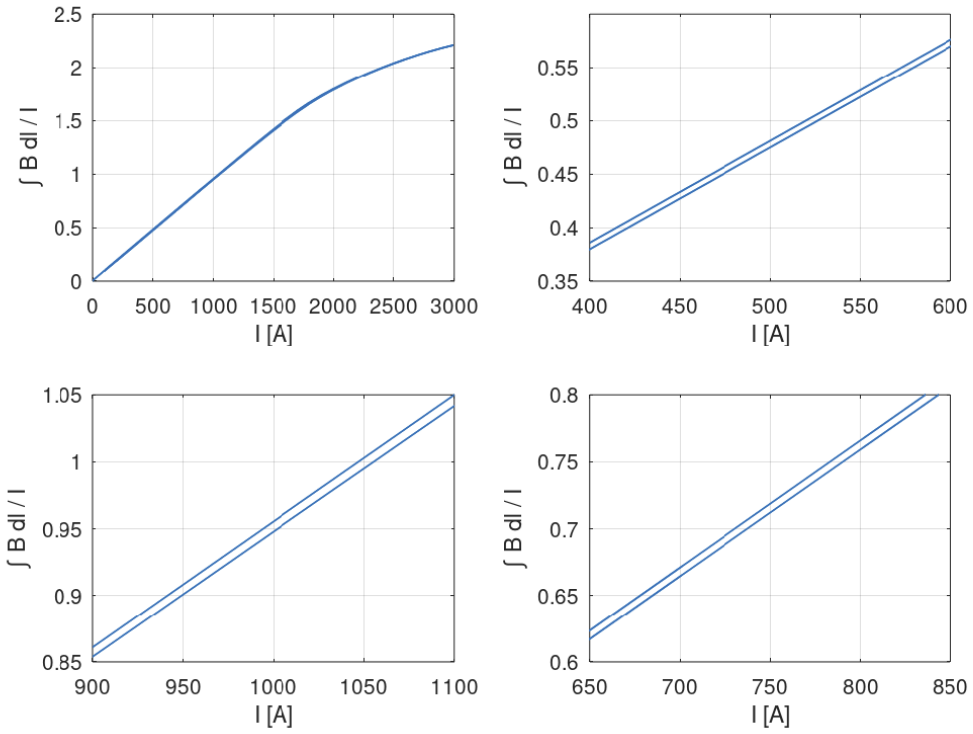

Figure 7.15: Integrated magnetic field as a function of current, measured with a rotating coil during current ramp-up (lower curves) and ramp-down (upper curves), showing the hysteresis effect. The zoomed panels show the deviations at different current settings.

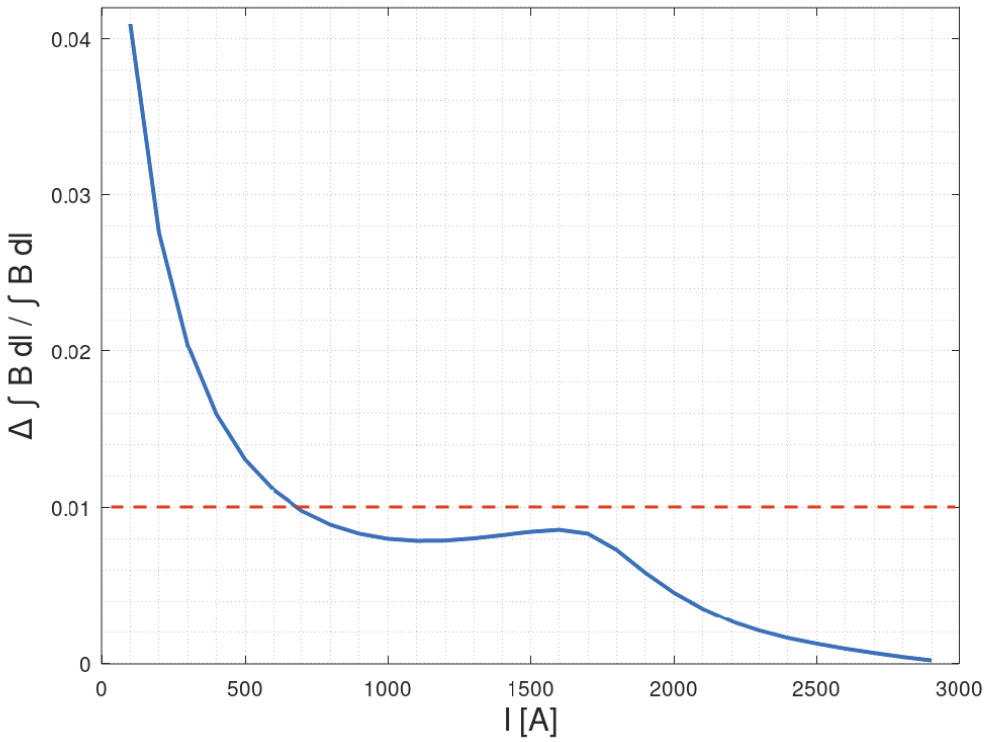

Figure 7.16: Relative difference between the integrated magnetic field during current ramp-up and ramp-down, calculated as $\left(\frac{\text{up}}{\text{down}} - 1\right)$. The red dashed line marks the assigned 1% systematic uncertainty within the NOvA magnet's operating range (500 A to 1250 A).

The resulting changes in reconstructed momentum were measured (Figure 7.17 and Table 7.2), and the mean shifts for each chamber were then combined in quadrature, resulting in a total uncertainty of approximately 11 MeV/c for the beamline momentum, this is equivalent to 2.2% at 500 MeV/c and 0.8% at 1250 MeV/c momentum. This value is taken as the systematic uncertainty associated with wire chamber position resolution and alignment. Similar to all beamline systematics, to study the effect on the energy response due to this systematic, single particle simulated samples were created by shifting the incident momentum up and down by 11 MeV/c [23].

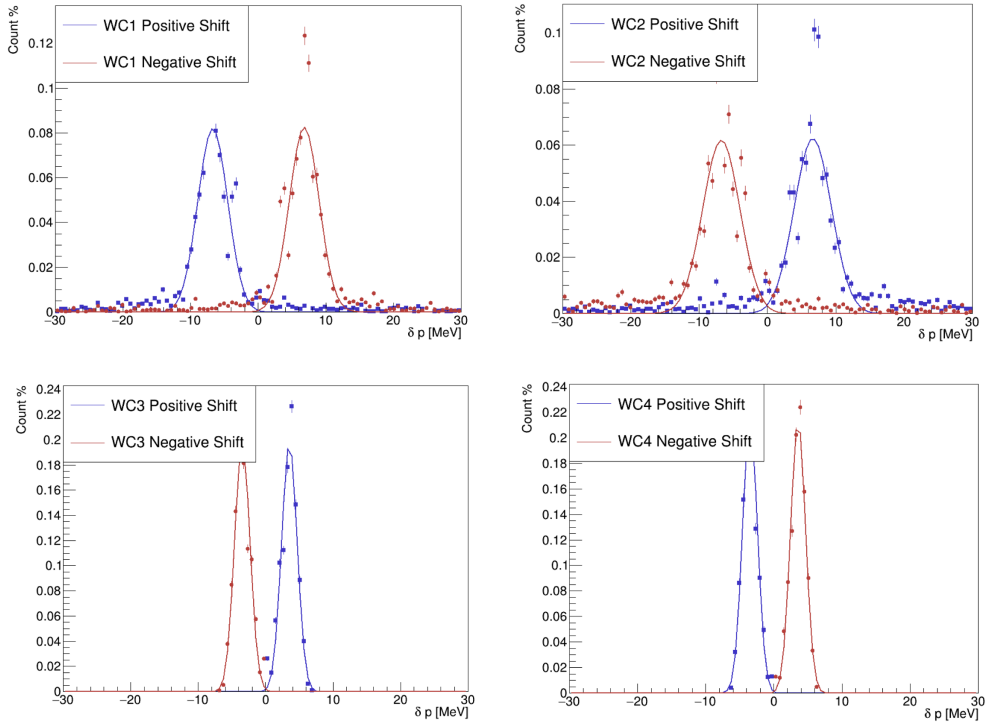

Figure 7.17: Impact of ±1 mm hit position shifts in each wire chamber (WC1–WC4) on the reconstructed beam momentum. The distributions show the momentum shift (δp) for positive (blue) and negative (red) displacements. The measured mean shifts were used to estimate a combined uncertainty of 11 MeV/c on the beamline momentum [23].

WC	Mean Positive Shift (MeV/c)	Mean Negative Shift (MeV/c)
1	6.826 ± 0.029	-6.760 ± 0.029
2	-6.801 ± 0.040	6.658 ± 0.040
3	3.538 ± 0.014	-3.510 ± 0.014
4	-3.545 ± 0.014	3.577 ± 0.014

Table 7.2: Mean momentum shifts for positive and negative hit displacements (±1 mm) of each wire chamber (WC1–WC4).

