FERMILAB-THESIS-2018-07

MEASUREMENT OF LONG BASELINE NEUTRINO OSCILLATIONS AND IMPROVEMENTS FROM DEEP LEARNING

Fernanda Psihas

Submitted to the faculty of the University Graduate School

in partial fulfillment of the requirements

for the degree

Doctor of Philosophy

in the Department of Physics,

Indiana University

February 19, 2018

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Doctoral Committee

Mark D. Messier

Jon Urheim

Rick Van Kooten

Herb Fertig

Defense Date February 19, 2018

DEDICATION

To God, who gave me five talents, which I continue to try and do my best with.

To my grandfather, whose example inspires me to work harder and do better.

To this country, which gives me the opportunity to do so.

FΨ

A Dios, quien me dio cinco talentos, los cuales continuo intentando usar al maximo.

A mi Noto, cuyo ejemplo me inspira a trabajar mas duro y mejor.

A este pais, que me da la oportunidad de hacerlo.

FΨ

ACKNOWLEDGMENTS

Thanks to **God**, for these years filled with many blessings, to **this country** for allowing me to follow my passion, and to **my family** for supporting my crazy idea to do science after all.

I would like to acknowledge my advisor, **Mark Messier**, for whom I have great admiration. While allowing me to explore every single tangent I wanted to take on this path, he still taught me the hard lessons about how to find the right direction. I have learned much from him about asking the right questions, about being productively thorough, about unbiased opinions, and about leadership based on integrity. To Mark and to the rest of my committee, **Jon Urheim**, **Rick Van Kooten**, and **Herb Fertig**, thank you for your patience, your time, and your advice.

Thank you, also, to my friend **Stan Wojcicki**, for listening to my stories and even my practice talk for defending this thesis, and for continuing to inspire me to work on problems which spark my passion for science.

During these years, I have had the privilege to work close to people who have shaped the course of both NOvA and my career, starting with **Gary Feldman** who asked me the first ever scary question about my work on NOvA, and steered our experiment towards success as well as my work in the best direction, and **Peter Shanahan**, who always made time to give me advice and make sure I knew I have a true voice on the experiment. Peter's impact in my career could never be duplicated.

I owe a big thank you to **Tricia Vahle** and **Mayly Sanchez** for their guidance showing me the value

of strength and perseverance by example, and to **Ryan Patterson** for his pointed insight.

From my first projects on NOvA to my latest, I have enjoyed the guidance of many people without whom my appreciation for the experiment would be incomplete. Thank you to **Alec Habig**, my advisor since Duluth, whose amazing support and no-nonsense attitude about physics and life continues to be refreshing. Also to **Rick Tesarek**, **Mat Muether**, **Jaroslav Zalesak**, **Andrew Norman**, **Jim Musser**, and **Jon Paley** for their guidance from my first lines of code and pieces of hardware to the work on this thesis.

Today I am one of the longest standing members of **Young NOvA**, their support made me carve a place for myself and a home in the experiment, thank you to all. Along the way, many of you have been my peers, postdocs or students, roles have changed back and forth, but my thanks for carving the path with me remain the same. Thank you to **Travis Olson**, **Kanika Sachdev**, **Chris Backhouse**, and **Minerba Betancourt**. Especially, I want to acknowledge **Evan Niner** for his support, his respect and also his patience, and to **Michael Baird** for his wisdom and example. Whatever stage of our careers Evan and Michael will continue to be my older brother examples to follow.

Naturally, a special thanks should go to my closest CVN companions, **Dominick Rocco** who academia misses terribly, **Alex Radovic** my best arguing buddy, and **Adam Aurisano** who is just super nice. More recently, to **Micah Groh**, who has taken on so much in such a short time, thank you for allowing me to get you started and for all the fun that followed. And to **Diana P. Mendez**, **Teresa Lackey**, **Ryan Murphy** and **Shiqi Yu** for taking the many batons I had been running with on NOvA. Science is a passion from which half the fun is talking about. Thank you for listening.

I was lucky to learn important lessons from many of my peers, for those lessons I thank **Erika Catano-Mur**, an unexpected companion at crunch time and an accomplice to always sneak in "just one more improvement", as well as my friends **Kuldeep Kaur Maan**, **Biswaranjan Behera**, and **Liudmila Kolu-**

paeva who are my examples of hard work and commitment and who make me always want to do better.

I would also like to acknowledge people outside of NOvA who gave me advice, friendship, and opportunities for growth, such as **Joseph Zennamo** with whom I am delighted to talk about any physics, any day. To **Sergei Gleyzer** for introducing me to new perspectives to problems and the whole field, and especially my good friend **Francesco Capozzi** with whom I endured the longest, best hikes ever, and who in turn endured with me the proof-reading of my theory chapter.

Lastly, a special acknowledgement to the people who have marked my PhD and my last few years, and potentially the few who have continued to read this far:

To my friend, colleague and co-watchdog leader **Louise Suter**, whom I love not only for her friendship but for her example of brilliance and integrity.

To **Gavin Davies** for being for me, like for NOvA, the string tying all loose ends together. Sometimes things seem to stand tall on their own, but there is always a support that keeps them from falling.

To **Justin Vasel**, who is my family away from family, my brother and partner in adventure. Thank you for all the proof-reads, all the encouragement, and for always believing in me. I would not be where I am now if I didn't have you in my life.

Thank you to Fermi National Accelerator Laboratory, a home away from home where I learned the true meaning and value of great American science. My research was supported by the Department of Energy.

The Deep Learning algorithms developed and optimized for this work used the resources from Indiana University's Big Red II cluster. This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute, and in part by the Indiana METACyt Initiative. The Indiana METACyt Initiative at IU was also supported in part by Lilly Endowment, Inc.

Fernanda Psihas

MEASUREMENT OF LONG BASELINE NEUTRINO OSCILLATIONS AND IMPROVEMENTS FROM DEEP LEARNING

NOvA is a long-baseline neutrino oscillation experiment which measures the oscillation of muon neutrinos from the NuMI beam at Fermilab after they travel through the Earth for 810 km. In this dissertation I describe the operations and monitoring of the detectors which make it possible to record over 98% of the delivered neutrino beam. I also present reconstruction and identification techniques using deep convolutional neural networks (CNNs), which are applicable to multiple analyses. Lastly, I detail the oscillation analyses in the $\nu_\mu \to \nu_\mu$ and $\nu_\mu \to \nu_e$ channels on a dataset of 8.85×10^{20} protons on target, with emphasis on the application of deep learning, energy reconstruction, and treatment of systematic uncertainties.

The CNN single particle identifier achieves 65%, 73%, 74%, 83%, and 45% efficiency and 81%, 71%, 86%, 73%, and 45% purity for electrons, photons, muons, protons, and pions, respectively, with no additional selection. The identification of signal events using CNNs achieves 93% purity and 31% efficiency for muon neutrinos and 66% purity and 64% efficiency for electron neutrinos, with no additional selection. The number of events in the energy region of interest is 66 over 20.5 expected background events for appeared electron neutrinos and 126 over 9.2 background events for surviving muon neutrinos, with leading uncertainties from cross sections, calibration, and near/far normalization. These results, fit to the three flavor model of neutrino oscillations with constraints from known oscillation parameters are consistent with values of $\sin^2\theta_{23} = 0.558^{+0.041}_{-0.033}$, $\Delta m^2_{32} = 2.444^{+0.079}_{-0.077} \times 10^{-3}$ eV^2, and $\delta_{CP} = 1.213\pi$ with a preference for the normal neutrino hierarchy at a significance level of 1.8σ.

CONTENTS

Chapter 1

Introduction

"I want to know how God created this world. I am not interested in this or that phenomenon, in the spectrum of this or that element. I want to know His thoughts, the rest are details."
–Albert Einstein

The field of experimental particle physics seeks to answer questions at the most elementary level. The particles which we study are the main protagonists in the story of how the universe came to be and the fundamental laws of nature. The nature of neutrinos and their masses is amongst the most interesting questions in physics today. Neutrinos are presently the only particles experimentally shown to be beyond the description provided by the Standard Model of particle physics. Furthermore, through the non-zero value of their masses they are linked in the theory to potential explanations for long sought after questions such as the matter anti-matter asymmetry in the universe.

Neutrino detection is a challenge both technological and scientific. Not only are neutrinos difficult to detect because they interact so rarely with matter, their interactions with nuclei are not well described by existing models. When they do interact and we detect them, there is still some uncertainty associated with how likely they were to interact in the first place. While these may be obstacles, they present great opportunities for collaboration and discovery. Such is the nature of scientific endeavors: every obstacle brings a lesson, every unexpected result a new adventure on the path to unveiling the mysteries that God may have left for us in nature.

Indeed, obstacles which turn into opportunities are at the heart of experimental neutrino physics. The first time we tried to measure the flux of solar neutrinos was in the 1960s, with Ray Davis' Homestake experiment. When Ray Davis saw but a third of the neutrinos from the Sun that John Bahcall had predicted [1], he might have known there was a mystery there to uncover. It is unlikely, however, that he anticipated the emergence of a new field once it was discovered that the mysterious effect was the mixing of neutrino flavors, also referred to as neutrino oscillations.

Little was known about the behavior of neutrinos back then. They were known to be a by-product of the nuclear processes which make stars like the Sun shine, but their rate had never been measured. The Homestake experiment made a measurement of incoming neutrinos by counting ^{37}Ar atoms, produced in the interaction of electron neutrinos (ν_e) with Cl atoms in a 100,000 gallon volume of cleaning fluid (Fig.1.1). For decades, the source of the discrepancy between the measurement and Bahcall's calculations, what we now call the solar neutrino problem, remained unproven. The knowledge that most of the ν_e's had undergone oscillations would explain this discrepancy, but this phenomenon would remain unobserved for decades.

Figure 1.1: Ray Davis (left) and John Bahcall (right) standing next to the C_2Cl_4 tank of the Homestake Experiment.

Since Ray Davis' experiment in the 1960's, many experiments have aimed to better understand the solar neutrino problem, now known to be explained by the phenomenon of neutrino oscillations: the fact that a neutrino, known to be one of the flavors ν_e, ν_μ, or ν_τ when produced, can be observed as a different one after some time. Neutrino oscillations are possible because the neutrino mass states that exist, ν_1, ν_2, and ν_3 do so as a superposition of the flavor states ν_e, ν_μ, and ν_τ.

Being in a superposition of the flavor states means that a neutrino is not one but all three flavors simultaneously, in a state governed by oscillating probabilities, a concept which Heisenberg described as *"a strange kind of physical reality, just in the middle between possibility and reality"* [2]. Only at their birth or death does one open the box, collapse the wave function, and resolve the neutrino's identity. The reality that we can only know the flavor of neutrinos at production and upon their interaction with matter influences much of our experimental methodology to measure their oscillations: choose a source of neutrinos whose initial flavor is known, allow them to travel for some distance, and then measure their flavor again to see if it changed.

This simple principle is the basis of neutrino oscillation experiments, but matters quickly get complicated. The fact that neutrino cross sections are extremely small, i.e., they interact with matter ever so rarely, requires an effort to provide them with as much matter as possible such that their chances of finding an atomic nucleus in their path and interacting with it are enhanced. Thus, neutrino detectors tend to be as massive or as dense as we can afford, a fascinating problem of engineering on its own. Another challenge is to probe all possible combinations of neutrino mixing, or at least as many as are needed in order to understand this phenomenon,[1] by measuring their flavor and energy.

Neutrino physicists have performed oscillation measurements using neutrinos from a variety of natural sources. Solar electron neutrinos, which brought the solar neutrino problem to light, have also been studied by experiments like Kamiokande [3], SAGE [4], Gallex [5], and SNO [6]. Atmospheric muon and electron neutrinos, which are produced by cosmic rays in the Earth's atmosphere, fueled the first observations of neutrino oscillations by the Super-Kamiokande [7] experiment in 1998 and

[1] While it is actually not the case that every single channel needs to be observed in order to understand as much as we can about neutrino oscillations, a large subset of them does need to be probed by experiments to complete the picture of probabilities for each single flavor to oscillate into another.

have since been observed by SNO [8], MACRO [9], and others. Neutrinos are also produced in the far reaches of space. Extragalactic muon neutrinos are probed by the IceCube [10] experiment in Antarctica.

On Earth, there are man-made sources of neutrinos such as nuclear reactors, which make electron anti-neutrinos as a byproduct of fission processes. The Daya Bay [11], RENO [12], and Double Chooz [13] experiments have performed oscillation measurements from reactor neutrinos. We also produce our own neutrinos in particle accelerators, with the advantage that the energies and fluxes can be engineered and controlled by design. NOvA is one of the experiments which measures the oscillations of accelerator neutrinos. Some accelerator neutrino experiments have measured oscillations close to the beam source, such as LSND [14] and MiniBooNE [15], while others like K2K [16], MINOS [17], and OPERA [18] have studied these neutrinos over longer baselines. Experiments such as T2K [19] and NOvA [20] continue to study accelerator neutrinos, and improvements are already underway. The next generation of neutrino experiments like DUNE [21] and HyperK [22] are currently in development and will be taking data in the next decade.

The techniques employed by these experiments have utilized and sometimes driven the advancements of new technologies. Detection techniques have been varied, from the counting of ions in a chlorine tank to the observation of Cherenkov radiation—cones of blue light emitted by particles going faster than light—in the ultra-pure ice of Antarctica. A number of techniques have also been developed and employed to make possible the detailed identification and study of neutrino interactions. While the first visualizations of neutrino interactions came from the photographic plates of spark chambers and bubble chambers in the 60s and 70s [23] (Fig.1.2), on NOvA, we employ deep learning image recognition techniques for the first time in the field to study these data.

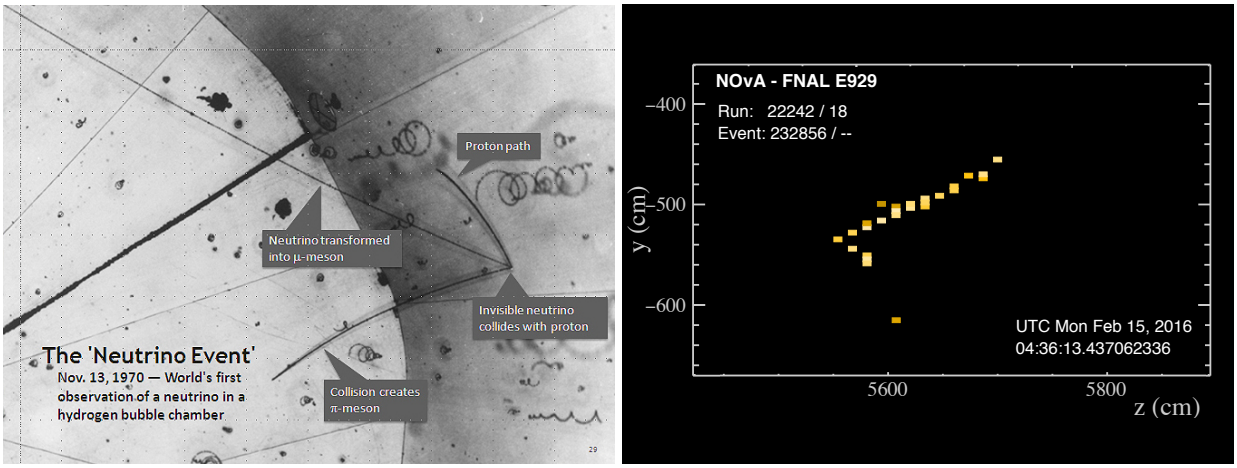

Figure 1.2: Left: neutrino interaction in a bubble chamber. Right: neutrino interaction in the NOvA detector.

We have learned much about neutrino oscillations in the past few decades. Precise measurements have been made in many oscillation channels and in some cases anomalies[2] have been found which are currently being investigated. However, some of the most interesting questions remain unanswered. The nature of neutrinos, specifically whether they are Majorana or Dirac particles (the same or different from their anti-particles), and whether they are connected to the fact that the universe is made up of mostly matter and not anti-matter, what we call the matter-antimatter asymmetry are both fundamental questions about nature itself and how things came to be what we observe in the universe today. In particular, two pieces of information are yet to be determined from neutrino oscillations which are related to these fundamental questions. One is the neutrino mass *ordering*, or mass *hierarchy*, a question of whether there is one heavier and two lighter neutrinos or vice-versa. The other is CP violation, the question of whether neutrinos violate the fundamental Charge-Parity symmetry. The mass hierarchy and CP violation can be determined by measuring ν_e appearance in a ν_μ source (and their anti-particles), after they have traveled through the Earth for a long distance (see Chapter 2). These oscillations, $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ are the main measurements we perform at NOvA.

[2] In particular, an excess of low energy, short-baseline ν_e's, [14,15] and the anomaly found in Gallium experiments [5].

Dissertation Overview

In this dissertation, I describe neutrino oscillation measurements in NOvA data for the $\nu_\mu \rightarrow \nu_e$ and $\nu_\mu \rightarrow \nu_\mu$ channels, as well as improvements from deep learning techniques to identification and reconstruction. I begin by motivating these measurements through a description of neutrinos and the phenomenon of oscillations in Chapter 2 and introducing the challenges regarding neutrino interactions and detection in Chapter 3.

In Chapter 4, I describe the NOvA experiment, detailing some of my contributions to commissioning, maintenance, and detector operations and monitoring in Secs. 4.2.2, 4.2.3, and 4.2.4. The reconstruction algorithms used for these analyses are described in Chapter 5, where my contributions to the application and optimization of the deep learning technique for event identification are described in Sec. 5.2.2, as well as the full development and application of the techniques in Secs. 5.2.3, 5.3.2, and 5.4.

I detail the methodology for the analyses of the ν_e appearance and ν_μ disappearance channels in Chapter 6, including the improved signal identification and reconstruction methods using deep learning techniques. The description in Chapter 6 pertains to the analysis made public by the collaboration in January 2018, to which I contributed to the development of the event selection for ν_μ signal events and systematic uncertainties for both analyses, in addition to the contributions detailed in Chapter 5. Finally, in Chapter 7, I show the resulting energy spectrum for each analysis, as well as the limits set on $\sin^2\theta_{23}$, Δm^2_{32} , and δ_{CP} from a fit of the oscillation probabilities to this result.

Chapter 2

The Nature of Neutrinos and Their Masses

"Neutrinos, they are very small..."

– John Updike

When Wolfgang Pauli proposed the existence of the neutrino in 1930, he had already come up with his famous exclusion principle, which governs the behavior of fermions, as well as the existence of a new quantum number, which he called Zweideutigkeit, nowadays known as spin. We knew little about the elementary particles at the time (only protons and electrons had been observed), but contributions like Pauli's would bring about an era of great advancement in our understanding of the fundamental constituents of nature. Neutrinos and their masses would be, and remain, one of the least understood pieces of the theory of elementary particles and fundamental forces.

Initially, the neutral particle proposed by Pauli in the famous letter to a conference [24] was intended to account for both the atomic integer spin—unaccounted for by electrons and protons—(the spin-statistics problem) and for the seemingly missing energy in β decay:

$$ {}^{A}_{Z}X \rightarrow {}^{A}_{Z-1}X' + e^{+} + \nu_e. \tag{2.1} $$

Eq. 2.1 shows the general form of β decay for which, without the detection of the neutrino in the final state, a fraction of the energy and momentum would appear to be missing.

In the end, Pauli effectively proposed two neutral particles which ultimately accounted for the missing pieces in both of these phenomena. The first was the neutron, discovered in 1932 [25] and found to resolve the spin-statistics problem, and the second was the neutrino, which was first observed in 1953 [26, 27] and was found to account for the missing energy in β decay.

As mentioned in Chapter 1, neutrinos are far from fully understood. In the following sections I detail our current understanding of the neutrino particles and their properties, give an overview of the remaining questions regarding their nature and their masses, and introduce the phenomenon of neutrino oscillations.

2.1 Neutrinos in the Standard Model

"Had I foreseen this, I would have gone into botany!"
– Wolfgang Pauli

In contrast with the picture of elementary particles from the early 1900s, which contained only what is found in the atom, it is now known that the number of elementary particles in nature is at least seventeen. Five of these particles, called bosons, are responsible for the fundamental forces which govern particle interactions. The remaining twelve are the fermions, which comprise all the matter we know to exist in the universe.

The Standard Model (SM), see Fig.2.1, is a quantum field theory that describes the fundamental particles and interactions through fields and symmetries and which describes the strong nuclear, electromagnetic, and weak forces. The model is characterized by the gauge symmetry SU(3) $\otimes$ SU(2) $\otimes$ U(1) and its dynamics are uniquely defined once its 19 free parameters (which include the masses of the fermions) are determined.

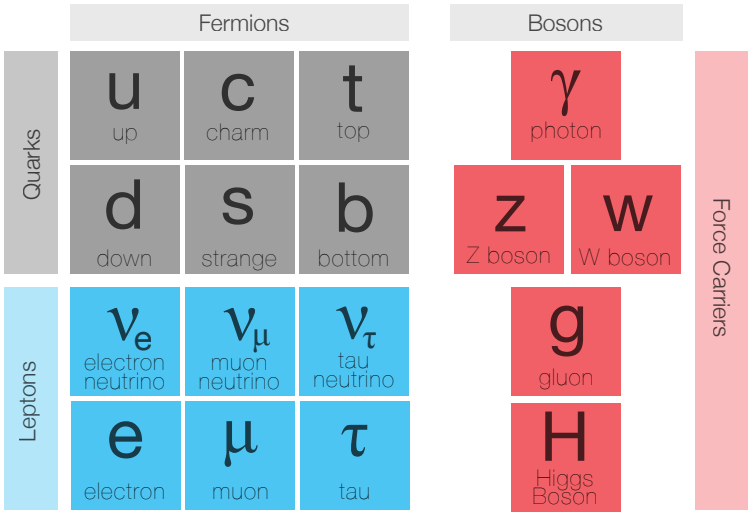

Figure 2.1: The Standard Model of particle physics.

The Standard Model is widely regarded as one of the most successful scientific theories in history. It contains a fundamental description of nature and makes accurate predictions of particles and associated observables that have been confirmed by experiments in the past several decades. With the discovery of the Higgs boson in 2012 [28], the full set of particles predicted by the model have been observed, although precision measurements of particles' characteristics and searches for other exotica are still underway [29].

The only observed discrepancy between experimental observations and what is predicted by the Standard Model is related to the behavior of neutrinos. The three neutrinos of flavors ν_e , ν_μ , and ν_τ are light fermions that interact only via the weak force, as they have no electric charge and no color charge. Neutrinos are produced in association with their corresponding heavy lepton, conserving lepton number.

While the Standard Model predicts the existence of neutrinos, it requires them to be massless. The non-zero mass of neutrinos and its implications in the theory are discussed in the following sections.

2.2 Neutrino Masses

"God used beautiful mathematics in creating the world."
–Paul Dirac

The non-zero value of the neutrino mass is the only laboratory based evidence of physics beyond the Standard Model to date, and it is an extremely meaningful deviation from the model because of its potentially large implications in the theory. To understand the consequences of non-zero neutrino masses, it is useful to discuss the possible modifications to the Standard Model that can generate them.

In the Standard Model, neutrino masses are not generated because there is no right-handed neutrino field. However, it is possible to introduce the neutrino masses via two different (yet not mutually exclusive) mechanisms: the Higgs mechanism and the Seesaw mechanism [30][1].

The Higgs mechanism [32] is responsible for the masses of all the fermions in the SM and can be used to give neutrinos their masses if a right-handed neutrino ν_R and corresponding left-handed anti-neutrino $\bar{\nu}_L$ are introduced. In that case, the Yukawa Lagrangian for the lepton masses is expanded to include a second term and is of the form:

$$\mathcal{L}_{H-\ell} = -\sum_{\alpha,\beta=e,\mu,\tau} \left(Y^{\ell}_{\alpha\beta} \overline{L_{\alpha L}} \Phi \, \ell_{\beta R} + Y^{\nu}_{\alpha\beta} \overline{L_{\alpha L}} \widetilde{\Phi} \nu_{\beta R} \right) + h.c. \, . \tag{2.2}$$

Here, Φ is the Higgs field, the $L_{\alpha L}$ are what generate the lepton fields:

$$L_{eL} = \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}, \quad L_{\mu L} = \begin{pmatrix} \nu_{\mu L} \\ \mu_L \end{pmatrix}, \quad L_{\tau L} = \begin{pmatrix} \nu_{\tau L} \\ \tau_L \end{pmatrix}, \tag{2.3}$$

and $Y^{\ell}_{\alpha\beta}$ and $Y^{\nu}_{\alpha\beta}$ are 3×3 Yukawa matrices, where the Yukawa couplings of the Higgs field generate the lepton ($\ell_{R,L}$) masses.

[1] What is presented here is not an exhaustive list, but only the most commonly discussed mechanisms, in their simplest form. A clear and detailed review of existing options can be found in [31] .

This model treats neutrinos as Dirac particles, but it allows for the generation of a mass term of the form:

$$\mathcal{L}_D = -m_D \left(\overline{\nu}_R \nu_L + \overline{\nu}_L \nu_R\right). \tag{2.4}$$

Alternatively to the Higgs-generated masses, the introduction of Majorana masses via the type I Seesaw mechanism is also possible for neutral particles, such as neutrinos. This extension to the SM introduces Majorana mass terms to the Lagrangian, which for the left and right handed neutrinos are of the form:

$$\mathcal{L}_{M_L} = -\frac{1}{2} m_L \left(\overline{\nu}_L \nu_L^c + \overline{\nu}_L^c \nu_L\right) \tag{2.5}$$

$$\mathcal{L}_{M_R} = -\frac{1}{2} m_R \left(\overline{\nu}_R \nu_R^c + \overline{\nu}_R^c \nu_R\right), \tag{2.6}$$

where ν^c is the charge-conjugated field. This is only allowed for $m_L = 0$ in the SM, but the right-handed Majorana mass term and the Dirac mass term can be combined into one Lagrangian of the form:

$$\mathcal{L}_{\mathcal{D}+\mathcal{M}} = -\frac{1}{2} \begin{pmatrix} \overline{\nu}_L & \overline{\nu}_R^c \end{pmatrix} \begin{pmatrix} 0 & m_D \\ m_D & m_M^R \end{pmatrix} \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + h.c. \tag{2.7}$$

The m_R mass scale is not linked to the SM scale, which means that its value is unconstrained and, in principle, can be much larger than the SM scale would require. With the assumption that $m_R \gg m_D$, the diagonalization of the mass matrix in $\mathcal{L}_{D+M}$ leads to two Majorana eigenstates:

$$m_{\text{light}} \simeq \frac{m_D^2}{m_R}, \qquad m_{\text{heavy}} \simeq m_R. \tag{2.8}$$

In this model, the light, active neutrino flavors observed in nature emerge: $m_{light} \approx 0.1$ eV, when $m_D \approx 100$ GeV, which is on the quark mass scale. On the other hand[2], the massive right-handed

[2] Pun not intended.

Majorana neutrinos acquire masses on the order of $m_M^R \approx 10^{14}$ GeV and are not observed, although they could have been found in the early universe due to the high energy density.

One of the questions that arises from the introduction of neutrino masses in the Standard Model is the Dirac or Majorana nature of neutrinos. Whereas Dirac particles are distinct from their anti-particles of the same helicity, as are the rest of the fermions, Majorana particles are effectively their own anti-particles when accounted for in lepton flavor conserving processes such as double beta decay.[3] The observation of processes unique to Majorana neutrinos, such as neutrino-less double beta decay ($0\nu\beta\beta$) [34], are the only probe on the nature of the neutrino masses. The contribution from $0\nu\beta\beta$ searches and other probes to our understanding of the neutrino masses will be mentioned in relation to the neutrino mass hierarchy in Sec. 2.3.2.

2.2.1 Neutrino Eigenstates

The term "neutrino masses" does not refer to the masses of the particles portrayed in the Standard Model diagram in Fig.2.1, but to the mass eigenstates of neutrinos: ν_1, ν_2, and ν_3. These three mass eigenstates are in turn quantum superpositions of the flavor eigenstates ν_e, ν_μ, and ν_τ. In fact, the neutrino eigenstates can be written in either the flavor or the mass eigenstate basis, which diagonalize the weak interaction and the free particle Hamiltonian, respectively. Thus, the mass eigenstates can be written as a linear combination of the flavor eigenstates (and vice-versa):

$$|\nu_k\rangle = \sum_\alpha U_{k\alpha} |\nu_\alpha\rangle \,, \tag{2.9}$$

[3] The statement that Majorana neutrinos are their own anti-particle is commonly used in this description, but this is not necessarily the case (see [33]). Two neutrinos may still have opposite helicity but be effectively identical, as far as the process of annihilation is concerned. This is different from the case of particles which are identical to their anti-particle in all quantum numbers.

where U is a unitary rotation matrix known as the Pontecorvo, Maki, Nakagawa, Sakata (PMNS) matrix [35, 36]. This matrix relates the $|\nu_\alpha\rangle$ flavor eigenstates, where $\alpha = e, \mu, \tau$ and the $|\nu_k\rangle$ mass eigenstates, where $k = 1, 2, 3$. The elements $U_{k\alpha}$ are amplitudes of the flavor eigenstates ν_α in mass state k.

In order to define neutrinos conceptually in terms of their unchanging identity and characteristic properties, it is necessary to state that neutrinos exist as the mass eigenstates, since it is the mass eigenstates that have unchanging identity and mass. However, it is impossible to observe or measure the properties of ν_1, ν_2, and ν_3 directly, since upon interacting, their wave function is collapsed into one of the flavor eigenstates. In other words, neutrino properties can only be measured at moments at which they are bound to be in a single flavor eigenstate by the known laws of conservation.

2.3 Neutrino Oscillations

"There is no quantum world. There is only an abstract quantum physical description. It is wrong to think that the task of physics is to find out how nature is. Physics concerns what we can say about nature..."
–Niels Bohr

Neutrinos are unambiguously in one flavor state $|\nu_\alpha\rangle$ only when they are produced and upon their interaction with matter. At any other moment, neutrinos are in a superposition of the flavor eigenstates which allows for the phenomenon of neutrino oscillations to occur. While the flavor of a neutrino may be known at creation, it propagates as one of the mass eigenstates $|\nu_k\rangle$ given by Eq. 2.9. Thus the probability that a neutrino is created as one flavor eigenstate but interacts as another is non-zero. This is known as neutrino flavor mixing or neutrino oscillations.

The probability to observe a neutrino of a given flavor can be obtained by propagating the initial flavor eigenstate from Eq. 2.9 in time:

$$|\nu_\alpha(t)\rangle = \sum_k e^{-iE_k t} U^*_{\alpha k} |\nu_k\rangle \tag{2.10}$$

where the e^{-iEt} is the time-evolution operator for the free particle Hamiltonian in natural units.

For a relativistic free particle, $E_k^2 = p^2 + m_k^2$ where we assume $p_k = p_l \equiv p$ [37], and since neutrino masses are very small with respect to their energies, $E_k \approx |p| + \dfrac{m_k^2}{2|p|}$. The probability of oscillation from flavor α to flavor β is given by

$$P_{\alpha\to\beta}(t) = |\left\langle \nu_\beta(t) \,|\nu_\alpha(t=0)\right|\nu_\beta(t)\,|\nu_\alpha(t=0)\rangle\,|^2 = |\sum_k U^*_{\alpha k} U_{\beta k} e^{-i\frac{m_k^2}{2|p|}t}|^2. \tag{2.11}$$

Under the assumption that the neutrinos are ultra-relativistic, we can use $t \approx L$ to write $t/|p|$ as L/E, where L and E are the distance traveled by the neutrino and the neutrino energy, respectively. L/E is typically expressed in units of km/GeV.

Taking advantage of the unitarity of U, Eq. 2.11 can be expanded into

$$P_{\alpha\to\beta} = \delta_{\alpha\beta} - 4\sum_{i>j} Re\left[U^*_{\alpha i}U_{\alpha j}U_{\beta i}U^*_{\beta j}\right]\sin^2\left(\frac{\Delta m^2_{ij}}{4E}L\right) + 2\sum_{i>j} Im\left[U^*_{\alpha i}U_{\alpha j}U_{\beta i}U^*_{\beta j}\right]\sin^2\left(\frac{\Delta m^2_{ij}}{2E}L\right) \tag{2.12}$$

which, for the survival probability ($\nu_\beta = \nu_\alpha$) reduces to:

$$P_{\nu_\alpha\to\nu_\alpha}(L,E) = 1 - 4\sum_{k>j}|U_{\alpha k}|^2|U_{\alpha j}|^2\sin^2\left(\frac{\Delta m^2_{kj}L}{4E}\right), \tag{2.13}$$

where $\Delta m^2_{ij} \equiv m^2_i - m^2_j$.

Neutrino oscillations are therefore governed not only by the elements of the matrix U, but also by the squared mass differences, Δm^2_{21}, Δm^2_{31}, and Δm^2_{32}.

A parametrization of the PMNS matrix in terms of rotation angles θ_{12}, θ_{13}, θ_{23}, and a complex phase δ_{CP} is convenient to describe neutrino oscillations and is given by:

$$U = \begin{bmatrix} c_{13}c_{12} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -c_{23}s_{12} - s_{13}s_{23}c_{12}e^{i\delta} & c_{23}c_{12} - s_{13}s_{23}s_{12}e^{i\delta} & c_{13}s_{23} \\ s_{23}s_{12} - s_{13}c_{23}c_{12}e^{i\delta} & -s_{23}c_{12} - s_{13}c_{23}s_{12}e^{i\delta} & c_{13}c_{23} \end{bmatrix}, \tag{2.14}$$

where $c_{ij} \equiv \cos\theta_{ij}$ and $s_{ij} \equiv \sin\theta_{ij}$. The s_{ij} terms in the matrix describe the amplitude of the oscillation probabilities in Eqs. 2.12 and 2.13 which, in some cases, also depend on the value of δ_{CP}.

It is useful to look at the expanded form of this matrix:

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix}\begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{bmatrix}\begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{2.15}$$

as it separates it into factors containing parameters that drive the oscillations in different channels.

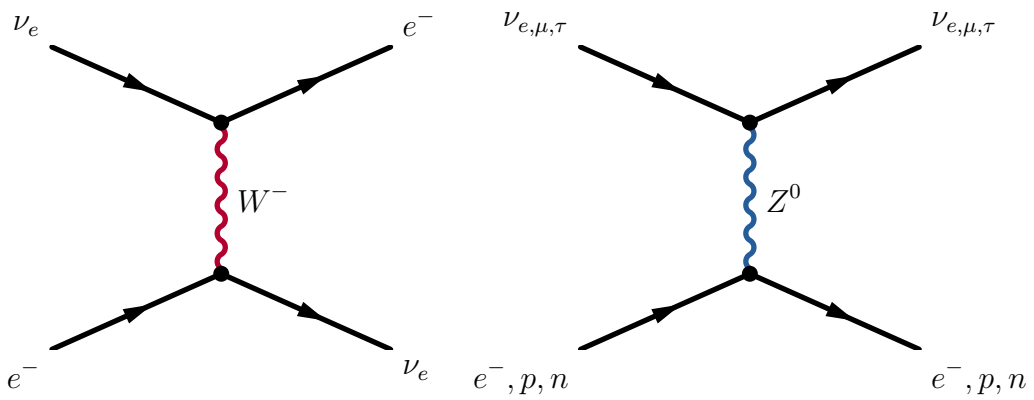

Figure 2.2: **Feynman diagrams for coherent forward scattering for neutrinos in matter (time goes from left to right). Left: an electron neutrino scattering via a charged-current (CC) interaction. Right: any neutrino flavor scattering via a neutral-current (NC) interaction.**

2.3.1 Matter Effects on Oscillations

When neutrinos travel through matter, an additional term is introduced to the Lagrangian to account for their interaction with particles in the medium they traverse.This effect, due to the presence of electrons, protons, and neutrons, changes the effective masses of the neutrinos in the medium due to coherent forward scattering of the neutrinos, an effect that is somewhat analogous to the refraction of light in a medium [38].

The presence of particles in the medium introduces potential terms in the Hamiltonian, $H_0 = U_{\beta i}\,\frac{m_i^2}{2E}\,U^\dagger_{i\alpha}$. As seen in Fig. 2.2, the potential coming from the presence of electrons will be felt only by the ν_e flavor eigenstates. This potential is called V_{CC} as it results from the charged current interactions ($W^\pm$ exchange) and is the expected potential for a cold electron gas, $V_{CC} = \sqrt{2}G_F N_e$. The effects from the NC potential—right panel of Fig.2.2—are felt by all neutrino flavors and through cancellation of the electron and proton contributions can be reduced just to the effect from the neutrons, $V_{NC} = \sqrt{2}G_F N_n$.[4] In turn, given that V_{NC} is felt equally by all flavors, it can be absorbed into a phase term

[4]This cancelation and more details of the calculation starting from a low-energy charged-current weak interaction Lagrangian can be found in Chapter 9 of [39], and a derivation analogous to refraction index is in the original paper [38].

and ignored as far as the effect in CC channel measurements is concerned.

The remaining potential, V_{CC}, is introduced in the Hamiltonian by adding an element, $A_{\alpha\beta} = \delta_{\alpha\beta}\delta_{\alpha e}V_{CC}$, such that the effective Hamiltonian is now given by:

$$H_{\alpha\beta} = \sum_i U_{\beta i}\,\frac{m_i^2}{2E}\,U_{i\alpha}^{\dagger} + A_{\alpha\beta}. \tag{2.16}$$

The Hamiltonian including matter effects is then:

$$H = \begin{pmatrix} U_{e1}+A & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \tag{2.17}$$

where $A \equiv \pm\sqrt{2}G_F N_e$, where the positive and negative sign apply to neutrinos and anti-neutrinos, respectively.

The matter effect can be seen in the leading term of the oscillation probability:

$$P_{\nu_\mu\to\nu_e} \approx \sin^2\theta_{23}\sin 2\theta_{13}^{\text{eff}}\sin^2\left(\frac{\Delta_{13}^{\text{eff}}L}{2}\right), \tag{2.18}$$

$$\sin 2\theta_{13}^{\text{eff}} = \frac{\Delta_{13}^2\sin 2\theta_{13}}{\left(\Delta_{13}^{\text{eff}}\right)^2}, \tag{2.19}$$

where

$$\Delta_{13}^{\text{eff}} = \sqrt{\left(\Delta_{13}\cos 2\theta_{13} - A\right)^2 + \Delta_{13}^2\sin^2 2\theta_{13}} \tag{2.20}$$

$$\Delta_{13} = \frac{\Delta m_{31}^2}{2E}. \tag{2.21}$$

Note that the first term inside the square root in Eq. 2.20 depends on the relative sign of Δm_{13}^2 and A. Thus, the oscillation probability $P(\nu_\mu \to \nu_e)$ is different for neutrinos versus antineutrinos as well as for normal versus inverted hierarchies.

When introducing matter effects the oscillation probability is given by:

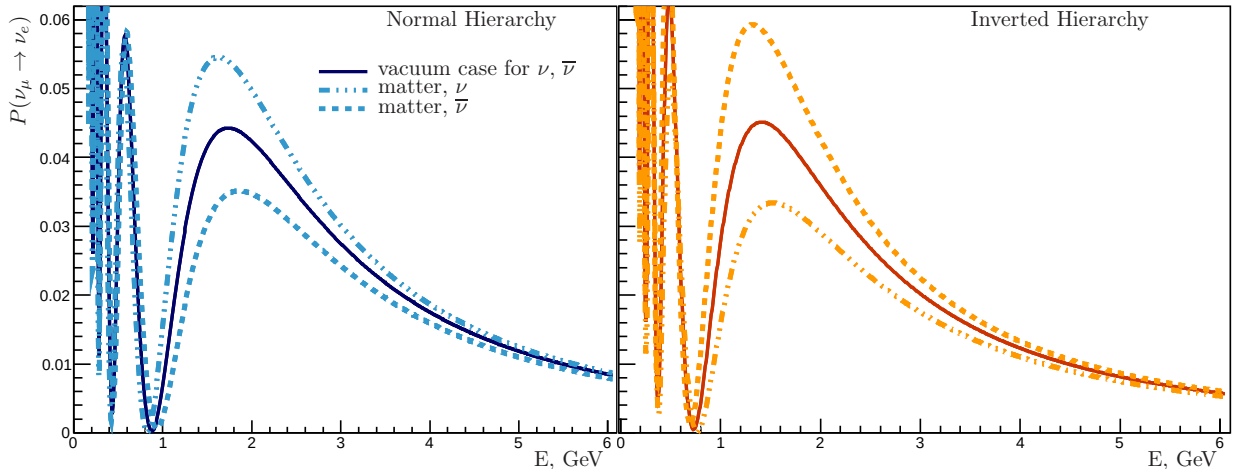

Figure 2.3: $\nu_\mu \to \nu_e$ **oscillation probabilities, from [40]. The orange and blue lines are probabilities for inverted and normal hierarchy, respectively. The dotted lines indicate that, as expected, the probability is enhanced (suppressed) for neutrinos (anti-neutrinos) in the normal hierarchy and enhanced (suppressed) for anti-neutrinos (neutrinos) in the inverted hierarchy.**

$$P_{\nu_\mu \to \nu_e} = P_{atm} + P_{sol} + 2\sqrt{P_{atm}}\sqrt{P_{sol}}\left(\cos\Delta m^2_{32}\cos\delta \mp \sin\Delta m^2_{32}\sin\delta\right), \tag{2.22}$$

where

$$P_{atm} = \sin^2\theta_{23}\sin^2 2\theta_{13}\frac{\sin^2\left(\Delta_{31} \mp AL\right)}{\left(\Delta_{31} \mp AL\right)^2}\Delta^2_{31}, \tag{2.23}$$

$$P_{sol} = \cos^2\theta_{23}\sin^2 2\theta_{12}\frac{\sin^2\left(\mp AL\right)}{\left(\mp AL\right)^2}\Delta^2_{21}, \tag{2.24}$$

with the positive and negative signs now applying to anti-neutrinos and neutrinos, respectively.