The uncertainties mentioned up to here have a direct impact on the reconstructed pion beamline momentum. The wire chamber uncertainty is the largest one up to ∼1 GeV in incident energy when comparing to the pion energy loss and magnet uncertainties. The impact on the energy response for positive and negative charge pions is shown in Figure 7.18. The ratio to nominal shows the error band associated to these systematics added in quadrature, with a maximum of 2.8% effect on lower energies decreasing up to a minimum of 1.3% for higher energies for positive pions and with a maximum of 2.3% and a minimum 1.2% for negative charge pions.

Muon Background Uncertainty

As mentioned in subsection 7.2.2, the muon background has a small impact on the energy response for lower incident energies, this background is subtracted bin-by-bin in each of the energy response distributions to isolate the pion sample.

Since the muon sample is obtained using the G4Beamline simulation, and there are no other control samples in data to perform validation of this component, a normalization uncertainty associated to the total number of expected muon backgrounds in the pion plus muon sample is implemented.

To achieve this, each simulated distribution (pions and muons) is scaled to the data counts, and the statistical error on the scaled muon background sample (nominal background) is then used as a normalization factor to scaled the nominal muon distributions up and down, the shifted muon distributions are then subtracted from nominal pion and muon samples and the residual effect of this calculation is taken as the systematic uncertainty.

Figure 7.19 Shows the effect of this systematic uncertainty on the energy response for positive and negative charge pions. The effect is energy dependent with a maximum rise of 0.4% for positive and 0.5% for negative charge pion, due to the absence of muon backgrounds at higher energies the uncertainty gets fully suppressed at incident energies beyond 1 GeV.

7.3.2 Detector Systematics

In subsection 4.1.4, we fully discussed the main detector systematic uncertainties implemented for the ND and FD for neutrino oscillation analysis. This includes energy response, light level, and Cherenkov efficiencies uncertainties.

For the Test Beam Experiment analysis, we include the light level and Cherenkov efficiencies uncertainties. The energy response uncertainty is not included here since one of the main purposes of the Test Beam experiment is to address the

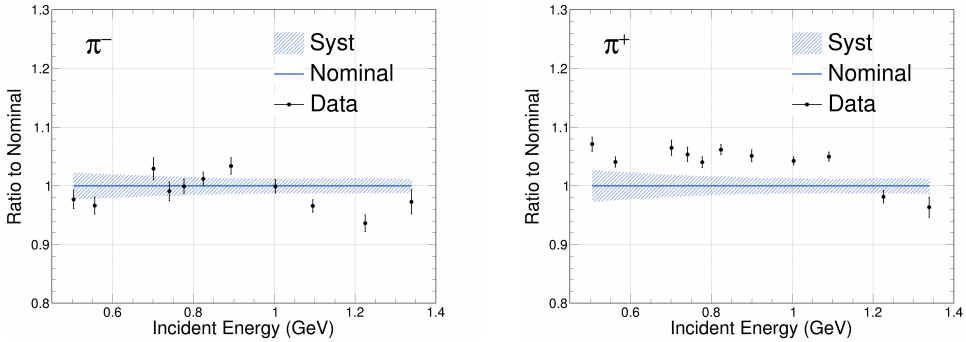

Figure 7.18: Impact of beamline-related systematics on the energy response for negative (left) and positive (right) pions. The shaded blue band shows the total uncertainty due to wire chamber misalignment, pion energy loss, and magnetic field variations, combined in quadrature. The ratio is shown with respect to the nominal response. The total effect reaches a maximum of 2.3% for π^- and 2.8% for π^+ at low energies, decreasing to 1.2% and 1.3% respectively at higher energies.

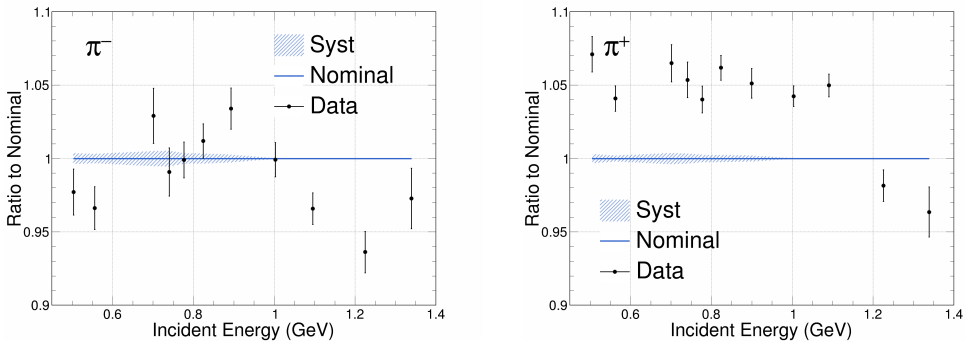

Figure 7.19: Impact of the muon background normalization uncertainty on the energy response for negative (left) and positive (right) pions. The shaded band shows the residual variation in the energy response after scaling the nominal muon background up and down using its statistical normalization uncertainty. The effect is more prominent at lower energies, reaching a maximum of 0.5% for π^- and 0.4% for π^+. The uncertainty vanishes above 1 GeV due to the absence of muon contamination in that region.

energy response uncertainties of single particles.

The light level systematic addresses the effects of thresholds (how much light is needed to produce a hit) in the detector, whereas the Cherenkov efficiency uncertainty focuses on how efficiently we can simulate the Cherenkov light produced by particles interacting in the detector. To estimate these uncertainties, the same procedure as for the ND and FD is used. This approach uses the light level tuning to extract the uncertainties. This leads to a 5% for the light level and approximately 6.2% for the Cherenkov uncertainty.

The effects of the Cherenkov efficiency uncertainty on the energy response for positive and negative charge pions can be seen in Figure 7.20. The impact of this uncertainty is relatively small, with a maximum of 0.4% effect. On the other hand, the impact of the light level systematic, as shown in Figure 7.21, is negligible, with a maximum of 0.06% effect on the energy response.

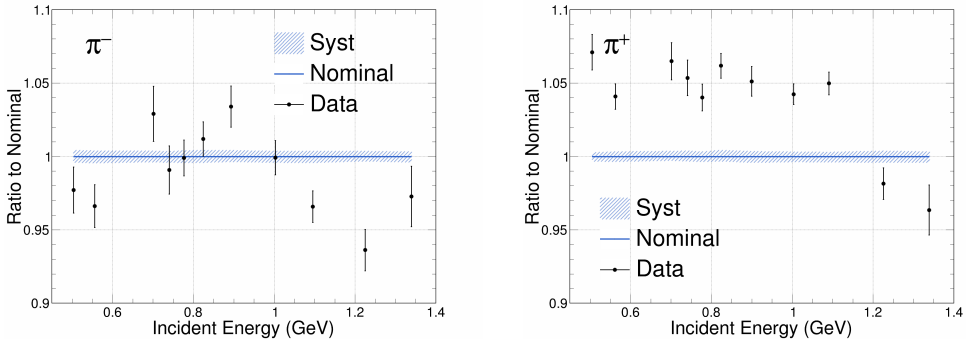

Figure 7.20: Impact of the Cherenkov efficiency uncertainty on the energy response for negative (left) and positive (right) pions. This systematic accounts for uncertainties in the modeling of Cherenkov light production and detection. The overall effect is small, reaching a maximum variation of 0.4%.

Other systematic uncertainties were studied, including the impact on energy response by changing the hadronic models in Geant4. The detector simulation

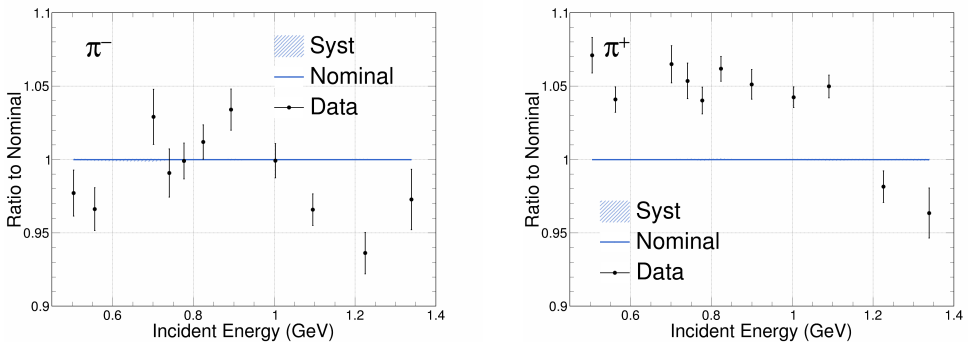

Figure 7.21: Impact of the light level systematic uncertainty on the energy response for negative (left) and positive (right) pions. This uncertainty reflects the effect of the detector threshold for producing a hit. The total impact on the energy response is negligible, with a maximum variation of 0.06%.

used the Bertini Cascade model (BERT) [56] as the nominal simulation for hadronic interactions. To study the effect of different models, simulations were generated using the Intra-Nuclear Cascade Model (INCLXX) [104] and Binary Cascade Model (BIC) [105].