The matter effects on the oscillation probability present an advantage for experiments seeking to determine the value of δ_{CP} and distinguish between the two neutrino mass hierarchies. In particular, the probability for NOvA's main channel, $P_{\nu_\mu \to \nu_e}$, is shown in Fig.2.3 for different values of the oscillation parameters it will measure. The influence of these effects on the design of the NOvA experiment will be explained in Chapter 4.

2.3.2 Oscillation Parameters and Neutrino Properties

Experiments have measured oscillations and constrained the values of the mixing angles and Δm^2_{ij} parameters to high precision, and phenomenological analyses exist that study their results and obtain global fits representing the current knowledge of neutrino oscillations in the field. The status of our knowledge about the oscillation parameters given by the global analysis described in [41] is summarized in the following table and depicted in Fig.2.4:

Parameter	Ordering	Best fit	1σ range	2σ range	3σ range
$\delta m^2/10^{-5}$ eV^2	NO, IO, Any	7.37	7.21 – 7.54	7.07 – 7.73	6.93 – 7.96
$\sin^2\theta_{12}/10^{-1}$	NO, IO, Any	2.97	2.81 – 3.14	2.65 – 3.34	2.50 – 3.54
$\lvert\Delta m^2\rvert/10^{-3}$ eV^2	NO	2.525	2.495 – 2.567	2.454 – 2.606	2.411 – 2.646
	IO	2.505	2.473 – 2.539	2.430 – 2.582	2.390 – 2.624
	Any	2.525	2.495 – 2.567	2.454 – 2.606	2.411 – 2.646
$\sin^2\theta_{13}/10^{-2}$	NO	2.15	2.08 – 2.22	1.99 – 2.31	1.90 – 2.40
	IO	2.16	2.07 – 2.24	1.98 – 2.33	1.90 – 2.42
	Any	2.15	2.08 – 2.22	1.99 – 2.31	1.90 – 2.40
$\sin^2\theta_{23}/10^{-1}$	NO	4.25	4.10 – 4.46	3.95 – 4.70	3.81 – 6.15
	IO	5.89	4.17 – 4.48 $\oplus$ 5.67 – 6.05	3.99 – 4.83 $\oplus$ 5.33 – 6.21	3.84 – 6.36
	Any	4.25	4.10 – 4.46	3.95 – 4.70 $\oplus$ 5.75 – 6.00	3.81 – 6.26
δ/π	NO	1.38	1.18 – 1.61	1.00 – 1.90	0 – 0.17 $\oplus$ 0.76 – 2
	IO	1.31	1.12 – 1.62	0.92 – 1.88	0 – 0.15 $\oplus$ 0.69 – 2
	Any	1.38	1.18 – 1.61	1.00 – 1.90	0 – 0.17 $\oplus$ 0.76 – 2

Table 2.1: Results of the global 3ν oscillation analysis, in terms of best-fit values for the mass-mixing parameters and associated $n\sigma$ ranges ($n = 1, 2, 3$), defined by $\chi^2 - \chi^2_{\text{min}} = n^2$ with respect to the separate minima in each mass ordering/hierarchy (NO, IO) and to the absolute minimum in any ordering. δm^2 is defined herein as $m^2_3 - (m^2_1 + m^2_2)/2$, and that δ is taken in the (cyclic) interval $\delta/\pi \in [0, 2]$. From [41].

As mentioned in Chapter 1, neutrinos are linked to fundamental questions through the non-zero value of their masses and the phenomenon of oscillations. In particular, the neutrino mass hierarchy, the octant of the angle θ_{23}, and the value of δ_{CP} are parameters which have yet to be determined, a task to which the NOvA experiment has been designed to contribute.

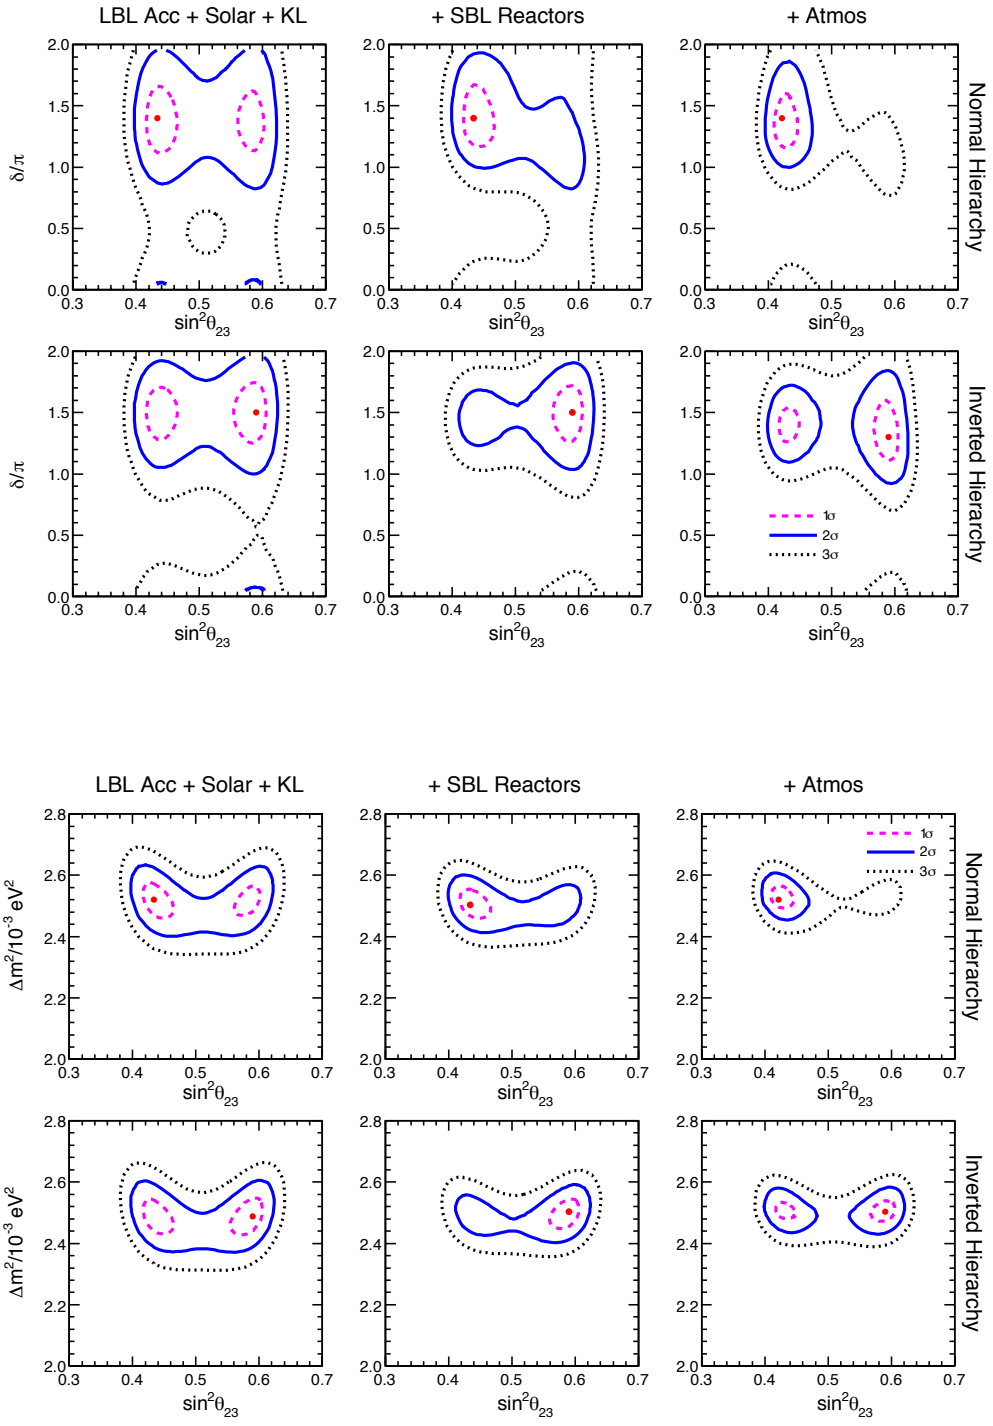

Figure 2.4: Covariance plot for $\sin^2\theta_{23}$, δ (top) and $\sin^2\theta_{23}$, $\sin^2\theta_{13}$ (bottom). From left to right, the regions allowed at Nσ = 1, 2 and 3 refer to the analysis of long-baseline accelerator + Solar + KamLAND data (left panels), plus short-baseline reactor data (middle panels), plus atmospheric data (right panels). Best fits are marked by dots. The three upper (lower) sub-panels refer to normal (inverted) ordering. From [41].

NEUTRINO MASS HIERARCHY

Neutrino masses are difficult to probe directly, given their small value as well as the small neutrino cross sections. However, indirect probes on the neutrino masses have been explored to set limits on related parameters, like the absolute effective neutrino mass $m_{\beta\beta}$, the lightest of the neutrino masses, and the neutrino mass ordering or mass hierarchy, the question of whether the mass of the ν_3 state is lighter or heavier than the states ν_1 and ν_2 (Fig.2.5).

The complementarity of different probes on the neutrino masses can be seen in Fig.2.6, which shows the current limits on $m_{\beta\beta}$ and $m_{lightest}$ set by different probes. Because the neutrino masses are related to structure formation in the early universe, measurements of the cosmic microwave background are sensitive to upper limits on the sum of the neutrino masses [42]. The absolute mass scale $m_{\beta\beta}$ can be estimated from searches for neutrino-less double beta decay ($0\nu\beta\beta$) and is given by:

$$m_{\beta\beta,NH} = \left| c_{12}^2 c_{13}^2 m_{\text{min}} + e^{\imath\alpha_1} s_{12}^2 c_{13}^2 \sqrt{m_{\text{min}}^2 + \Delta m^2} + e^{\imath\alpha_2} s_{13}^2 \sqrt{m_{\text{min}}^2 + \Delta m^2 + \frac{\Delta m^2}{2}} \right|, \tag{2.25}$$

$$m_{\beta\beta,IH} = \left| c_{12}^2 c_{13}^2 \sqrt{m_{\text{min}}^2 + \Delta m^2 - \frac{\Delta m^2}{2}} + e^{\imath\alpha_1} s_{12}^2 c_{13}^2 \sqrt{m_{\text{min}}^2 + \Delta m^2 + \frac{\Delta m^2}{2}} + e^{\imath\alpha_2} s_{13}^2 m_{\text{min}} \right|, \tag{2.26}$$

where $\alpha_1 = 2\lambda_1$ and $\alpha_2 = 2(\lambda_2 - \delta)$, and $\lambda_{1,2}$ are the Majorana CP phases [43].

The neutrino mass scale can also be expressed in terms of the measurable half-life of $0\nu\beta\beta$ decay $T_{1/2}^{0\nu}$ for a given isotope:

$$m_{\beta\beta} = (G^{0\nu})^{-1} |M^{0\nu}|^{-2} T_{1/2}^{0\nu} \tag{2.27}$$

where $G^{0\nu}$ is the phase space factor and $M^{0\nu}$ is the nuclear matrix element.

Currently the best upper limit on $m_{\beta\beta}$, from the Kamland-ZEN experiment [44], is set at (61-165) meV, although the uncertainties on $M^{0\nu}$ are large, which impacts the sensitivity of these measure-

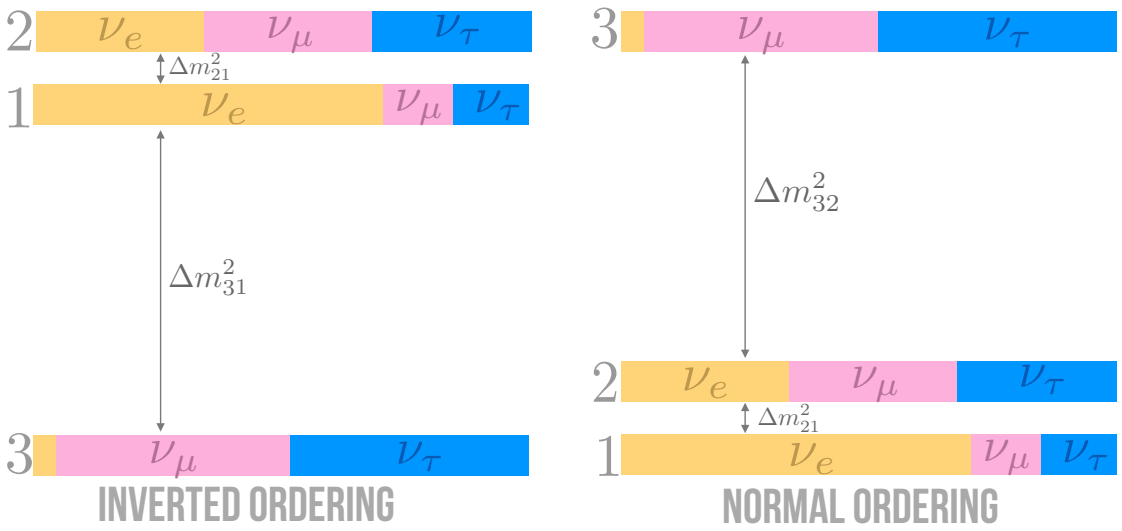

Figure 2.5: This diagram depicts the two possible mass orderings. The three horizontal bars on each column represent ν_1, ν_2, and ν_3. The colors represent the combination of states ν_e, ν_μ, and ν_τ in each mass state. The terms Δm^2_{ij} are the squared differences between two masses.

ments [45].

The relation between the different limits seen in Fig.2.6 could be interpreted as an indication that the hierarchy can be determined from a combination of constraints from cosmology and $0\nu\beta\beta$ experiments. However, the correlation depicted only holds for the type I Seesaw model presented in Sec. 2.2 and, thus, an unambiguous determination of the origin of the neutrino masses would be necessary for such a claim. The measurements of neutrino oscillation probabilities are the only unambiguous probe for the neutrino mass hierarchy regardless of the mass scale and, in turn, can constrain the values of $m_{\beta\beta}$ and the lightest neutrino mass for type I Seesaw.

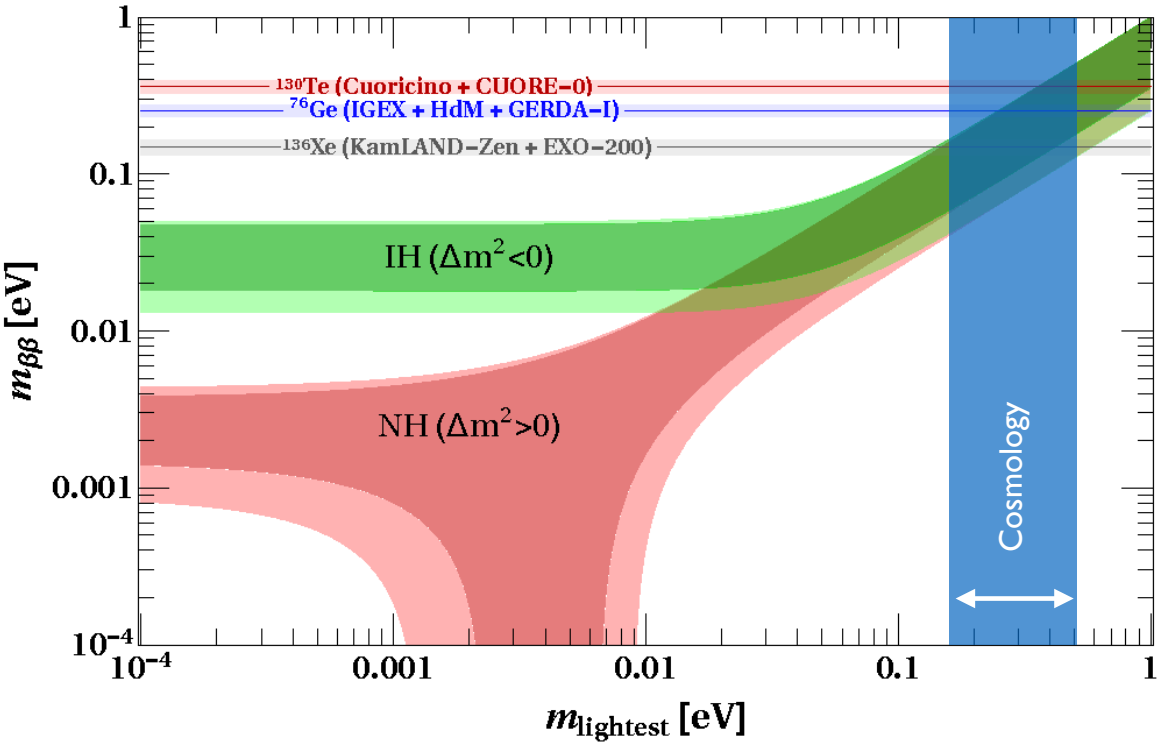

Figure 2.6: Current limits on $m_{\beta\beta}$ and the lightest neutrino mass. The green and red bands show the currently allowed regions for inverted and normal hierarchy, respectively. The three horizontal bands are the upper limits to $m_{\beta\beta}$ set by $0\nu\beta\beta$ searches in different isotopes [46]. The blue, vertical band indicates the upper limit to the lightest neutrino mass from cosmological constraints [47], in the style of [48].

THE OCTANT OF θ_{23}

The allowed values of the amplitude $\sin^2\theta_{23}$ at 3σ (from Table 2.1) are in the range 0.384–0.636. Thus, the question of whether this mixing is maximal is yet to be answered. The diagram in Fig.2.7 depicts the relation of the quantity $\sin^2\theta_{23}$ to the composition of the state ν_3. A maximal value of this mixing angle would imply equal fractions of ν_μ and ν_τ in the state ν_3 and maximize the amplitude of the corresponding oscillation probabilities. Otherwise, the octant of this angle would imply which flavor eigenstate is more largely represented in ν_3.

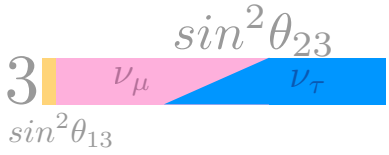

Figure 2.7: A diagram of the flavor eigenstate composition of the mass eigenstate ν_3.

THE δ_{CP} PHASE

The δ_{CP} phase is also not currently determined, as seen in Table 2.1. This parameter indicates the measure to which the assumed Charge-Parity symmetry in the lepton sector is violated, which is one of the conditions for connecting neutrinos to the question of baryogenesis[5]. The values of δ_{CP} that imply CP conservation are 0 and π, all others mean CP is violated, maximally so for δ_{CP} values of $\pi/2$ and $3\pi/2$.

The separation between the curves depicted in Fig. 2.8 allows one to translate the oscillation probability for neutrinos and anti-neutrinos into a constraint for the value of the δ_{CP} phase as well as a potential determination of the mass ordering, for some values of the $P(\nu_\mu \to \nu_e)$ $P(\bar{\nu}_\mu \to \bar{\nu}_e)$ phase space.

[5] Technically, it is the Majorana phases in Eq. 2.26 that are connected to leptogenesis, according to Sakharov's conditions [49]. However, the only probe on those phases is a combination of precision measurements of both the oscillation parameters and $0\nu\beta\beta$. Evidence of CP violation in oscillations would be the first indication of this possibility in the lepton sector.

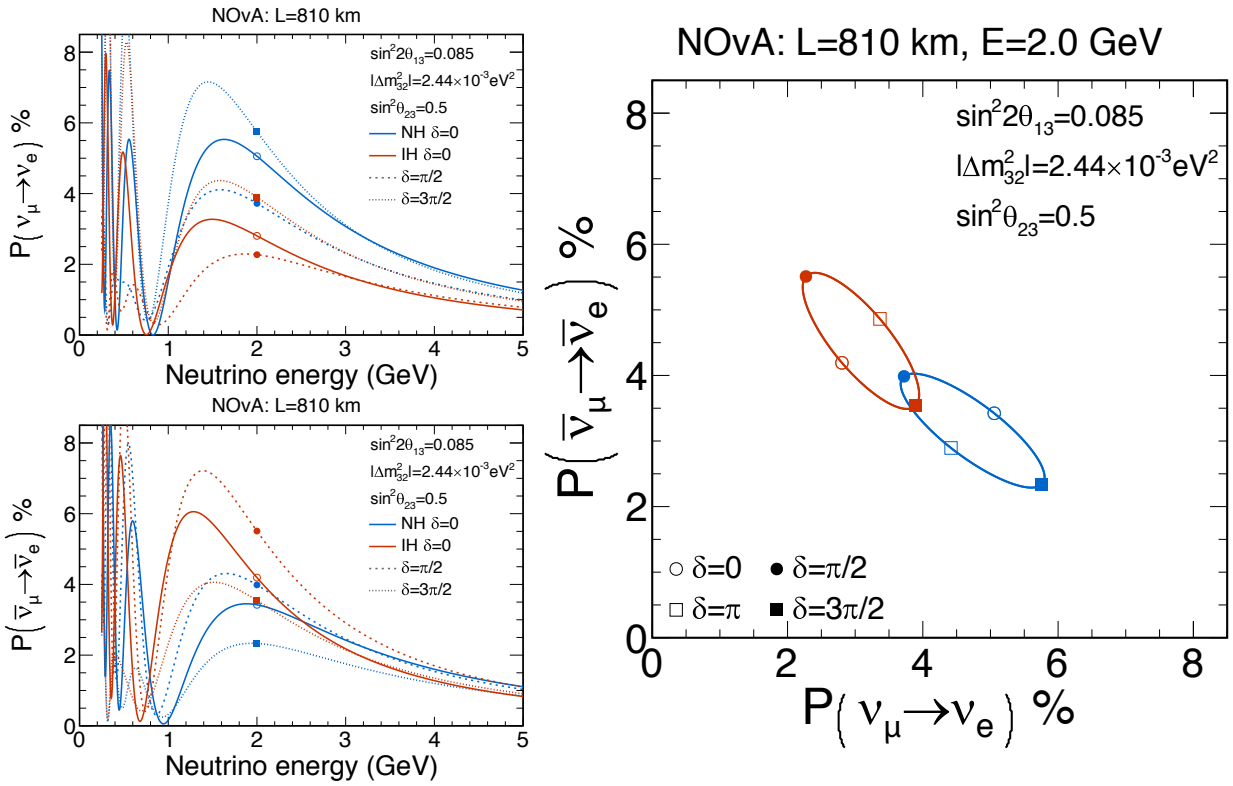

Figure 2.8: Oscillation probability curves for an 810 km baseline for $\nu_\mu \rightarrow \nu_e$ oscillations, assuming $\sin^2 2\theta_{13} = 0.085$, Δm^2_{32} = 2.44×10^-3 eV^2, and $\sin^2 \theta_{23} = 0.5$. Left: Oscillation probability for neutrinos (top) and anti-neutrinos (bottom). Right: Bi-probability plot. The curves on this plot show varying points of δ_{CP} for one point in $\sin^2\theta_{23}$ from the probability distributions on the left. Every curve in this plot corresponds to a point at $E_\nu = 2$ GeV, the rest of the curve shows the variation over δ_{CP}. The colors in all figures are for the inverted (orange) and normal (blue) hierarchies.

Chapter 3

Neutrino Detection

"We have to remember that what we observe is not nature herself,
but nature exposed to our method of questioning."
– Werner Heisenberg

The motivation behind particle detectors of any kind is to observe identifiable behaviors of the particles in which we are interested, through their direct interaction with some material. The technologies used for this purpose have evolved, from marks in photographic plates to electronics specially designed for the signal events expected, and are of varied type and complexity according to the needs of the experimental array in question, as are the analysis techniques that follow. The basic questions that drive the R&D behind particle detectors, however, are common across the board: How do particles interact in different environments? And how can we engineer media in which the interaction of the particles of interest (signal) is distinguishable from the interactions of other particles (background)? For this reason, a discussion of particle interactions and detection techniques is relevant in order to motivate the challenges present in experiment design in general, and to motivate the choices behind the details of the NOvA design in particular.

In the case of neutrino oscillation measurements, the metrics of interest are the flavor and the energy of the incoming neutrino. Thus, a knowledge of the expected energy spectrum and the particle interactions with matter in this context is necessary. Neutrinos interact solely via exchanges of W or Z

bosons, yielding a lepton of the same flavor and one or multiple hadrons. Neutrino detection is, thus, done through the detection of the outgoing particles from such interactions, whose characteristics are the real drivers of the technology of choice. In particular, the identification of the lepton flavor is of particular importance for analyses regarding oscillations, cross sections measurements, or any flavor dependent effects.

This chapter describes the main characteristics of the interaction of neutrinos and their daughter particles with matter, as well as the methods used by some neutrino experiments to detect them. The following sections will detail neutrino interactions with matter through which they produce leptons and hadrons in our detectors. The types of neutrino interactions are shown in Sec. 3.1. The characteristic energy depositions of leptons and hadrons as a handle to identification and energy reconstruction are presented in Sec. 3.2 and a motivation for the design of the NOvA detectors in the context of other technologies is introduced in Sec. 3.3.

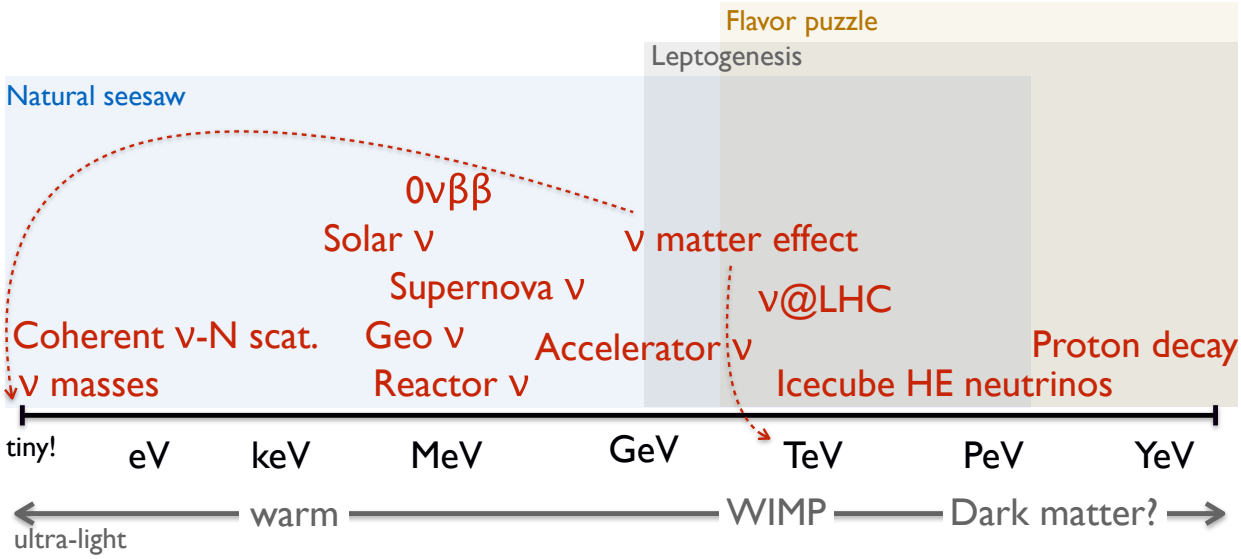

Figure 3.1: A diagram of the broad energy range of processes for neutrino production and their relation to the questions presented in Sec. 2.3.2.

3.1 Neutrino Interactions

The energy range in which neutrino sources are available, illustrated in Fig.3.1, is broad and rich in probes of a variety of physical phenomena, from nuclear fusion to supernovae. However, their measurement is a grand challenge in experimental particle physics.

It was thought even by Pauli himself that neutrinos were not only neutral but also extremely light and thus impossible to detect. Indeed, for most particle physics experiments and for collider experiments in particular, neutrinos appear merely as missing energy in known decays or interactions and their detection remains to this day a challenging physics problem on its own, due to their small interaction cross sections (less than 10^{-30} cm^2) over a large energy range, as shown in Fig.3.2.

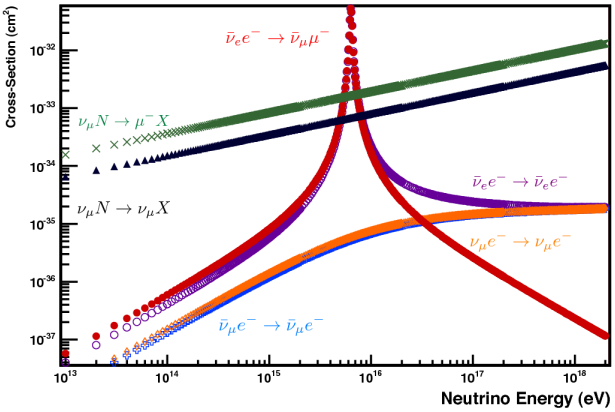

Figure 3.2: Neutrino cross sections at large energies. Note that even the largest cross-sections, coming from $\nu_e\, e \to \nu_\mu\, \mu$ (red) are still only 10^{-31} cm^2. Plot from [50].

Neutrino interactions with matter occur only via the weak force by the exchange of any of the weak force carriers, the Z or $W^\pm$ bosons. Upon this interaction with a nucleon the total number of leptons from the same family (where the number is negative for anti-leptons) will be the same in the final state. This is called lepton-flavor conservation, as seen in Fig.3.3. The interactions of neutrinos with nucleons are known to be lepton flavor conserving and they can be categorized by the boson that is exchanged as well as by the type of interaction with the nucleon.

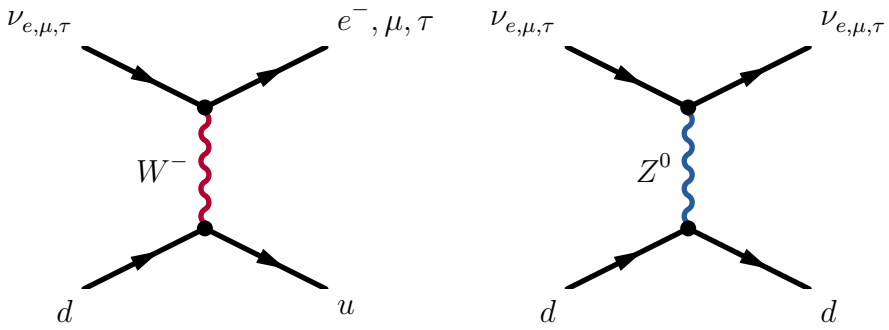

Figure 3.3: Feynman diagrams for charged current and neutral current interactions.

In order to detect neutrinos and measure their properties, there are two main obstacles to overcome: the small statistics due to the aforementioned small interaction cross sections, and the uncertainties in their modeled interactions with atomic nuclei. Since the number of interactions seen is the product of both the neutrino flux and the cross sections, this challenge calls for massive detectors as well as large neutrino fluxes, the likes of which we can produce in accelerators. The NOvA experiment employs both a massive detector and an intense beam of neutrinos to maximize signal statistics, as detailed in Chapter 4.

A more in-depth classification of neutrino interactions can be made beyond the charged current (CC) and neutral current (NC) interactions shown in Fig.3.3.

The interactions of the incoming neutrinos with atomic nuclei can be further classified into Quasi-Elastic (QE), Deep Inelastic Scattering (DIS), Resonant Pion Production (RES), and Coherent Scattering (COH), as seen in Fig.3.4 and other less prominent ones for the energy ranges relevant to NOvA[1]. However, while existing nuclear models can separate the contributions of each of these interaction types, they do not map uniquely into final states, i.e., the same final state particles can be found for more than one interaction type. This leads to an uncertainty in the expected energy distribution of neutrino events in our detectors shown in Fig.3.5.

[1] For instance, there are also interactions with the atomic electrons with even smaller cross sections.

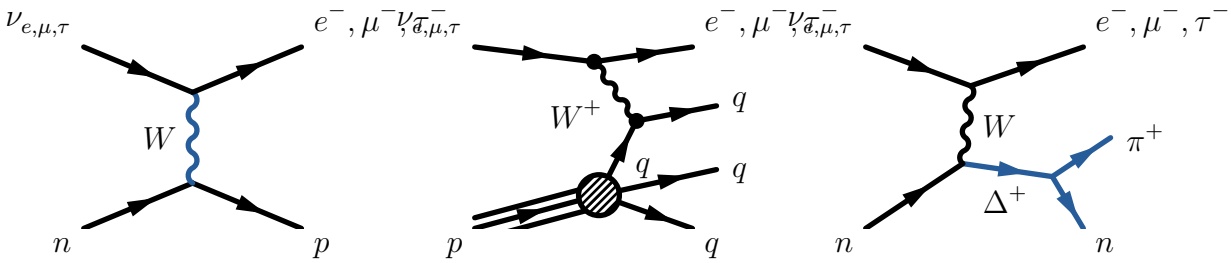

Figure 3.4: Neutrino interaction types by process: Quasi-Elastic (CC), Deep Inelastic Scattering (DIS), Resonant Pion Production (RES).

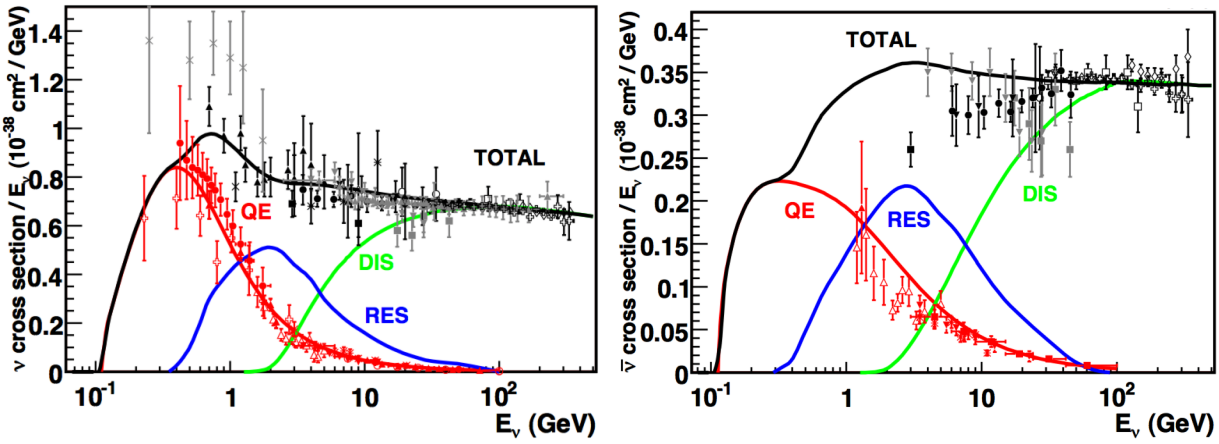

Figure 3.5: Neutrino Cross Sections in the energy range relevant to NOvA for neutrinos (left) and antineutrinos (right). From [50].

The study of neutrino cross sections is in itself a large effort in the field of neutrino physics. While not the topic of this dissertation, it is important to mention that our lack of complete understanding of neutrino cross sections is the largest systematic uncertainty in our analyses. The details of the effect of uncertainties in the cross-sections in our measurements are further explored in Sec. 6.6.

3.2 Handles on Identification and Energy

In order to design experiments that can detect and reconstruct the observables of interest in neutrino physics, it is important to understand the processes by which we may gather information from the interactions we study. The types of interactions that particles undergo and through which they de-

posit energy depend mainly on the energy of the particles and the characteristics of the active detector material they traverse. The circumstances (energy, material characteristics, etc.) that determine the predominant types of interactions are also different for each type of particle.

While mechanisms for energy losses vary greatly for different particles, a useful categorization to make broad distinctions is the following: Electromagnetic (mostly electrons and photons) and hadronic. This difference is further explored in section Sec. 5.3.2. The energy of neutrino interactions goes predominantly into charged heavy leptons[2], which upon interacting with matter lose their energy mostly through radiative losses or ionization. For the energy range encountered in NOvA, the main losses for the different types of leptons are dominated by the following:

Muon stopping power. The stopping power of muons and the dominant process for their energy losses is described by the curve in Fig.3.6. In the case of muons from neutrino interactions, the typical energies (constrained by the energy of the incoming neutrinos) usually start at the hundreds of MeV. Thus, their energy losses are typically in the Bethe-Bloch or the radiative regions, as defined in the figure.

Electron stopping power. In the case of electrons, the amount of energy lost to radiative processes is significant in comparison to their rest energy. Above some critical energy, E_c, their energy loss is primarily through Bremsstrahlung.

Compton scattering is also an important effect in the behavior of these particles as they pass through matter. However, the relative magnitude of this effect is dependent on the energy range of the particles in question as well. In the case of muons, pions and protons, scattering in the typical energy range of NOvA this effect will become important and useful for identification and reconstruc-

[2] Typically this refers to muons and electrons only, as the lifetime of ν_τ's is small. This causes them to decay (usually into lighter leptons)

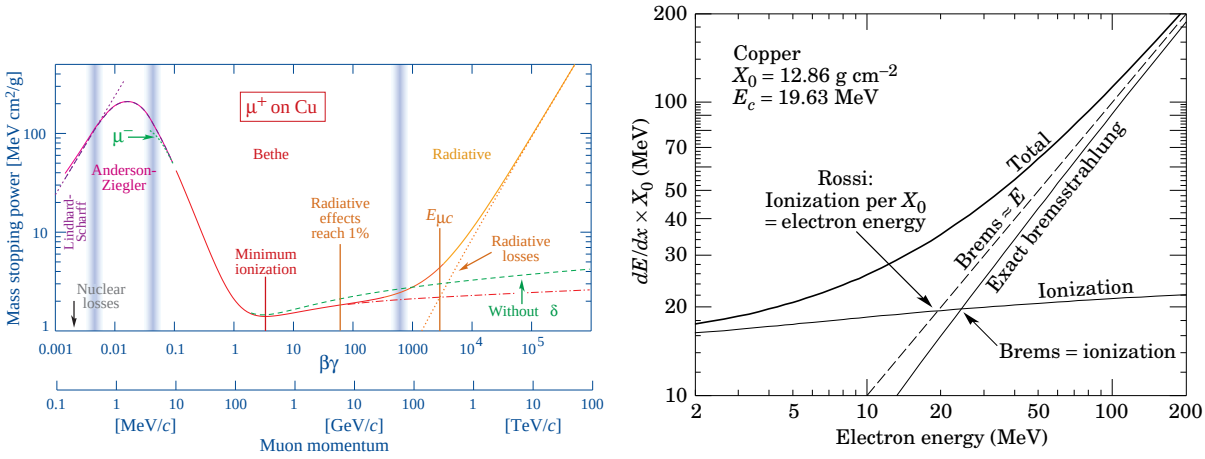

Figure 3.6: Stopping power of muons (left) and electrons (right).

tion. See Sec. 5.1.4.

3.3 Neutrino Detectors

While specific detector designs can vary significantly, most neutrino detectors can be classified into one or several groups based on the interaction medium. Most current and planned neutrino detectors fall into one or two of these groups: water Cherenkov, scintillator, high-Z, and noble element. These materials serve as the interaction medium for the neutrinos, but they must also be coupled with sensors and other electronics to perform the detection.

Cherenkov detectors employ water or oil as an interaction medium. Charged particles produced from neutrino interactions traveling faster than the speed of light in the medium produce a Cherenkov ring of light that falls incident upon light sensors, typically Photo-Multiplier Tubes (PMTs). These detectors are well-suited to measure lower-energy neutrinos in the MeV range, like solar and supernova neutrinos, as well as ultra-high energy neutrinos, such as astrophysical neutrinos.

Noble element detectors, particularly those of liquid argon outfitted with a Time Projection Cham-

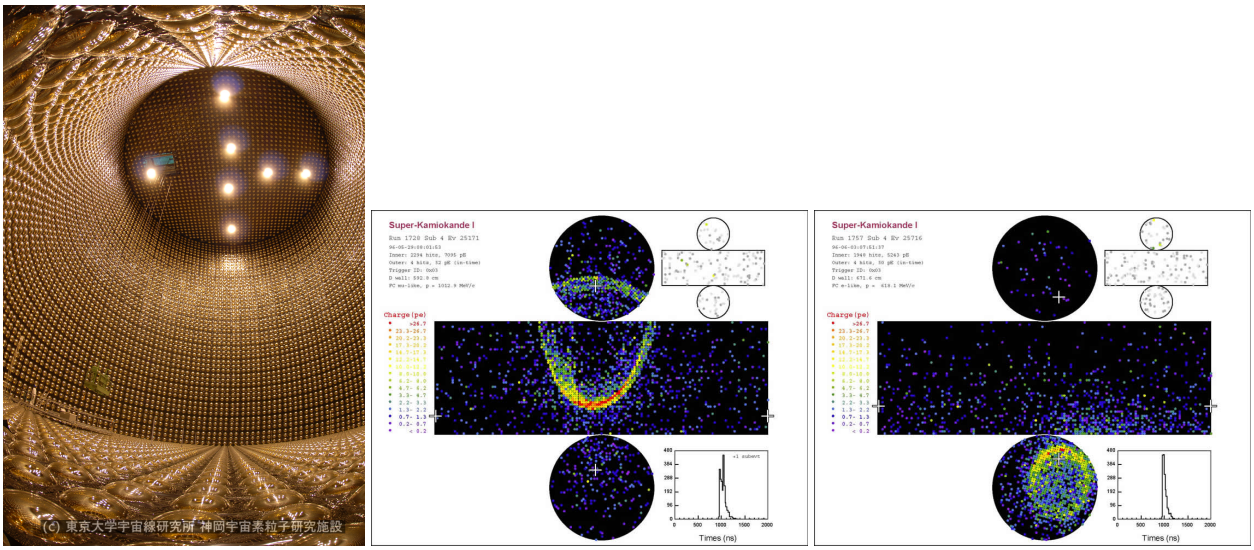

Figure 3.7: The Super-K detector and Event Displays. From [51].

ber (TPC) apparatus, are increasingly popular for neutrino detection. The liquid argon provides a dense medium for neutrino interactions whose charged products create electrons though ionization, the process in which energy is transferred to molecules in the medium, creating ions and ionization electrons. The electrons are forced to drift in a strong electric field towards a collection plane. This delayed signal, combined with a prompt optical signal gives these detectors excellent spatial and energy resolution.

Scintillator detectors, like MINOS and NOvA, rely on the light produced by the passage of charged particles through the detector, by exciting atomic electrons which de-excite by photon emission. In some cases, PMTs are used to collect the scintillation light, but in the case of NOvA the detector is outfitted instead with Avalanche Photo Diodes (APDs). The reasons for this choice will be detailed in Chapter 4.

The design and technology of choice varies widely experiment to experiment and is coupled to other constraints like the available resources and the expected reconstruction capabilities. The re-

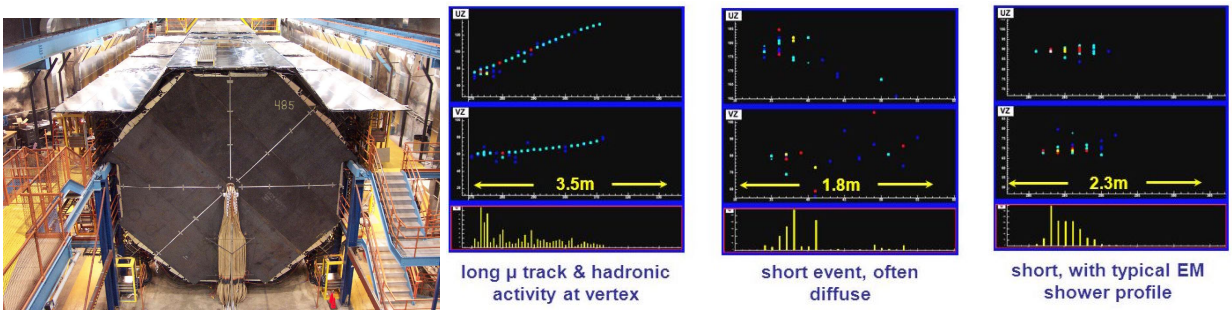

Figure 3.8: The MINOS detector and Event Displays. From [52].

quirements are driven by characteristics of the expected signal such as type of particle and energy of the signal and expected background, as well as the type of measurement which is to be performed and the sensitivity required. A good example of this is the comparison between the technology in the water Cherenkov detector of Super-Kamiokande (Fig.3.7) compared to the steel and scintillator technology in MINOS (Fig.3.8). The Super-Kamiokande detector's expected signal events are cones of blue light that should be distinguishable between muons and electrons. On the other hand, the MINOS detector was built to look for muon neutrino disappearance in an accelerator neutrino beam. Identifying and containing the muon energy was of paramount importance to their measurement. As seen in Fig.3.8, electron neutrinos are more difficult to identify in a detector like MINOS, but the combination of steel and scintillator and its magnetic field facilitate the containment, identification and reconstruction of muon neutrino events.

One of the design choices made for NOvA is an improvement on the capabilities of the MINOS detector, specifically the ability to identify and reconstruct electron neutrino interactions. The NOvA design is detailed in Chapter 4.

Chapter 4

The NOvA Experiment

"A scientist in his laboratory is not only a technician: he is also a child placed before natural phenomena which impress him like a fairy tale."
–Marie Curie

The NuMI Off-Axis ν_e Appearance (NOvA) experiment is a two detector neutrino oscillation experiment designed to look primarily for the appearance of ν_e's in a beam of predominantly ν_μ's. The choice of the $\nu_\mu \rightarrow \nu_e$ and $\nu_\mu \rightarrow \nu_\mu$ channels and the design of the experimental array are motivated by the following physics goals:

Precision measurement of $\sin^2\theta_{23}$. As seen in Sec. 2.3, the mixing angle θ_{23} contains information about the fraction of ν_μ and ν_τ present in the mass eigenstate ν_3.

Determination of the Neutrino Mass Hierarchy or Neutrino Mass ordering, as depicted in Fig.2.6.

Measurement of the δ_{CP} parameter. Specifically, the goal is to constrain possible CP conserving values in order to set a limit on the possibility for CP violation in the lepton sector.

The design of the NOvA experiment was drawn from the optimization and careful choice of detector and data acquisition technologies, as well as planned improvements to the neutrino beam given these physics goals and our current understanding of neutrino oscillations, as detailed in Chapter 2. The effect of neutrino passage through matter on the oscillation probabilities, discussed in Sec. 2.3, is particularly important and calls for a long baseline. The small cross sections, discussed in Sec. 3.1, call for an effort to maximize the neutrino flux and a massive detector.

The task of operating the experiment for optimal data-taking is also of crucial importance to the final result. Optimal operation and monitoring ensure maximal uptime synchronized with the uptime of the accelerator array, guarantee a long life for the electronics, and ensure optimal quality for the data that we collect. This chapter further motivates and details the experimental configuration in Sec. 4.1 and discusses the operation and monitoring of the experiment, from the detectors to the handling of the data, in Sec. 4.2.

4.1 Experiment Design

"We've got no money, so we've got to think."
– Ernest Rutherford

Experimental design is an optimization problem of accomplishing specific physics goals as precisely as possible, as soon as possible, and within realistic funding constraints. NOvA's design addresses the following challenges:

Maximizing the signal. Given the small value of the neutrino interaction cross-sections discussed in Sec. 3.1, large statistical uncertainties are expected for these measurements. This challenge is addressed on two fronts: Increasing the intensity of the neutrino beam, and maximizing detectors' active mass.

Flavor Identification. As discussed in Sec. 3.1, identification of the lepton is the only handle on identification of the signal events. The choice of material and detector structure was optimized, in particular, for the identification of electrons.

Reducing background. Suppressing the number of background events is a significant challenge, one which I will discuss in this chapter as well as in the reconstruction and analysis (Chapters 5 and 6). The experimental design impacts background reduction through accurate timing and the chosen beam energy.

GENERAL OVERVIEW OF THE NOVA EXPERIMENT:

A proton beam is accelerated and delivered by Fermilab's Main Injector [53] into a graphite target. The collision of the protons with the target produces a beam of neutrinos. The neutrinos are sampled at two locations, near the source with a near detector and some distance away from the source with a far detector, both of which are placed off axis relative to the beam. The interactions of particles in the detectors are triggered and recorded by a time-synchronized data acquisition system (DAQ).

The following subsections will describe the NOvA experimental setup in detail, starting with the neutrino beam in 4.1.1, and explaining the choice of off-axis location in 4.1.2. The details of the two detectors are presented in 4.1.3, and the detection mechanism, including the electronics involved, is described in 4.1.4. Finally, the details of the DAQ and the timing system are shown in Secs. 4.1.5 and 4.1.6.

4.1.1 NOvA's Neutrino Beam

NOvA's neutrinos come from the beam of Neutrinos at the Main Injector (NuMI) at Fermilab. This beam is produced by colliding protons, which are first accelerated at energies up to 120 GeV, into a graphite target 1.2 meters in length. This collision is known to produce charged mesons, specifically kaons and pions of both positive and negative charge, which predominantly decay via the following channels:

$$\pi^{\pm} \rightarrow \mu^{\pm} + \overset{(-)}{\nu}_{\mu}, \tag{4.1}$$

$$K^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}. \tag{4.2}$$

The NuMI beam itself is structured into packets or clusters of 4.5×10^{12} protons called bunches. By limiting the length and frequency of the proton bunches the NuMI beam can produce a pulsed stream

of neutrinos— also called spills—which are 10 μs streams separated by 1.3 sec intervals and containing 12 bunches each. The purpose of this segmented beam is to localize the neutrino signal in time, which is useful to minimize cosmogenic backgrounds.

The characteristics of the neutrino beam are largely consequential to the experimental sensitivities of NOvA. Not only are oscillation probabilities dependent on neutrino energy, the signal statistics are also a function of the beam's intensity. NuMI has undergone improvements starting before NOvA began taking data and occasionally since then. The original design intensity of NuMI is 700 kW, which was reached in 2017 (See Fig.4.1). NuMI is currently the most intense neutrino source ever built.

The total flux of neutrinos delivered by NuMI is dependent on the proton beam intensity as well as the processes in Eq. 4.1. The beam exposure is reported in the unambiguous quantity POT (protons on target), the number of protons which have been accelerated into the target, which is later translated into an expected flux.

The flux of neutrinos from the pion and kaon decays on a cross sectional area A at a distance z from the source is given by:

$$\Phi = \frac{2\gamma}{1+\gamma^2\theta^2}\frac{A}{4\pi z^2} \tag{4.3}$$

where θ is the angle between the meson and the neutrino direction and $\gamma = E_{(\pi,K)}/m_{(\pi,K)}$.

While this flux is dominated by ν_μ's, there is a secondary decay of interest: $\mu \rightarrow e + \bar{\nu}_e$. This decay introduces some contamination of ν_e's in the predominantly ν_μ beam. The composition of this neutrino beam has been measured by both the MINOS and the MINERvA experiments, which are also situated close the NuMI beam, and is simulated for NOvA as shown in Fig.4.2.

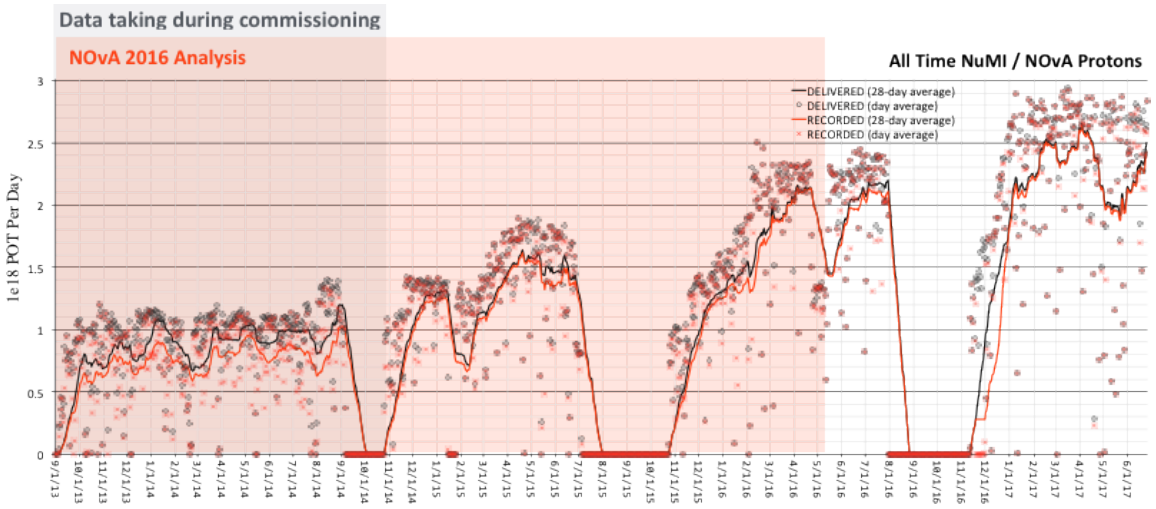

Figure 4.1: Delivered and recorded Protons on Target (POT) over NOvA's entire data taking period. The points are daily recorded values and the lines are traced over a 28-day average. NOvA's uptime efficiency, the fraction of time the detectors are running and recording beam data, is above 95% overall and above 98% after the detector was commissioned.

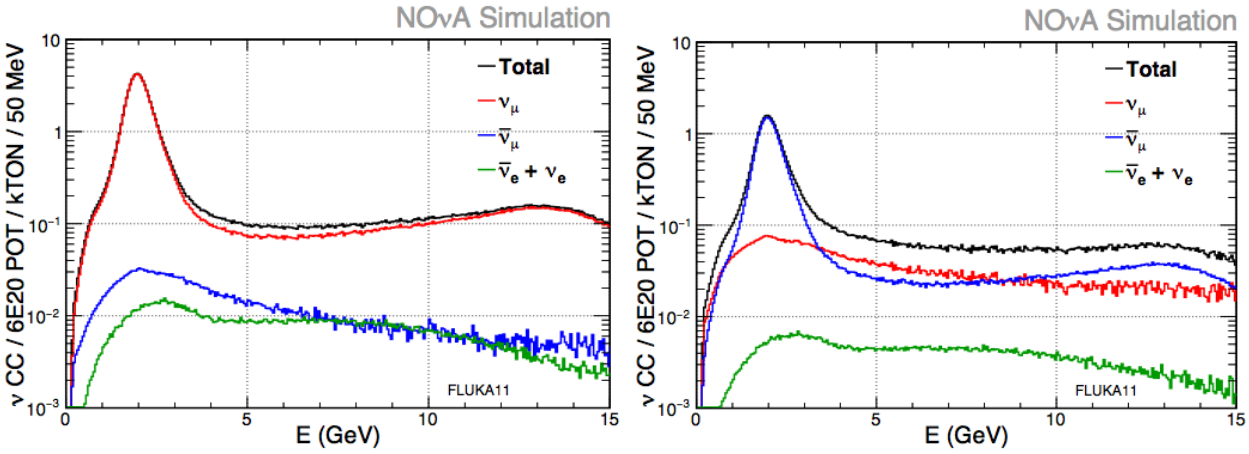

Figure 4.2: Simulated NuMI flux×cross-section in neutrino (left) and antineutrino (right) mode. The visible difference in wrong-sign contamination is an effect of mostly the difference in cross sections between neutrinos and antineutrinos and also by charge conservation from the initial proton-on-carbon collision. The simulation on this plot is made [54] using Flugg [55] and Fluka [56].

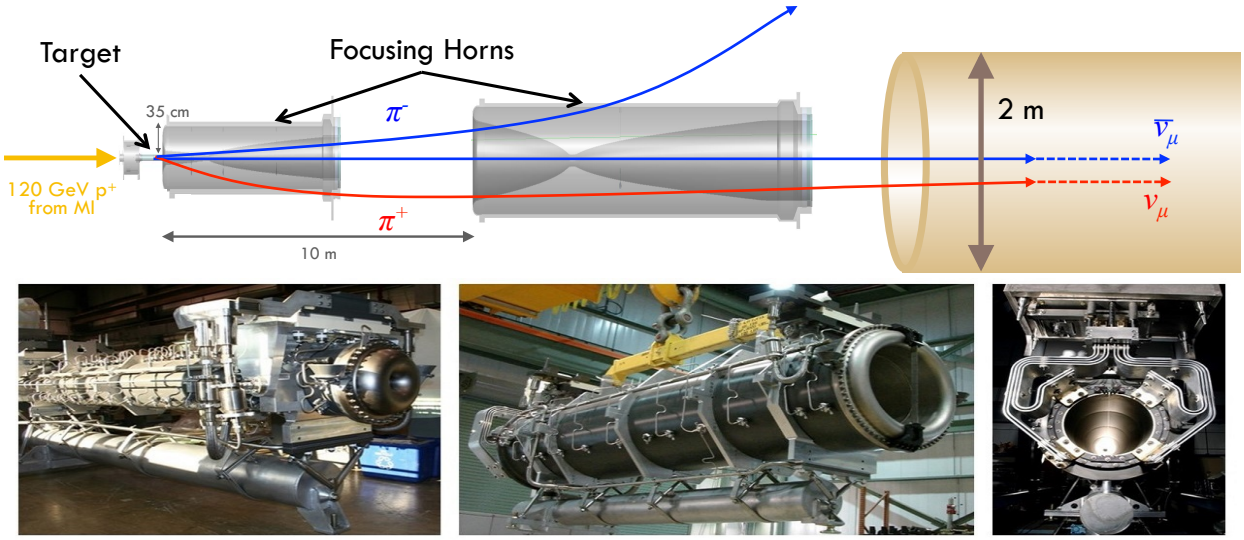

Figure 4.3: A diagram (top) and photographs (bottom) of the NuMI focusing horn. Protons coming from the accelerator (left of the page) impact a graphite target. The outgoing π^+ and K^+ mesons (only π^+ depicted) get focused by the field induced by a 100kA current in horn 1, and into horn 2, 10 meters away. The same current de-focuses negative π^- and K^- mesons. Reversing the polarity of horn 1 has the opposite effect, a design choice that allows for running in neutrino and antineutrino modes interchangeably. [53]

The measurement of oscillations requires not only flavor identification, but also the distinction between neutrinos and antineutrinos for reasons explained in Sec. 2.3. In order to isolate each contribution as much as possible, focusing horns are used to select the charged mesons produced by the proton interaction with the target. The focusing horns depicted in Fig.4.3 allow for two running modes, with neutrinos or antineutrinos for the main composition of the beam, respectively. Reversing the current on the horns will cause the opposite effect, producing a beam of predominantly antineutrinos.

4.1.2 Off-axis Spectrum and NC Backgrounds

As described in Sec. 3.1, neutral current (NC) events are the main backgrounds to the $\nu_\mu \rightarrow \nu_e$ electron neutrino signal coming from the NuMI beam. As the outgoing neutrino in these events carries with it some of the incoming neutrino energy, NC events with energies higher than that of signal events can mimic them if misidentified. Thus, constraining the energy of the incoming neutrinos is beneficial to the goal of background rejection. In addition to reducing the neutral current backgrounds, con-

straining the incoming neutrino energy E_ν distribution can be used in combination with the baseline L to constrain the mean energy of the expected signal events at a value consistent with the expected maximum of the oscillation probability.

In the case of NOvA, E_ν is constrained by placing the detector off-axis with respect to the NuMI beam direction. This takes advantage of decay kinematics, given that the angle of the outgoing neutrino from the meson decay is related to its energy. Thus, the mean E_ν increases as the angle with respect to the beam increases. The NOvA far detector (FD) was built at a location 14 mrad off-axis with respect to the direction of the NuMI beam. As seen in Fig.4.4, this effectively constrains the flux of incoming neutrinos to a narrow beam spectrum peaking close to 1.9 GeV.

Constraining the incoming neutrino energy represents an additional advantage to the measurement of oscillation probability. As seen in Eq. 2.22, this probability depends on the ratio L/E, where L is the baseline in km and E is the incoming neutrino energy E_ν. Fig.4.4 shows the expected probability of oscillations for an assumption of matrix parameters compatible with current best values. This shows the oscillation maximum expected for $\nu_\mu \rightarrow \nu_e$ at $L/E \sim 400$ km/GeV. With a baseline of 810 km, NOvA's L/E is 426.3 km/GeV at the energy peak, precisely chosen to maximize the sensitivity to these oscillations.

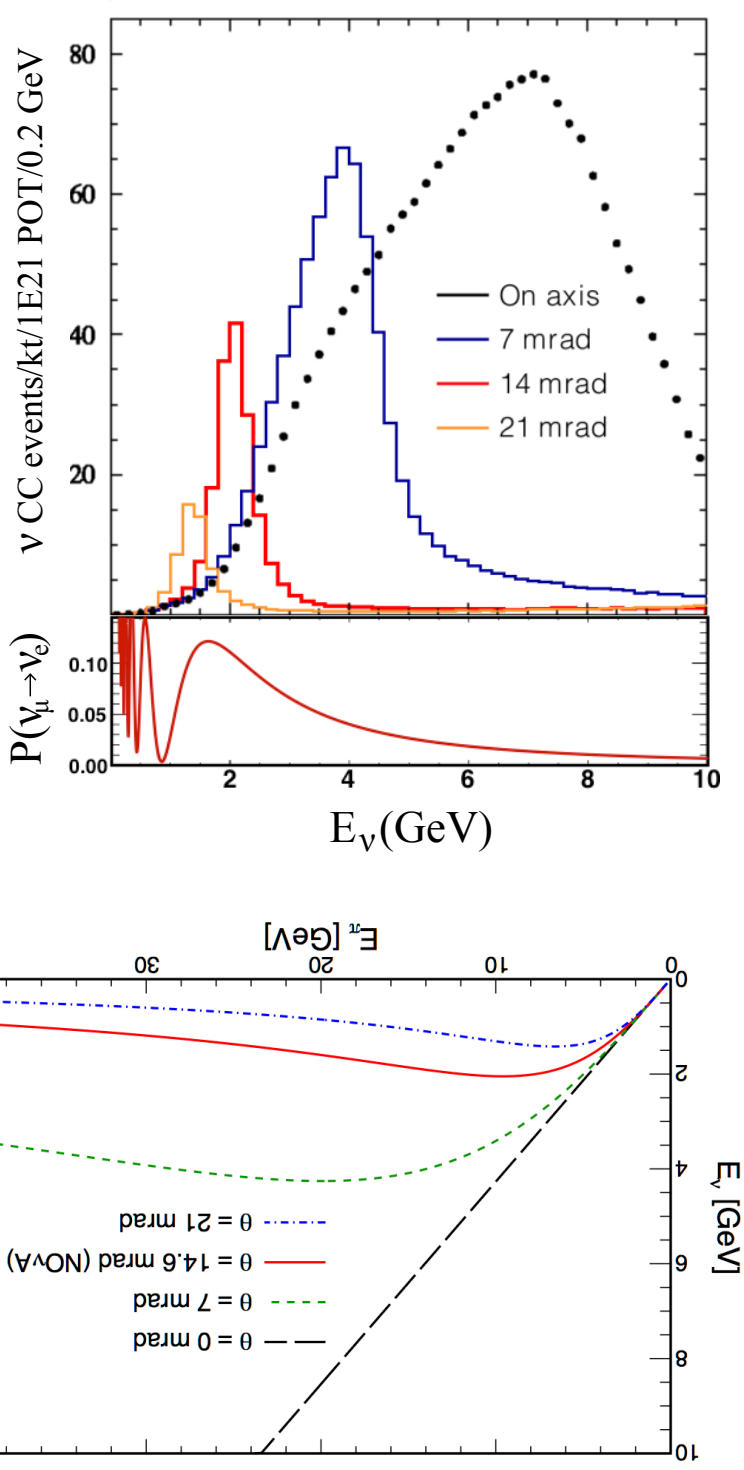

Figure 4.4: NOvA's Off-Axis Energy. Bottom: the neutrino energy as a function of the energy of the mother π for different angles. Top: the energy spectrum of the NuMI neutrinos on axis (black dotted) and for different angles off-axis. The bottom panel shows the oscillation probability as a function of E_ν for L=800 km which is approximately NOvA's baseline. This figure illustrates the advantage of the 14 mrad design choice, which places the peak energy of the neutrinos at the far detector close to the expected oscillation maximum at 1.9 GeV.

4.1.3 Detectors

Figure 4.5: Map of the NOvA experiment layout. The NuMI beam located at Fermilab makes a beam of neutrinos aimed at the location where the MINOS detector was once located. The NOvA far detector sits close to the Canadian border, 810 km away from the NuMI beam.

NOvA's experimental array, shown in Fig.4.5 consist of two detectors, a near detector (ND) $\sim$1 km from the NuMI beam and a far detector (FD) 810 km away, in Ash River, MN. The ND is located 105 meters underground, within the MINOS underground area at Fermilab. The FD is located on the surface, shielded only by the building which contains it in a 6-inch layer of barite overburden placed on the roof of the building, providing some shielding from cosmic ray particles entering from the atmosphere.

The NOvA detectors are made up of PVC modules, which are filled with liquid scintillator, and optimized for electron identification. The PVC modules are 15.8 meter long extrusions (3.84 m for the near detector) that are segmented into 3.5×5.6 cm cells, as shown in Fig.4.6. These modules of 32 cells each are then joined side-by-side to form detector planes which are 15.8×15.8 meters for the far detector, and up to 3.84×3.84 meters for the near detector. The full detector volume is made up of these planes, positioned front-to-back in alternating vertical and horizontal orientations. This yields two orthogonal views of the energy depositions, as shown in Fig.4.6.

While treatment of systematic uncertainties will not be detailed until Chapter 6.6, it is important to note that the methodology employed in their treatment is greatly impacted by the design of

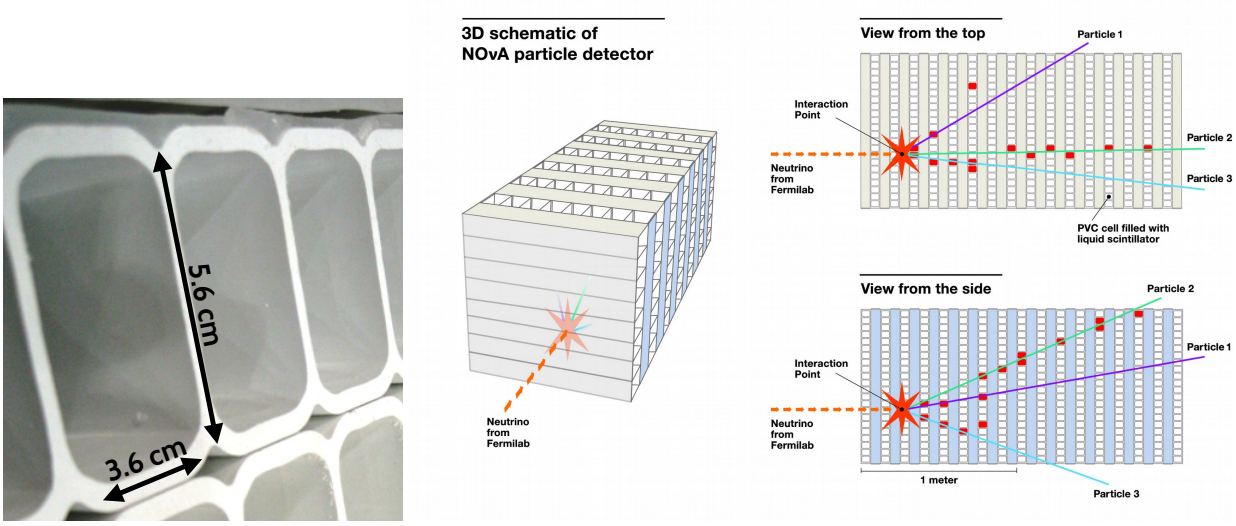

Figure 4.6: Detector structure. Left: A photograph of the PVC extrusions that form the NOvA detector planes. Right: A structural diagram of the detector plane orientation, from [57]. Detector readout is done in vertical (side view) and horizontal (top view) planes independently. This design allows for 3D reconstruction of particle trajectories.

the detectors and their similarities (or lack thereof). For instance, some systematic uncertainties are common to both detectors and are expected to cancel upon a comparison of measurements between the two, if the uncertainties apply identically. This is the motivation behind the choice to build both detectors as functionally equivalent in technology as possible.

The near detector is an array of 192 alternating horizontal and vertical planes, the last 22 of which alternate a steel plane with each vertical-horizontal plane as shown in Fig.4.7. This section of the ND is called the *muon range* . The purpose of the higher density of the steel planes is to absorb a larger fraction of the energy of passing muons, in order for them to stop before escaping the detector.

The far detector is an array of 896 planes, put together in sets of 64 planes called diblocks[1]. There are a total of 14 diblocks in the far detector, making up a total active mass of 14 ktons, 35% of which is PVC and 65% of which is scintillator.

[1]Technically, the far detector is divided into sections of 32 planes called blocks. However, given that blocks were joined in sets of two through their interplay with common pieces of electronics diblocks are usually taken as one unit for simplicity.

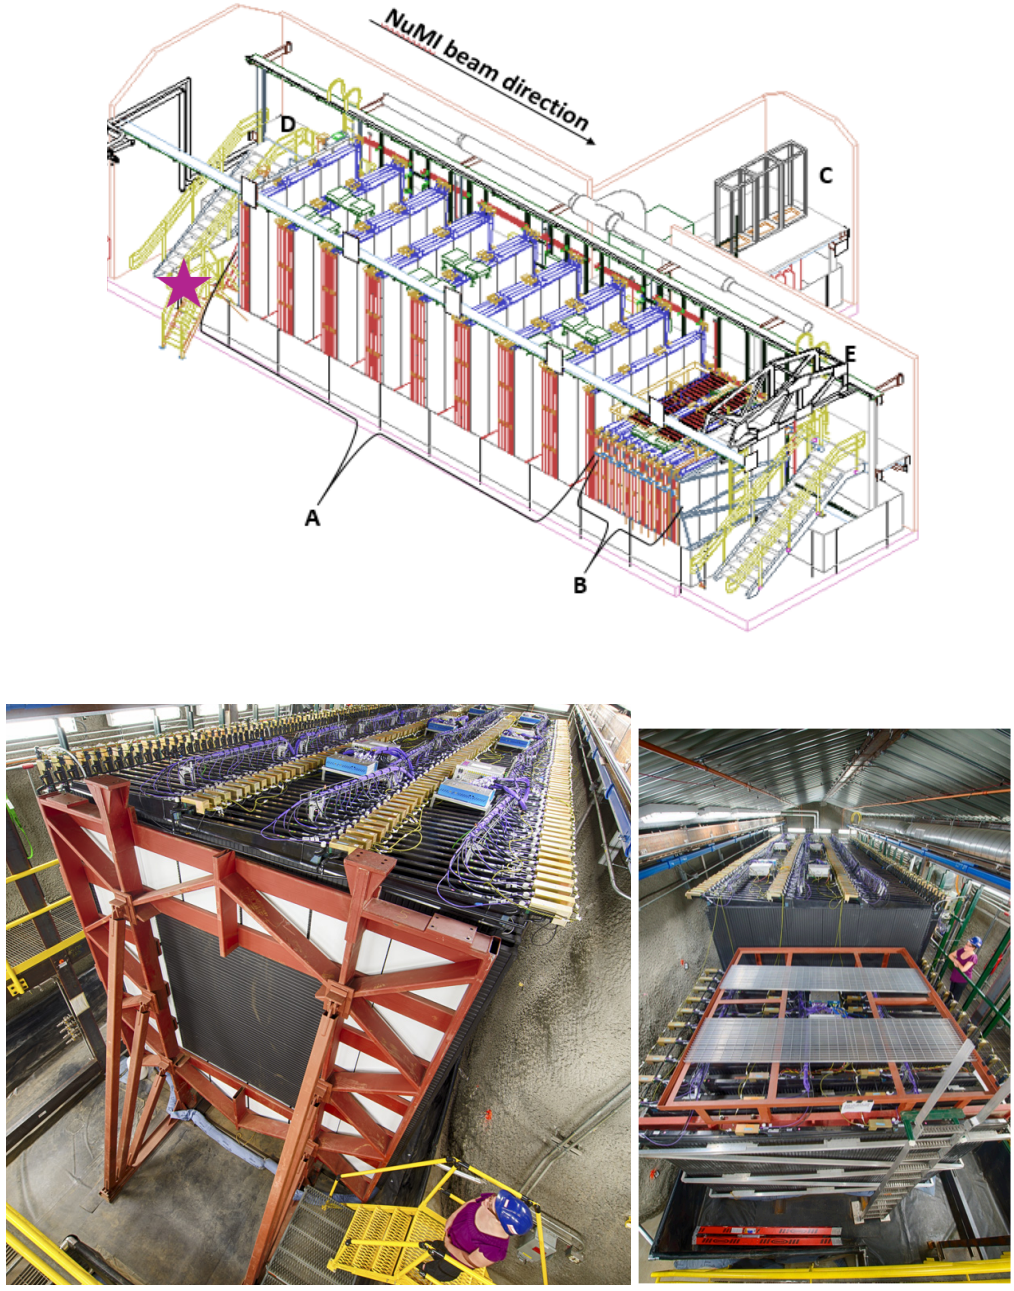

Figure 4.7: Top: A diagram of the near detector. Region A on this diagram represents the main body of the detector made of PVC+scintillator modules. The muon range (alternating active modules with steel planes) is labeled as region B. Left: A photograph of the near detector inside the detector hall underground at Fermilab and my friend Louise Suter who is represented by a star on the diagram on the top. Right: A photograph from the back of the near detector, the shorter planes at the front of the image form the muon range.

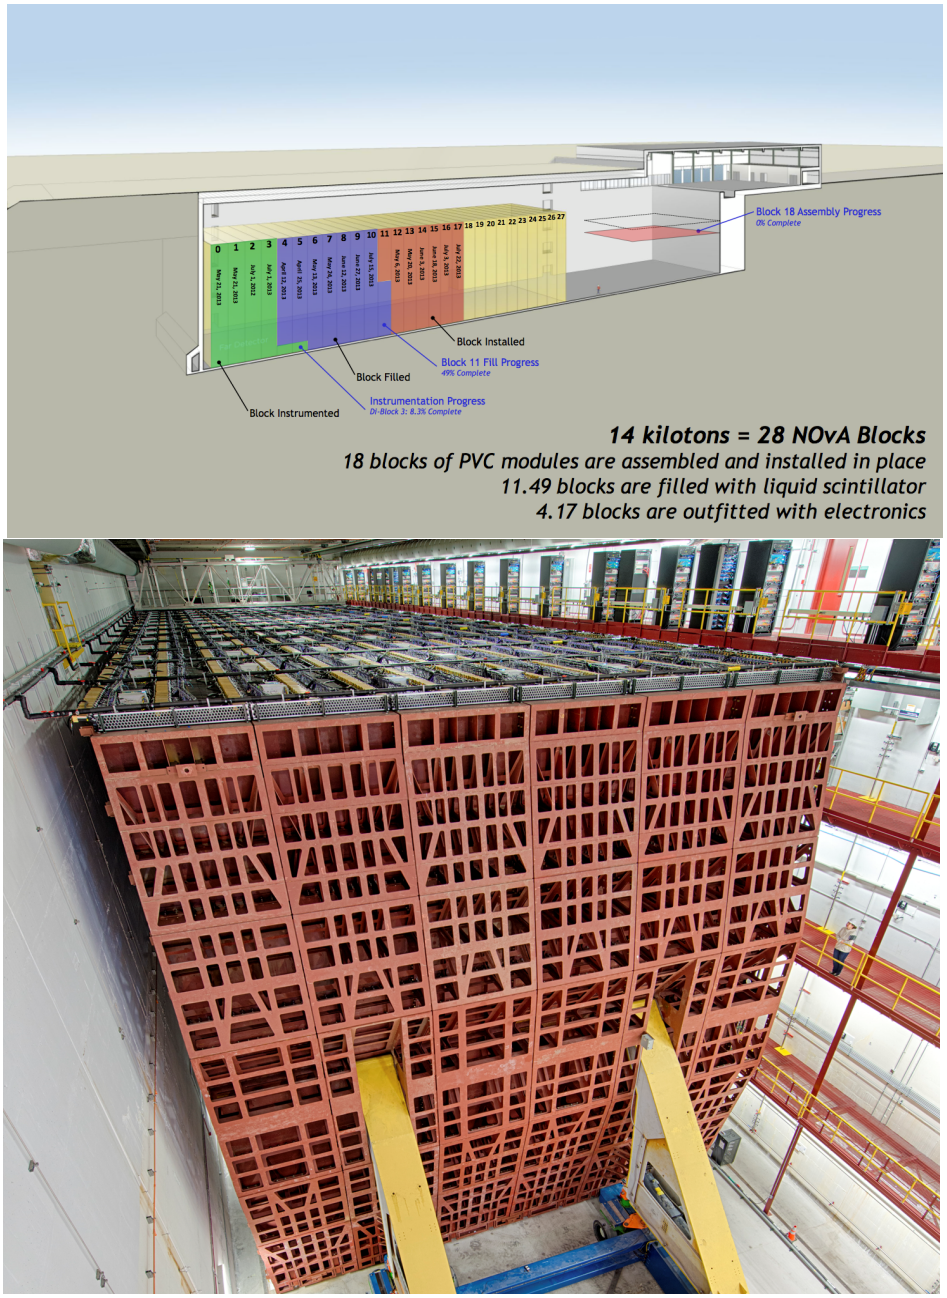

Figure 4.8: Top: A diagram of the construction and commissioning stages of the far detector. During construction, the pivoter platform (which now sits at the end of the detector on the right-hand side photograph) would serve as an assembly area for detector blocks and would then transport them to the end of the building. Once in place, modules were filled, equipped with electronics, and started taking data while the rest of the commissioning continued. Bottom: A photograph of the NOvA far detector, a $15.8 \times 15.8 \times 62$ meter structure of planes which are functionally equivalent to the near detector planes. The NOvA far detector is perhaps the largest standing plastic structure ever built.

NOvA data-taking began in the summer of 2014, overlapping with installation and commissioning of the far detector. We took advantage of the segmented structure of the readout to begin taking data in advance of full completion of the detector construction. Once the initial four diblocks had been commissioned, the completed region of the detector would take neutrino data, as the rest of the diblocks were being moved into place in the detector hall, filled, and fitted with readout electronics (see Fig.4.8). In future chapters, numbers for POT are reported as "POT equivalent" which implies the conversion to equivalent POT for a complete 14 kton active volume, scaling for periods of data-taking with the partially completed detector.

4.1.4 Signature Process

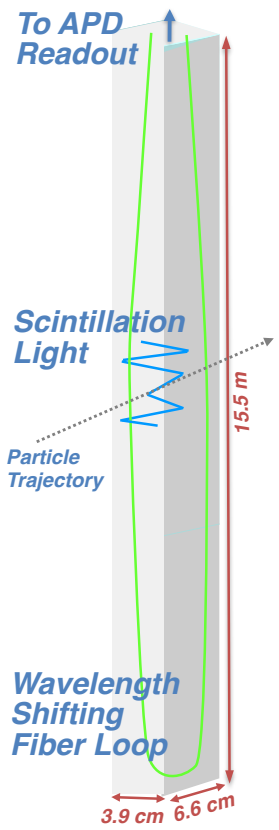

Figure 4.9: NOvA detector element.

Detection of charged particles is achieved through the scintillation light they produce as they travel through the detectors. This is made possible by the composition of the liquid scintillator with which the PVC modules are filled. This liquid is made of mostly mineral oil to which the following substances have been added to make scintillation possible while retaining the highest possible transparency of the material:

- **Pseduocumene (4.1%)** - Scintillant that emits light at wavelengths of 360–390 nm.
- **PPO and bis-MSB (9200 ppm)** - Additives that shift the wavelength to 450 nm.
- **Stadis-425 (0.3 ppm)** - Anti-static agent that prevents charge build-up.

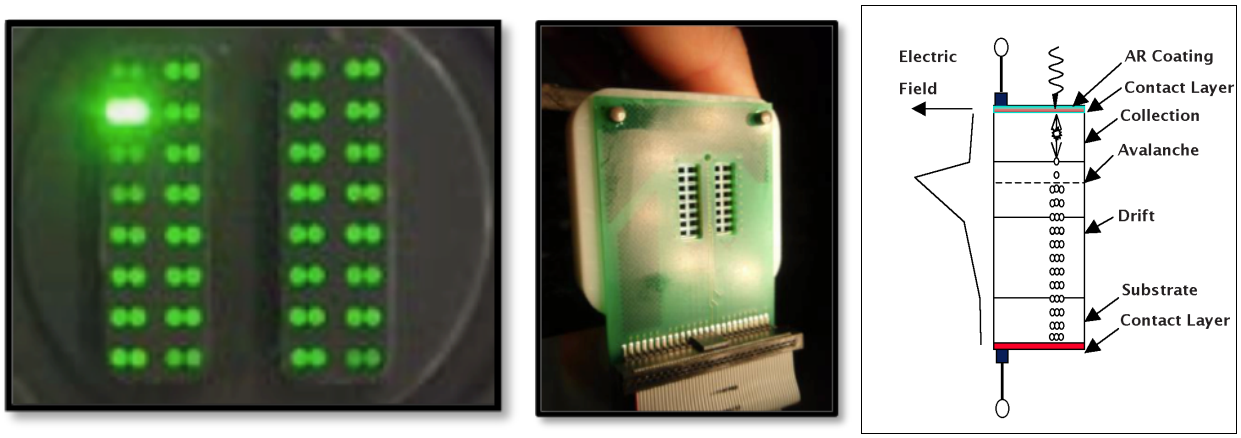

Figure 4.10: Left: The endcap of modules (planes of 32 channels like Fig.4.9) which bundles the wave-shifting fibers and couples to an Avalanche Photo Diode (APD). Middle: A custom NOvA APD. Right: APD schematic corresponding to one of the 64 wire end ports. The photons from the wire come from the top at the contact layer and through photoelectric effect they make photo-electrons. These get drifted by the existing field in the APD that causes the avalanche through the material as they travel to the bottom contact layer.