The results of switching the hadronic physics list in Geant4 from BERT to BIC are shown in Figure 7.22. The effect on the energy response is negligible. The nominal simulation and the BIC simulation lay on top of each other. This shows that the BIC model is not able to capture the largest differences observed in data, especially in the positive charge pion sample. In contrast, we can observe in Figure 7.23 that the INCLXX has large variation across the incident energy, with a smaller response at lower energies and a large variation at higher energies. Even though the variations are large, this model does not provide a good description of the data.

The study of these two alternative hadronic models provides insight into the limitations of both the BERT model and the alternatives in accurately re-

producing the pion energy response observed in data. These studies are not intended to estimate systematic uncertainties, but rather to evaluate how well each model performs in this analysis.

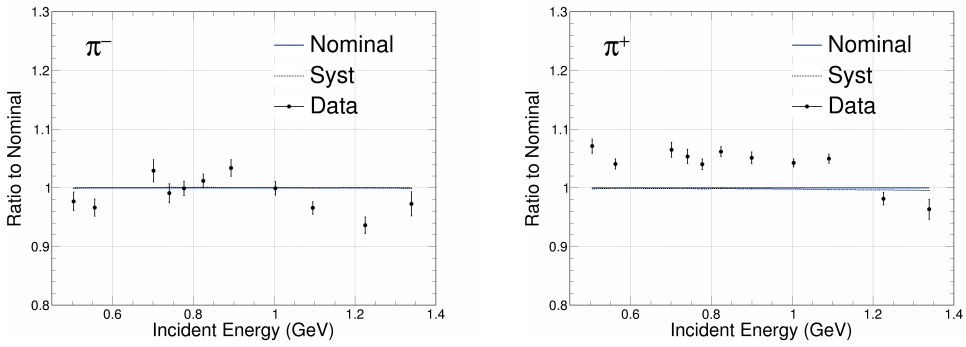

Figure 7.22: Ratio to nominal energy response for negatively (left) and positively (right) charged pions when switching the hadronic model in Geant4 from BERT to BIC. The negligible deviation between the two simulations indicates that BIC reproduces the BERT behavior closely and does not capture the data–simulation discrepancy, particularly in the π^+ sample

Summary on Systematic Uncertainties

A summary of the impact on the energy response from all systematic uncertainties is shown in Table 7.3. For the beamline systematics, the largest contribution comes from the combined effect of the pion energy loss correction, dipole magnet uncertainty, and wire chamber alignment.

This combination results in an energy-dependent uncertainty that reaches up to 2.8% for π^+ and 2.3% for π^- at lower energies and decreases to 1.3% and 1.2%, respectively, at higher energies. The muon background is treated separately from the rest of the beamline systematics since it acts as a background component rather than modifying the beam momentum. Its effect is also con-

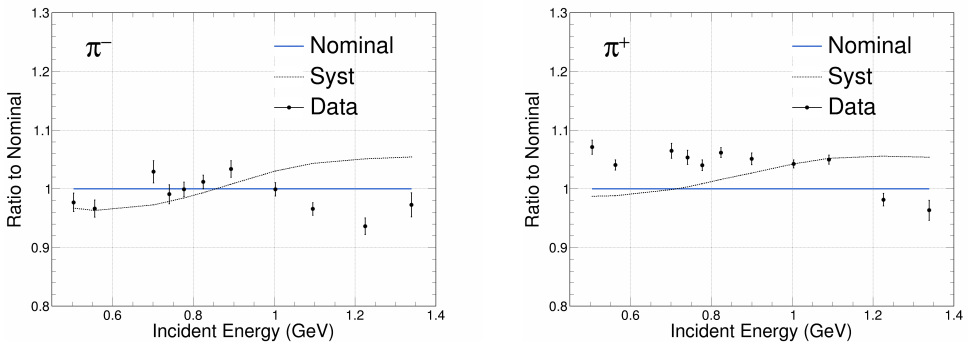

Figure 7.23: Ratio to nominal energy response for negatively (left) and positively (right) charged pions when switching from the BERT model to INCLXX. A pronounced deviation from the nominal simulation is observed, especially at higher energies, yet the INCLXX model fails to improve agreement with the data.

fined to low energies, with a maximum impact of 0.4-0.5%, and it becomes negligible above 1 GeV.

On the detector side, the Cherenkov efficiency introduces a small variation in the energy response, with an effect up to 0.4%. The light level uncertainty, which accounts for possible mismodeling of how much light is needed to produce a hit in the detector, has a negligible impact, reaching only 0.06%. These detector-related systematics are estimated using the same approach developed for the NOvA Near and Far Detectors.

7.4 MINERvA Energy Response Comparisons

The MINERvA collaboration ran a Test Beam experiment at the FTBF from 2010 to 2012 [24]. It was designed to tune the MINERvA simulation and

Source	π^- Impact	π^+ Impact	Notes
Beamline systematics	2.3% → 1.2%	2.8% → 1.3%	Combined effect from wire chamber, dipole magnet, and pion energy loss uncertainties; decreases with incident energy
Muon background	0.5% → 0%	0.4% → 0%	Beamline-related background; effect only present below 1 GeV
Cherenkov efficiency	0.4%	0.4%	Uncertainty from Cherenkov light efficiency
Light level	0.06%	0.06%	Minimal impact from light-level threshold systematic

Table 7.3: Summary of systematic uncertainties affecting the pion energy response for negative and positive charge pions. The beamline-related sources are combined in quadrature and dominate at low energies, while detector-related uncertainties are small across all energies.

evaluate systematic uncertainties of single charged particles, such as protons, pions, and electrons in the beamline momentum range of 0.35 to 2 GeV/c. This was in support of their neutrino cross-section analyses.

A comparison between the NOvA and MINERvA of the energy response for charged pions can be conducted given that the incident particle momentum ranges are similar between the two experiments. However, one should note that the detector technology and design are different. While the NOvA Test Beam detector is composed mainly of PVC planes filled with liquid scintillator, the MINERvA Test Beam detector ran in two different configurations that included plastic scintillator and absorber planes that were interleaved by placing the absorber plane upstream of each scintillator plane.

In the first configuration, half of the detector contained lead planes as absorbers to study electromagnetic interactions (ECAL region), and in the second half of the detector, iron planes were used as absorbers to study hadronic

interactions (HCAL region). In the second configuration, the first half of the detector was composed only of plastic scintillator strips (Tracking region), followed by an ECAL region in the second half of the detector. These two configurations were called EH and TE, respectively.

The energy response results of the MINERvA Test Beam experiment for positive and negatively charged pions in the EH configuration are shown in Figure 7.24.

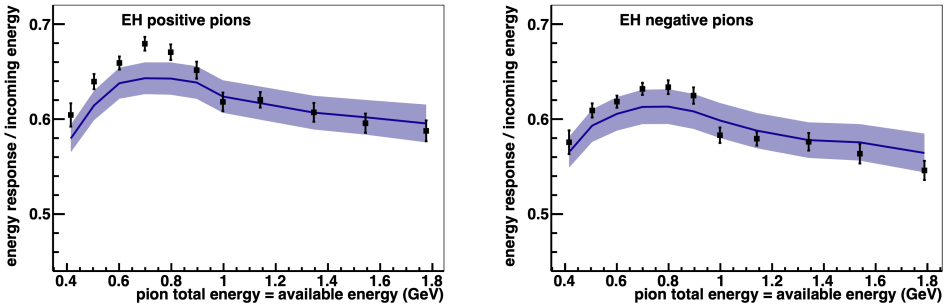

Figure 7.24: MINERvA mean energy response results for positive (right) and negative (left), charged pions in the EH detector configuration. The data (in black) is plotted with statistical uncertainties. The MC line and the systematic uncertainty band are represented in purple [24].

When comparing these results to the ones obtained in the NOvA Test Beam Experiment (depicted in Figure 7.9). It can be noticed that the mean energy response results obtained by NOvA are significantly lower compare to the MINERvA results. These differences may be explain by the different calibration procedures and detector configurations. As described in section 7.1, the pion energy response analysis is conducted using the visible energy in the detector without taking into account the energy loss in the PVC cells.

On the other hand, MINERvA implemented corrections for energy losses in

the absorbers and the inactive fraction regions of the detector to their analysis using the following expression [106]:

$$E_{\text{cal}} \equiv \sum_i C_i E_i. \tag{7.2}$$

Here, E_{cal} is the calorimetric energy after energy loss corrections. The summation is performed over the hits in the detector and C_i is the calorimetric constant that depends on the region of the detector (tracker, ECAL or HCAL). The calorimetric constants are calculated from the energy loss per unit length, dE/dx, of a minimum ionizing particle at normal incidence that includes the active fraction of the scintillator plane and is defined as:

$$C = \frac{E_{\text{abs}} + E_{\text{sc}}}{f \times E_{\text{sc}}}, \tag{7.3}$$

Where E_{abs} is the energy loss in one absorber plane, E_{sc} is the energy loss in one scintillator plane, and f is the active fraction of the scintillator plane. Since the tracking regions do not contain any absorber material, $E_{\text{abs}} = 0$, and the active fraction, $f = 0.8185$, gives $C_{\text{tracker}} = 1/f = 1.222$. For comparison, $C_{\text{ECAL}} = 2.013$ and $C_{\text{HCAL}} = 10.314$.

To better compare the energy response between NOvA and MINERvA, the same approach can be applied to the NOvA results. The NOvA detectors do not contain any absorber ($E_{\text{abs}} = 0$); thus, the calorimetric constant is $C_{\text{NOvA}} = 1/f$. The active fraction of the NOvA detectors is known to be 0.63 (63%) [107], which yields $C_{\text{NOvA}} = 1.59$. After applying these corrections to

the energy response for pions, the comparisons were made and are shown in Figure 7.25.

The comparisons show a higher energy response for the NOvA Test Beam Experiment, this could be explain by the differences between the liquid scintillator (NOvA) and plastic scintillator plus absorbers (MINERvA), since the liquid scintillator tends to have a greater light yield compare to the plastic scintillator strips in the MINERvA detector.

One can also observe similar data vs simulation discrepancies measured by the two experiments. Specifically, the positively charged pion sample shows a larger energy response in data compared to simulation in the incident energies between 0.5 GeV to 0.8 GeV. Here, the mean energy response exhibits the same level of discrepancy for the two experiments.

It can also be noted that for the positively charged pions, the NOvA measurements have discrepancies in incident energies beyond 0.8 GeV. For the negatively charged pions, data and simulation have a better agreement, except that in the NOvA results, the data tends to have a lower energy response in data compare to simulation for higher incident energies.

7.5 Pion Energy Resolution

The energy resolution is the precision with which we measure the energy in the calorimeter. This quantity along with the energy response are considered essential parameters in detector performance. The energy resolution is typi-

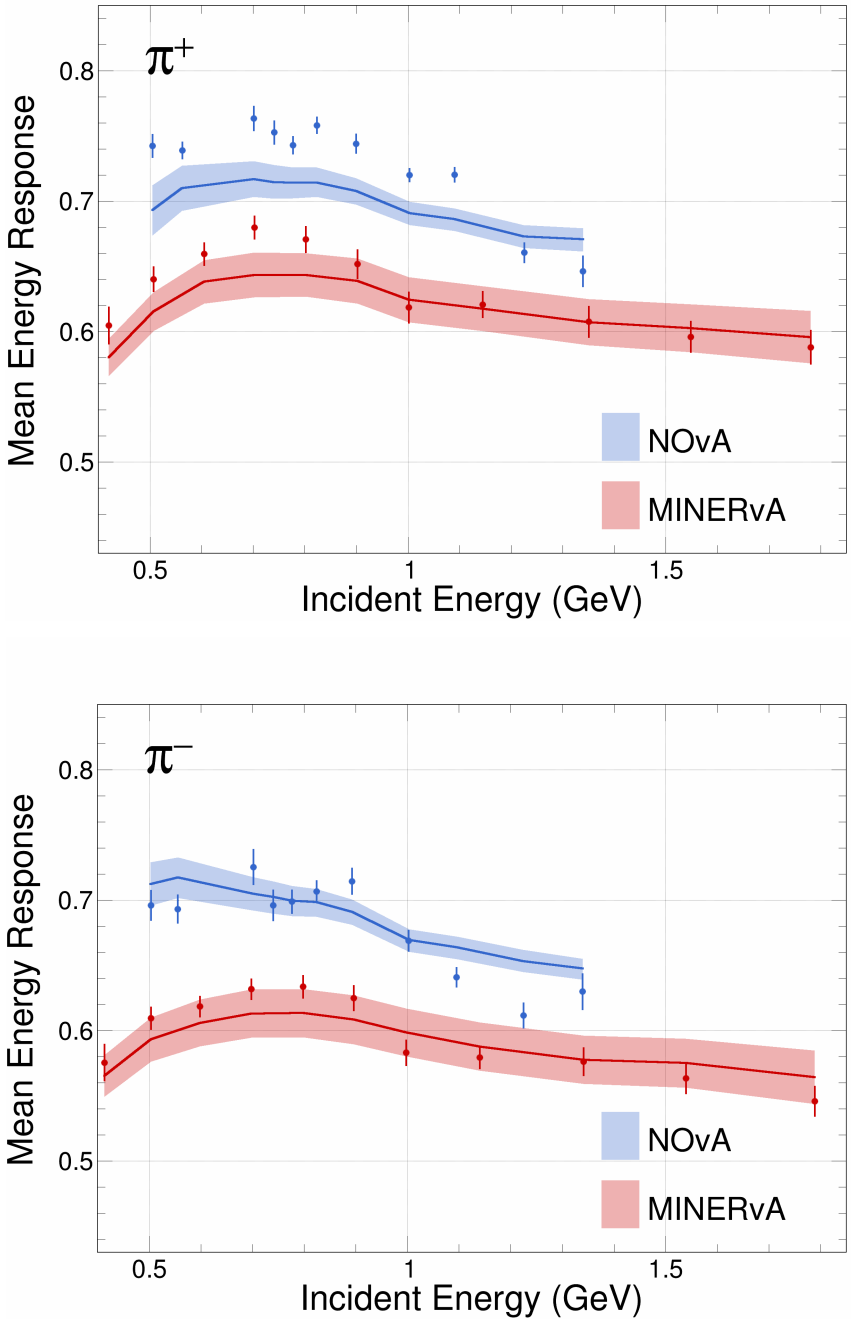

Figure 7.25: Comparison of the mean energy response between the NOvA (blue) and MINERvA (red) Test Beam Experiments for positively (top) and negatively (bottom) charged pions.

cally defined as the relative width (σ/μ) of the energy distribution measured in the detector. In addition, the energy resolution is energy dependent, and can be parameterized as:

$$\frac{\sigma}{\mu} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E}. \tag{7.4}$$

Here, E is the incident energy of the particle, $\oplus$ implies that the terms on the right hand side of the equation are added in quadrature, and a and b are the fit parameters.

The fit results for the energy resolution using Equation 7.4 for positively and negatively charged pions are shown in Figure 7.26.

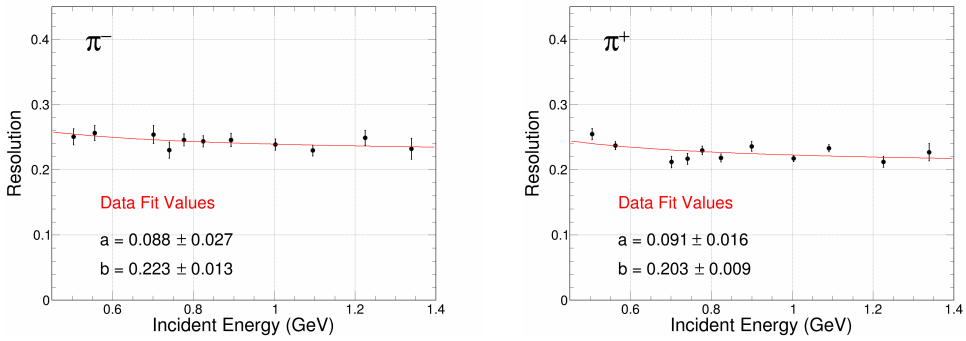

Figure 7.26: Energy resolution for negatively (left) and positively (right) charged pions.

These results show that the energy resolution ranges from $\sim 25\%$) to $\sim 22\%$) for both positively and negatively charged pions. The fit for the resolution parameters a and b are shown in Table 7.4.

Particle	a	b
π^-	0.088 ± 0.027	0.223 ± 0.013
π^+	0.091 ± 0.016	0.203 ± 0.009

Table 7.4: Best-fit parameters for the energy resolution parametrization $\frac{\sigma}{\mu} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E}$ for π^- and π^+.

7.6 Energy per Number of Hits Discrepancy

In the sterile neutrino analyses presented in section 4.2, a discrepancy exists between data and simulation for the number of hits divided by the reconstructed energy distribution, as shown in Figure 7.27. The discrepancy is primarily observed around the ∼0.23 energy per hit region for the NC-selected events in the ND. This discrepancy is not well understood and requires further investigation (more discussion in Chapter 8).

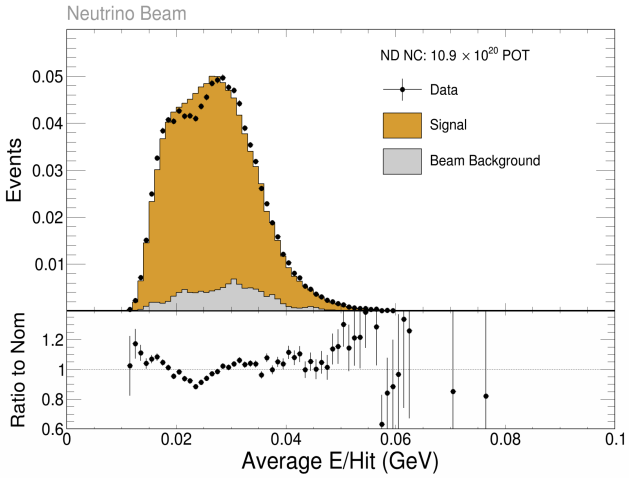

Figure 7.27: Energy per hit distribution for neutrino NC interactions in the ND.