Before being filled with liquid scintillator, each PVC cell was equipped with a loop of wavelength shifting fiber. This fiber carries the scintillation light out to the readout electronics, while also shifting its wavelength from the blue 400–450 nm to the green 490–550 nm which the readout electronics expect as input. The process from scintillation to detection is shown in Fig.4.9.

The green light from the fibers is detected by avalanche photodiodes (APDs), depicted in Fig.4.10. These were chosen in part for having a high quantum efficiency of 85% in this wavelength range as shown in Fig.4.11. These APDs absorb the light from the endpoints of 32 fibers, one from each cell in the detector, and output a current. Operationally this occurs when the absorbed photons cause electron emission through the photoelectric effect and the emitted electrons are drifted by the voltage applied to the APD, causing an electron avalanche.

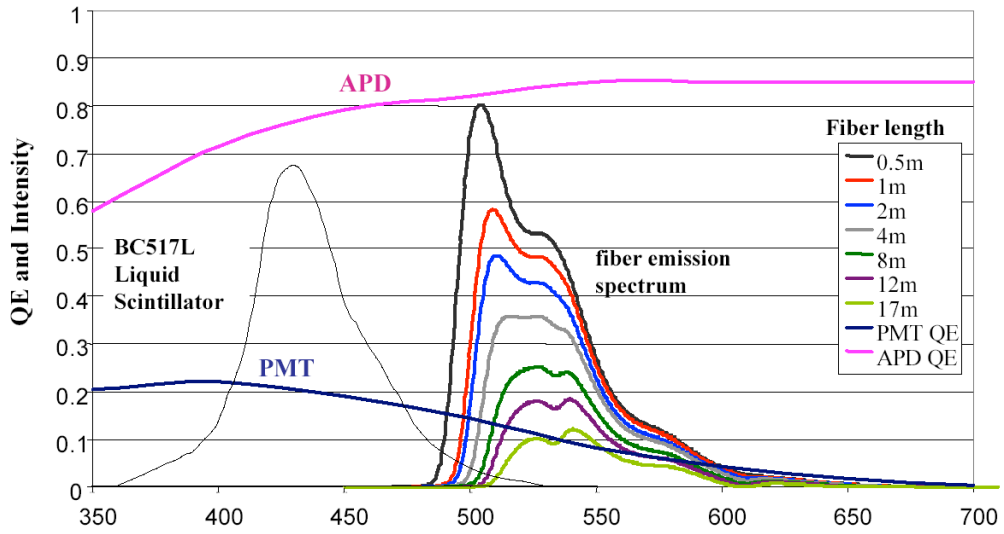

Figure 4.11: This figure shows the quantum efficiency of APDs and PMTs (photo-multiplier tubes)—a common alternative technology—as a function of incoming photon wavelength. It also shows the wavelength spectrum for different fiber lengths. For reference, the far detector planes are 15.8 meters long and 3.84 meters long for the near detector.

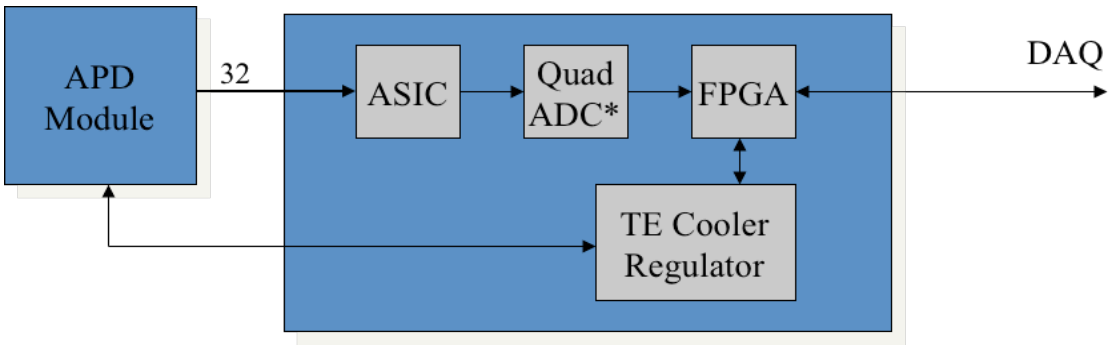

Figure 4.12: This diagram highlights the components of the front-end board (FEB) architecture. After the signal is received from any of the 32 channels of the avalanche photo-diode (APD), the ASIC (application specific integrated circuit) amplifies and shapes the signal in one of its 32 channels and feeds the resulting signal to the Quad ADC, which digitizes it at a rate of 2 million samples per second. The FPGA (field-programable gate array) then receives the signal in dual-correlated sampling mode and extracts a pulse height as well as timing information by looking for signals above a threshold set for each particular channel.

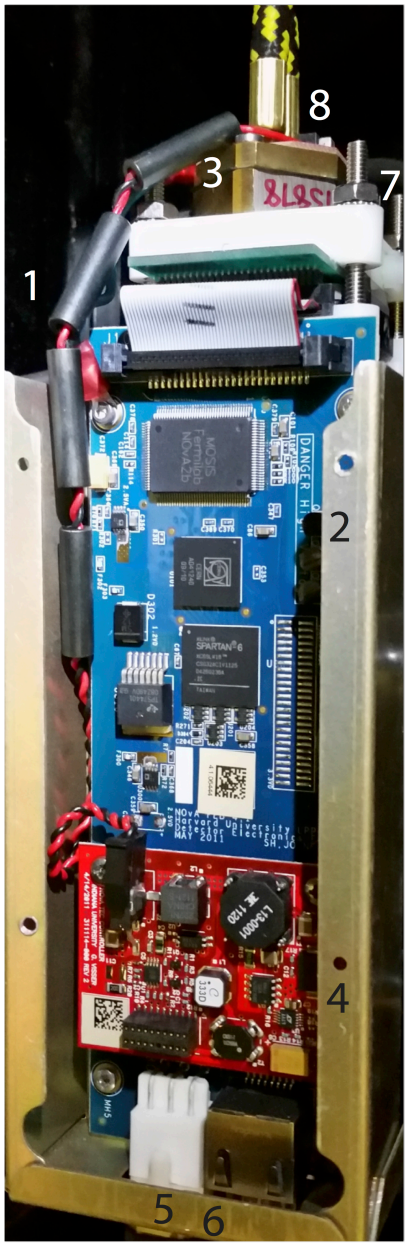

Figure 4.13: Electronics box. 1. APD connected through a ribbon cable to 2. FEB 3. Thermo-electroc cooler (TEC). 4. TEC controller (TECC). 5. Network cable. 6. Data cable to DCM. 7. Dry gas ports and hoses. 8. Cooling water ports and hoses.

This signal is then read, shaped, and digitized by a Front-End Board (FEB), whose functionality is further described in Fig.4.12. Finally, the digital signal is sent as a data packet to a data concentrator module (DCM), which aggregates data from multiple FEBs at a time.

The NOvA detectors operate with a total of 10,752 electronics boxes for the FD (and 329 for the ND) like the one depicted in Fig.4.13, which read-out 32 detector channels each. Each box contains an avalanche photodiode, a front end board—which processes the signal—as well as a thermo-electric cooler and controller for the water cooling system. The connections to water and dry gas are also located in this box, as are the connections to power supplies and a network connection to the data concentrator module.Their location on the FD is shown in Fig.4.14. These units operate at specific temperature and voltage conditions and the suite of on-line monitoring systems described in Sec. 4.2 ensure that the behavior of all these components is optimal for data taking.

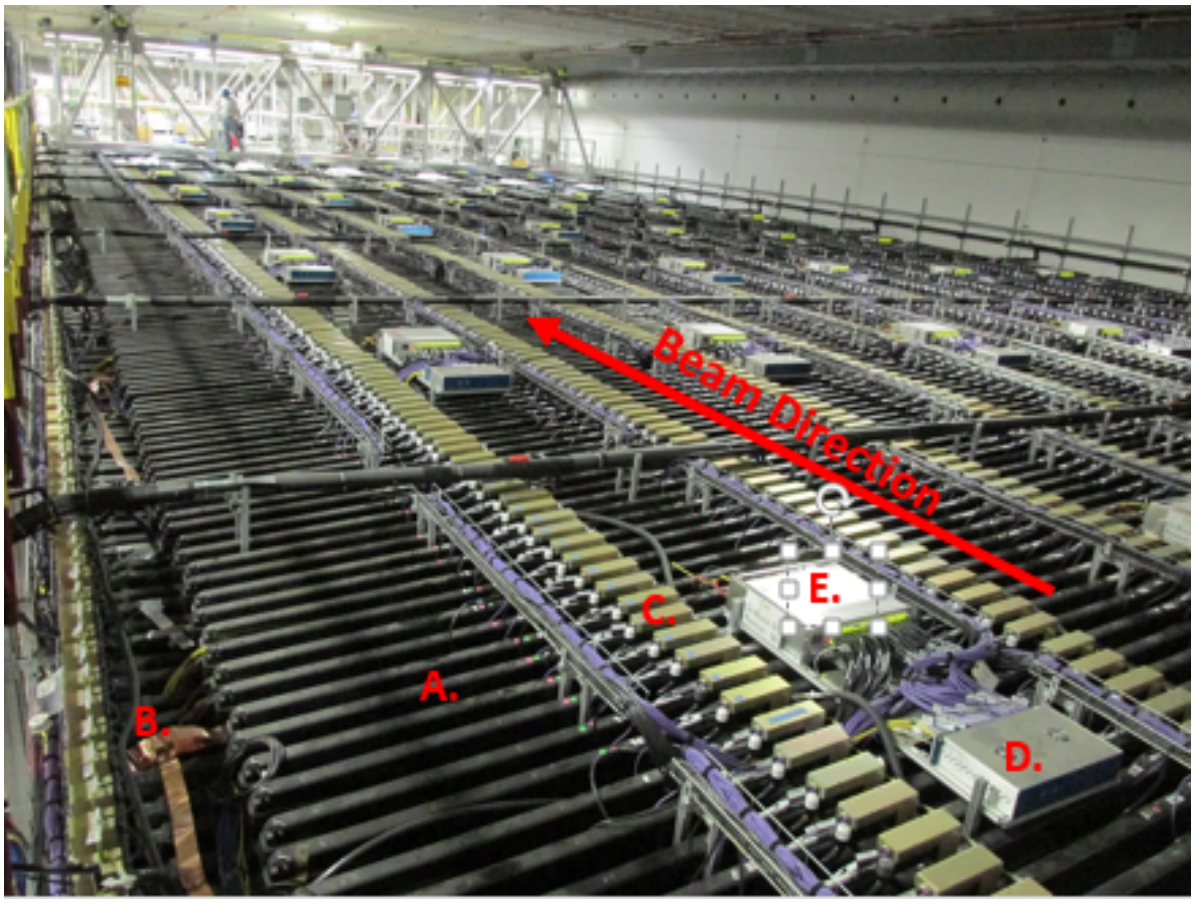

Figure 4.14: View of the top of the far detector. A. Manifold enclosing the fiber from each cell. B. Grounding strip. C. Electronics box (see figure 4.13) containing an APD, an FEB, and cooling system boards. D. Data Concentrator Module (DCM) E. High voltage power supply which delivers 425 V to APDs and 24 V to the DCMs.

4.1.5 DAQ and Timing

The electronics described above allow for a readout system with 100% live time for data-taking from the 344,064 (20,192) channels in the far (near) detector. The data acquisition (DAQ) systems are responsible for the delicate next steps. These include aggregating all data from the 32 FEBs per module, triggering—which is the decision process to either store or discard data, and writing out information to disk.

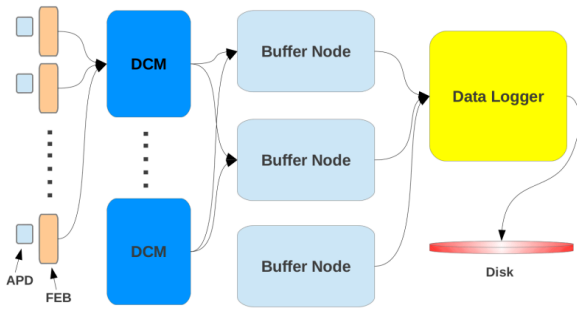

Figure 4.15: Data flow from the readout electronics to data stored on disk.

Each DCM receives signals from 2,048 channels, which it aggregates and feeds into a circular buffer. This buffer is comprised of a farm of machines that store all data for up to 16 minutes before it is discarded. During that time, the DAQ awaits triggers to either store or ignore data from the buffer. Fig.4.15 shows the elements of the readout and DAQ systems, as well as the progression of signals through the buffer and into files on disk in those cases when a trigger is received.

There are three types of triggers for NOvA data which are in use within the DAQ: Clock triggers, signal triggers, and data-driven triggers. *Clock triggers* are those for which there is a specific time interval at which data is stored. *Signal triggers* are those for which the system receives an external signal upon which data is stored for a period of time which is pre-defined for the trigger in question. *Data-driven triggers (DDTs)* are driven by specific conditions in the data which, if encountered, will prompt the DAQ to store the data from the time period during which such behavior was observed. All triggers on NOvA cause data for given time intervals in multiples of 50 μsec to be stored. The following is a partial list triggers that are currently active in the DAQ.

- **Cosmic Data Trigger:** Timed to store readouts at 10 Hz in the far detector. These are 550 μsec long readouts. These data are used as a minimum bias sample for detector calibration as well as to estimate the cosmic backgrounds in the analyses (Fig.4.19).
- **SNEWS Trigger:** Takes long readout (on the order of seconds) upon receiving an alert from the SuperNova Early Warning System [58], a world-wide network of neutrino detectors which distributes an alert upon detection of potential galactic supernovae signals (Fig.4.17).
- **Supernova Trigger:** Data-driven trigger that performs basic reconstruction on the data in real time to estimate a number of supernova-like candidates. It triggers a readout on the order of seconds when the number of candidates exceeds some pre-set threshold [59]. A simulation of supernova events is shown in Fig.4.18.
- **NuMI Data Trigger:** Stores a readout upon receiving a signal (in this case a GPS timestamp from the accelerator). Readouts are 550 μsec long. We expect our neutrino signal to be contained in these data, in addition to the characteristic cosmic signal in the far detector. The 550 μsec readout is centered on the 10 μsec beam spill time window, given by the accelerator complex as detailed in Sec. 4.1.1 (Fig.4.19).
- **DDEnergy Trigger:** DDT which stores readouts upon encountering 50 μsec of data with total deposited charge in the detector exceeding some threshold. Readouts are 50 μsec long or consecutive if the condition is continually filled (Fig.4.20).
- **DDActivity:** DDT which stores readouts upon any activity being seen in the near detector. These are 550 μsec long readouts. These readouts are used to calibrate the near detector as well as to evaluate running conditions independent of beam activity.

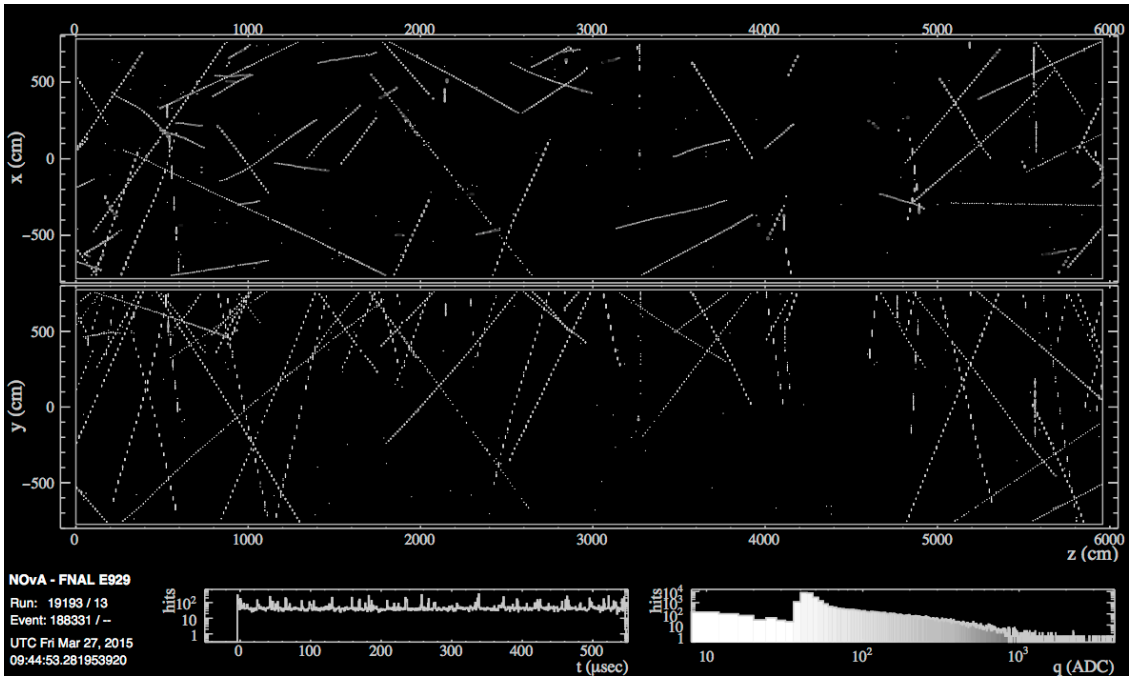

Figure 4.16: Readout for the cosmic pulser trigger.

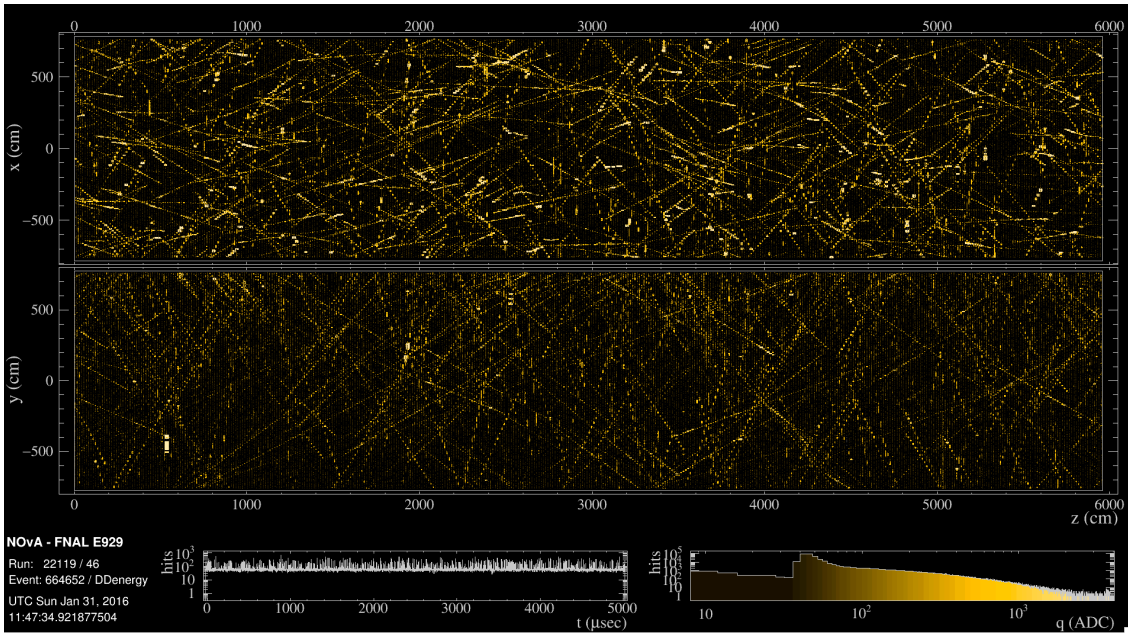

Figure 4.17: Readout for the SNEWS trigger.

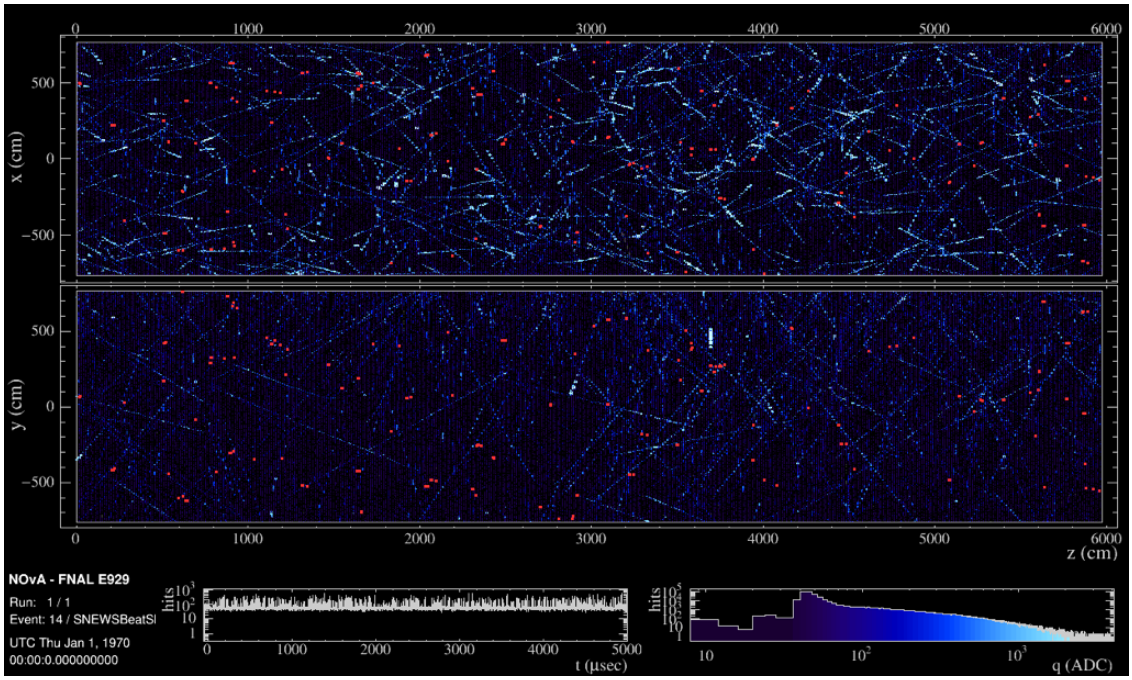

Figure 4.18: Readout for a simulated 10 kpc supernova over a period of 5 ms. Red: Simulated interactions from supernova neutrinos. Blue: minimum-bias overlay.

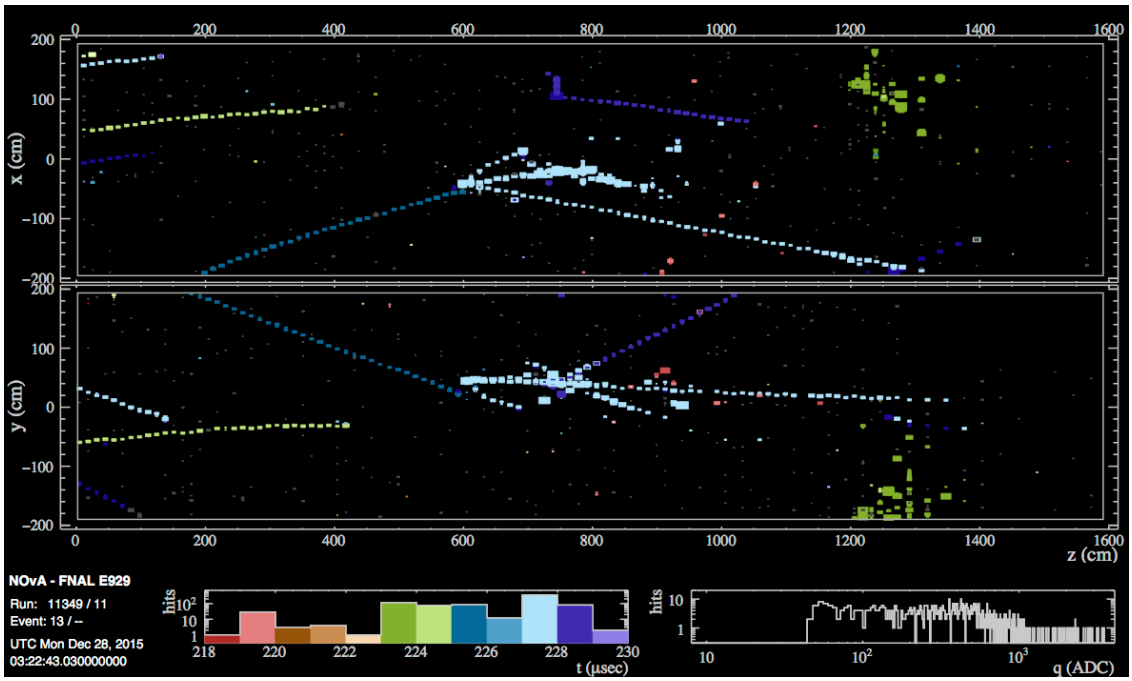

Figure 4.19: Readout for the NuMI trigger (550 μsec).

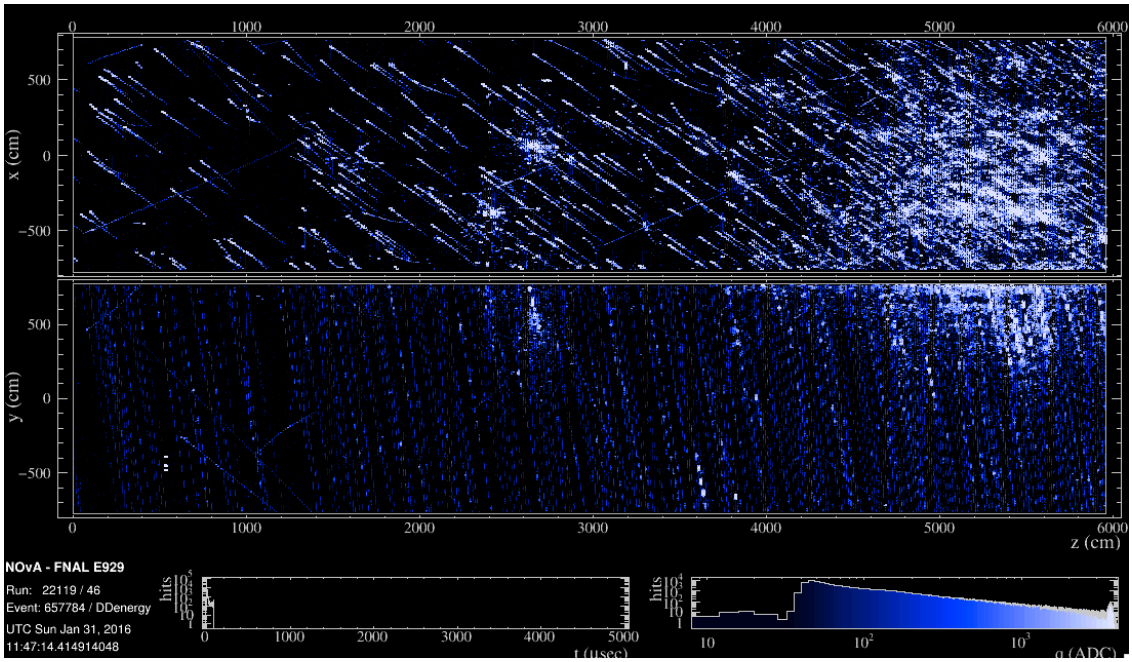

Figure 4.20: Readout for the DDEnergy trigger. The pattern in this event is a large energy deposition from air showers incoming from the top corner of the detector.

4.1.6 Timing

The timing system is a crucial component of the data acquisition system. It ensures accurate time-stamping of signals from the accelerator and the timing synchronization of detector components. Every channel in the detector is synchronized to the same internal time. This is achieved through a system controlled by a series of timing distribution units (TDUs), which ensure that all of the DCMs are simultaneously reporting consistent time. TDUs also maintain the global detector time in sync with an external "wall time" from GPS satellites, which is consistent with both the accelerator time and the second detector's time. Another component of the system is called the spill server.

The accelerator will send a signal to the NOvA data acquisition system in order to prompt the storage of data at the time at which neutrinos are expected as described in Sec. 4.1.5. The spill server is responsible for time-stamping signals from the accelerator and delivering that information to a

global trigger system on both detectors, which triggers the data readout. The entire timing system ensures that readouts are taken exactly around the beam window, and that the time reported by all channels is consistent, which is necessary for reconstruction downstream.

4.2 Experiment Operations

The ongoing task of operating an experiment is rich in opportunities to reexamine one's assumptions about the expected behavior of the experimental array and the data-taking procedures. Continuous evaluation of the expected conditions is required, as well as assessment of new, unexpected effects. Collecting data is an active task that directly impacts the uncertainties in our measurement. Specifically, statistical uncertainties are impacted by the amount of quality data collected, which requires reliable understanding of the performance of the apparatus and the ability to quantify it, as well as ongoing assessments of optimal operating conditions. For low statistics measurements, such as detection of neutrino interactions, the statistical uncertainties dominate over systematics in the sensitivity of the results, which makes optimal data taking efficiency a high priority. The task of ensuring optimal data quality goes beyond the design and commissioning of the experiment and continues onto new challenges—expected and otherwise—throughout its running life. Variations in hardware performance, understood fluctuations, unexpected damage, cascading system failures, and even weather conditions are examples of possible causes for detector downtime which must be understood and anticipated when possible, and addressed when necessary.

On NOvA, a continuous effort exists to understand these effects as well as identify new ones promptly. This effort has been carried out by experts over the years and is continuously under improvement and further development. Monitoring detector operations starts with immediate inspection of live data and running conditions, and it continues onto several failure identification systems from detection electronics and downstream. This section details the broad operation procedures and monitoring of the NOvA systems. Operation and monitoring of the DAQ and support systems is discussed in 4.2.1. Data monitoring in real-time and close-to-real-time is detailed in 4.2.2 and 4.2.3.

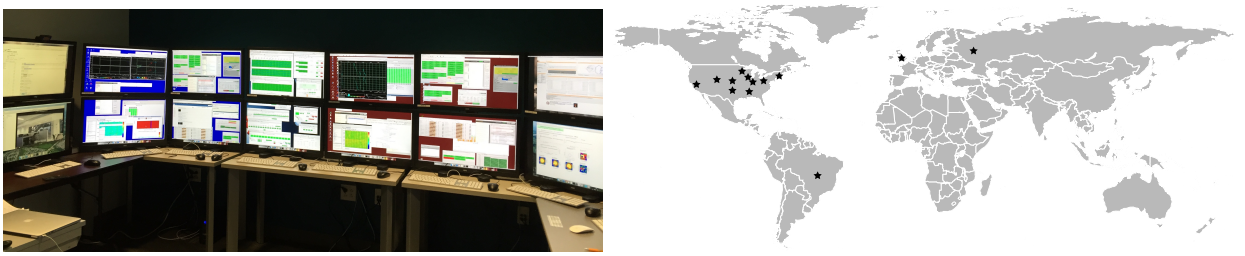

Figure 4.21: Left: Photograph of a NOvA Remote Operations Center (ROC). Controls for detector operations and monitoring are available to collaborators from 20 institutions, from 4 countries, as depicted on the maps in the middle and right panels.

Finally, the hardware monitoring and maintenance prioritization system is described in 4.2.4.

The ability to monitor data as it is being recorded gives an immediate handle on data quality, running conditions and detector performance, which allows us to address and correct issues as they arise. It is difficult to quantify the exact statistical impact of all the tools implemented for monitoring of detector operations. However, it is clear that this effort plays a large role in the detector up-time of NOvA and in the fraction of optimal quality data recorded, which at the time of writing are over 98% and 97%, respectively.

4.2.1 Monitoring DAQ and Support Systems

As is common for most high energy physics experiments, the multiple DAQ and support systems for NOvA are under constant monitoring by one or multiple collaborators on rotating shifts which cover 24 hours a day, every day. This is done at one of the multiple remote operations centers set up by experts in the twenty locations shown in Fig.4.21. While the NuMI beam is active, this task is to prioritize detector uptime while documenting running conditions and consulting experts on how to address unexpected incidents. While the beam is not operating, the shifter's tasks also include aiding experts in documenting the effects of changes and improvements to the systems and of general maintenance.

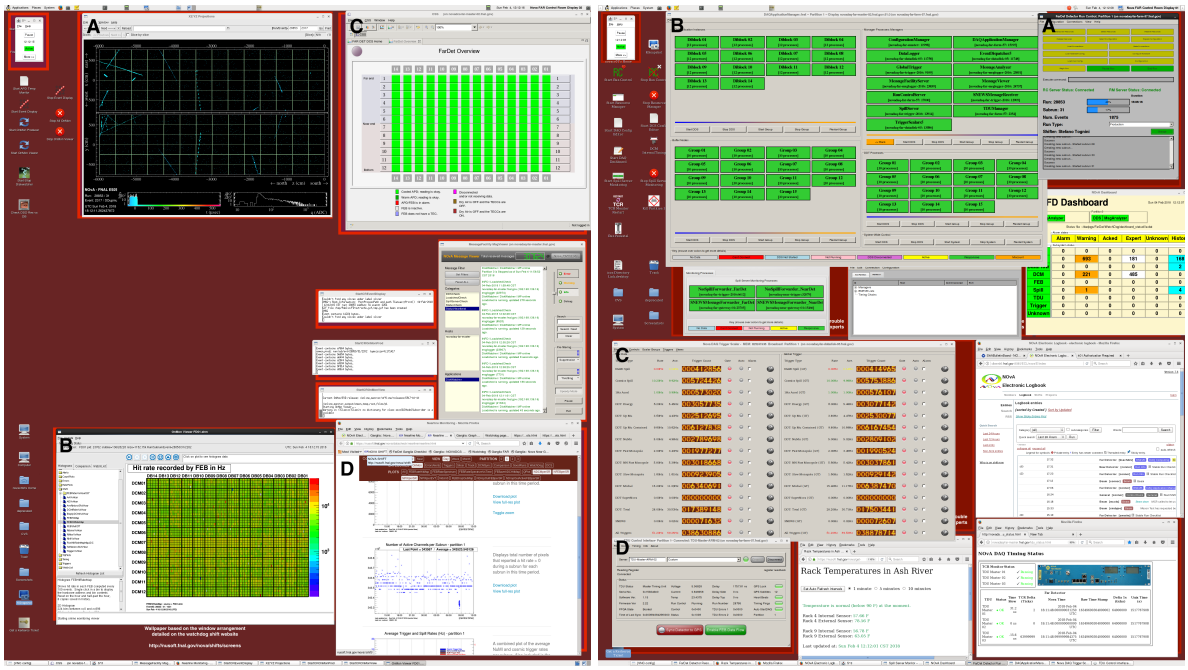

Figure 4.22: Operations screens for data acquisition for FD (ND screens are equivalent). Left: CR01 Control. A. Run control, starts and stops data taking. B. DAQ Application manager: Runs and monitors connection to the detector electronics. C. Trigger Scalars: Counts of triggers issued by the DAQ and events recorded. D. Timing system monitors for the TDUs and the spill server. Right: CR02 monitoring. A. Event display B. OnMon viewer. C. APD/FEB Temperature monitor. D. Nearline monitor.

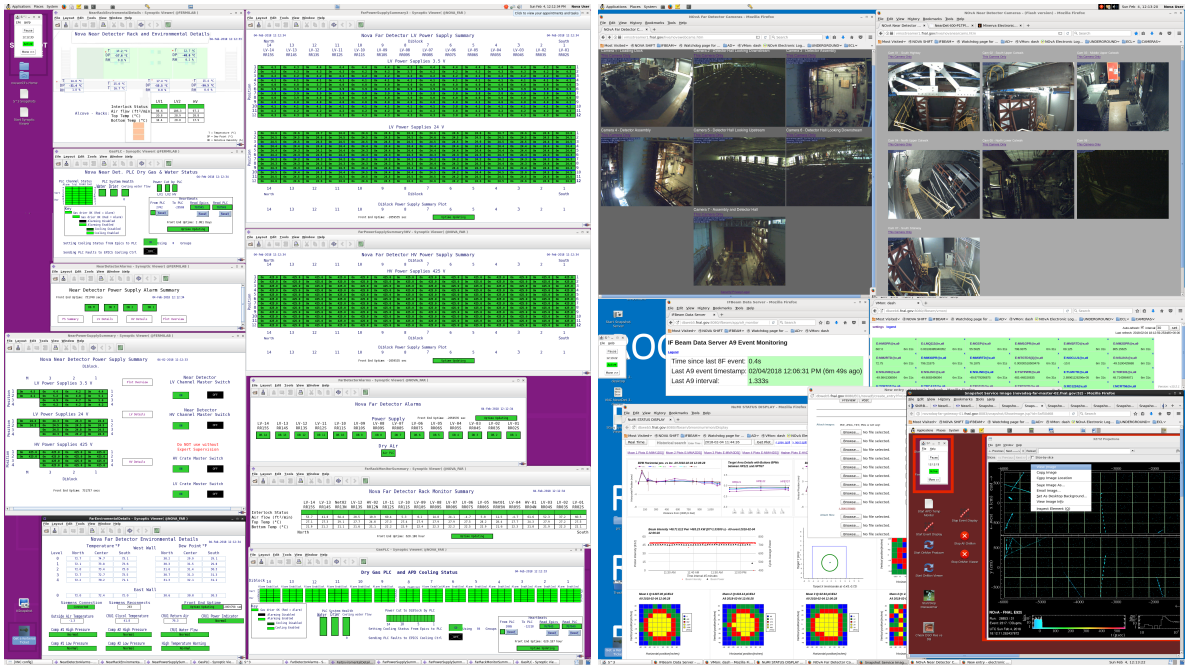

Figure 4.23: Operations screens for external monitors: Left: Synoptic displays with environmental details from both detectors and voltage settings for electronics. Right: Beam monitoring and detector cameras.

The standard tools available for these tasks, shown in Fig.4.22 and Fig.4.23 include monitoring for hardware conditions, data acquisition systems, and data quality. Hardware operating conditions are monitored by the systems described in the following subsections.

4.2.2 Online Monitoring

The Online Monitoring system provides the shifter with metrics calculated from the data as it is being read out by the DAQ. The system computes and constantly updates a set of metrics by sampling different triggered readouts while they are still available in the buffer depicted in Fig.4.24. Given the rate at which different readouts are taken this live system primarily samples two types of triggers, one providing a regular stream of data (the cosmic trigger) and another providing beam signaled data (the NuMI trigger). The raw data consists of ADC counts with time and location information. This raw data provides us the "what", "when", and "where" of detector activity. The charge deposited, given in ADC counts can be monitored for hardware thresholds and readout issues, the time of the hits can be converted to a rate whose normal values are known, and the location of the hits allows for correlation of the time synchronization over all the electronics.

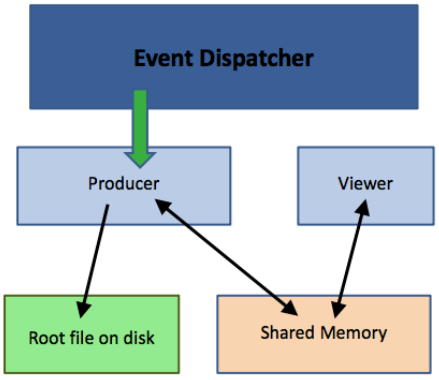

Figure 4.24: Data flow from the readout electronics to data stored on disk.

The Online Monitoring system has a backend that produces the metrics and a frontend GUI that displays them. The backend, called OnMon Producer, samples live data by subscribing to a shared memory segment from the event dispatcher (a component of the DAQ) for a set of chosen triggers and computes metrics in the following categories:

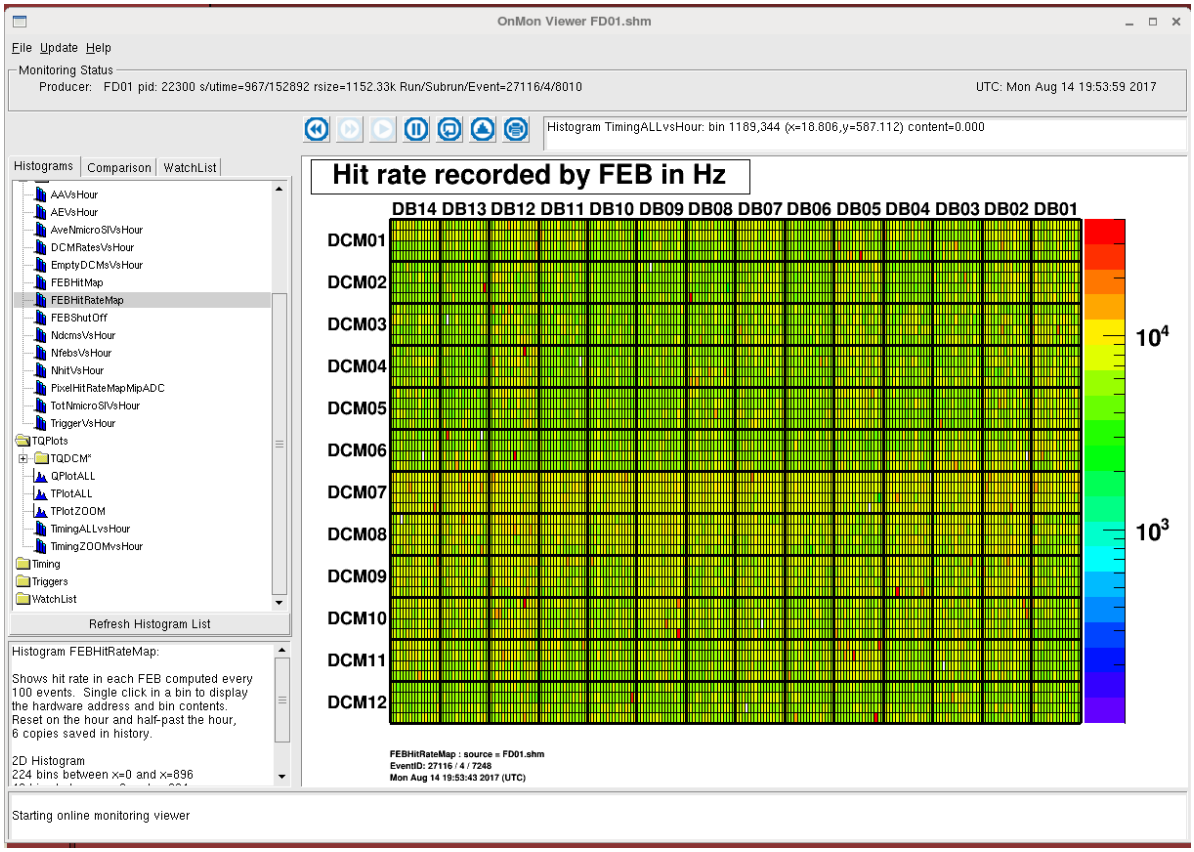

Figure 4.25: OnMon Viewer. A: Plot region. B. Information Box. C: Status Bar. D: Print Menu. E: Plot directories. Plot: The hit rate seen by every APD in the Far Detector. The normal rate are between 10 Hz and 5 kHz.

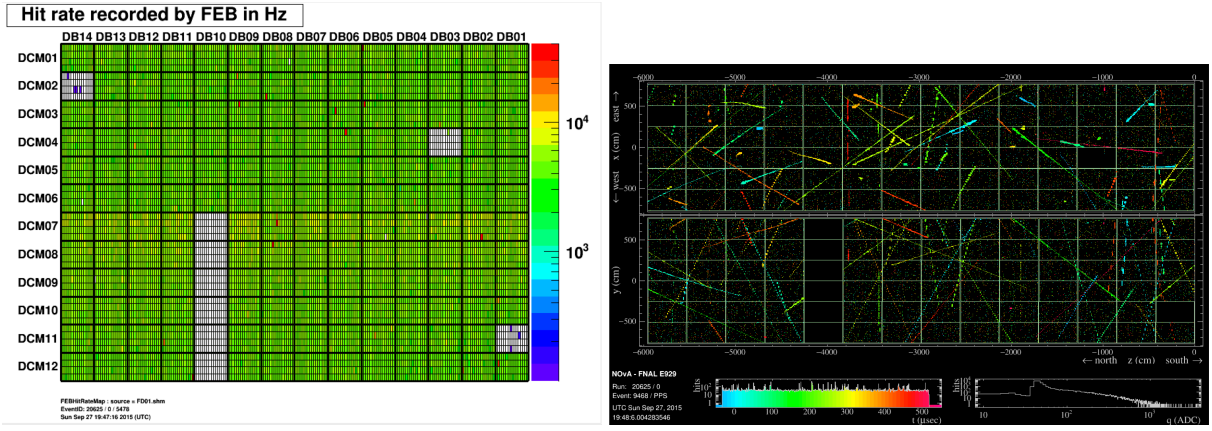

Figure 4.26: DCM synchronization issue seen on the hit rate map (left) and the event display (right)

Trigger metrics: These provide confirmation that triggers are being issued at expected rates.

Timing metrics: These allow for monitoring of the proper synchronization of the electronics as well as time structure of the neutrino beam spills.

Hardware activity: These metrics serve as indicators of hardware performing as expected.

The front end of the system, called the OnMon Viewer is a ROOT [60] GUI which incorporates formatting, plot information and comparison capabilities for further analysis. While the OnMon Producer aggregates live data to the summary metrics, the OnMon Viewer displays them on 30-minute summary plots or time series plots. The OnMon Viewer and the OnMon Producer communicate via a shared memory segment where the information is being modified by the Producer and accessed by the Viewer in real time. The OnMon Viewer is one of the main systems running in the control room monitors for the detectors as shown in Fig.4.22.

Accurate timing is an important part of optimal data-taking, and is important to monitor. The correct recording of triggers, the synchronized response of the hardware and the time of the NuMI beam spills can all be confirmed on these plots. For example, Fig.4.26 presents examples of useful online monitoring metrics to diagnose hardware issues and Fig.4.28 shows beam and timing system monitoring available within the same system.

An example of a check that can impact our signal statistics directly is the TDU offset metric, shown in Fig.4.28. While the NOvA NuMI triggers are in-time with the beam spills, as discussed in Sec. 4.1.5, it is possible for the timing system to be in a state of error, which causes an additional delay with respect to this time. This failure mode would be difficult to identify if present in one detector only, such that all metrics concerning triggers, stored events, and timing would appear normal.

However, it is possible to identify this issue from the delay in the 1 Hz trigger (a trigger signal received from the accelerator system and independent of the NOvA timing). The 1 Hz trigger records

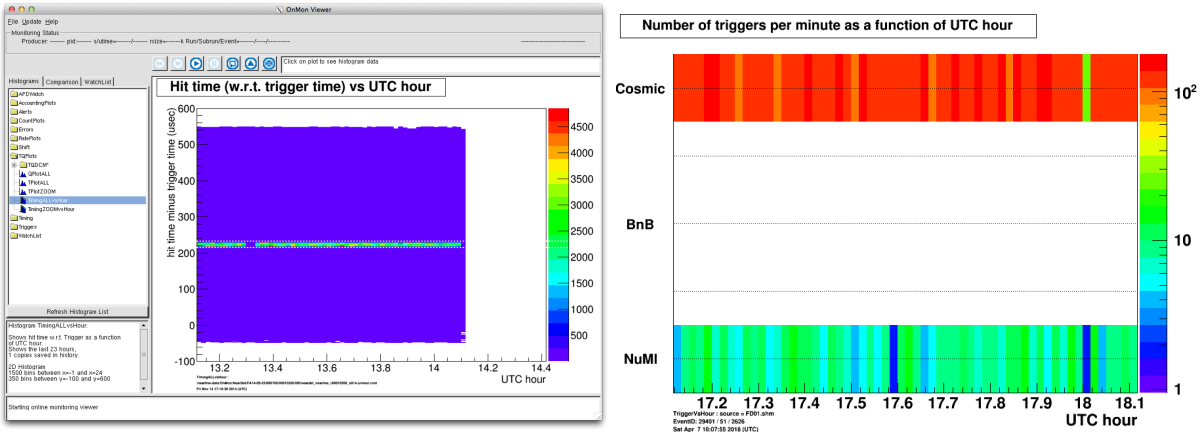

Figure 4.27: Left: Time structure of the NuMI beam seen in NOvA readouts which are 550 μsec long, centered around the 10 μsec beam spill. Right: The approximate number of triggers recorded as a fuction of UTC time.

an event on the tick every second. There is an expected 21 μsec offset from the time it takes the signal to be acknowledged by the DAQ. The failure is an additional 64 μsec delay of all signal triggers. The metric shown in Fig.4.28 shows the TDU offset calculated as:

$$\text{ND Trigger delay } = T_s * 64 \times 10^6 \tag{4.4}$$

and

$$\text{FD Trigger delay} = T_s * 64 \times 10^6 - ToF - \Delta_{MI}, \tag{4.5}$$

where T_s is the trigger start time from the DAQ, 64×10^6 is the number of DAQ clock ticks per second, ToF is the neutrino time of flight and Δ_{MI} is the time shift with respect to the Main Injector.

The trigger delay metric was incorporated to look for data which had been taken with a 64 μsec offset, a failure which occurred for some time between 2014 and 2015, leaving NuMI data within the recorded events, but away from the region of interest where the beam spill is expected. A total of 2.74×10^{20} POT was analyzed with this offset and two of the selected ν_e events were recovered by adding the second time window.

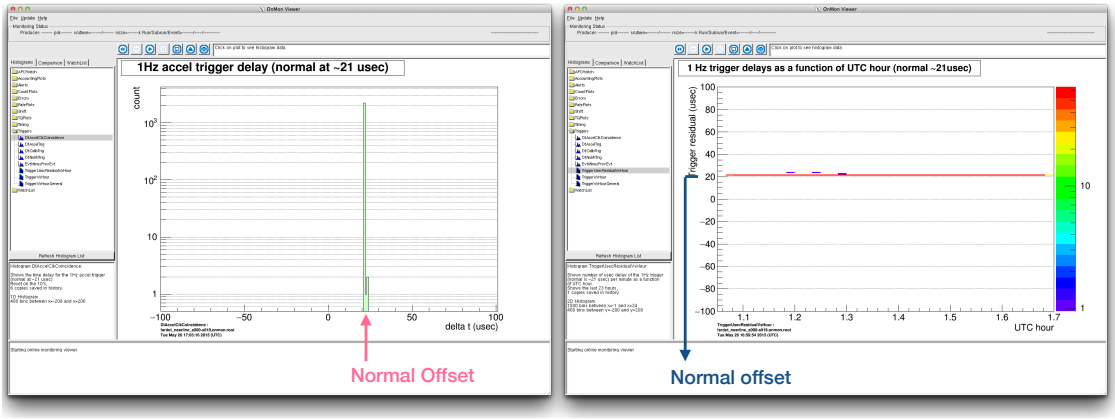

Figure 4.28: The signal travel time between the accelerator and the NOVA Timing Distribution Units detailed in Section 4.1.6. There is an expected 21 μsec offset from this time, which the NOvA DAQ corrects for at the trigger level.

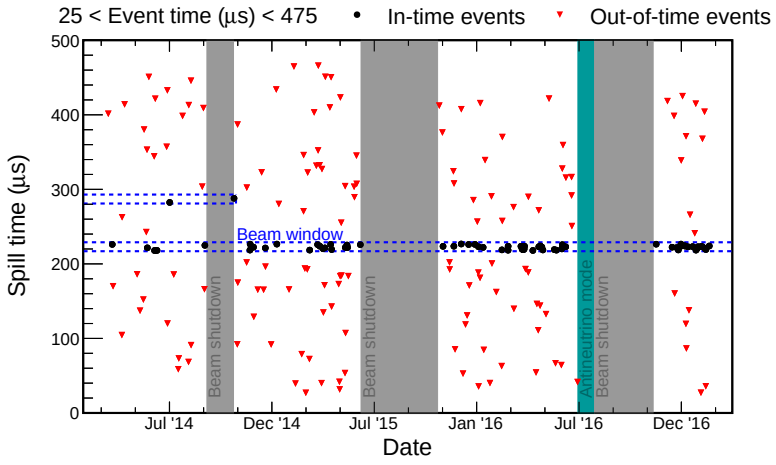

Figure 4.29: Event times for ν_e selected events. The time windows used for the analysis are contoured by blue dashed lines. Two events were found in the second time window.

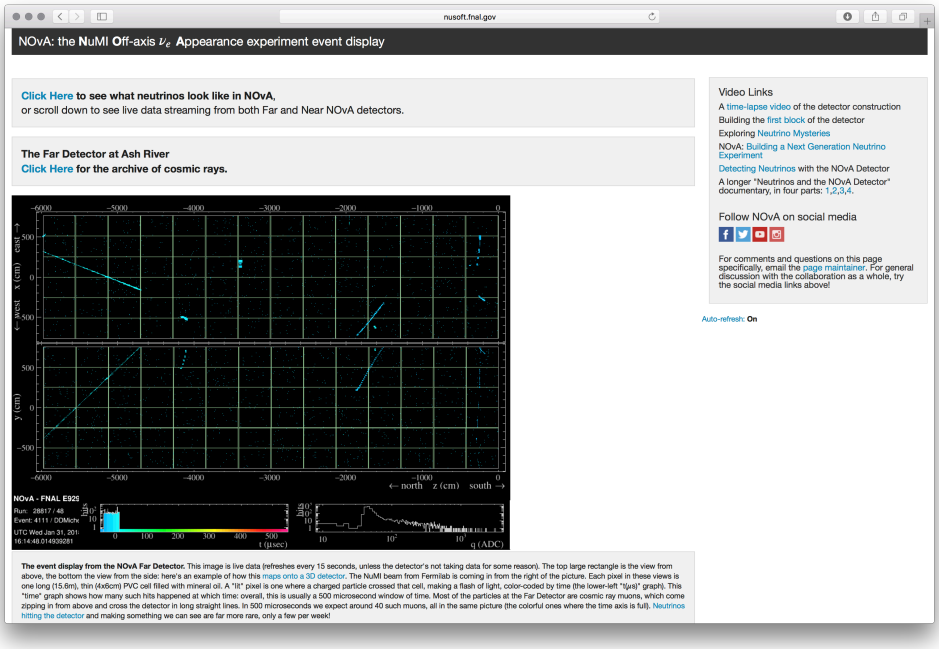

Figure 4.30: Live event display public webpage.

The live Event Display is another component of the online monitoring tools available in the control room as well as online, as shown in Fig.4.30. Like the OnMon producer, the live Event Display subscribes to the event dispatcher and displays the raw data in a map of the detector. Fig.4.26 shows an example of the live event display outlining the locations of individual DCMs, mapped in the OnMon viewer, which correspond to geographical locations in the detector. An example condition where hardware has fallen out of sync is shown in Fig.4.26, where issues in data from DCMs can be seen in coincidence on both the OnMon Viewer and the Live Event Display.

4.2.3 Nearline Monitoring

While the online monitoring provides a fast look at the raw detector data, a second "nearline" system provides monitoring of raw and reconstructed information over long periods.

The *Nearline* monitoring system if a set of algorithms that computes relevant summaries from

multiple operational metrics. These summaries are made for periods of 24 hours, seven days and 30 days worth of data, as shown in 4.31. In addition to producing the live summary metrics, the OnMon Producer also outputs summary files which are one of the inputs to the Nearline monitoring system. The Nearline output consists of a set of plots which either aggregate metrics or show their evolution over time. These metrics include the following categories:

OnMon metrics: Hit-rates, active/reporting hardware status, issued triggers, trigger delay, among others. Computed from the OnMon summary files. These give a handle on optimal hardware performance and hardware issues which evolve over time, as well as the correct recording of triggers by the DAQ.

Detector and Environmental metrics: From sensors in the FD and ND halls, these are temperature and dew point summaries for different locations. These are important to monitor given the dependence to temperature and dew point of the noise seen in our electronics due to hardware operating temperatures or condensation. The environmental metrics also include running voltages at the power supply sources and the DCMs.

Data and OnMon files: Monitoring of data file transfers from the DAQ disks into permanent storage, file processing and OnMon summary file production.

Beam metrics: From a combination of the Fermilab's accelerator database and the NOvA triggers, we can monitor the recording of beam events as well as the optimal use of scheduled detector downtime in coincidence with beam downtime, as seen in Fig.4.32.

A few examples and use-cases of nearline metrics can be seen in Fig.4.32.

The Nearline backend monitors the creation of new raw data files by the DAQ and executes an offline version of the OnMon Producer, which outputs one summary file per data file. When a new file is closed, it re-computes all metrics for the 24 hour period of time starting with the newest summary

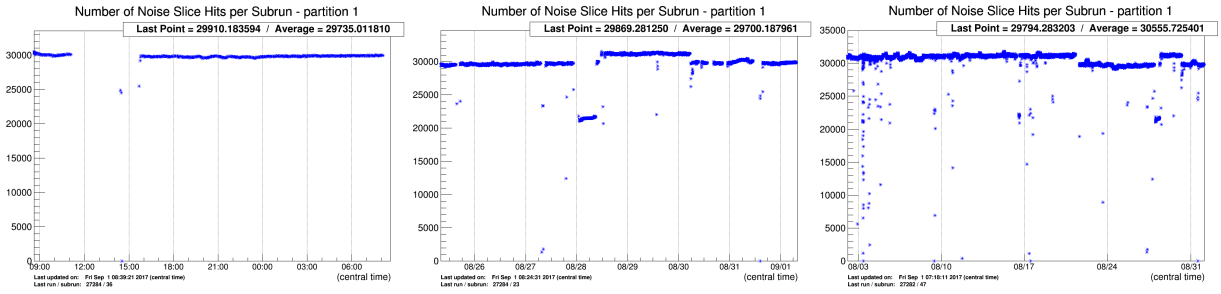

Figure 4.31: Evolution of noise hits over 24hrs, one week and 30 days.

file. The production of plots with Nearline metrics has a delay with respect to live data of approximately 20 minutes for the far detector and on the order of one hour for the near detector. This delay is driven by the frequency with which summary files are produced, which can only be computed for the finite time intervals of readout determined by the DAQ.

The Nearline front-end is a series of webpages which display a subset of plots selected by the user and holds all summary plots for both detectors calculated for the time intervals mentioned above. In addition to this main display, a summary of the most important metrics is available through a front page which shows a small subset of what are considered to be the most relevant plots for the shifters to actively monitor. Additionally, the system alerts users of known, time-sensitive error conditions in metrics as they are computed and displays alerts accordingly. The nearline front page even detects when the user has not interacted with the system for a period of time and issues a reminder to do so. Fig.4.33 shows an example of the nearline front page and Fig.4.34 shows the full content of the nearline webpage.

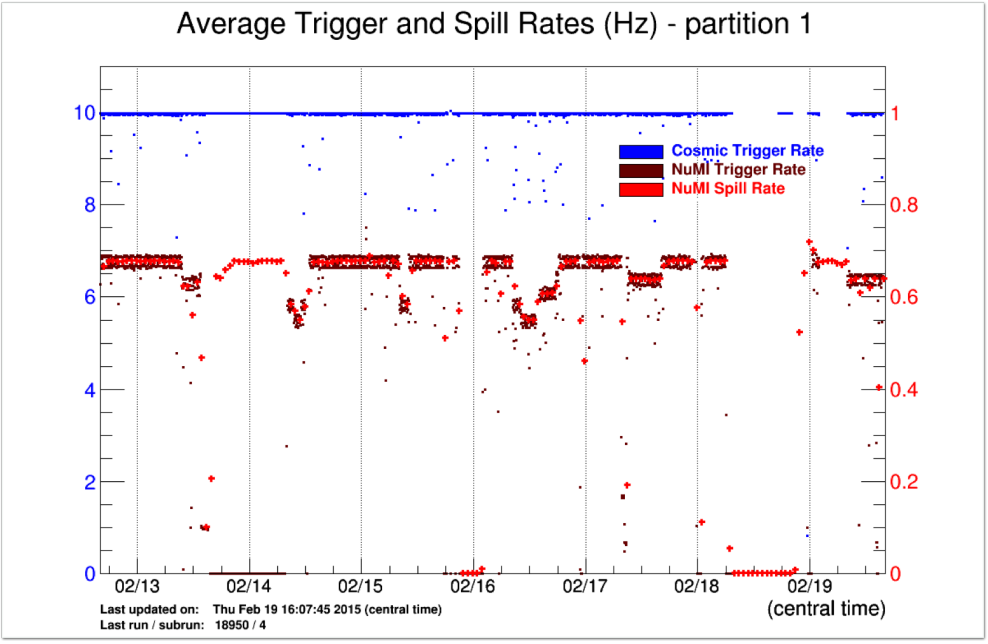

Figure 4.32: Nearline Trigger Metrics. The first two metrics are rates of collected readouts, calculated from readouts in data files. Blue: Pulser Cosmic trigger rate gives a measure of detector uptime. Maroon: Rate of collected NuMI trigger readouts. Red: Rate at which neutrino spills are delivered from the accelerator. This plot contains useful uptime information like how much of the NuMI beam spill available we record (when the maroon and red lines match), errors of triggering on the beam spill (when the detector is running but we fail to record beam spills marked A) and detector downtime taken for maintenance in coincidence with beam downtime marked B.

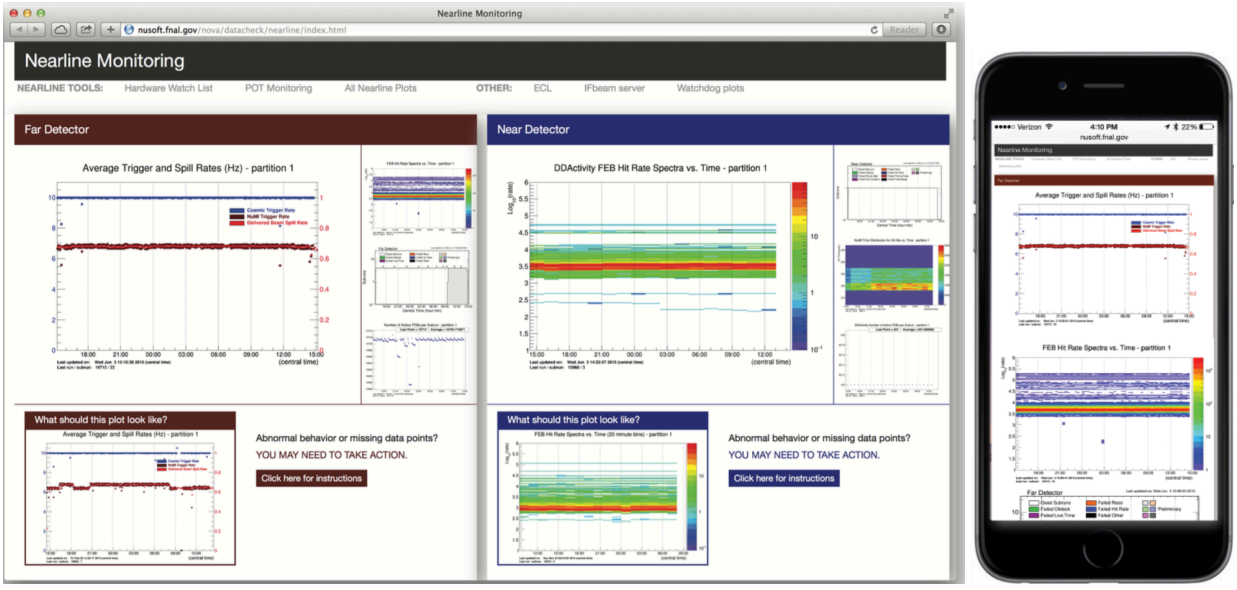

Figure 4.33: Nearline webpage accessible for expert remote monitoring of all systems.

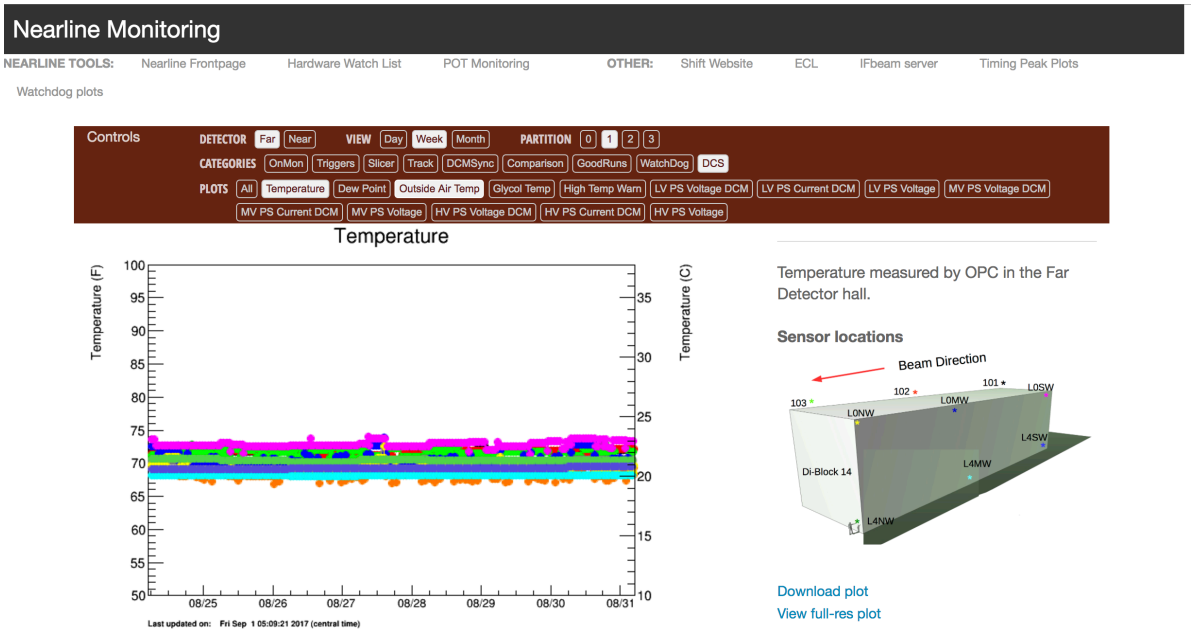

Figure 4.34: Nearline environmental metrics webpage. Left: The temperature from sensors inside the detector hall, markers on the right show the physical location of the sensors.

4.2.4 Hardware maintenance

Ensuring the continuing high quality of the data for the duration of the experiment requires prompt and effective maintenance to be done on the hardware by expert technicians. Routine testing for failures of the almost 11,000 FEBs, APDs, TECs, etc. would be impossible. However, using the online and nearline summaries we can identify problems when they occur and categorize them enabling a more efficient response from the experiment technical staff.

Element	Constraint	Normal Values
APD	Voltage	25 V
APD	Temperature	-15 C
APD	Hit Rate	$\sim$500 Hz
FEB	Temperature	35 C

Table 4.1: Hardware Operating Conditions.

In some cases, running conditions may deviate from the nominal values listed in Table 4.1, but do not represent a threat to the electronics. However, running within these limits is necessary for optimal data taking and in some cases known effects on the quality of the recorded data occur when falling outside of the normal values. For example, elevated operating temperature on APDs is correlated with noise and with higher saturation rates from large charge depositions. However, some operating conditions have been studied over time and adjusted to optimal values with positive effects on the data. For example, the voltage gain changes in the APDs which will be discussed in Sec. 5.3.1 as well as different algorithm settings within the FEB firmware.

The hardware watch system was designed to assess the stability of the detectors, identify features in the data which can correlate to specific hardware failures, and prioritize maintenance tasks optimizing for detector performance. This process begins by identifying the pieces of hardware that show

failures or unexpected behaviors and other features over long periods of time. Once identified, the maintenance needs are prioritized by severity and frequency of the issues. Such prioritization is done by using the data summary outputs from the online monitoring covered in Sec. 4.2.2 for each of the channels in the detectors over large periods of time. The most indicative metrics the hardware watch system employs are the average noise rates over week-long periods of time, minimum or maximum reported ADC and the thresholds set by the system. The system identifies issues which are defined as values outside the expected noise rate across the detector shown in Fig.4.36.

The hardware watch system was not part of the original monitoring design. The concept and development of this tool happened during detector commissioning after I employed a simpler first iteration to assess a manufacturing defect of the first set of electronics installed in the far detector. At the time, the first two diblocks of the detector had been completed and instrumented with the Hamamatsu APDs shown in Fig.4.10, with a custom paraffin coating to avoid condensation issues from the cooling system.

The coated APDs developed an incremental noise profile over time, as charge would accumulate on the coating. This issue was characterized and potential solutions were discarded through studies of the hit rate evolution over long periods of time. In the end, this study helped to determine that no operational modification would resolve the issue, after which the design of the APDs was changed and the more than 1500 pieces of hardware replaced.

Fig.4.35 shows the evolution of the noise incidents in the APDs, where incidents were defined as:

$$\text{Incident rate} = \frac{\text{Number of noise incidents}}{\text{Number of data files sampled}} \tag{4.6}$$

where incidents are defined as hit rates exceeding an allowed maximum of $Rate_{max} = 500$ Hz.

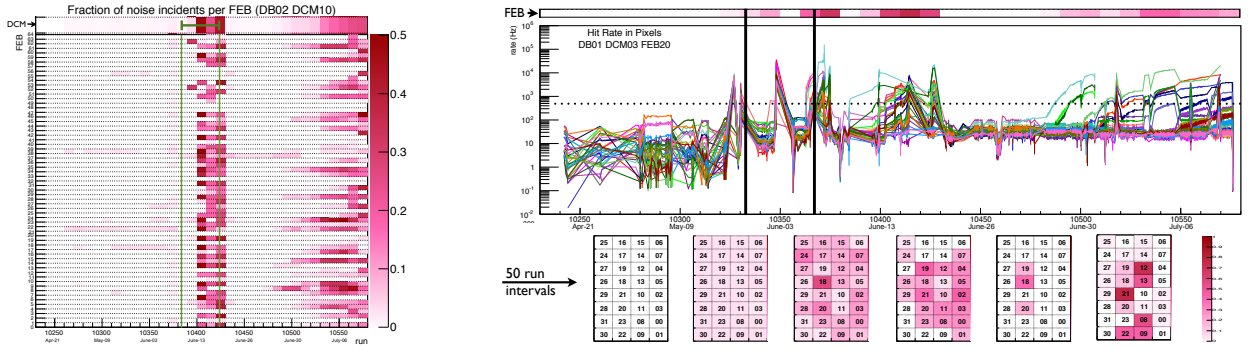

Figure 4.35: Hardware noise reduction diagnostics. Left: incident rate as defined in 4.6 is shown in color over a period of 3 months (on the x axis) for each FAB/APD set (on the y axis) on an example DCM. The issue rate for the DCM is displayed on the top bar. Right: an example noise profile for one FEB. Top: the hit-rate of each of the 32 channels read at that location. Bottom: the issue rate for each channel over six time-periods, displayed as ordered on the APD input face. Note that the x axis on both

This methodology inspired the design of the system I developed, called the Hardware Watch List, to monitor the hardware. It operates as follows:

Every week, an *issue rate* is calculated for every channel in the detector (i.e. one fiber output from Fig.4.9). Issue rates are reported as an average over the 32 channels readout in a single location corresponding to an electronics box like the one shown in Fig.4.13. The score is computed in the following way:

$$\text{Issue rate} = \frac{\text{Number of failure instances}}{\text{Number of data files sampled}} \tag{4.7}$$

where all data files from the current week are sampled, except for those corresponding to period of time with documented issues in running conditions.

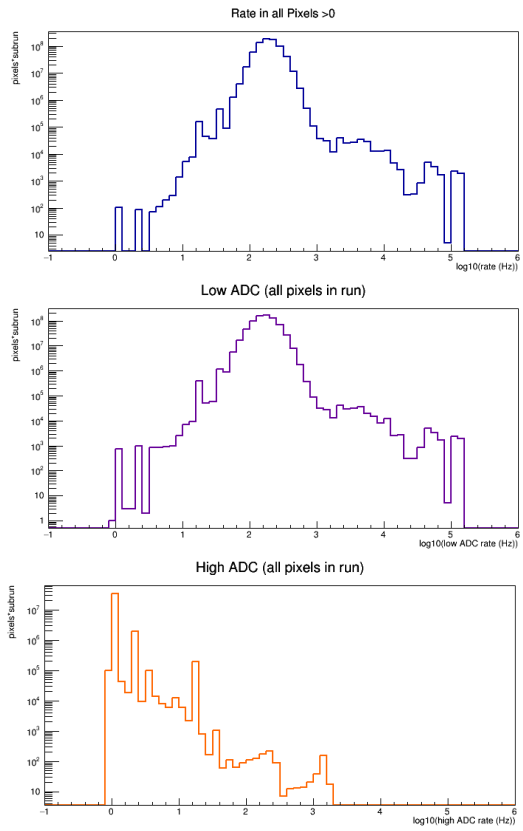

Figure 4.36: Hit rates in cosmic triggers for all APDs in the FD for one week. Top: Hit rate from all hits. Middle: Rate of hits with ADC bellow the MIP range. Bottom: Rate of hits with ADC above the MIP range. Lines indicate the thresholds for issue tracking. Note that the x axis is not uniform in units of time, as it is computed by run numbers.

Failure instances may fall into one of six categories:

Noise failure. The hit rates exceed an allowed maximum of $Rate_{max} = 500$ Hz. Figure 4.36 shows this metric and the failure rate is displayed in pink on Fig.4.37.

Quiet failure. The rates are below an allowed minimum of $Rate_{min} = 10$ Hz. Figure 4.36 shows this metric and the failure rate is displayed in blue on Fig.4.37.

Loud noise failure at high ADC counts. Noise failure for rates of hits above a threshold ADC_{max} Figure 4.36 shows this metric and the failure rate is displayed in orange on Fig.4.37.

Loud noise failure at low ADC counts. Noise failure for rates of hits below a threshold ADC_{min} Figure 4.36 shows this metric and the failure rate is displayed in purple on Fig.4.37.

Non-reporting channels. The 32 channels report null rates for the duration of that file. This failure rate is displayed in gray on Fig.4.37.

Channel drop out. Null rates for portions of time larger than $0.1/Rate_{min}$.

The issue rate uses mostly hit rate information and can be interpreted as the fraction of time$\times$channels with issues within the seven day period. The list of maintenance requirements is available to the tech-

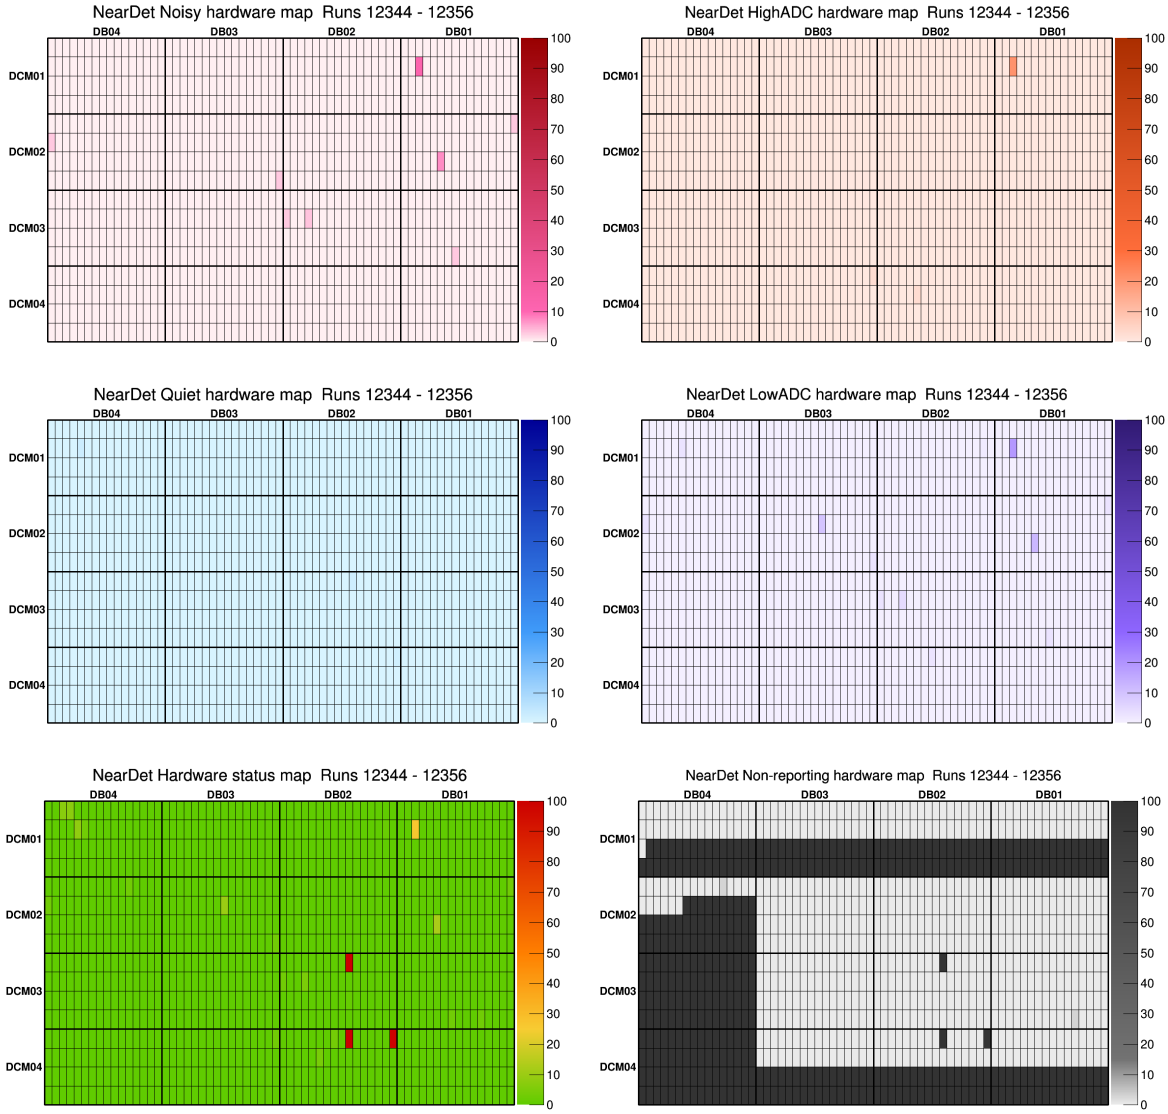

Figure 4.37: Bottom-left Panel: Issue rate on every location at the near detector. Remaining panels: Components of the issue rate.

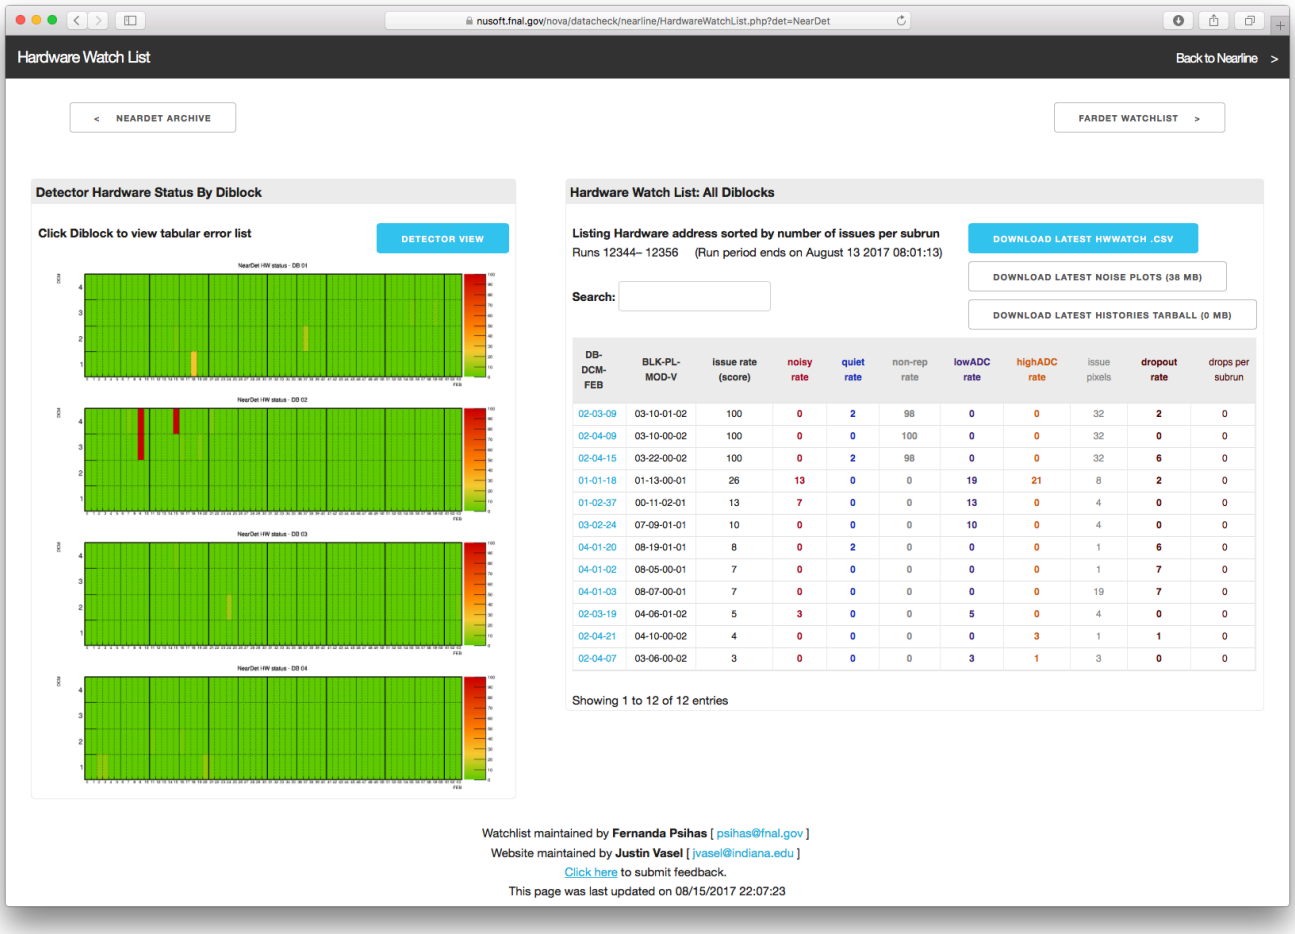

Figure 4.38: Hardware watch list front end web archive.

nicians as a summary of the computed metrics which in turn relate to issues with a specific piece of the hardware (or multiple) in the box shown in Fig.4.13.