Since neutrino NC interactions produce hadrons in their final state in the

NOvA detectors, such as pions and protons. It is of great interest to study the average energy per hit distributions for pions using the NOvA Test Beam experiment results.

Before discussing the results of this study. It is important to mention that the energy per hit results for charged pions in the NOvA Test Beam detector cannot be directly compared to the results for the neutrino NC interactions. This is due to the fact that the energy of the neutrinos is extracted using energy estimators, as discussed in section 7.1, and the fact that the NC events contain other hadrons. However, we can make use the pion samples to verify for possible disagreements between data and simulation in the energy per hit variable.

Figure 7.29 shows the energy per hit study results for negatively charged pions across three different energy regions: low, medium, and high. The left column illustrates the number of hits, the middle column shows the reconstructed energy in the detector, and the right column presents the energy per hit variable. It can be observed that the energy per hit variable exhibits shape discrepancies similar to those seen in the NC case, particularly in the intermediate energy region (0.755–0.8 GeV).

In this region, both the number of hits and the reconstructed energy in the detector appear reasonable. Therefore, it is difficult to determine the cause of the disagreement between data and simulation in the energy per hit distribution. In contrast, the energy per hit variable shows better agreement in the low and high energy regions. However, in these two cases, the number of hits distribution disagrees from the simulation.

Similar to the negatively charged pion case, Figure 7.29 shows energy per hit study results for positively charged pions across three different energy regions. In this case, we observed data-simulation disagreements on the three energy regions for the energy per hit variable. These results maybe explain by the differences observed in the reconstructed energy distributions, where the energy is higher compare to simulation.

The interpretation of these results remains inconclusive, as we cannot directly attribute the differences observed in the energy per hit distribution to the individual discrepancies in the reconstructed energy or the number of hits distributions. Further investigation is necessary to better understand the source of these differences.

Similar to the energy response discrepancies observed in previous studies, it would be beneficial to perform reweighing studies in which the simulation is tuned to better match the data. Once this adjustment is made, the analysis can be repeated to verify whether the disagreement in the energy per hit distribution still persists.

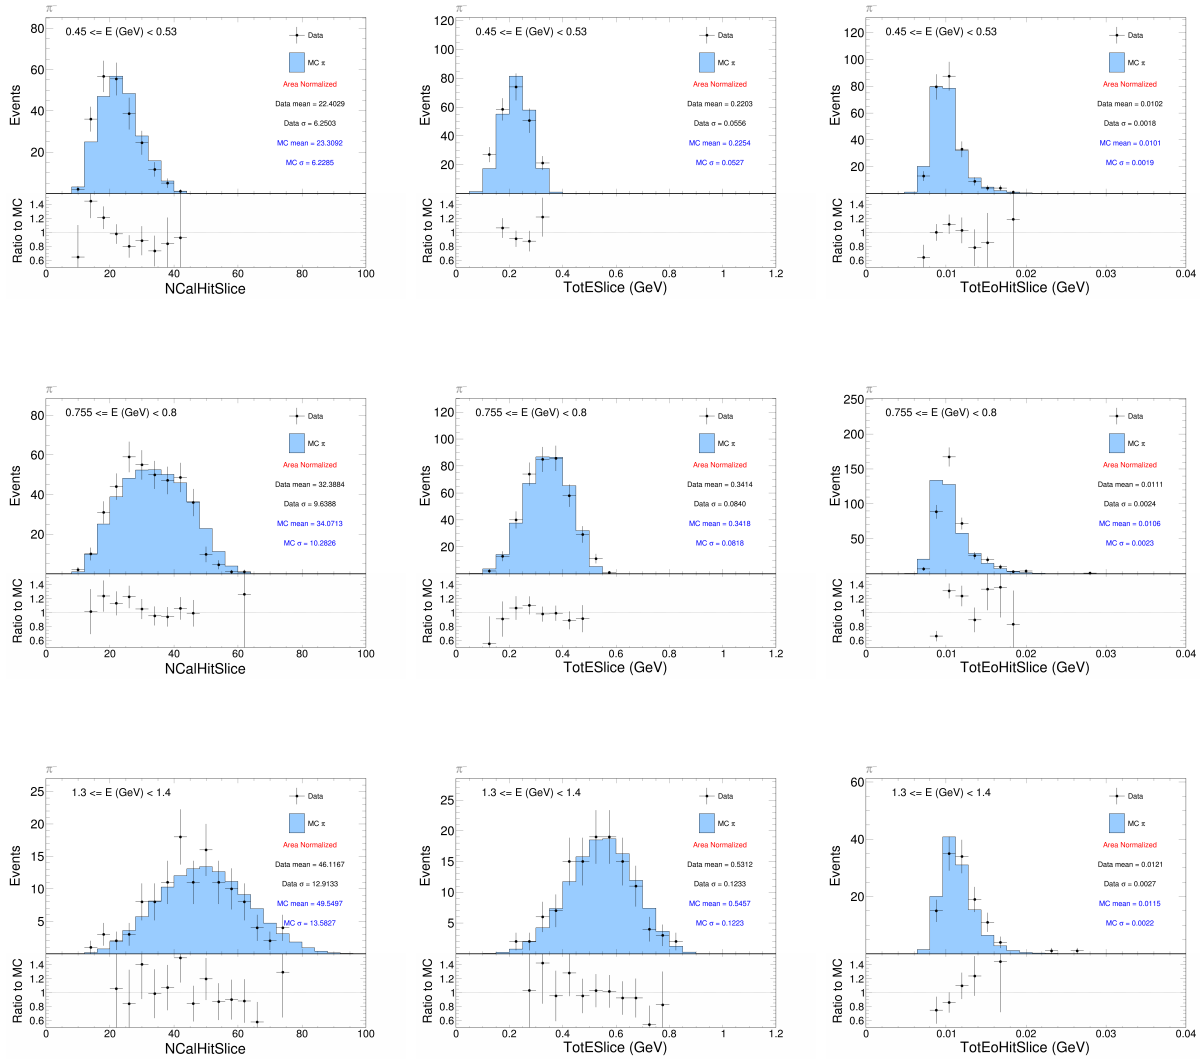

Figure 7.28: Distributions of number of hits (left), reconstructed energy (middle), and energy per hit (right) for negatively charged pions in three energy regions: low, medium and high.