The front end of the hardware watch system is an web-based list available to the technicians and is shown in Fig.4.38. This tool displays the current list of maintenance needs as well as plots of each element of the issue rate shown with respect to the physical location on the detector.

In addition to computing the priority maintenance list weekly for technicians, the system keeps

the history of the behavior for all pieces of hardware, computed for up to four months after an issue has disappeared. This is done by computing the noise rate history per channel over a period of several months and by keeping a summary of previous activity at that location. Fig.4.39 shows examples of the noise history for three locations, before and after maintenance. This information allows technicians to verify the stable behavior of channels after maintenance has been performed.

The effectiveness of the maintenance over larger periods of time can be computed by keeping track of resolved error conditions and correlating with the maintenance records kept by technicians in a database over periods of several months. Fig.4.40 shows the percentage of channels which operate normally after maintenance in three categories corresponding to the rate of issues prior to maintenance. These saved channels are a metric for the success rate of the maintenance performed, which improves significantly and remains above 95% after the detector was fully commissioned. This is a metric of the stability of our hardware over time.

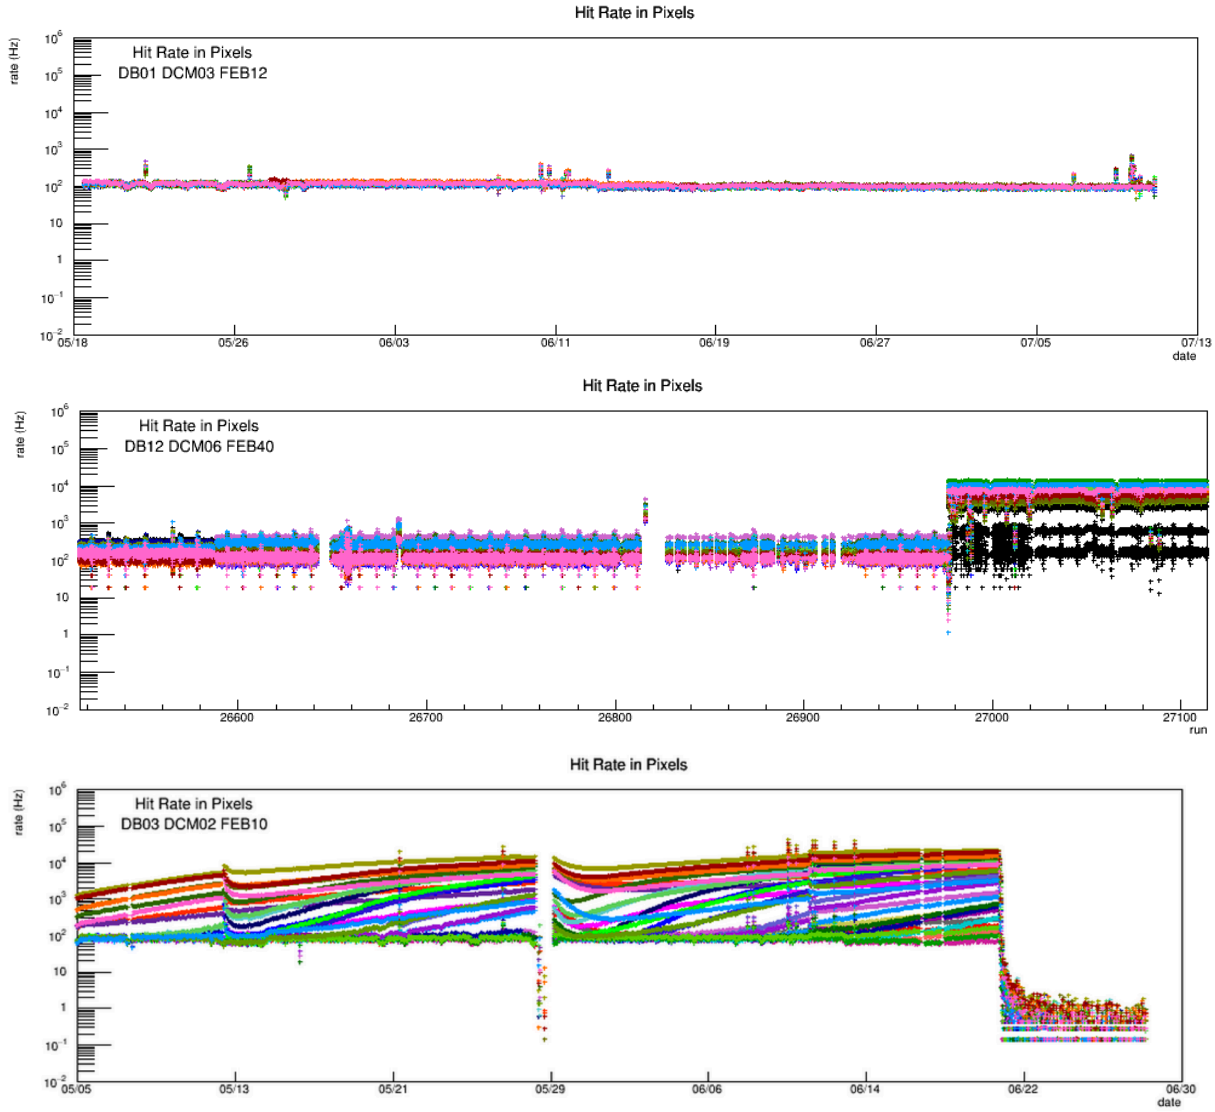

Figure 4.39: Noise History in three locations.Top: normal rates in all pixels. Center:

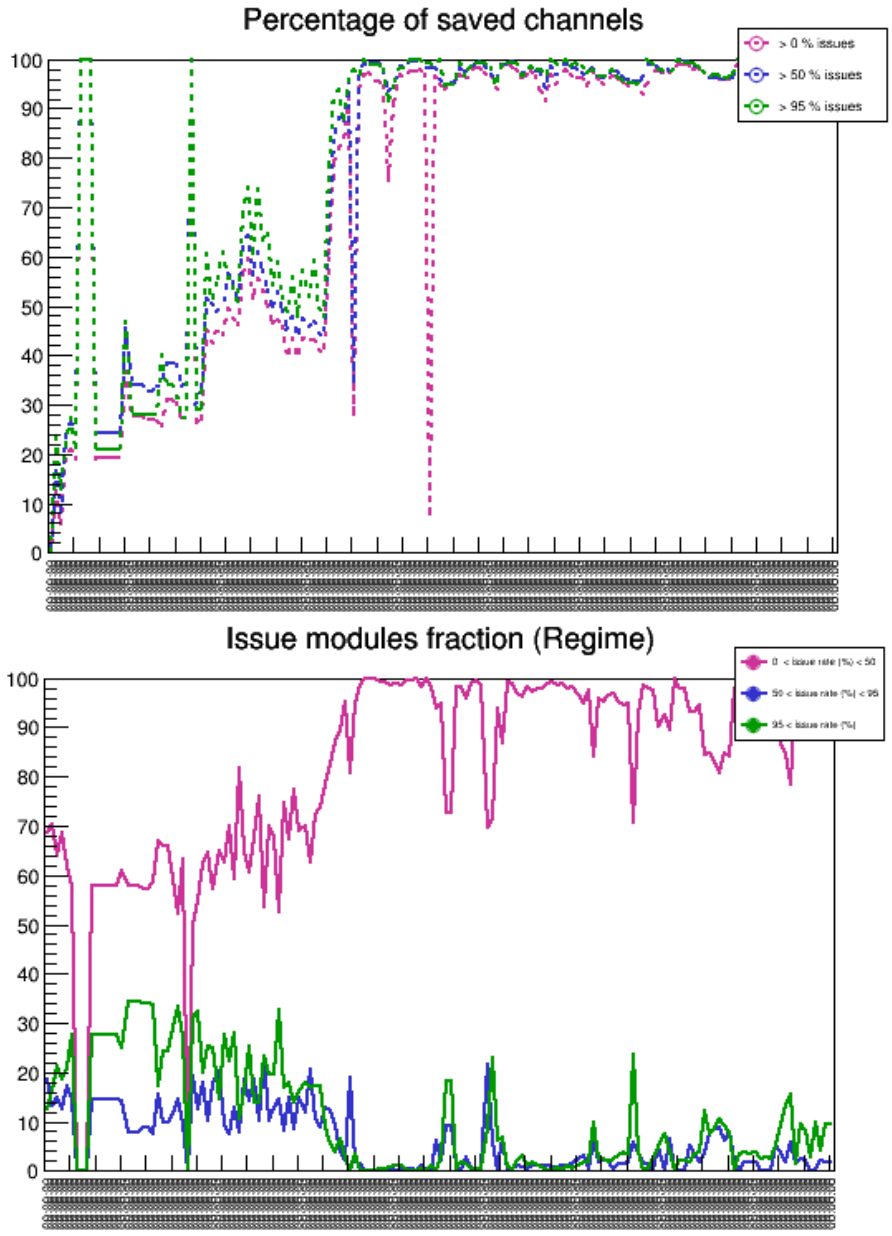

Figure 4.40: Top:Successful maintenance over a one year period of time. Bottom: Maintenance effectiveness over time (left) and hardware stability over time (right). The three populations shown correspond to issue rates over one week periods of time being non-zero, above 50%, and above 90% in pink, blue, and green, respectively.

Chapter 5

Deep Learning and Event Reconstruction

Event reconstruction uses calibrated detector readouts, consisting of hit times and charges, to extract the physical meaning of the particle interactions which produced them. The goal of reconstruction for neutrino events is to separate the neutrino interactions from other activity in the detector, to identify the neutrino flavor in the charged-current interactions , and to estimate the energy and direction of the neutrino through the particles produced in its interaction.

This chapter details the multiple approaches through which we reconstruct NOvA events. The next sections describe algorithms which are commonly used in particle physics experiments, beginning with traditional reconstruction and followed by the newest deep learning algorithmic implementations we have developed in the last couple of years [61].

The clustering, vertexing and tracking algorithms used by NOvA were inspired by existing implementations from other high energy physics experiments. Sec. 5.1 describes how these algorithms have been adapted to the geometry and readout of the NOvA detectors. For categorization both of the interaction as a whole and of the individual particles produced in it we employ deep learning techniques which are new to physics experiments. These algorithms were inspired by implementations for

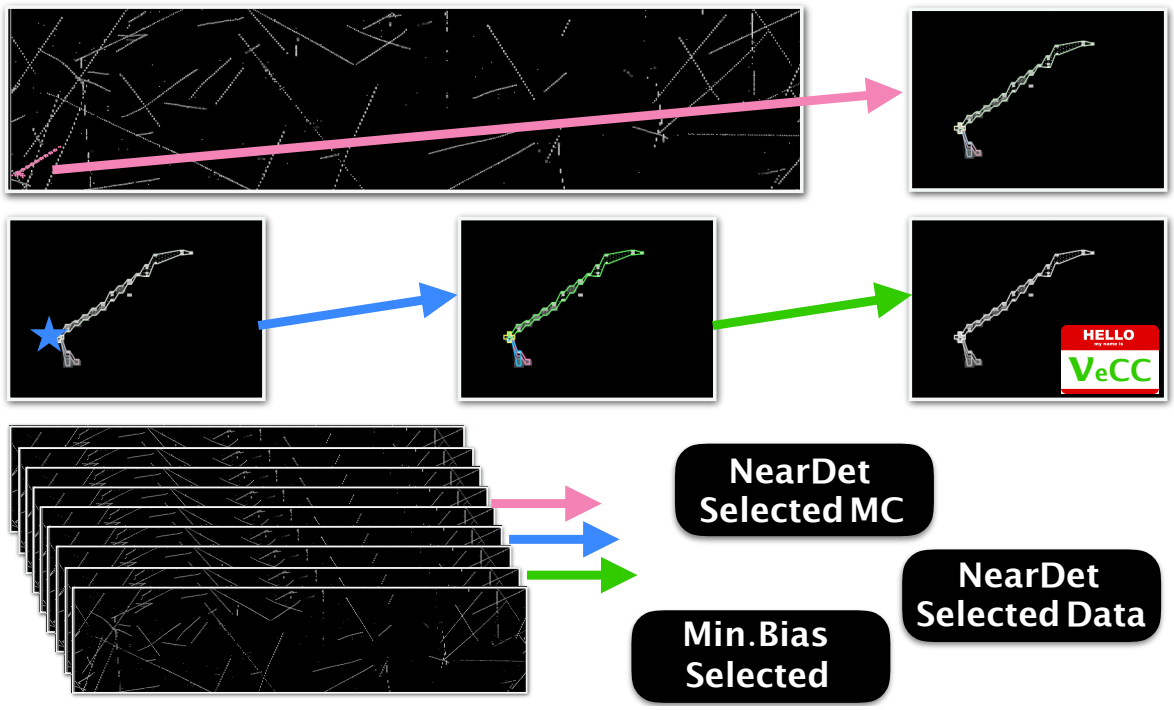

Figure 5.1: Diagram of the NOvA reconstruction chain. Top: Isolating an interaction from the 550 μsec readout. Middle: Event reconstruction: Vertex location (left), single particle clustering (center) and identification (right). Bottom: The same reconstruction algorithms are run on minimum bias data and ND MC for background estimation, then on real data.

image recognition and were adapted to the NOvA calibrated readout. The diagram in Fig.5.1 shows the main steps of event reconstruction:

Isolating Interactions: Separating each neutrino interaction or cosmic ray contribution from the rest of the activity within the 550 μsec readout. This is detailed in Sec. 5.1.1

Event Reconstruction: Extract physics information from the topology of the event and the characteristic energy depositions of different particles discussed in Sec. 5.1.2.

Event Identification. Classify different types of events in order to separate signal events from the different backgrounds as detailed in Sec. 5.2.2.

Energy Reconstruction. By the use of algorithms detailed in Sec. 5.4, use the topological features and

energy depositions in the event to map to the incoming neutrino energy.

Data/MC Comparisons Evaluate these algorithms in MC and data for the near detector and understand the discrepancies if any are found.

Data Processing. Only after the preceding steps have been developed and studied run the analysis on far detector data to compute the oscillation results.

It is crucial to point out that no one type of algorithm is an improvement over another generically, and that all techniques should be approached with the same amount of rigor when trying to understand their inner-workings. Above all else, the greatest improvements coming from different types of algorithms, machine learning or otherwise, are consistently made by careful study of the processes involved and, most importantly, by thoroughly applying the knowledge and intuition which help us relate back to the physics principles which govern the task at hand.

5.1 Event Reconstruction

While the methodology is mostly analogous, there are important differences between the ν_e appearance and the ν_μ disappearance channels. It is useful to keep the topology of signal events in mind throughout the discussion of the reconstruction chain. Characteristic signal events are shown in Fig.5.2. As discussed in Chapter 4, the detectors are optimized for electron identification. Electromagnetic depositions make distinctive showers of a handful of cells in width transversely, as shown in the middle and bottom panels in Fig.5.2. Muons, on the other hand, leave long tracks of minimally ionizing depositions, as discussed in Sec. 3.2.

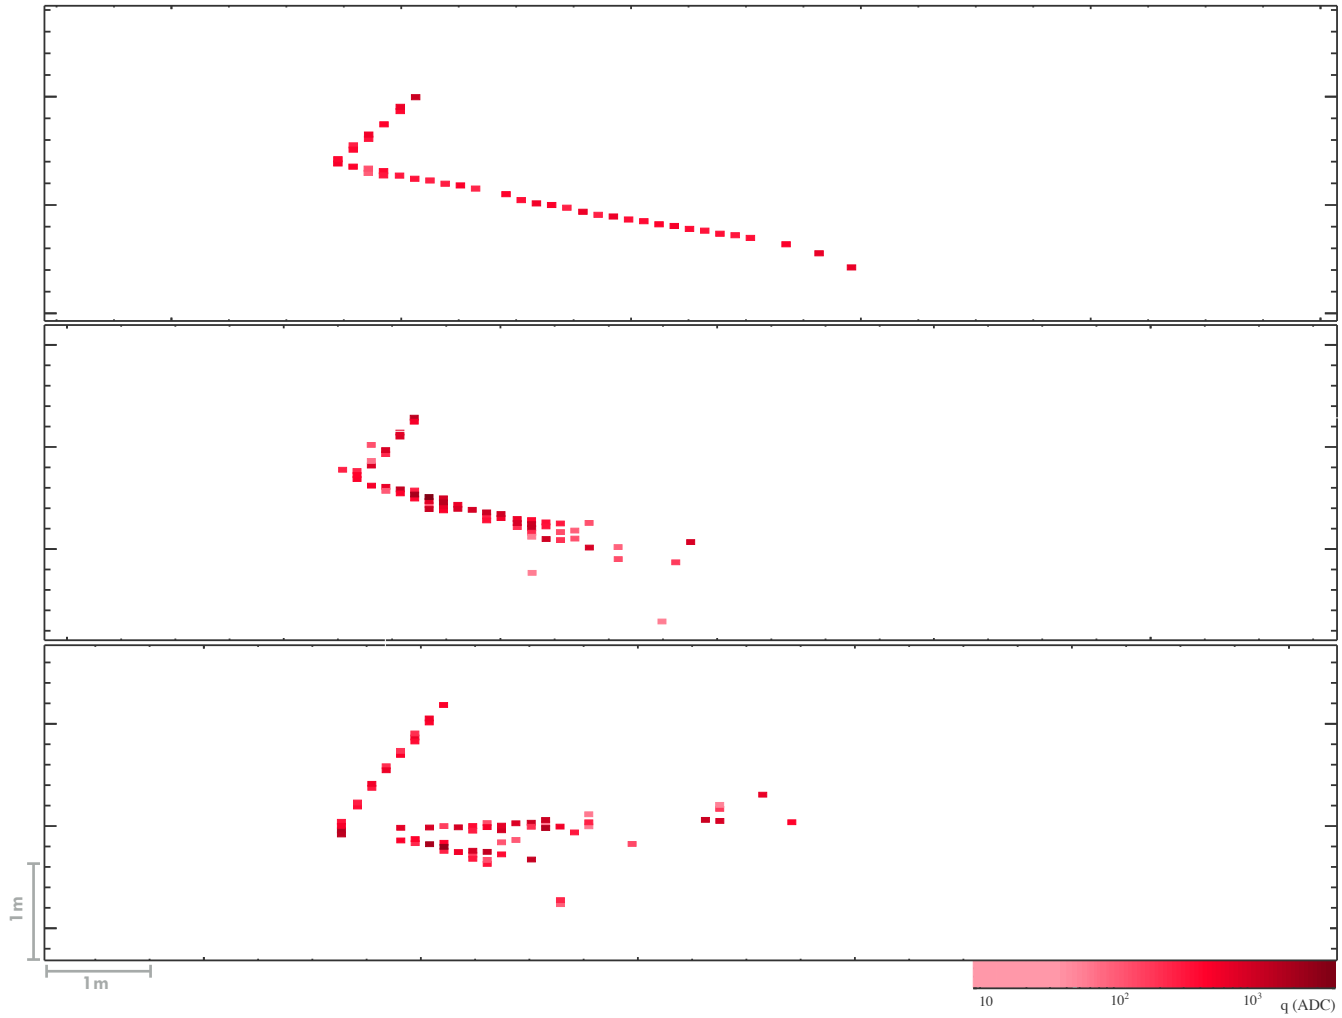

Figure 5.2: Signature event topologies for the ν_e CC and ν_μ CC signal events as well as a characteristic neutral current event. Top: ν_μ CC event with short hadron track and a long, minimally ionizing muon track. Middle: ν_e CC event with short hadronic track and characteristic electromagnetic shower corresponding to outgoing electron. Bottom: NC event with a hadronic track and two electromagnetic showers from $\pi^0 \rightarrow \gamma\gamma$.

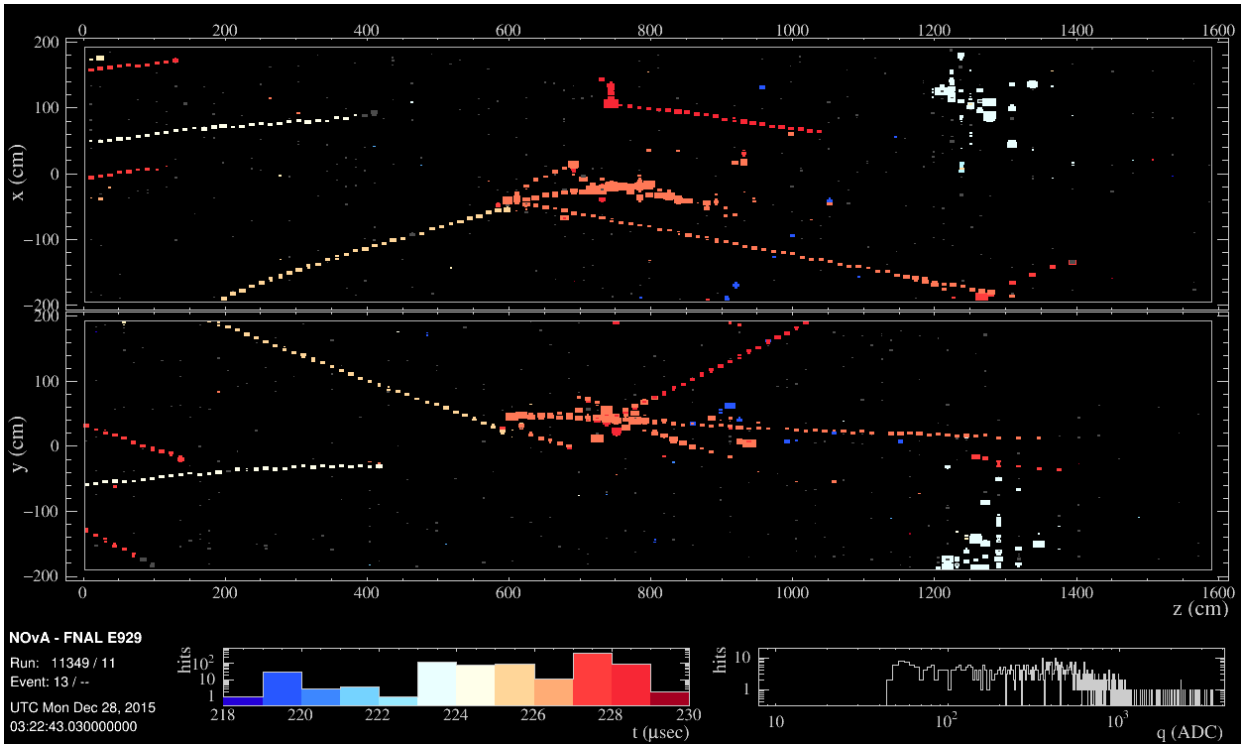

Figure 5.3: Time separated neutrino interactions in ND NuMI data. Color indicates time.

5.1.1 Isolating Particle Interactions

The NOvA DAQ records all detector activity within a single 10 μsec beam spill window of time as shown in Sec. 4.1.5. In the case of NuMI triggers, a single neutrino interaction can exist in one 550 μsec readout, along with additional detector activity. In the case of the near detector multiple neutrino interactions are usually present, and for far detector readouts the additional activity comes from the interaction of cosmic ray particles. The goal of the first stage in our reconstruction is to cluster the hits from single interactions, of neutrinos or cosmic rays, while separating them from other detector activity.

As discussed in Sec. 4.1.5, timing resolution is an important asset of our experimental design. As is evident by the readout depicted in Fig.5.3, we can take advantage of the timing resolution and

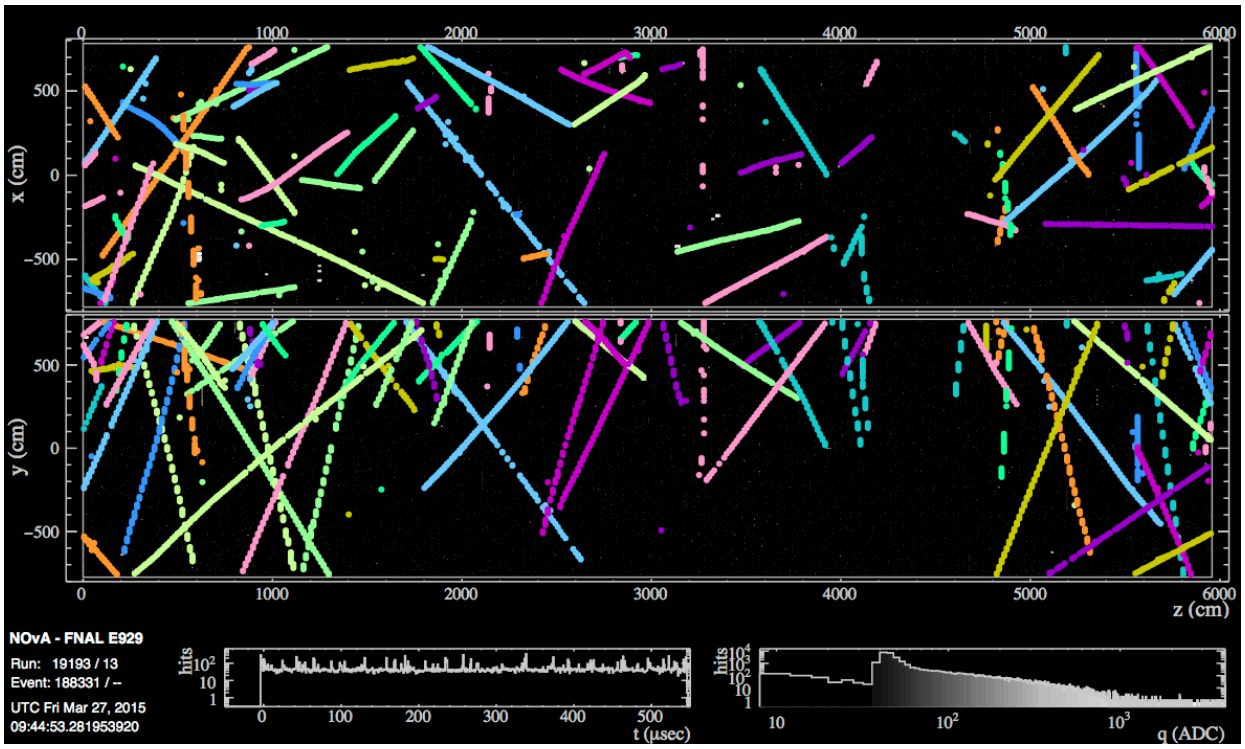

Figure 5.4: Slices in a FD data readout, marked by different colors.

granularity of the detector to isolate interactions by their time and space separation alone. Separating contributions by time only, as is implicitly done in the figure, already shows promising efficiency for separation of individual interactions, to first order. The algorithm used on NOvA for separating neutrino interactions into clusters within one readout is called *Slicer4D* and was developed by M.Baird [62] based on the DBSCAN clustering algorithm [63]. In this section the term *slices* refers to the clusters of hits being built by the Slicer4D algorithm.

Slicer4D uses two main metrics: the separation in time and space between hits (Δs), and the minimum density of hits which is considered a *core cluster*

$$d_{core\ min} = N_{hits\ min}/\Delta s_{max}. \tag{5.1}$$

Pairs of hits for which $\Delta s < \Delta s_{max}$ are considered *neighbors*. Once Δs is calculated for every pair of

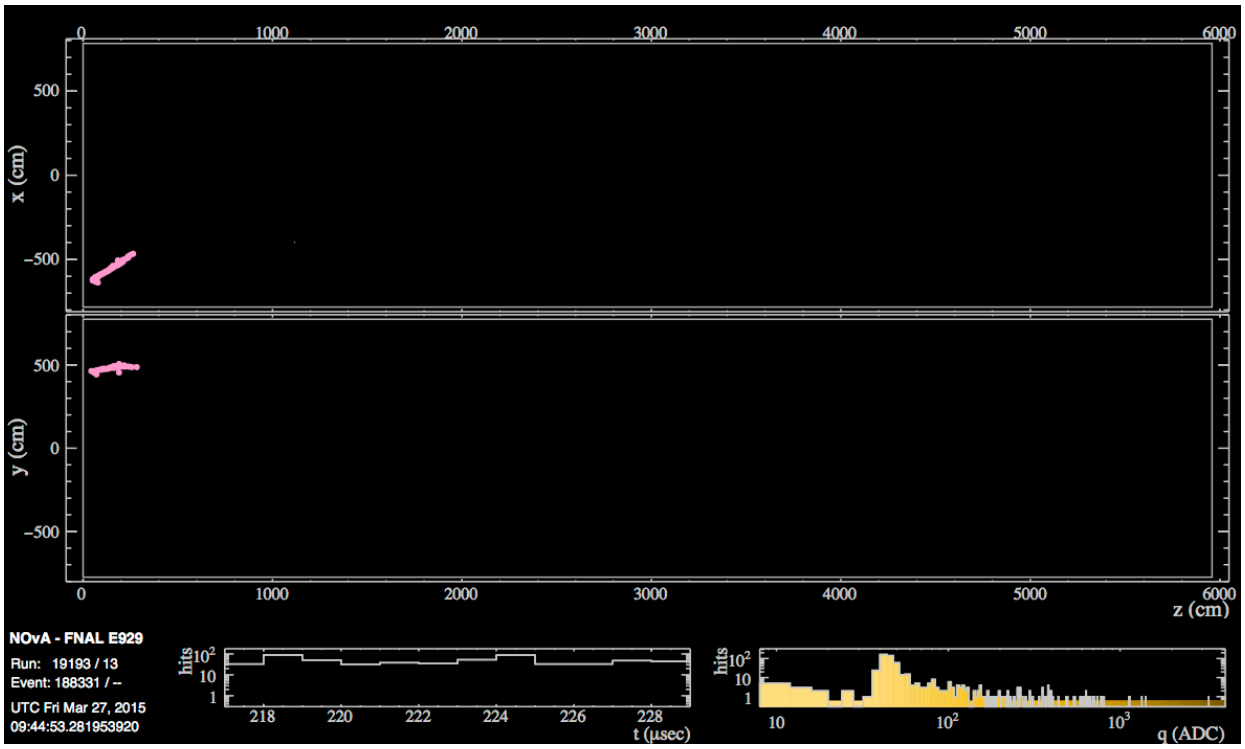

Figure 5.5: A single slice within the FD readout.

hits in the readout [1], hits belonging to regions of density above $d_{core\ min}$, called *core hits* are clustered together in a same slice. Hits outside the core cluster but which are *neighbors* to *core hits* are included in the slice as well.

$$\Delta s = \left(\frac{|\Delta t| - |\Delta \overrightarrow{r}|/c}{T_{res}}\right)^2 + \left(\frac{\Delta z}{D_{pen}(1-\delta_{v_i v_j} V_{pen})}\right)^2 + \left(1-\delta_{v_i v_j}\right)\left(\frac{\Delta x_{or} y}{D_{pen}}\right)^2, \qquad (5.2)$$

is the time-space separation used by Slicer4D, where T_{res} is the timing resolution of the hits, PE is the total number of photoelectrons for both hits, and D_{pen} and V_{pen} are penalty terms for hits separated by a large distance and hits on different views (indexed v) respectively. These are additional parameters considered for the specific geometry of NOvA events and capabilities of the detectors as well as other geometrical considerations.

[1] For the far detector running time optimizations involve not computing Δs for every pair of hits in the readout, of course, but rather skipping the rest of the calculation for hits which are too far away in time.

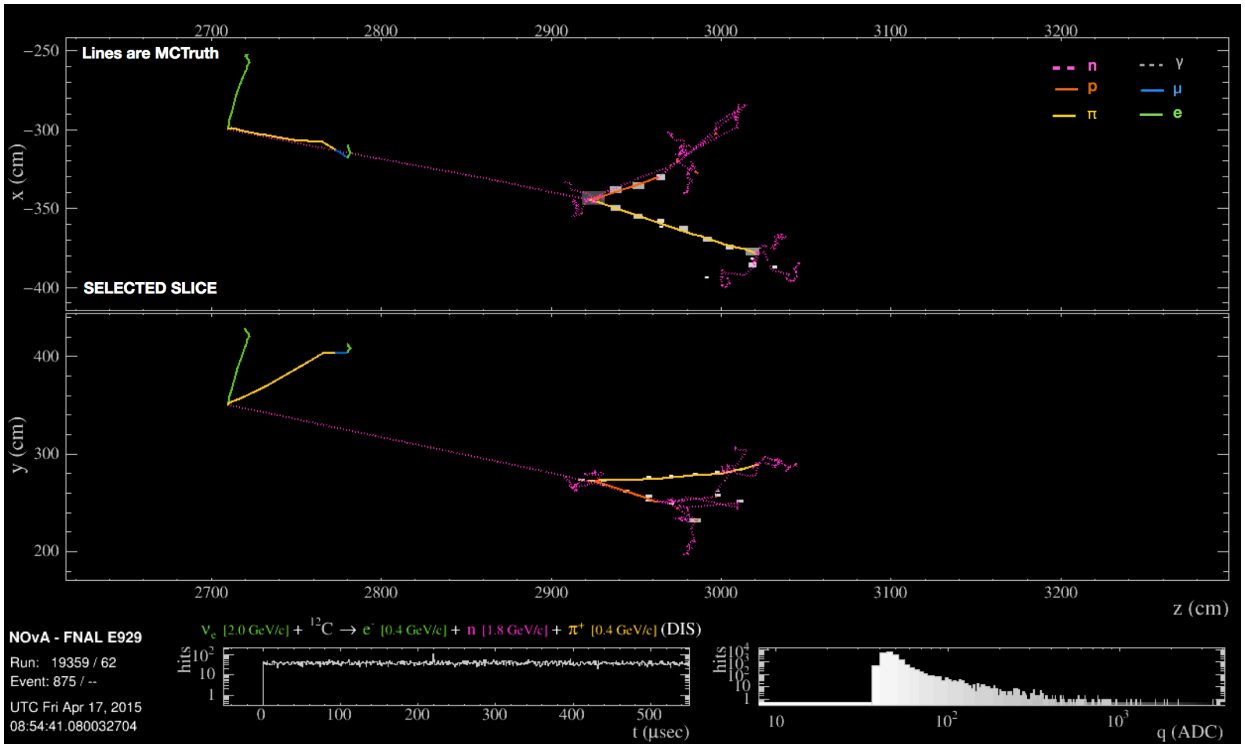

Figure 5.6: Incomplete slice. The hits shown along the path of the proton (orange) and pion (yellow) are clustered together, but hits corresponding to the electron and primary pion are not in this slice, as they are significantly separated in time.

The parameters in Eq. 5.2 were varied to maximize a figure of merit (FOM). The FOM chosen for Slicer4D was the number of slices with purity and completeness above 90% within the 550 μsec readout. Purity and completeness are calculated with respect to the true energy[2] of the simulated interactions as follows:

$$\text{Purity} = \frac{\nu \text{ Interaction true energy in this slice}}{\text{Total true energy in this slice}}, \tag{5.3}$$

$$\text{Completeness} = \frac{\nu \text{ Interaction true energy in this slice}}{\text{Total true energy of the } \nu \text{ Interaction}}. \tag{5.4}$$

One of the known inefficiencies to note comes from the fact that some processes lay beyond the

[2] True refers to the quantities known from the simulation.

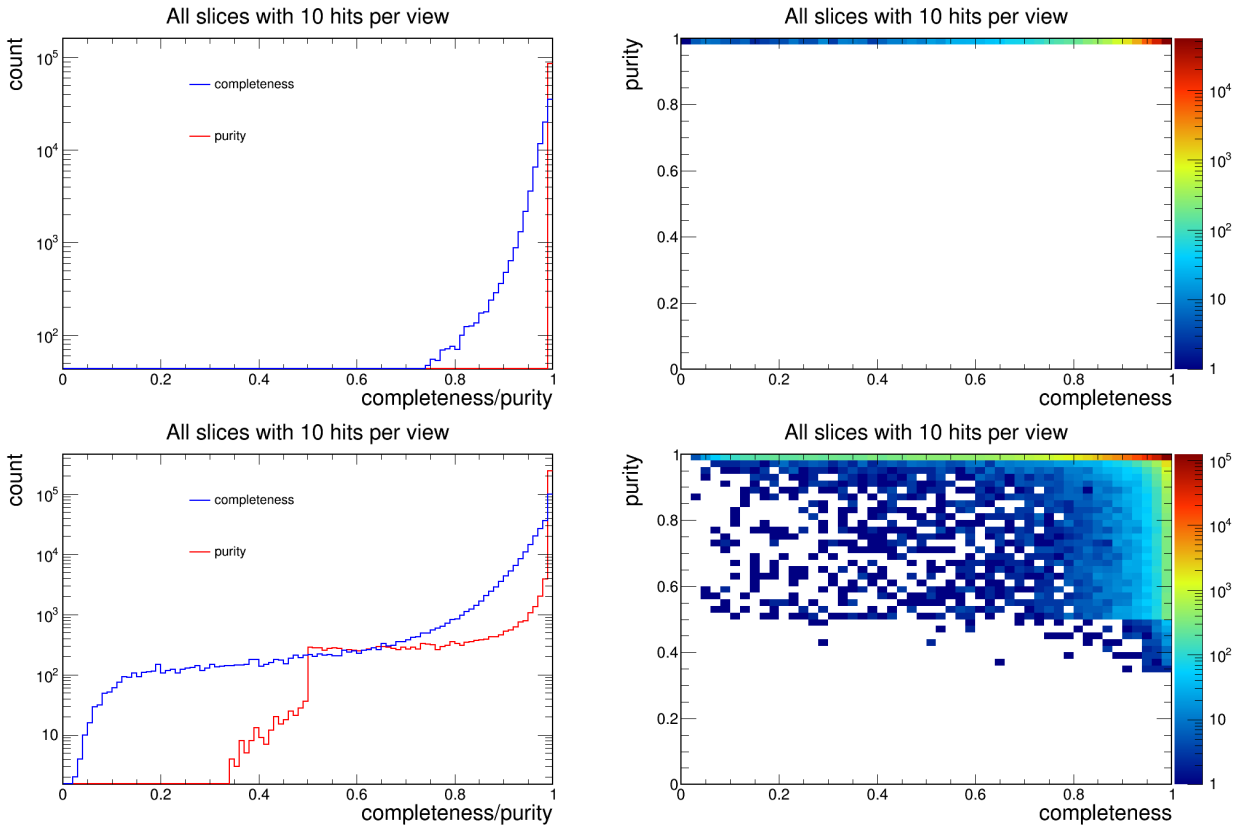

Figure 5.7: Slicer performance. Completeness and purity of slices in the far detector (top) and the near detector (bottom). Note that the differences come from the pileup of events which happens in the near detector given the multiplicity of interactions within the 10 μsec window of the neutrino spill.

acceptable Δt of the Slicer4D algorithm. An example of one such process is shown in Fig.5.6, where a neutron, product of the neutrino interaction, will travel for some distance in the positive z direction before it interacts with the detector. In this case the hits from the pion and proton from the neutron interaction will be clustered together, but will be so separated from the neutrino interaction vertex that they will be clustered as a different slice.

The performance of Slicer4D is shown in Fig.5.7 in terms of purity and completeness for events in each detector.

From this section onward the term “event” and “slice” will be use interchangeably, referring to the

cluster which is expected to contain mostly hits common to the same interaction.

5.1.2 Finding the Interaction Point

To reconstruct the original momentum four-vector for the incoming neutrino in our signal events it would be necessary to extract information from the particles in the final state. Postponing the question of identification for discussion in Sec. 5.2, the following sections will focus on the algorithms designed to separate contributions of individual particles to the event, starting with geometrical approaches that incorporate the physics knowledge of how the particles are expected to interact with the detector, as covered in Sec. 3.2.

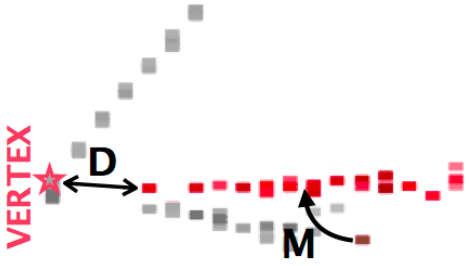

Figure 5.8: Clustering parameters.

The first task is to cluster groups of hits associated with the same primary particle from the neutrino interaction. This is done in two steps, vertexing and clustering. It is useful to begin by reconstructing the location of the interaction point or vertex, and then cluster hits assuming that starting point for the contribution of every primary particle.

The NOvA vertexing algorithm is called *Elastic Arms*. It was developed by M. Messier [61, 64, 65] based on other similar algorithms employing Hough transforms [66]. The progression of the Elastic Arms algorithm chain is the following:

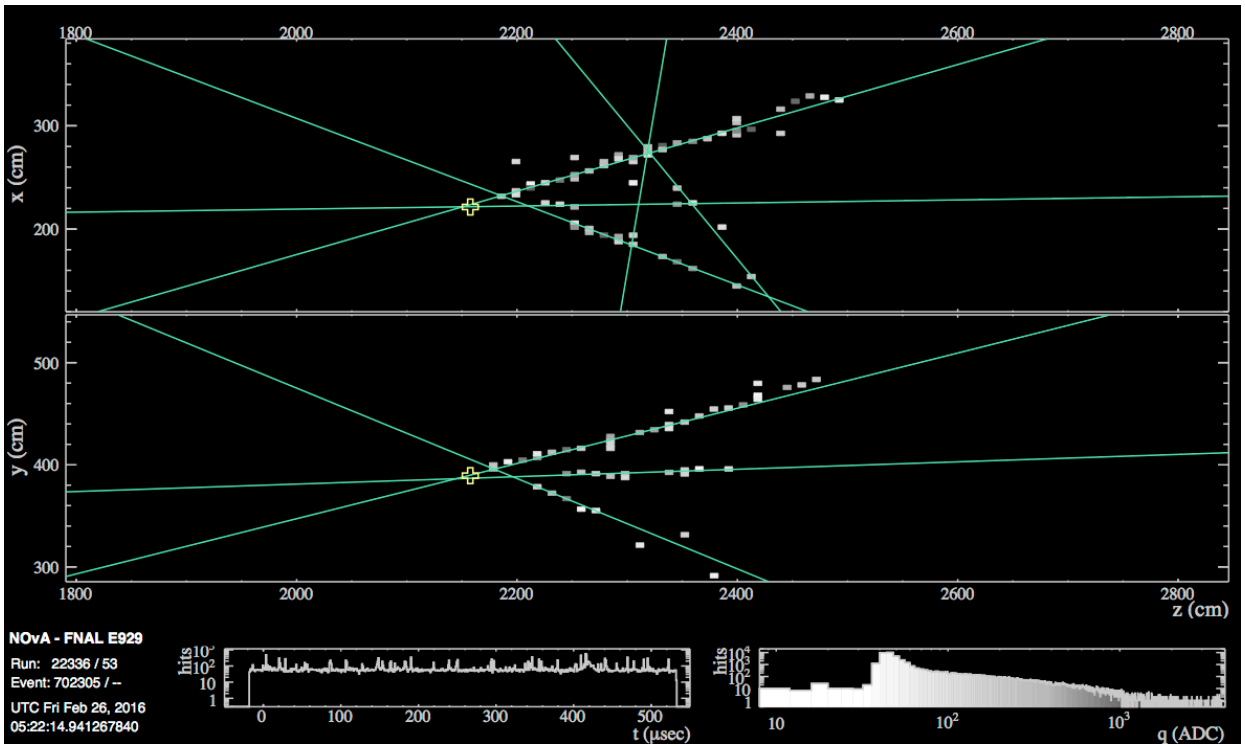

Figure 5.9: The Hough lines (green) and reconstructed elastic arms vertex (yellow) for one neutrino interaction.

Making Hough lines: Draw a line through each pair of hits separated at least some threshold distance apart, and describe it in polar coordinates θ, ρ.

Selecting Hough lines: The predominant line directions will appear as peaks in Hough (θ, ρ) space and will correspond to the main lines in the event.

Elastic Arms vertex: Through a process of minimization of Eq. 5.5 with the main hough lines as seeds, a main vertex is selected for the event. This is done through minimizing the energy function of the distance and directions to the arms (lines in 3D space), and the location to the vertex. The energy function is given by:

$$E = \sum_{i=1}^{N}\sum_{a=1}^{M} V_{ia} M_{ia} + \lambda \sum_{i=1}^{N}\left(\sum_{a=1}^{M} V_{ia} - 1\right) + \frac{2}{\lambda_\nu}\sum_{a=1}^{M} D_a. \tag{5.5}$$

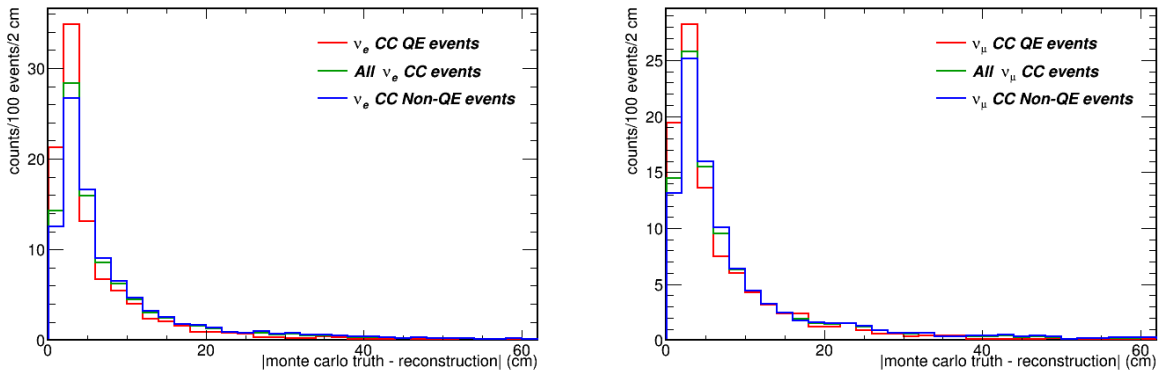

Figure 5.10: Elastic Arms performance. Distance of the reconstructed vertex to the true vertex for simulated ν_eCC events (left) and ν_μCC events (right). Most of the reconstructed vertices lay less than 10 cm away from the true vertex. The performance is better for the ν_e CC events, which is expected given their simpler topology.

M_{ia} is the distance of the hit to the arm, and D_a is the distance from the start to the vertex, as seen in Fig.5.8. V_{ia} measures the strength of the association between the ith hit and the ath arm and λ is a penalty for the hits with no associations to arms.

The performance of the vertexing algorithm is shown in Fig.5.10. More than 90% of the vertices are placed within 20 cm of the true interaction points, peaking below 5 cm. The best performance is observed on quasi-elastic events, whose topology is simpler than non quasi-elastic events. The reconstructed interaction vertex is used as input in the steps of the reconstruction which follow.

5.1.3 Clustering Particle Contributions

In order to group hits into clusters corresponding to contributions from different particles, we use a geometrical algorithm which takes in the reconstructed vertex location and the location of all hits as input. NOvA's *FuzzyK* clustering algorithm, developed by E. Niner [67], is a modified k-means clustering algorithm [68, 69] that incorporates a possibilistic clustering [70] approach, i.e., it allows unconstrained values for hit *membership* over all clusters in order to ensure that hits can belong to multiple clusters. Clusters are formed by minimizing the function:

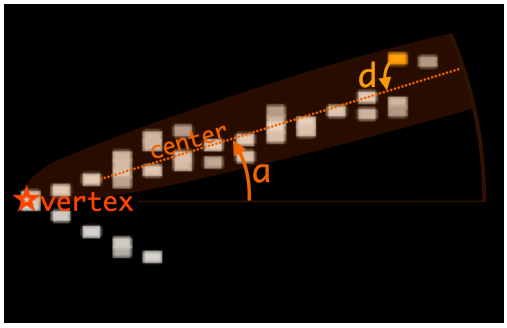

Figure 5.11: Clustering parameters.

$$J = -\frac{\beta}{m^2\sqrt{c}} \sum_{i=0}^{c} \sum_{j=0}^{h} \exp\left\{ \left(\frac{m\sqrt{a_i}\, d_{ij}}{\beta} \right), \right\} \tag{5.6}$$

where a and d are angular distances as shown in Fig.5.11, m is the allowed total membership of hits in clusters, and β is a measure of how wide the clusters can be, in units of standard deviations. The indices i and j are the corresponding cluster and hit, respectively.

Clustering is done separately for each view and then clusters are matched into 3D clusters through a KS test [71]. This involves a comparison of the deposited energy profile per unit length along a track in the direction of the cluster. This comparison is used for view matching, as shown in Fig.5.13 for the corresponding prongs in Fig.5.12.

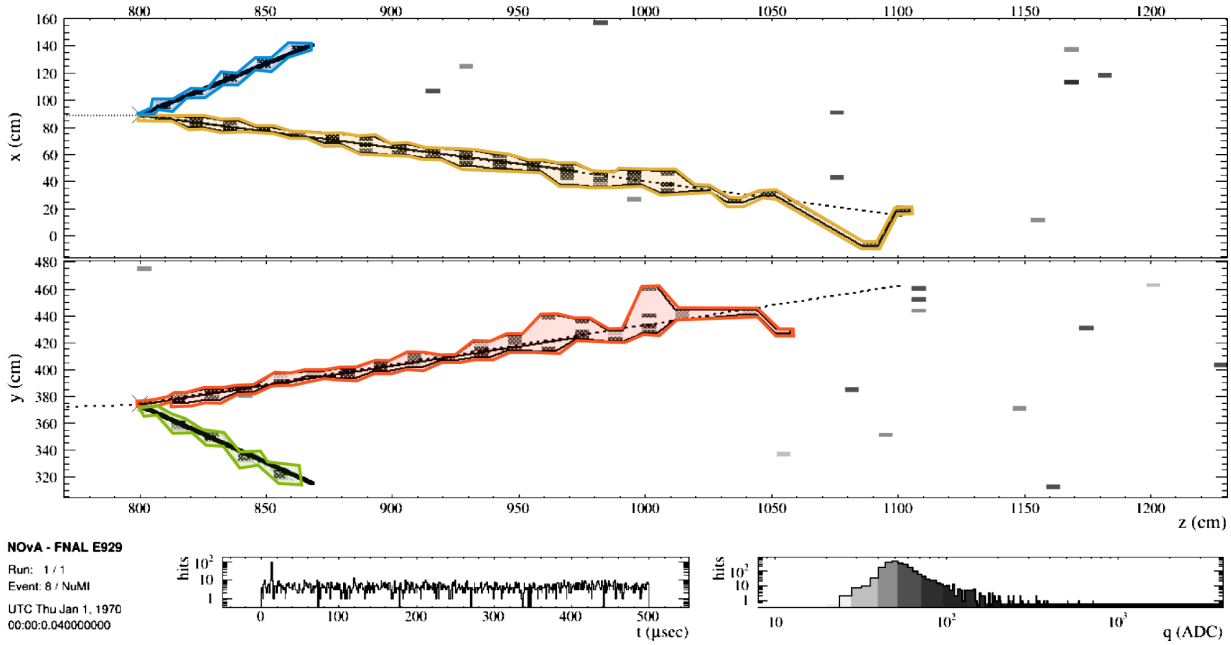

Figure 5.12: A simulated ν_eCC interaction in the Far Detector with completed 3D prong reconstruction from the fuzzy-k algorithm. From [67].

5.1.4 Tracking and Energy

The term tracking usually refers to reconstructing the trajectory a particle followed as it traveled through the detector. Tracking is a task which requires knowledge of the particle's interactions in the detector with the same considerations discussed in Sec. 3.2. Typically, on NOvA, we are interested in the details of the trajectory of muons and hadrons with higher importance, since they do not induce showers, as electrons and photons do at these energies. Our NOvA tracking algorithm, which also incorporates momentum reconstruction is called *Break Point Fitter* or BPF. It approaches the problem of tracking as different for protons, muons and pions by employing the understanding of multiple scattering, and the fact that it is different for each of these three. This tracking algorithm was developed by M.Baird [72] based on the tracking described in [66], following my initial investigations of muon multiple scattering for energy reconstruction on NOvA [73].

As a particle travels through the detector material it undergoes multiple Coulomb scattering.

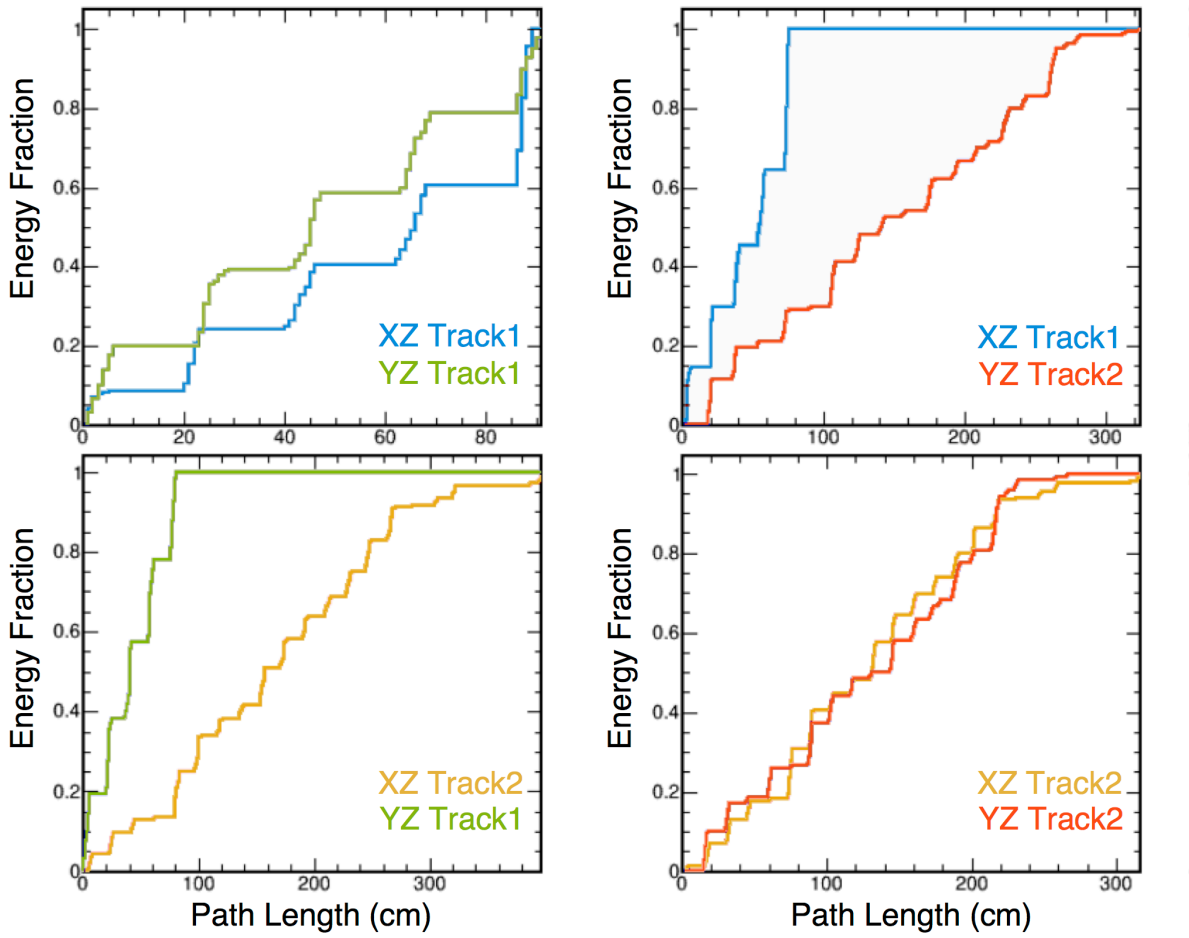

Figure 5.13: Cumulative energy profile as a function of path length along a prong for 3D match candidates, with the preferred being the upper-left and lower-right panels (matched due to having the smallest gap). From [67].

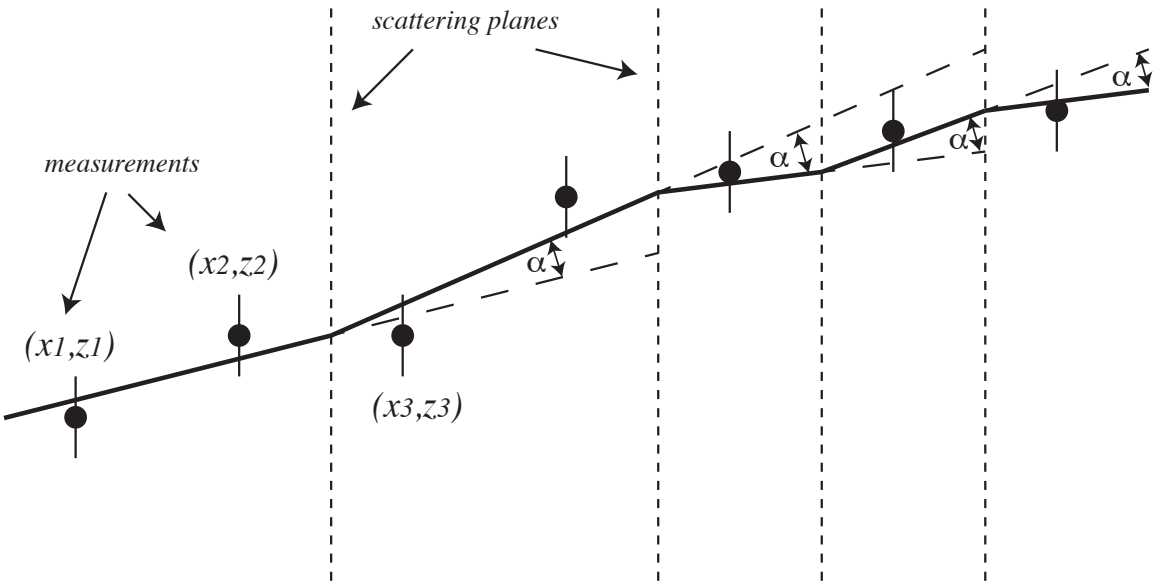

Figure 5.14: Break Point Fitter Track Building Diagram

Specifically, after traveling a distance s, in units of radiation length X_0, it is expected to scatter by an angle θ given by:

$$\theta_0 \approx \frac{13.6}{(pc\beta)}\sqrt{\frac{x}{X_0}} \times \left[1 + 0.038 \ln\left(\frac{x}{X_0}\right)\right]. \tag{5.7}$$

Note that this also depends on the momentum of the particle and that the angle is assumed to be gaussian-distributed around θ_0.

The procedure followed for each 3D cluster, constructed as detailed in Sec. 5.1.3, is depicted in Fig.5.14. It begins by counting from the end of the track back, and computing the distance traveled in units of radiation lengths X_0 as well as the deposited energy and the fraction of material (PVC or scintillator) which the particle has crossed.[3] The transverse location ξ_i of the particle's trajectory at

[3] Given the cell structure of the detectors shown in Fig.4.9, the fraction of crossed scintillator versus PVC is a function of the direction and angle of the trajectory at that point.

the z location of the i^{th} measurement can be expressed as

$$\xi_i = a + bz_i + \sum_{J=1}^{M} \alpha_J (z_i - Z_J)\Theta(z_i - Z_J), \tag{5.8}$$

where a and b are the intercept and slope of the initial track direction, α_J is the scattering angle at the J^{th} scattering plane (assumed to be small), and the Heaviside step function is used to account only for scattering planes upstream of the measurement plane. Here, the measurements shown in Fig.5.14 are the positions of hits in the detector.

Eq. 5.8 gives a measure of the change in momentum of the particle as it travels through the material, which is different for muons, protons, and pions, which are the main contributors to the event energy apart from electromagnetic energy. After these ξ_i are computed from the end of the track onward, the algorithm places an "allowed plane of scattering" perpendicular to the direction of the track whenever one of the following conditions are met:

- The predicted transverse scattering distance, $d \times \theta$, exceeds 3 mm
- d becomes greater than 100 cm or greater than $2X_0$

These choices are to optimize the location of the scattering planes. Once the scattering planes have been set, the trajectory is built, allowing the trajectory to be deflected by an expected scattering angle at each scattering plane. Because the fraction of the lost momentum given the distance traveled is different for protons, pions, and muons, so is the location of the scattering planes. Thus, BPF goes through the process three times, assuming each of the particle types. It produces a proton track, a pion track, and a muon track, each made under the corresponding assumption. In addition to producing tracks, BPF also computes a momentum estimate for each particle assumption, based on the integrated dE/dx along the reconstructed track.

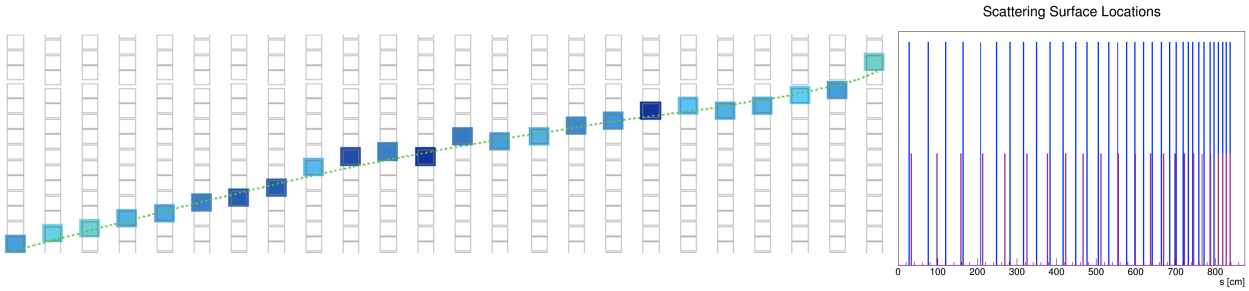

Figure 5.15: Break Point Fitter Planes

5.2 Identification with Deep Learning

Image identification is a well studied problem in the field of image processing, and is adaptable to other classification problems, such as event identification in physics experiments. Whether it is an object represented in a picture or a particle track in a projection of our detector, these techniques are designed to exploit topological correlations in objects of the same kind. Thus, the same set of tools are natural options for identification of particle detector data and other reconstruction tasks, provided that the detector readout as inputed by the algorithm retains the existing topological correlations.

In NOvA, we employ deep learning techniques for two tasks in the reconstruction and analysis. We use convolutional neural networks for event classification into different classes of neutrino interactions or cosmic rays and for particle classification starting from individual clusters. The latter is later used in energy reconstruction (see Sec. 5.3). This section contains a basic description of Convolutional Neural Networks or CNNs, as well as a description of our implementations on NOvA events.

5.2.1 Convolutional Neural Networks

The principle behind convolutional neural networks is to take advantage of the fact that images can be mapped onto matrices of numbers corresponding to pixel contents. Specifically, images are treated as $w \times h \times d$ arrays where the width w and the height h are the dimensions of the image in pixels and the depth d is the number of channels, commonly colors in RGB space. The content of each matrix element corresponds to a color in the image, which preserves the features and spatial correlations for analysis.

IMAGES AND PIXEL MAPS

In the context of CNN applications in particle physics, the term *image* is sometimes used loosely and there are many ways in which data from a physics detector can be mapped unto an N×N matrix of values, or as it is typically done, three matrices of values between 0 and 255 in order to mimic the RGB input channels for which these types of networks have been originally designed as implied in Fig.5.16.

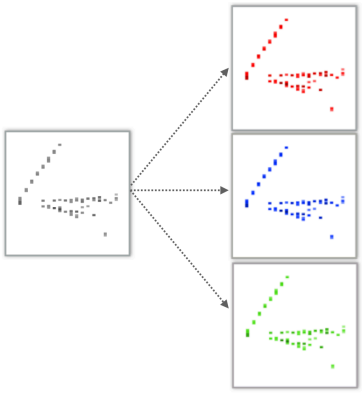

Figure 5.16: Cartoon of RGB channels.

While this constraint is not necessary for implementation, it allows for simpler adaptations of existing open-source toolsets and for comparisons with network performance on benchmark datasets in the image vision community like the well known MNIST dataset [74]. Many physics experiments have implemented CNNs for their own classification problems and interestingly different ways to map detector raw data have emerged, like the mapping onto a θ vs η coordinate plane for LHC jets and the mapping of PMT signals onto a 2D projection for Daya Bay events, both shown in Fig.5.17.

The three channels used in place of the RBG channels vary for different implementations and can sometimes combine similar topological information, or other correlated data. In the case of NOvA events, it is straightforward to map directly onto matrices called *pixel maps* corresponding to cell locations in XZ or YZ planes, as depicted in the event displays shown throughout this document. For the conversion to pixel maps, one for each detector view, the cell and plane location of the hits correspond to one element of a 80×200 matrix, mimicking an image of these dimensions. Fig.5.18 shows an example of these pixel maps where the values for each pixel are calibrated hit energies converted to a scale of 256 discrete values.

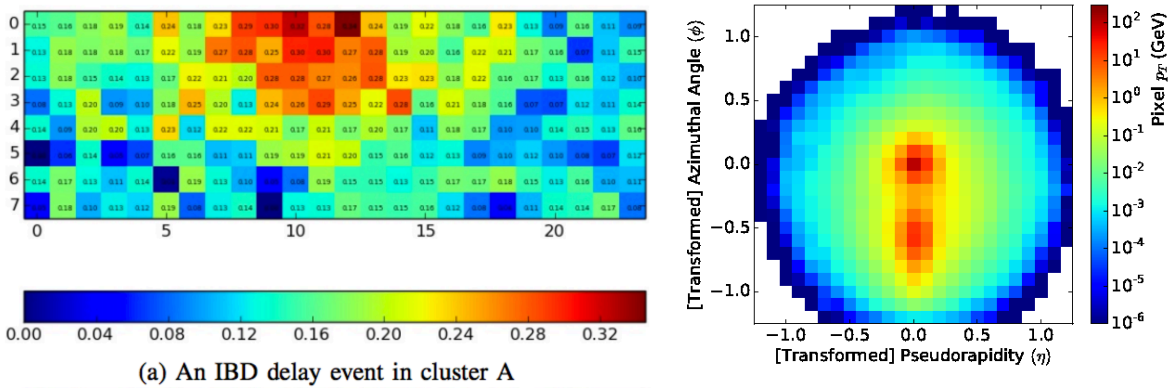

(a) An IBD delay event in cluster A

Figure 5.17: Other approlaches at making pixel maps. Left: Daya Bay "pixel map" is a mapping of charge deposits in PMTs along a cylinder unwrapped into a 2D (8 "ring" x 24 "column") pixel image. [75] Right: ATLAS maps jet structures into a binned azimuth θ vs pseudo-rapidity η [76] converted so all jet structures point in the same direction in θ.

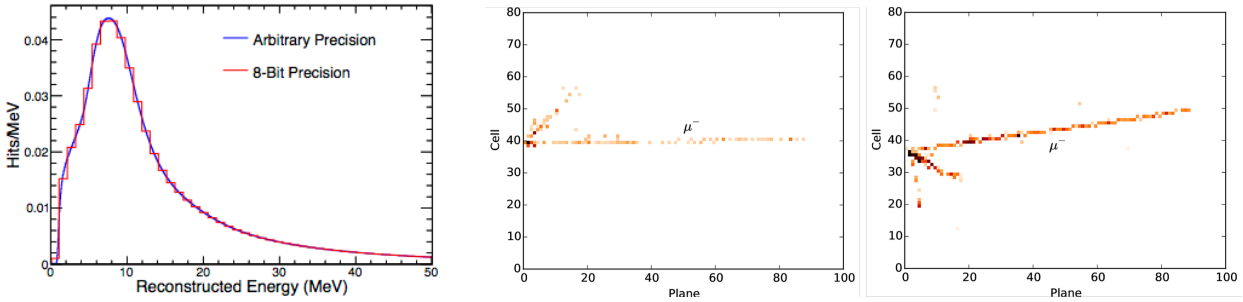

Figure 5.18: Left: The deposited energy (in MeV) in hits converted to the 256 scale values (8 bit precision). Center and Right: Examples of Pixel Maps for a ν_μ simulated event from the top and side views, respectively.

FEATURES AND FEATURE EXTRACTION

CNNs analyze an image by applying matrix filter operations across it. These matrix filters, also called *kernels*, are different matrix operations which when applied over the image extract features consistent with their size. *Convolutions*, the application of an image kernel across the image (typically from left to right repeatedly by rows as depicted in Fig.5.19), are uniquely defined by the operation they perform as well as three parameters: kernel size, stride, and pad which define the dimensions of the kernel matrix, the number of pixels to skip between operations, and the number of pixels to skip at the edges of the image, respectively.

In a way, *feature extraction* is the additional step to what is done by commonly used fully connected NNs like the ones detailed in Appendix A. These simpler networks will typically take in a vector of values as input, where the values are usually parameters known to be correlated with the target task. Extracted features, on the other hand, are learned rather than set a priori. Nevertheless, Deep CNNs and traditional NNs approach learning tasks in fundamentally the same way and there is great commonality between the structure and the learning process of the two types of networks. Fig.5.20 describes the basic components of a CNN but a more complete glossary of common NN concepts and CNN particularities is included in Appendix A.

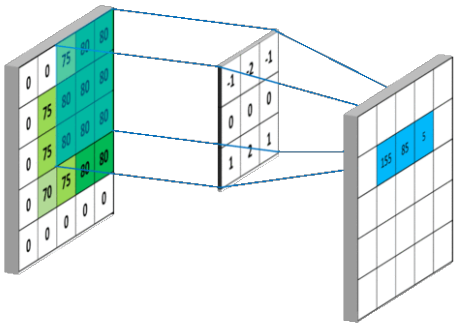

Figure 5.19: Convolutions: The application of an image kernel across an image.

One of the important subtleties in the use of CNNs is that not only are the kernels not pre-defined by the user but they also evolve throughout the training. This means that the set of features that the network uses to classify objects is learned rather than limited to the extracted physical quantities which

are commonly used on KNNs or other fully connected networks. The process of feature extraction is, thus, decoupled from traditional reconstruction and, as such, unaffected by the inefficiencies of every step of the reconstruction. Contrasting the advantages of decoupling from pre-defined reconstructed features is the fact that there is no knowledge of what physics quantities are used for this classification a-priori. It falls on the analyzer to perform the relevant investigations to understand the inefficiencies of the algorithm and to introduce physics knowledge as necessary.

NETWORK COMPONENTS

The process of training networks for classification requires understanding of the typical components of a CNN, shown in the diagram in Fig.5.20.

Convolutional Layers: These layers apply learned kernels of fixed dimensions to the image, as shown in Fig.5.19. Their output is of the same dimensions as the input image.

Pooling Layers: These layers are functionally similar to convolutional layers but focus on reducing the size of the output. Adding or averaging pieces of the image are examples of simple pooling operations.

Fully Connected Layer: Typically CNNs include a fully connected layer at a final stage. At this point the features extracted by the previous layers have been turned into sets of variables and weights which can be fed to a traditional network.

Learning Minimization Function: One important component not depicted in diagram but visible in real network architectures like the one in Fig.5.21 is the minimization function. This function is customizable and the choice influences network performance and propensities to reach local versus global minima, overtraining, etc.

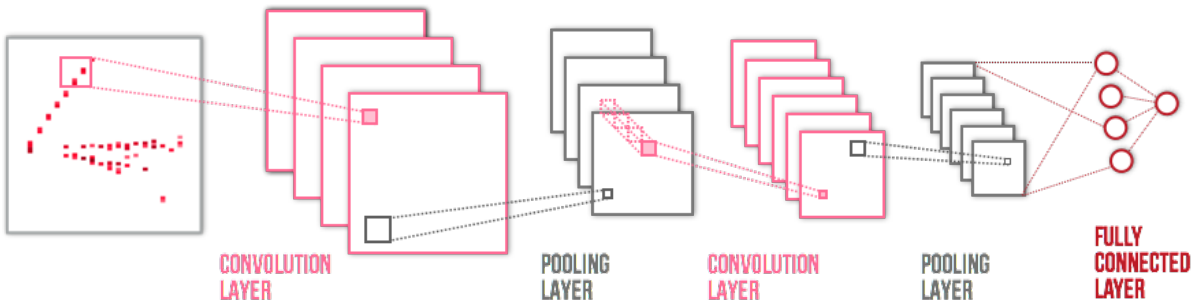

Figure 5.20: Basic Structure of a Convolutional Neural Network: Convolutional Layers (pink). Pooling Layers (gray).Fully Connected Layer (red).

The number of convolutional layers, pooling layers and their characteristics, as well as the learning rates and minimization functions vary for different implementations. These *hyperparameters* can be optimized for networks solving particular problems for images of different sizes, as well as the specific target the network is being trained for [4], number of categories to classify, etc. The network architectures used for classification on NOvA employ two key elements which serve our image classification task well: the siamese architecture structure and the concept of inception module.

The Inception module: In an implementation first employed by the GoogLeNet [77], these modules are layers that not only perform one convolution over the image but combine many convolutions of different sizes in a network-in-network structure. This allows for the extraction and merging of features of different dimensions at the same step.

Siamese structure: Typically for image recognition a same set of operations is performed on all three channels, R, G and B. In a siamese tower structure like the one depicted in Fig. 5.21, the network learns from each channel independently, thus, learning different operations for each one.

[4] While the implementations shown in this dissertation focus on classification of particles and events it is possible to specify different targets as the output of the network. For instance, energy reconstruction and vertex finding are some of the applications which we are currently exploring.

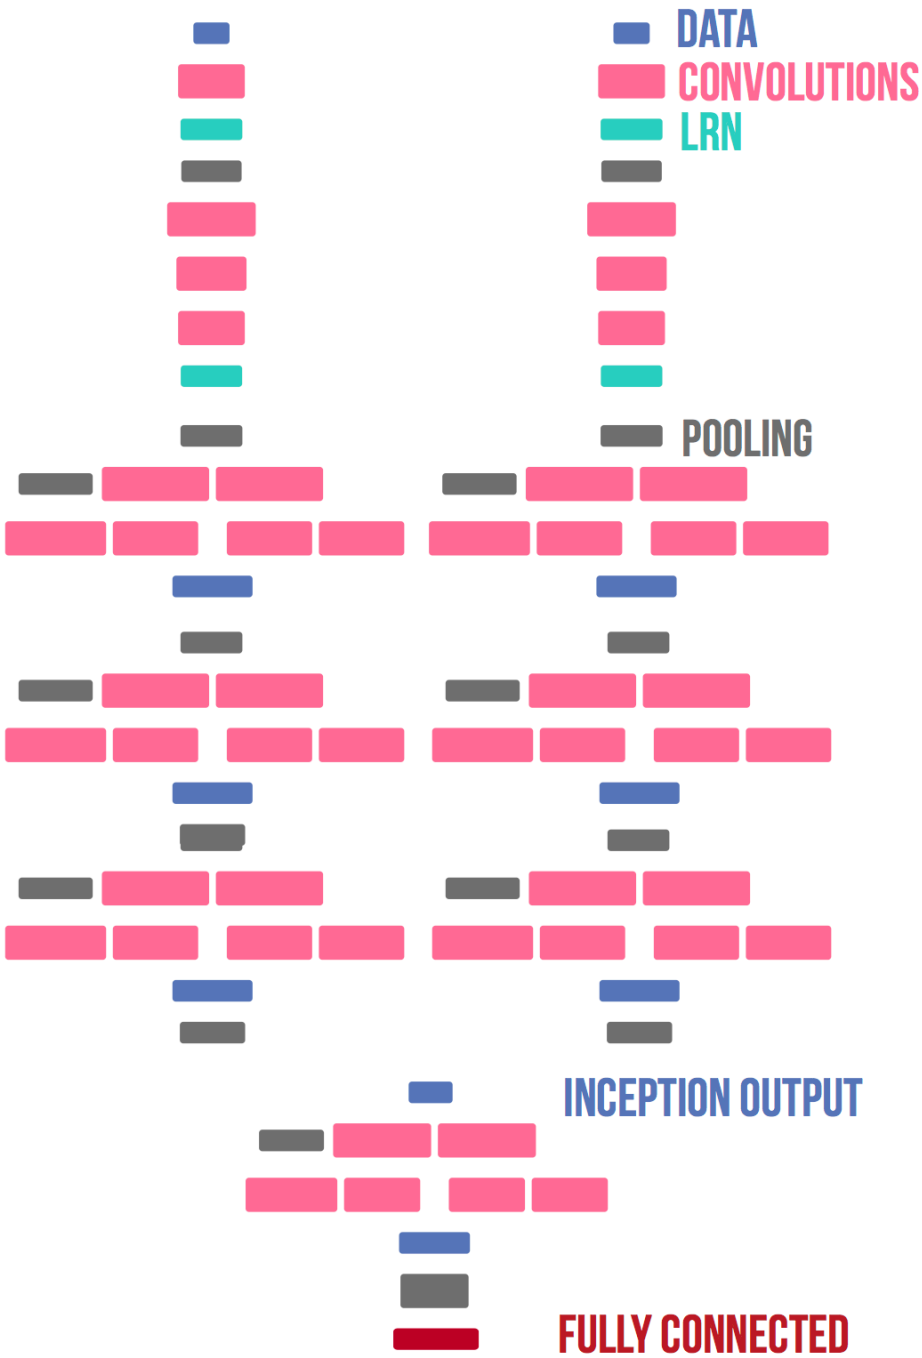

Figure 5.21: Event CVN Network Structure Diagram. The siamese (two-branch) structure receives the XZ (top) and YZ (side) views as independent inputs (top blue boxes). The fully connected layer computes softmax output (bottom red box). Full details of the architecture are available in Appendix A.

It may seem natural to use the two views on NOvA events as different channels of an image, and there are some implementations which attempt this approach [78]. However, the lack of spatial correlation between the XZ and YZ views motivates the choice of a two-pronged network architecture instead. The input to the network (shown in Fig. 5.23) consists of the two views separately, such that each view is operated on by one of the siamese towers throughout most of the learning process. This means that part of the network architecture is divided into two main branches, which encourages the learning of features in each view independently before combining the features at later steps. The intention of using a siamese architecture like the one depicted in Fig.5.21 is both to maximize the amount of useful information extracted from each view and to utilize information uniquely available to one view like the incoming angle of cosmic ray muons.

Figure 5.22: Event CVN Input: Two views of the full event and a particle label.

5.2.2 Event Classification with CVN

The event identification network, which we call *EventCVN* for Convolutional Visual Network, is a two-prong siamese architecture and it is depicted in Fig.5.21. The full details of the architecture, including the structure of each of the layers, is included in Appendix A. The EventCVN network was inspired by the GoogLeNet network architecture for image recognition [77]. Our network was designed and trained using the open-source Caffe [79] framework, a set of libraries and methods to train CNNs. Given the computationally intensive task that convolutions on millions of images represent, these networks are more efficiently handled by graphics processing units or GPUs. The event classification network was trained in the Wilson GPU Cluster at Fermilab and further optimized at the BigRedII supercomputer cluster at Indiana University.

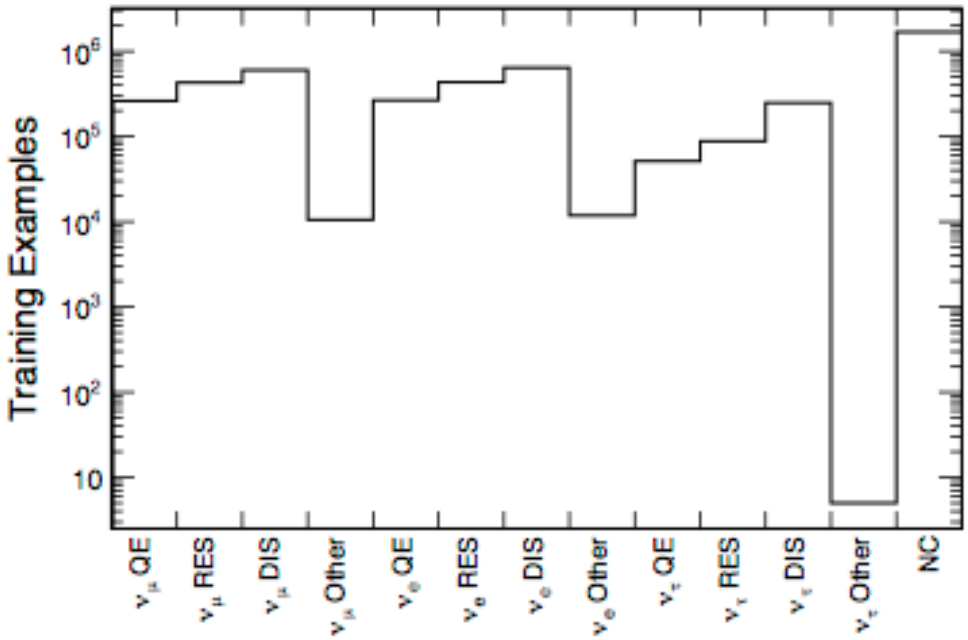

Figure 5.23: EventCVN training sample composition.

EventCVN was trained on approximately 4.7 million events distributed as indicated in Fig.5.23 where all the neutrino interactions come from simulated events and all cosmic ray events come from minimum bias data taken at the far detector. The input for the network (per event) are two pixel maps (one for the top view, one for the side view) and a label in one of the categories listed in Fig.5.23.

This work was implemented for the first time in a high energy physics result on NOvA's 2016 ν_e appearance analysis [80]. More details of the event classifier can be found in our publication, "A Convolutional Visual Network Neutrino Event Classifier" [81], and an implementation for the ν_μ disappearance analysis is explored by D. Rocco [82].