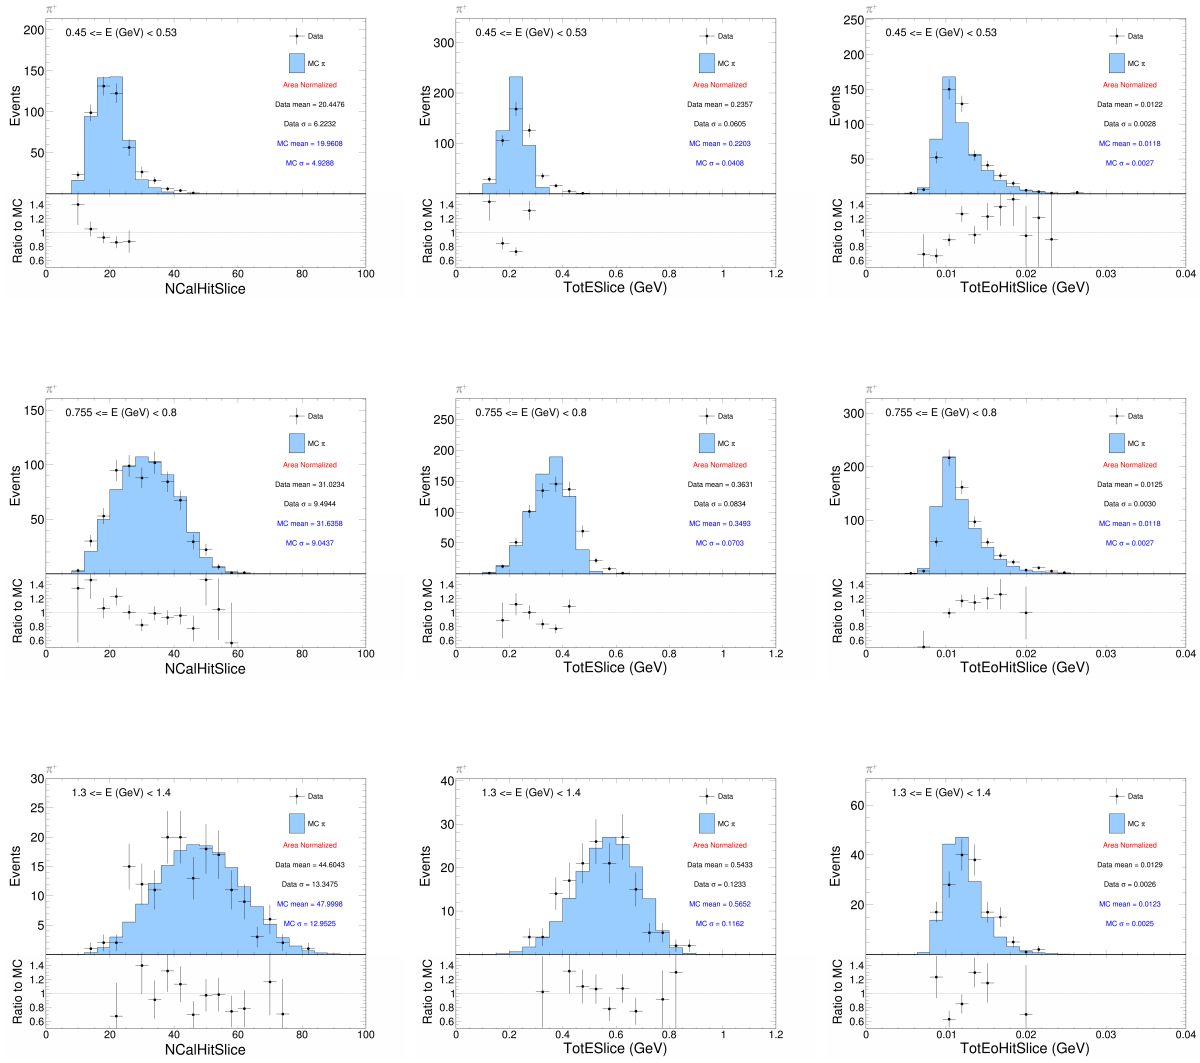

Figure 7.29: Distributions of number of hits (left), reconstructed energy (middle), and energy per hit (right) for positively charged pions in three energy regions: low, medium and high.

7.7 Conclusions on Pion Energy Response and Resolution

This section summarizes the main results of the pion energy response and resolution measurements in the NOvA Test Beam Detector. These studies were performed using a dedicated tertiary beam setup that allowed for a clean separation of charged pion and muon events. By carefully simulating the beam in G4Beamline, the muon background sample was estimated and subtracted from the samples. The analysis compared the measured energy response for pions in data to the detailed Geant4-based single-particle simulation.

For negatively charged pions, the agreement between data and simulation is generally good. The energy response remains consistent across most of the energy range, with some disagreements at higher incident energies.

In contrast, the response for positively charged pions shows a more persistent discrepancy. The data shows a consistently higher energy response compared to simulation across nearly the full incident energy range, even after the muon background has been subtracted. These differences are not explained by the systematic uncertainties evaluated in this analysis, which include beamline and detector systematics.

The Test Beam measurements show a direct evidence of mismodeling in the hadronic response for pions, particularly for the positively charged case. This discrepancy is clear in this controlled Test Beam environment, where the incident energy and particle type are well known.

While the precise origin of the discrepancy is not fully understood, the results suggest that the current simulation fails to accurately model certain pion interaction processes in the detector. These results will be used to guide the development of improved modeling strategies for NOvA, including possible reweighing schemes in the oscillation analysis. In addition, these results should translate in a reduction of the systematic uncertainty associated with hadronic response, which is currently 5%. Thus, the energy response studies presented here represent an important result of this work.

Following the energy response analysis, the energy resolution was also studied for negatively and positively charged pions. The results show that the energy resolution is nearly flat in the range of approximately 25% to 22% for both positively and negatively charged pions across the studied energy range.

Finally, the discrepancies between data and simulation in the energy per hit distribution of the neutrino NC event sample were also studied using the Test Beam pion data. It can be concluded that the origin of this discrepancy remains unclear, as discrepancies observed in the energy per hit variable for the charged pion samples cannot be directly linked to the differences observed in the reconstructed energy or number of hits distributions.

Chapter 8

Conclusions and Future Work

8.1 Summary of Results

This thesis presents two analyses developed within the NOvA experiment. The first is a search for sterile neutrino oscillations in the 3+1 framework using ν_μ-CC and NC samples using the Near and Far Detectors. The second is a measurement of the pion energy response using data from the NOvA Test Beam detector.

The sterile neutrino analysis used a two-detector fit technique developed by the UC neutrino group called PISCES, this sophisticated technique was used to search for sterile neutrinos in both detectors. Contributions to this effort included performing data quality checks on the selected NC samples and comparing data and simulation.

Validation studies were also performed using Asimov datasets and fluctuated pseudo-experiments to evaluate the stability of the fit under statistical and systematic fluctuations. Additional studies examined the impact of individual uncertainties, tested different sample configurations, and performed cross-checks used in the interpretation of the results. These efforts supported the results for the final limits imposed on sterile mixing parameters by the NOvA experiment.

The pion energy response analysis showed reasonable agreement between data and simulation for π^-, although some discrepancies were observed, particularly at high incident energies. For π^+, the simulation underestimated the energy response across most of the studied incident energy range.

This difference was not explained by the systematic uncertainties considered or by testing alternative Geant4 hadronic models. These results may help inform future updates to the pion energy response calibration in NOvA and could serve as input for tuning the simulation.

Both analyses, the sterile neutrino search and the pion energy response analysis were developed independently. The sterile neutrino study tested deviations from the three-flavor oscillation model. The Test Beam study provided a direct measurement of the detector's response to charged pions and produced results that can support future improvements to the modeling of hadronic energy in NOvA and reduction of systematic uncertainties.

8.2 Future Work

As the results of the pion energy response analysis indicate clear mismodeling in the simulation, future studies could be conducted to reduce the observed discrepancies between data and simulation. This includes tuning the pion cross sections in the simulation using G4Reweight to better match the data. If better agreement is achieved, it is recommended that the updated cross-section central value replace the one currently used in NOvA analyses.

These studies could also shed light on the discrepancy observed in the energy-per-hit variable for NC events in the sterile neutrino search. To directly compare the energy per hit between NC and pion events, it would be essential to remove the effects of the energy estimators applied in the NC events. This would allow a more accurate comparison of the energy per hit regions studied in the two analyses. Additionally, exploring the possibility of separating the NC sample into pion-only events could further clarify the origin of the observed discrepancies.

Other studies related to machine learning could also be conducted. In particular, the use of prongCVN on the Test Beam data sample (see subsection 3.3.5) aims to investigate whether the discrepancies observed in the energy response could lead to misidentification of pions in the Near and Far Detectors.

Finally, as discussed in subsection 6.3.2, the pion energy response analysis currently uses BPF tracks for selecting pion candidates. The BPF tracking algorithm reconstructs particle tracks in slices with 10 or more hits. To avoid this limitation, it is suggested that the analysis be re-run using prongs, which

do not have this hit threshold. This change could help recover lower-energy pion events that are currently excluded from the sample.

Bibliography

[1] F. Reines, C. L. Cowan, F. B. Harrison, A. D. McGuire, and H. W. Kruse. Detection of the free neutrino: A confirmation. *Physical Review*, 117. doi:10.1103/PhysRev.117.159.

[2] CERN. Cern70: A gargantuan discovery, 2024. URL: https://home.cern/news/series/cern70/cern70-gargantuan-discovery.

[3] ALEPH, DELPHI, L3, OPAL, and SLD Collaborations. Precision electroweak measurements on the z resonance. *Physics Reports*, 2006. URL: https://arxiv.org/abs/hep-ex/0509008.

[4] J. N. Bahcall and M. H. Pinsonneault. What do we (not) know theoretically about solar neutrino fluxes? *Physical Review Letters*, 2004. URL: https://doi.org/10.1103/PhysRevLett.92.121301.

[5] J. N. Bahcall and C. Peña-Garay. Solar models and solar neutrino oscillations. *New Journal of Physics*, 2004. URL: https://doi.org/10.1088/1367-2630/6/1/063.

[6] B. T. Cleveland et al. Measurement of the solar electron neutrino flux with the homestake chlorine detector. *Astrophysical Journal*, 496. URL: https://doi.org/10.1086/305343.

[7] M. Altmann et al. Complete results for five years of gno solar neutrino observations. *Physics Letters B*, 2005. URL: https://doi.org/10.1016/j.physletb.2005.04.068.

[8] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. iii: Results for the 2002–2007 data-taking period. *Physical Review C*, 2009. URL: https://doi.org/10.1103/PhysRevC.80.015807.

[9] J. Hosaka et al. Solar neutrino measurements in super-kamiokande-i. *Physical Review D*, 2006. URL: https://arxiv.org/abs/hep-ex/0508053.

[10] J. N. Bahcall, M. H. Pinsonneault, and S. Basu. Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties. *The Astrophysical Journal*, 2001. URL: https://doi.org/10.1086/321493.

[11] A. Bellerive et al. The sudbury neutrino observatory. *arXiv preprint*, 2016. URL: https://arxiv.org/abs/1602.02469.

[12] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007. URL: https://global.oup.com/academic/product/fundamentals-of-neutrino-physics-and-astrophysics-9780198508717.

[13] Y. Fukuda et al. (Super-Kamiokande Collaboration). Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 1998. URL: https://arxiv.org/abs/hep-ex/9807003.

[14] M. Sanchez et al. (Soudan-2 Collaboration). Observation of atmospheric neutrino oscillations in soudan 2. *arXiv preprint*, 2003. URL: https://arxiv.org/pdf/hep-ex/0307069.

[15] A. C. Booth. *Electron neutrino appearance at the NOvA experiment*. PhD thesis, Ph.D. Thesis, 2021. $\nu_\mu\,(\bar{\nu}_\mu) \rightarrow \nu_e\,(\bar{\nu}_e)$.

[16] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Physical Review D*, 64:112007, 2001. URL: https://arxiv.org/abs/hep-ex/0104049, arXiv:hep-ex/0104049, doi:10.1103/PhysRevD.64.112007.

[17] A. Aguilar-Arevalo et al. Significant excess of electronlike events in the miniboone short-baseline neutrino experiment. *Physical Review Letters*, 121:221801, 2018. URL: https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.121.221801, doi:10.1103/PhysRevLett.121.221801.

[18] Fermi National Accelerator Laboratory. Fermilab's accelerator complex. URL: https://www.fnal.gov/pub/science/particle-accelerators/accelerator-complex.html.

[19] P. Adamson et al. The numi neutrino beam. 2015. URL: https://arxiv.org/abs/1507.06690, arXiv:1507.06690.

[20] D. Ayres et al. The nova technical design report. Technical Report FERMILAB-DESIGN-2007-01, Fermi National Accelerator Laboratory, 2007. NOvA Collaboration. URL: https://inspirehep.net/literature/774999.

[21] A. Norman. The nova data acquisition system. In *Journal of Physics: Conference Series*, 2012. URL: https://iopscience.iop.org/article/10.1088/1742-6596/396/1/012035.

[22] Teresa Lackey. *Proton Scattering in NOvA Test Beam*. Ph.d. thesis, Indiana University, 2022. URL: https://inspirehep.net/files/30bb9db0965af7173ec957bc7cb500b1.

[23] Elizabeth Scheuermann. Uncertainty analysis of the nova test beam wirechambers. Internal document, NOvA-doc-59536, 2023.

[24] L. Aliaga et al. Minerva neutrino detector response measured with test beam data. 2015. URL: https://arxiv.org/abs/1501.06431.

[25] Wolfgang Pauli. Letter to the tübingen conference, 1930. URL: https://cds.cern.ch/record/83282.

[26] James Chadwick. Possible existence of a neutron. *Nature*, 1932. URL: https://doi.org/10.1038/129312a0.

[27] Enrico Fermi. Tentativo di una teoria dell'emissione dei raggi beta. *Il Nuovo Cimento*, 1934. URL: https://doi.org/10.1007/BF02959820.

[28] G. Danby, J.-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Physical Review Letters*, 1962. doi:10.1103/PhysRevLett.9.36.

[29] F. J. Hasert et al. Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment. *Physics Letters B*, 1973. URL: https://doi.org/10.1016/0370-2693(73)90499-1.

[30] F. J. Hasert et al. Search for elastic muon-neutrino electron scattering. *Physics Letters B*, 1973. URL: https://doi.org/10.1016/0370-2693(73)90494-2.

[31] Abdus Salam. Weak and electromagnetic interactions. In N. Svartholm, editor, *Elementary Particle Theory: Relativistic Groups and Analyticity (Nobel Symposium No. 8)*. Almqvist and Wiksell, 1968. URL: https://www.nobelprize.org/uploads/2018/06/salam-lecture.pdf.

[32] Sheldon L. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 1961. URL: https://doi.org/10.1016/0029-5582(61)90469-2.

[33] Steven Weinberg. A model of leptons. *Physical Review Letters*, 1967. URL: https://doi.org/10.1103/PhysRevLett.19.1264.

[34] G. Breit and E. Wigner. Capture of slow neutrons. *Physical Review*, 1936. URL: https://journals.aps.org/pr/abstract/10.1103/PhysRev.49.519.

[35] M. L. Perl et al. Evidence for anomalous lepton production in e^+e^- annihilation. *Physical Review Letters*, 1975. URL: https://doi.org/10.1103/PhysRevLett.35.1489.

[36] K. Kodama et al. Observation of tau neutrino interactions. *Physics Letters B*, 2001. URL: https://doi.org/10.1016/S0370-2693(01)00307-0.

[37] K. Kodama et al. Final tau-neutrino results from the donut experiment. *Physical Review D*, 2008. URL: https://doi.org/10.1103/PhysRevD.78.052002.

[38] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013. URL: https://www.cambridge.org/highereducation/books/modern-particle-physics/CDFEBC9AE513DA60AA12DE015181A948.

[39] P. Anselmann et al. Solar neutrinos observed by gallex at gran sasso. *Physics Letters B*, 1992. URL: https://doi.org/10.1016/0370-2693(92)91452-Z.

[40] S. P. Mikheev and A. Yu. Smirnov. Resonance amplification of oscillations in matter and solar neutrino spectroscopy. *Soviet Journal of Experimental and Theoretical Physics*, 1986. URL: https://link.springer.com/article/10.1007/BF02508049.

[41] M. Maltoni and A. Y. Smirnov. Solar neutrinos and neutrino physics. *European Physical Journal A*, 2016. URL: https://arxiv.org/abs/1507.05287.

[42] L. Wolfenstein. Neutrino oscillations in matter. *Physical Review D*, 1978. URL: https://doi.org/10.1103/PhysRevD.17.2369.

[43] A. Yu. Smirnov. Solar neutrinos: Oscillations or no-oscillations? *arXiv preprint*, 2016. URL: https://arxiv.org/abs/1609.02386.

[44] D. Casper et al. Measurement of atmospheric neutrino composition with imb-3. *Physical Review Letters*, 1991. URL: https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.66.2561.

[45] B. Pontecorvo. Mesonium and anti-mesonium. *Soviet Physics JETP*, 1957. URL: http://jetp.ras.ru/cgi-bin/e/index/e/6/2/p429?a=list.

[46] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Soviet Physics JETP*, 1958. URL: http://jetp.ras.ru/cgi-bin/e/index/e/7/1/p172?a=list.

[47] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 1962. URL: https://moscow.sci-hub.se/2711/871ae19d1cf0c63bd0a58a2074f95fbf/maki1962.pdf.

[48] K. Abe and et al. Solar neutrino measurements using the full data period of super-kamiokande-iv. *Phys. Rev. D*, 2024. URL: https://journals.aps.org/prd/abstract/10.1103/PhysRevD.109.092001.

[49] F. P. An and et al. Precision measurement of reactor antineutrino oscillation at kilometer-scale baselines by daya bay. *Phys. Rev. Lett.*, 130:161802, 2023. URL: https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.130.161802.

[50] J. Wolcott. New nova results with 10 years of data. Presented at Neutrino 2024 Conference, Milano, Italy, 2024. Conference presentation. URL: https://agenda.infn.it/event/37867/contributions/233955/attachments/121832/177712/2024-06-17%20Wolcott%20NOvA%202024%20results%20-%20NEUTRINO.pdf.

[51] Artur Sztuc and T2K Collaboration. Recent t2k neutrino oscillation results. *J. Phys.: Conf. Ser.*, 2020. URL: https://iopscience.iop.org/article/10.1088/1742-6596/1468/1/012133.

[52] Accelerator Division. *Concepts Rookie Book*. FNAL Operations Department, Batavia, IL, 2020. URL: https://operations.fnal.gov/rookie_books/concepts.pdf.

[53] Leonidas Aliaga Soplin. *Neutrino Flux Prediction for the NuMI Beamline*. Ph.d. thesis, College of William & Mary, 2016. URL: https://lss.fnal.gov/archive/thesis/2000/fermilab-thesis-2016-03.pdf.

[54] S. Agostinelli et al. Geant4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 2003. URL: https://doi.org/10.1016/S0168-9002(03)01368-8.

[55] V.V. Uzhinsky. The fritiof (ftf) model in geant4. In *Proceedings of Calorimetry for High Energy Frontiers (CHEF 2013)*, 2013. URL: https://s3.cern.ch/inspire-prod-files-d/d72cea91048dd2561e9cfe15fdb72ac3.

[56] D. H. Wright and M. H. Kelsey. The geant4 bertini cascade. *Nucl. Instrum. Meth. A*, 2015. URL: https://doi.org/10.1016/j.nima.2015.09.058.

[57] L. Aliaga et al. Neutrino flux predictions for the numi beam. *Phys. Rev. D*, 2016. URL: https://journals.aps.org/prd/abstract/10.1103/PhysRevD.94.092005.

[58] C. Andreopoulos et al. The genie neutrino monte carlo generator. *Nucl. Instrum. Meth. A*, 2010. URL: https://doi.org/10.1016/j.nima.2009.12.009.

[59] J. Nieves, J. E. Amaro, and M. Valverde. Inclusive quasi-elastic neutrino reactions. *Phys. Rev. C*, 2004. URL: https://journals.aps.org/prc/abstract/10.1103/PhysRevC.70.055503.

[60] R. Gran, J. Nieves, F. Sánchez, and M. J. Vicente Vacas. Neutrino-nucleus quasielastic and 2p2h interactions up to 10 gev. *Phys. Rev. D*, 2013. URL: https://journals.aps.org/prd/abstract/10.1103/PhysRevD.88.113007.

[61] Ch. Berger and L. M. Sehgal. Lepton mass effects in single pion production by neutrinos. *Phys. Rev. D*, 2007. URL: https://doi.org/10.1103/PhysRevD.76.113004.

[62] A. Bodek and U. K. Yang. Modeling deep inelastic cross-sections in the few gev region. *Nucl. Phys. B Proc. Suppl.*, 2002. URL: https://doi.org/10.1016/S0920-5632(02)01755-3.

[63] Steven Dytman. Final state interactions in neutrino-nucleus experiments. *Acta Physica Polonica B*, 2009. URL: https://www.actaphys.uj.edu.pl/fulltext?series=Reg&vol=40&page=2445.

[64] A. Aurisano, C. Backhouse, R. Hatcher, N. Mayer, J. Musser, R. Patterson, R. Schroeter, and A. Sousa. The nova simulation chain. *Journal of Physics: Conference Series*, 2015. URL: https://iopscience.iop.org/article/10.1088/1742-6596/664/7/072002.

[65] G. Folger and J. P. Wellisch. String parton models in geant4. *arXiv preprint nucl-th/0306007*, 2003. URL: https://arxiv.org/abs/nucl-th/0306007.

[66] John Apostolakis et al. Improvements in the geant4 hadronic physics. In *Journal of Physics: Conference Series*, 2011. URL: https://doi.org/10.1088/1742-6596/331/3/032002.

[67] E. Mendoza, D. Cano-Ott, T. Koi, and C. Guerrero. New standard evaluated neutron cross section libraries for the geant4 code and first verification. *IEEE Transactions on Nuclear Science*, 2014. URL: https://ieeexplore.ieee.org/document/6869047.

[68] V. Ivanchenko, S. Incerti, and M. Novak. Geant4 electromagnetic physics progress. Presentation at CHEP 2019, 2019. URL: https://indico.cern.ch/event/773049/contributions/3474766/attachments/1937617/3211571/G4EmCHEP2019_v6.pdf.

[69] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. 2014. URL: https://doi.org/10.1126/science.1242072.

[70] R.C. Prim. Shortest connection networks and some generalizations. *Bell System Technical Journal*, 1957. URL: https://doi.org/10.1002/j.1538-7305.1957.tb01515.x.

[71] Leandro A. F. Fernandes and Manuel M. Oliveira. Real-time line detection through an improved hough transform voting scheme. 2008. URL: https://doi.org/10.1016/j.patcog.2007.04.003.

[72] Mattias Ohlsson, Carsten Peterson, and Alan L. Yuille. Track finding with deformable templates–the elastic arms approach. *Computer Physics Communications*, 1992. URL: https://doi.org/10.1016/0010-4655(92)90074-9.

[73] M. Baird, J. Bian, M. Messier, E. Niner, D. Rocco, and K. Sachdev. Event reconstruction techniques in nova. *Journal of Physics: Conference*

Series, 2015. URL: https://iopscience.iop.org/article/10.1088/1742-6596/664/7/072035.

[74] J. C. Dunn. A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters. 1973. URL: https://doi.org/10.1080/01969727308546046.

[75] N. H. Kuiper. Tests concerning random points on a circle. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen, Series A*, 1960. URL: https://doi.org/10.1016/S1385-7258(60)50006-0.

[76] B. Rebel. A window tracking algorithm for cosmic ray muons. Internal document, NOvA-doc-15977, 2016.

[77] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 1960. URL: https://www.unitedthc.com/DSP/Kalman1960.pdf.

[78] N. Raddatz. Kalmantrack technical note. Internal document, NOvA-doc-13545, 2015.

[79] G. Lutz. Optimum track fitting in the presence of multiple scattering. *Nuclear Instruments and Methods in Physics Research Section A*, 1988. URL: https://www.sciencedirect.com/science/article/pii/0168900288908364.

[80] Michael David Baird. *An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter*. Ph.d. thesis, Indiana University, 2015. URL: https://lss.fnal.gov/archive/thesis/2000/fermilab-thesis-2015-24.pdf.

[81] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle. A convolutional neural network neutrino event classifier. *Journal of Instrumentation*, 2016. URL: https://doi.org/10.1088/1748-0221/11/09/P09001.

[82] F. Psihas et al. Context-enriched identification of particles with a convolutional network for neutrino events. *Physical Review D*, 2019. URL: https://doi.org/10.1103/PhysRevD.100.073005.

[83] M. A. Acero et al. An improved measurement of neutrino oscillation parameters by the nova experiment. *Phys. Rev. D.* URL: https://arxiv.org/pdf/2108.08219.

[84] Aaron Mislivec. Near-to-far extrapolation in pt for the nova 2020 3-flavor analysis. Internal document, NOvA-doc-44401, 2018.

[85] B. Jargowsky. PCA TNs: Flux, Genie, & Geant4. Internal document, NOvA-doc-62099, 2024.

[86] P. Désesquelles et al. Cross talk and diaphony in neutron detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1991. URL: https://www.sciencedirect.com/science/article/pii/0168900291902066.

[87] M. Rabelhofe. Menate neutron file systematic. Internal document, NOvA-doc-59999, 2024.

[88] J. Calcutt and et al. Geant4reweight: A framework for evaluating and propagating hadronic interaction uncertainties in geant4. *Journal of*

Instrumentation, 2021. URL: https://iopscience.iop.org/article/10.1088/1748-0221/16/08/P08042.

[89] G. J. Feldman and R. D. Cousins. Unified approach to the classical statistical analysis of small signals. *Physical Review D*, 1998. URL: https://journals.aps.org/prd/abstract/10.1103/PhysRevD.57.3873.

[90] Ben Jargowsky. *A Measurement of ν_e Appearance and ν_μ Disappearance Using 10 Years of Data from the NOvA Experiment*. Ph.d. thesis, University of California, Irvine, 2024. URL: https://escholarship.org/content/qt8t19p34n/qt8t19p34n.pdf?t=sljr7o.

[91] P. Adamson and et al. Search for active-sterile neutrino mixing using neutral-current interactions in nova. *Physical Review D*, 2017. URL: https://journals.aps.org/prd/abstract/10.1103/PhysRevD.96.072006.

[92] M.A. Acero et al. Search for active-sterile antineutrino mixing using neutral-current interactions with the nova experiment. *Phys. Rev. Lett.*, 2021. URL: https://inspirehep.net/literature/1867930.

[93] M.A. Acero et al. Dual-baseline search for active-to-sterile neutrino oscillations in nova. *Phys. Rev. Lett.*, 2025. URL: https://inspirehep.net/literature/2826226.

[94] V Hewes. Parameter inference with systematic covariance and exact statistics. Internal document, NOvA-doc-45977, 2022.

[95] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 1963. URL: https://doi.org/10.1137/0111030.

[96] R. Acciarri and et al. The liquid argon in a testbeam (lariat) experiment. *Journal of Instrumentation*, 2020. URL: https://arxiv.org/abs/1911.10379.

[97] Michael Hahsler, Matthew Piekenbrock, and Derek Doran. dbscan: Fast density-based clustering with R. 2019. URL: https://www.jstatsoft.org/article/view/v091i01.

[98] A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 1964. URL: https://pubs.acs.org/doi/abs/10.1021/ac60214a047.

[99] Robert Kralik. *Measuring the Muon Neutrino Magnetic Moment in the NOvA Near Detector*. Ph.d. thesis, University of Sussex, UK, 2024. URL: https://sussex.figshare.com/articles/thesis/Measuring_the_muon_neutrino_magnetic_moment_in_the_NOvA_near_detector/28623461?file=53102570.

[100] Inc. Muons. G4beamline – a geant4-based simulation program for beamlines and accelerators. URL: https://www.muonsinc.com/Website1/G4beamline.

[101] Fernanda Psihas. *Measurement of Long Baseline Neutrino Oscillations and Improvements from Deep Learning*. Ph.d. thesis, Indiana University, 2018. URL: https://lss.fnal.gov/archive/thesis/2000/fermilab-thesis-2018-07.pdf.

[102] Harry R. Hausner. *Sterile Neutrino Search with the NOvA Detectors*. Ph.d. thesis, University of Wisconsin–Madison, 2022. URL: https://inspirehep.net/files/c32974bc1b6a4e8964d07f858e21b678.

[103] Yagmur Torun and Brinden Carlson. M1 magnet. Internal document, NOvA-doc-32880, 2018.

[104] Pekka Kaitaniemi, Alain Boudard, Sylvie Leray, Joseph Cugnon, and Davide Mancusi. Incl intra-nuclear cascade and abla de-excitation models in geant4. *Progress in Nuclear Science and Technology*, 2011. URL: https://inspirehep.net/files/2c50b67853a5f61522b95c1fd4f87fc0.

[105] G. Folger and J. P. Wellisch. The geant4 binary cascade. *European Physical Journal A*, 2004. URL: https://cds.cern.ch/record/865824/files/p313.pdf.

[106] L. Aliaga et al. Design, calibration, and performance of the minerva detector. 2014. URL: https://arxiv.org/abs/1305.5199.

[107] M. A. Acero et al. Supernova neutrino detection in nova. 2020. URL: https://arxiv.org/abs/2005.07155.