NETWORK PERFORMANCE

Many approaches can be taken to evaluate the performance of the classification network, either during training or at the analysis stages. A typical metric for the performance of the network, regardless of the task it is designed for, is to look at the evolution of losses and overall accuracy as a function of training iterations. As detailed in Appendix A, the *loss* is a metric of how closely the current performance of the network maps the real behavior of a known test sample. More details of the loss function can be found in Appendix A but it is sufficient to understand it conceptually as a difference between the probability assigned to the data for the true class they belong to and the true probability to be in that class. In other words, the distance between the approximate classification score function built by the network and the true one from simulated data. Losses are computed with some set frequency in terms of number of training iterations as well as for the test sample, independently. Fig.5.24 shows the evolution of the loss and accuracy during the training of the event classifier.

In general, the loss and accuracy are monitored during training to detect simple errors like over-training or abnormal behavior of the losses, which could indicate either a poorly chosen loss rate or an unevenly populated training sample. However useful, these metrics are only measures of the network's performance at learning a given task. The physics implications and the performance differences for different categories must be explored in more detail post-training.

The output of EventCVN is a vector of scores which can be interpreted as a probability for each category. This is called *softmax* output and is described in Appendix A. This output vector contains a score between 0 and 1 for each of the categories trained for, in this case those shown in Fig.5.23. Being normalized to add up to 1, the scores can also be interpreted as PID values the same way the output of other NNs is typically treated. Furthermore, the assignment of scores beyond the highest one may

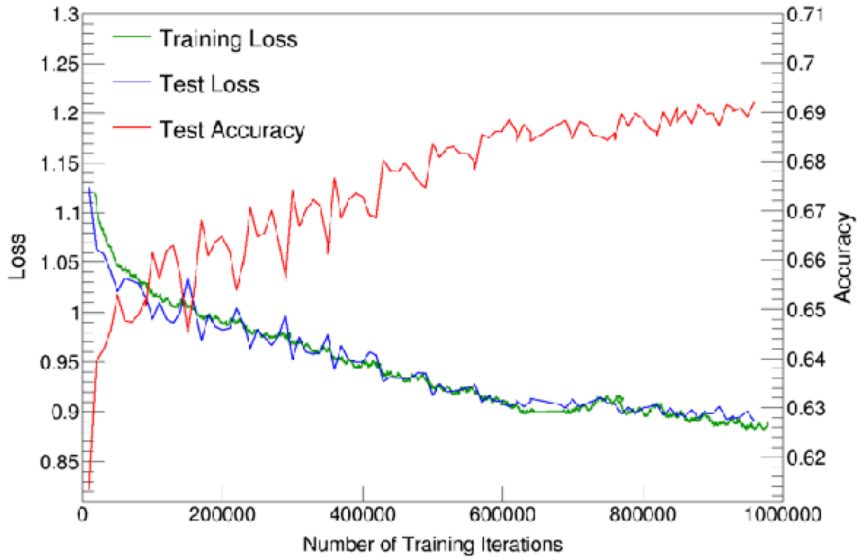

Figure 5.24: Loss and Accuracy plot

provide additional information about the event. The two main approaches to use the softmax output are to take the largest score in the vector as the selected class, as is done in Fig.5.25, or to use only the value of interest for specific signal events, as is done in Fig.5.26.

Classification matrices, like the ones shown in Fig. 5.25 are useful to evaluate efficiency or purity (depending on how they are normalized) for a simple selection based on the largest score, and they are also used as a first order check of the algorithm's performance compared with one's intuition. For instance, Fig. 5.25 left shows the best performance on the simplest topologies like cosmics and CCQE events, and Fig. 5.25 right shows poor performance for ν_τ CC events, mostly being misidentified as NC. This is expected feature given that these events are above the τ production energy threshold and the τ lifetime. ν_τ CC events mimic ν_μ events given the rapid τ decay into a muon, and they generally display large fractions of hadronic activity, which is consistent with NC events.

Another subtlety of the treatment of scores for analysis is that the classes chosen for training subdi-

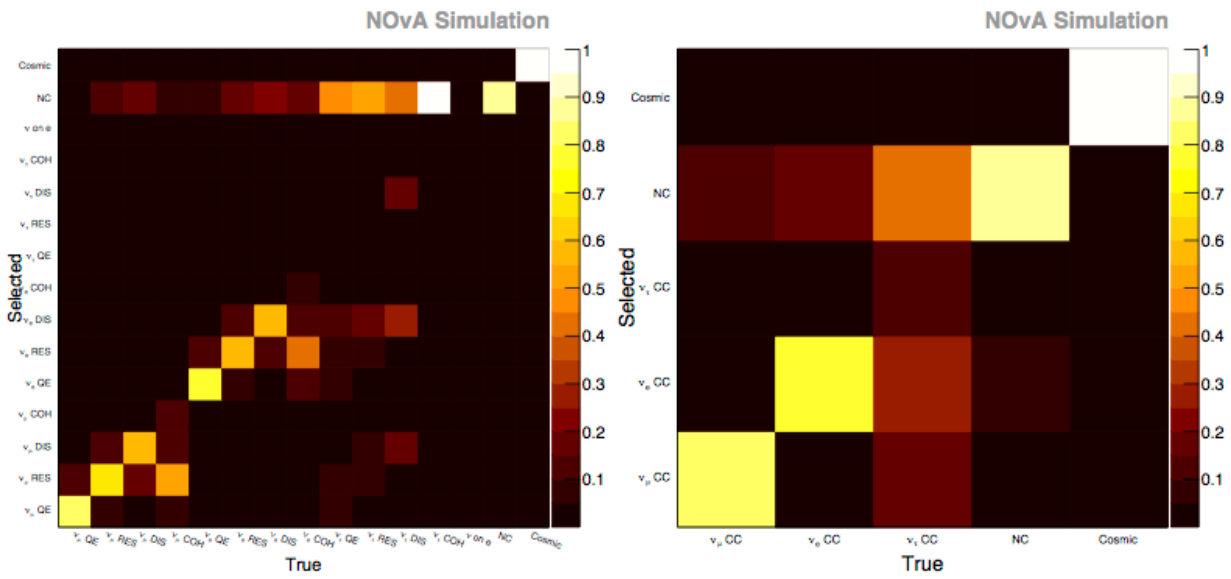

Figure 5.25: Classification Matrix. This plot shows the true class of events (on the x axis) and the class selected by the network on the y axis. In this case the selection is done by choosing the largest value from the output softmax vector. That is, the selected class is whichever class got the highest score on the output vector. Left: All the categories trained for. Right: Reduced categories used for analysis.

vide the categories of interest (ν_e CC, ν_μ CC, ν_τ CC, NC and Cosmic backgrounds) into more specific categories, in this case associated to the interaction label given by the simulator. The GENIE simulation assigns the interaction types Quasi-elastic, Resonance, Deep Inelastic Scattering for the processes discussed in Sec. 3.1, however, not all of these result in topological differences in the final state. Thus, the scores for subcategories have been added together, which preserves the normalization and reduces the set of scores to those shown on the right side of Fig.5.25.

The performance of the network was also evaluated both in the figures of merit for the main analyses shown in Table 5.1 and Table 5.2 as well as accuracy for the top category which is common in network training of this kind. The results shown in the tables correspond to a simple selection used for comparisons between different optimizations. For the optimized analysis selections, the use of EventCVN represented a 30% increase in exposure for ν_e CC signal events back when first imple-

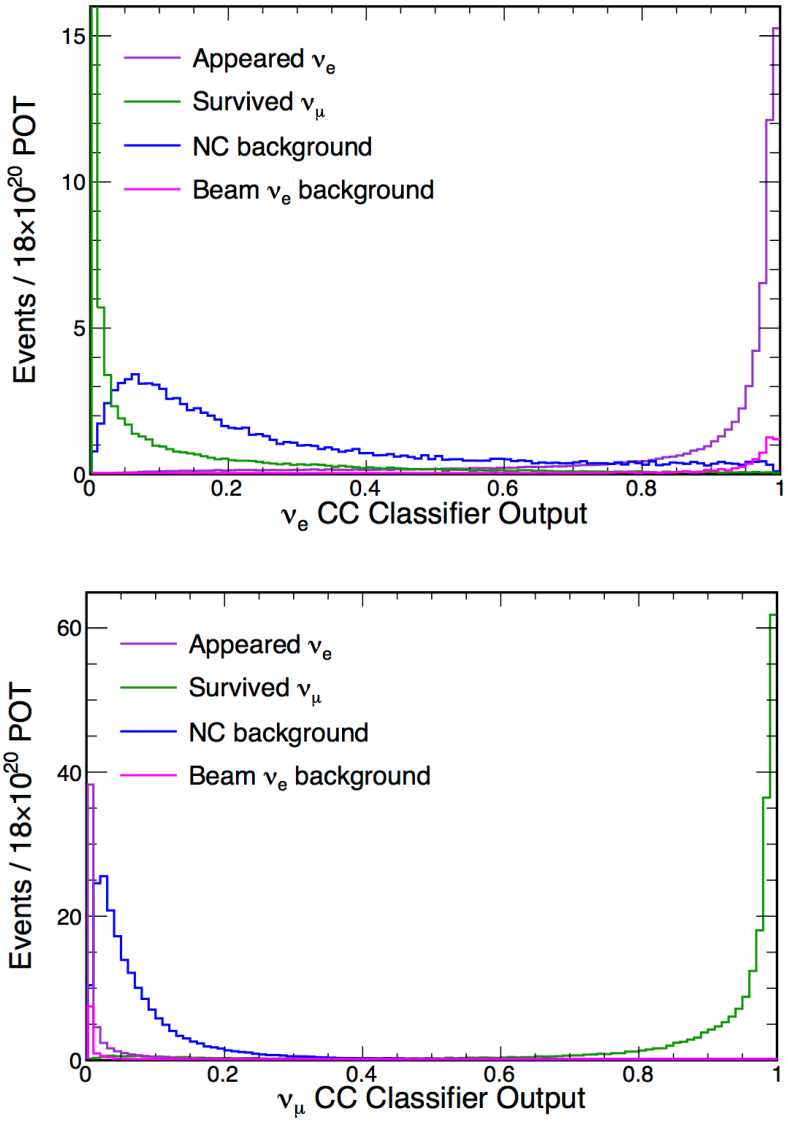

Figure 5.26: PID values for main analyses signal and background. In this case only one score from the output vector is used for selection. The ν_e CC on the left and ν_μ CC on the right. This comparison contains simulated events only, for which the most prominent inefficiency comes from the neutral current (blue) and beam ν_e (pink) backgrounds for the ν_e CC events, as expected given the topology of these events discussed in Section 6.1

mented [80], with respect to a selection based on an electron dE/dx likelihood identifier [65]. The introduction of CVN for selection of ν_μ CC events represents a 10% increase in exposure for the analysis presented here, with respect to a muon identification kNN-based selection utilizing muon-like track characteristics. The details of this selection will be discussed in Chapter 6.

	CVN Selection Value	ν_e sig	Tot bkg	NC	ν_μ CC	Beam ν_e	Signal Efficiency	Purity
Contained Events	–	88.4	509.0	344.8	132.1	32.1	–	14.8%
$s/\sqrt{b}$ opt	0.94	43.4	6.7	2.1	0.4	4.3	49.1%	86.6%
$s/\sqrt{s+b}$ opt	0.72	58.8	18.6	10.3	2.1	6.1	66.4%	76.0%

Table 5.1: This table shows relative selected event numbers for the various components of the NuMI beam, efficiency, and purity for two different optimizations for the selection of appearing electron neutrino CC interactions. Efficiency is shown here relative to the true contained signal. The numbers are scaled to an exposure of 18×10^{20} protons on target, full 14-kton Far Detector.

	CVN Selection Value	ν_μ sig	Tot bkg	NC	Appeared ν_e	Beam ν_e	Signal Efficiency	Purity
Contained Events	–	355.5	1269.8	1099.7	135.7	34.4	–	21.9%
$s/\sqrt{b}$ opt	0.99	61.8	0.1	0.1	0.0	0.0	17.4%	99.9%
$s/\sqrt{s+b}$ opt	0.45	206.8	7.6	6.8	0.7	0.1	58.2%	96.4%

Table 5.2: This table shows relative selected event numbers for the various components of the NuMI beam, efficiency, and purity for two different optimizations for the selection of surviving muon neutrino CC interactions. Efficiency here is shown here relative to the pre selected sample. The numbers are scaled to an exposure of 18×10^{20} protons on target, full 14-kton Far Detector.

The performance of the network was also evaluated for an improved architecture (shown in Appendix A) earlier this year. This short network (half in weight depth) was trained independently for forward and reverse horn current MC (where FHC is mostly neutrino and RHC is mostly antineutrino events). The network shows comparable purity and efficiency on FHC and up to 10% improvement in efficiency for the RHC trained network.

In addition to the information that can be extracted from the features seen in Fig.5.26 and Fig.5.25, an interesting tool to visualize the performance of these networks is the t-SNE [83]. The t-SNE

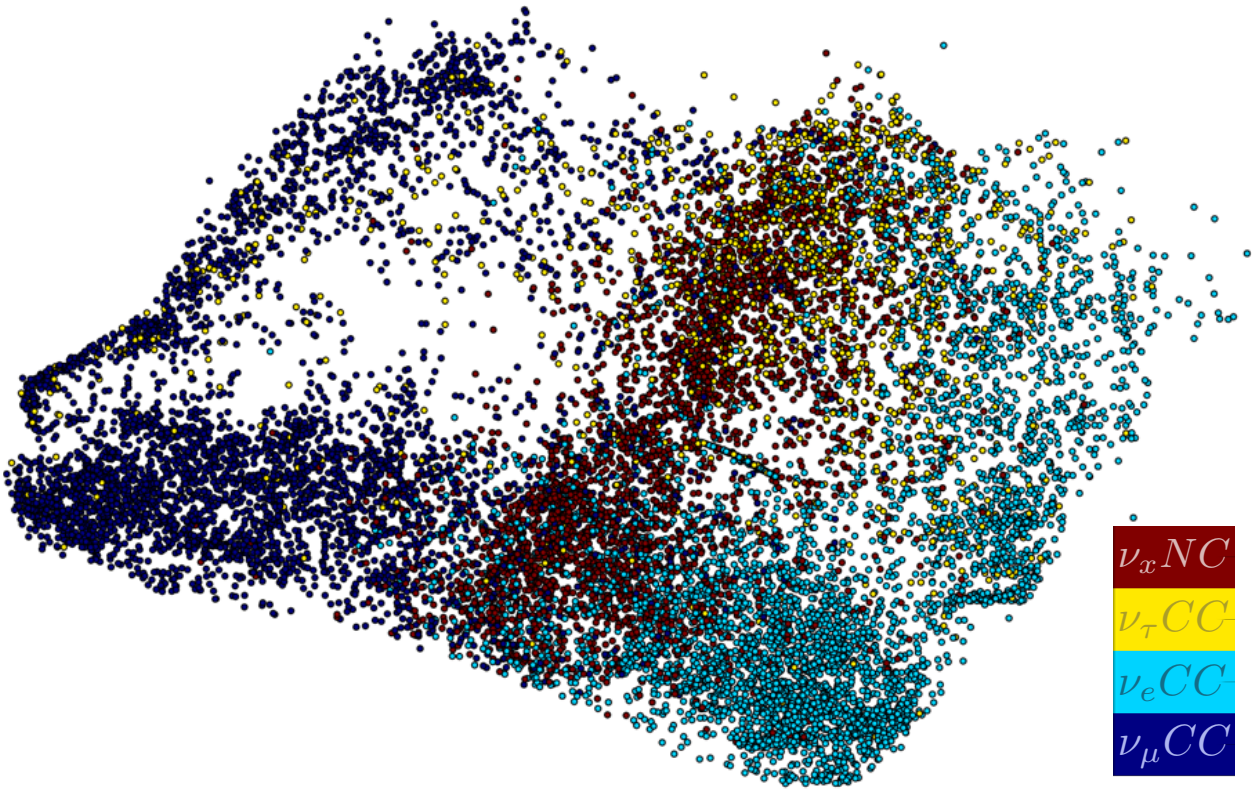

Figure 5.27: A t-SNE visualization of EventCVN. Each color represents one of the main categories.

method (which stands for t-distributed Stochastic Neighbor Embedding) is a method of dimensionality reduction which converts high dimensional Euclidean distances (like the distance between two classes of events in a classifier) into two dimensional probabilities which represent similarities. In essence, this method displays a set of classified events in a two dimensional space where the 2D separation between points corresponds to the ability of the network to separate them in a higher dimensional space. A t-SNE graphic for the event classifier can be seen in Fig.5.27, where some overlap is observed in expected categories, such as ν_e CC (light blue) and NC (brown), and good separation is observed by topologically distinct events, such as ν_e CC and ν_μ CC (dark blue).

5.2.3 Particle Classification with CVN

Another task for which we employ CNNs on NOvA is that of classifying individual particle contributions to the event. This process is especially important in order to be able to use other reconstruction objects like BPF energy estimates and especially for cross sections analyses searching for specific final states. Applications for energy reconstruction in particular are detailed in Sec. 5.4 and 5.5 and studies of event selection for cross sections analyses are ongoing within the collaboration.

The implementation of the particle classifier, called ParticleCVN or ProngCVN, is similar to the EventCVN described in the previous section, with the following improvements:

Multi-prong siamese architecture. While it is the individual particle contribution which will be classified by this network, the complete neutrino interaction is provided to the network as well, as it contains useful information which serves as context for the individual particles. The particle classification network uses four views in total, two of the full event and two of the individual prong with the additional hits removed, as shown in Fig.5.28. This requires an extension of the siamese architecture into a network with two branches depicted in Fig.5.29.

Further Network Optimizations. The details of the four branch network trained for ParticleCVN can be found in Appendix A. The main optimizations are a reduction of the number of inception modules as well as a reduction of the number of weights, making it a shallower network than that used by EventCVN.

Coupling to existing reconstruction. ParticleCVN is not completely decoupled from traditional reconstruction. It employs the FuzzyK clusters described in Sec. 5.1.3. These clusters are constructed and matched between views by traditional reconstruction upstream of the classifier.

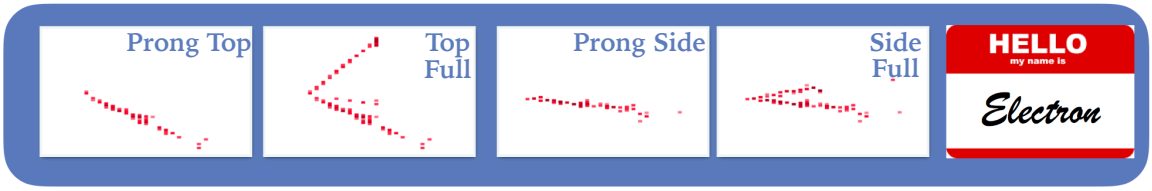

Figure 5.28: ProngCVN Input: Two views of the full event, two views of the FuzzyK prong and a particle label.

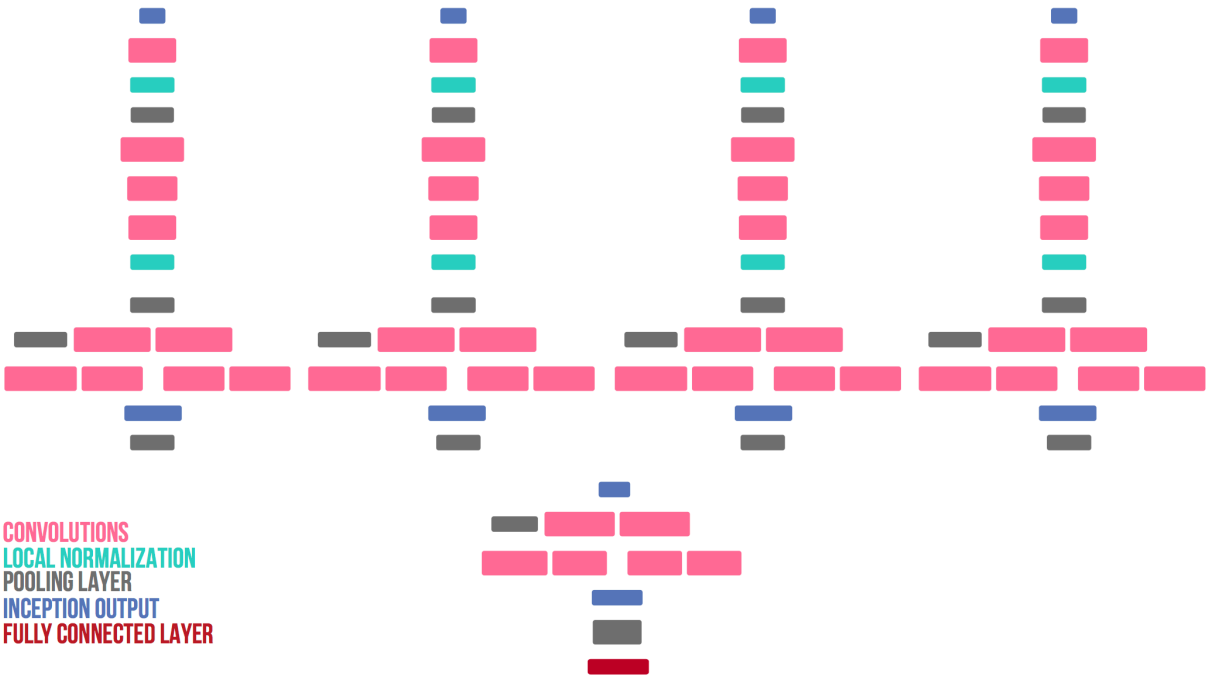

Figure 5.29: ProngCVN Network: A four branch siamese-type network with one inception module per branch. Full architecture details are in Appendix A.

The main impact of coupling to upstream reconstruction comes from the contamination that exists in prongs, that is, hits which do not correspond to the same particle. Fig.5.30 shows a zoomed-in version of one such case, where prongs are close enough to the vertex that it is difficult to disentangle individual components. This will cause prongs to be incomplete (not contain all the hits from one particle) or impure (contaminated with hits from other particles).

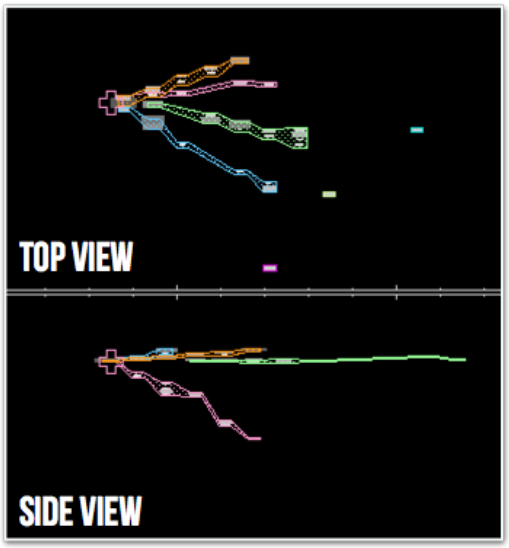

Figure 5.30: Overlap in FuzzyK prongs.

ParticleCVN classifies individual prongs by the type of particle which produced most of the hits in it. Fig.5.31 shows the prong purity, which is defined as the largest fraction of energy in a prong corresponding to the same particle. From the figure it is clear that protons and pions will be the least well defined, which is expected given that the short tracks they make will stay close to the vertex, where most of the inefficiencies are concentrated. Interestingly, the distributions for pion and proton prong purity in Fig.5.31 peak below 50%, a feature that arises from either incorrect matching of 2D prongs into a 3D object or—most prominently— contamination from other particle contributions to the prong, both of which are more common near the event vertex.

The selection of prongs used to train the network passed the following cuts:

Minimum prong purity: This cut selects only prongs above some threshold purity which is chosen to be inclusive of a similar fraction of the total number of prongs in each category. This is 35% for protons and pions, and 50% for electrons, muons, and photons.

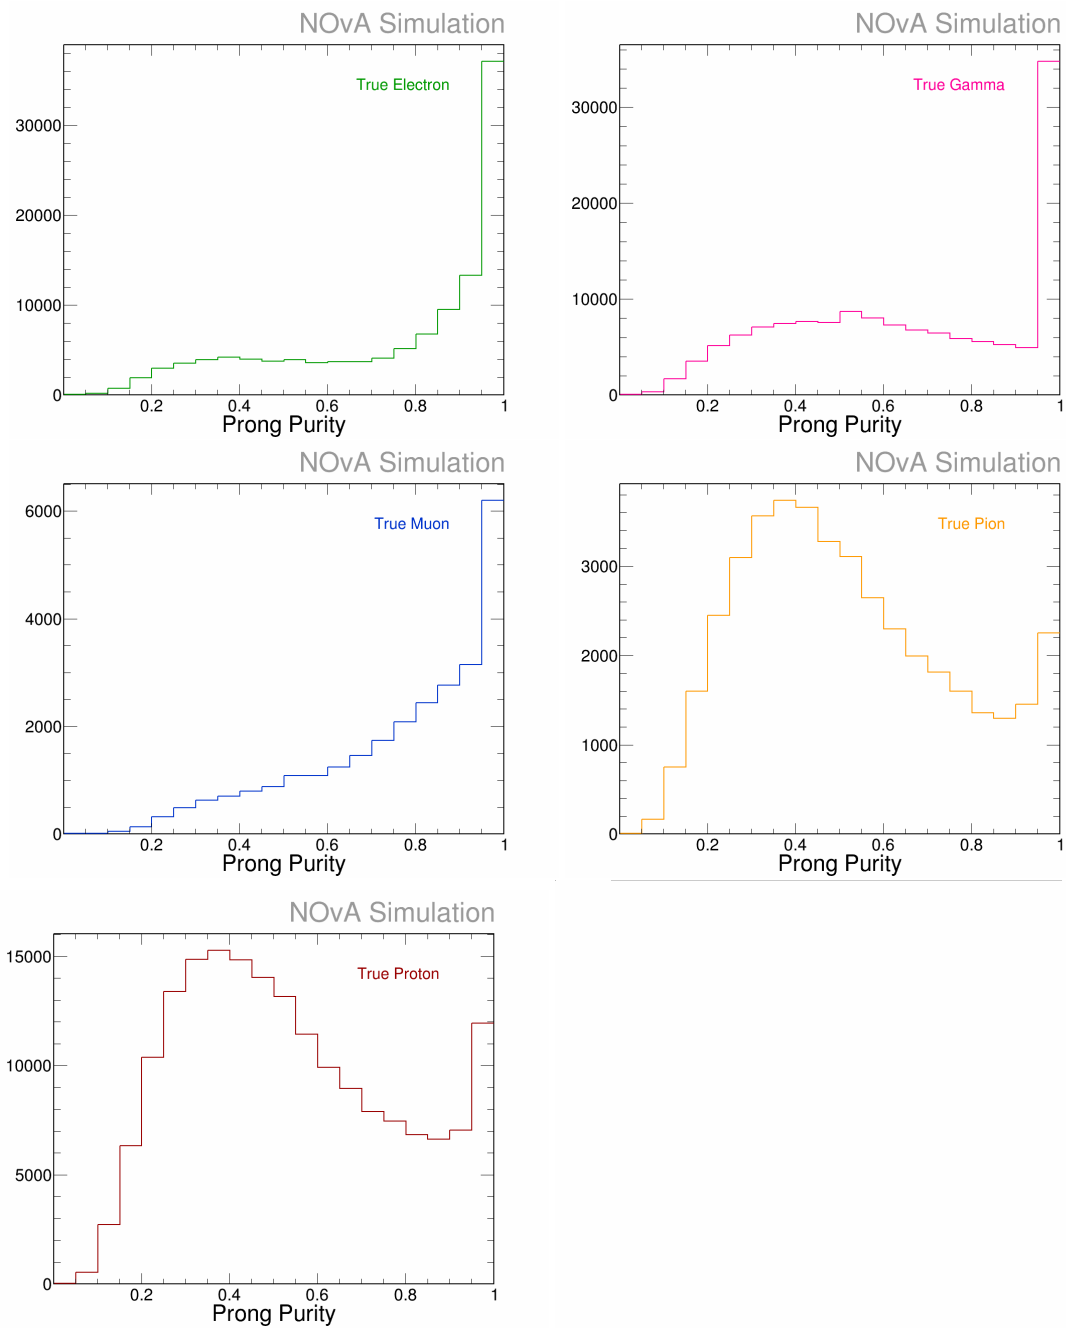

Figure 5.31: FuzzyK prong purity by particle type.

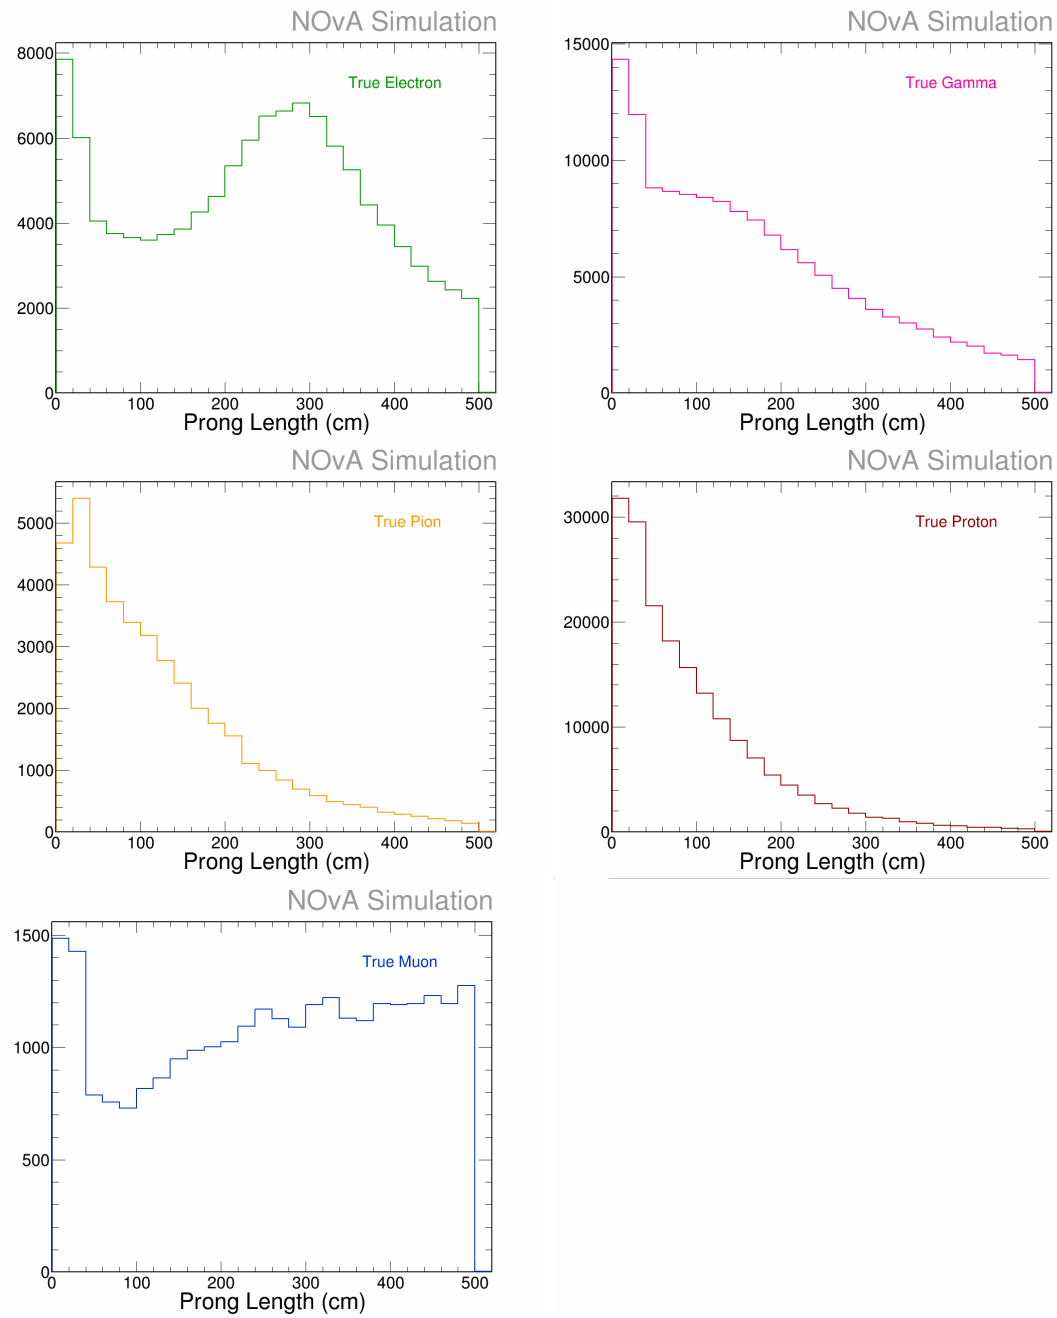

Figure 5.32: FuzzyK prong length by particle type.

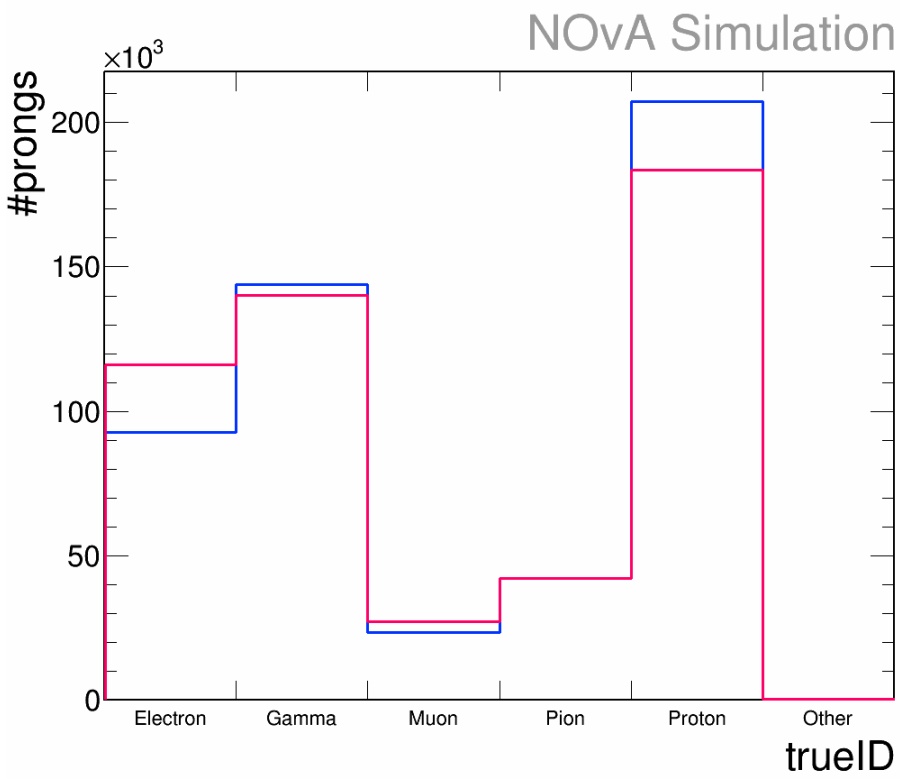

Figure 5.33: Input sample composition for Prong CVN (pink) and output score composition (blue).

Short prongs: Prong length below 500cm. This leaves out all long muon tracks, which have extremely low backgrounds from protons and pions, as seen in Fig.5.32.

ParticleCVN was trained on approximately 3.5 million prongs distributed into categories as shown in Fig.5.33 and coming from neutrino interactions from simulated events. Only prongs with better purity than a set threshold were used for training, where the threshold was chosen to retain a similar fraction of prongs of each type, with respect to the total. The network was trained and optimized on the Big Red II cluster at Indiana University.

NETWORK PERFORMANCE

While the network was trained on the five categories shown in Fig.5.34, the first implementations for energy estimation require only a distinction between electromagnetic and hadronic activity. Thus, the performance of ParticleCVN is evaluated from two slightly different standpoints; as a three-class network and a five-class network. The classification matrices and PID distributions shown in Figures 5.34 and 5.35 for the five categories, and in Figures 5.37 and 5.36 for the reduced categories, give an overview of the performance of the network for the different categories. Events are distributed in bins of true versus selected ID, where the largest of the scores has been chosen for the selection. The top and bottom matrices in each figure contain information from the same events, but are normalized to show efficiency and purity, respectively. Thus, the numbers in the top matrix are fractions of the true particles indicated for each column, while the numbers of the bottom (purity) matrices are fractions of the events selected as indicated for each row.

ParticleCVN is the first single particle classifier developed for NOvA events and, as such, comparisons with other algorithms are not available. However, performance for single electron identification compared with a technique based on likelihoods for transverse and longitudinal depositions (LID) is possible. The matrices shown in Fig.5.34, the performance for electrons is 81% purity and 65% efficiency, which yields a 7% increase in purity for the same efficiency from LID. The classification matrices also show other interesting features. For instance, Fig. 5.34 top shows misidentification at 10% for true electrons as identified as photons and 10% true muons as charged pions, both expected from the known dE/dx profiles of these particles, discussed in Sec. 3.2. The misidentification rate of charged pions as protons is 35%, which is not inconsistent with the lesser quality of the prongs corresponding to these particles, as discussed in the previous section. Despite the lesser performance for

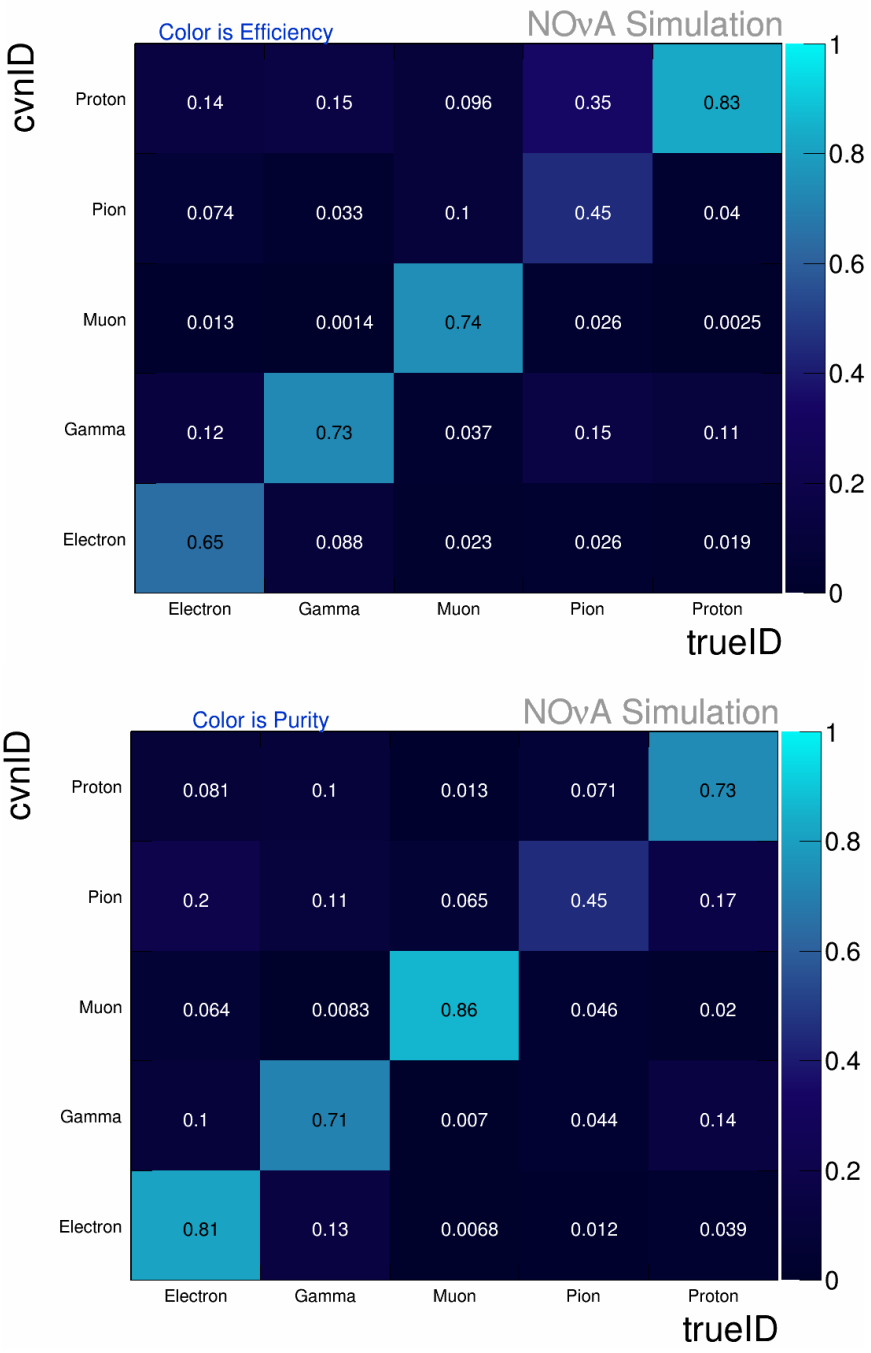

Figure 5.34: Classification Matrix for all classes trained for. Top: The diagonal elements are efficiency. Bottom: The diagonal elements are purity. The sample is FD MC neutrino events.

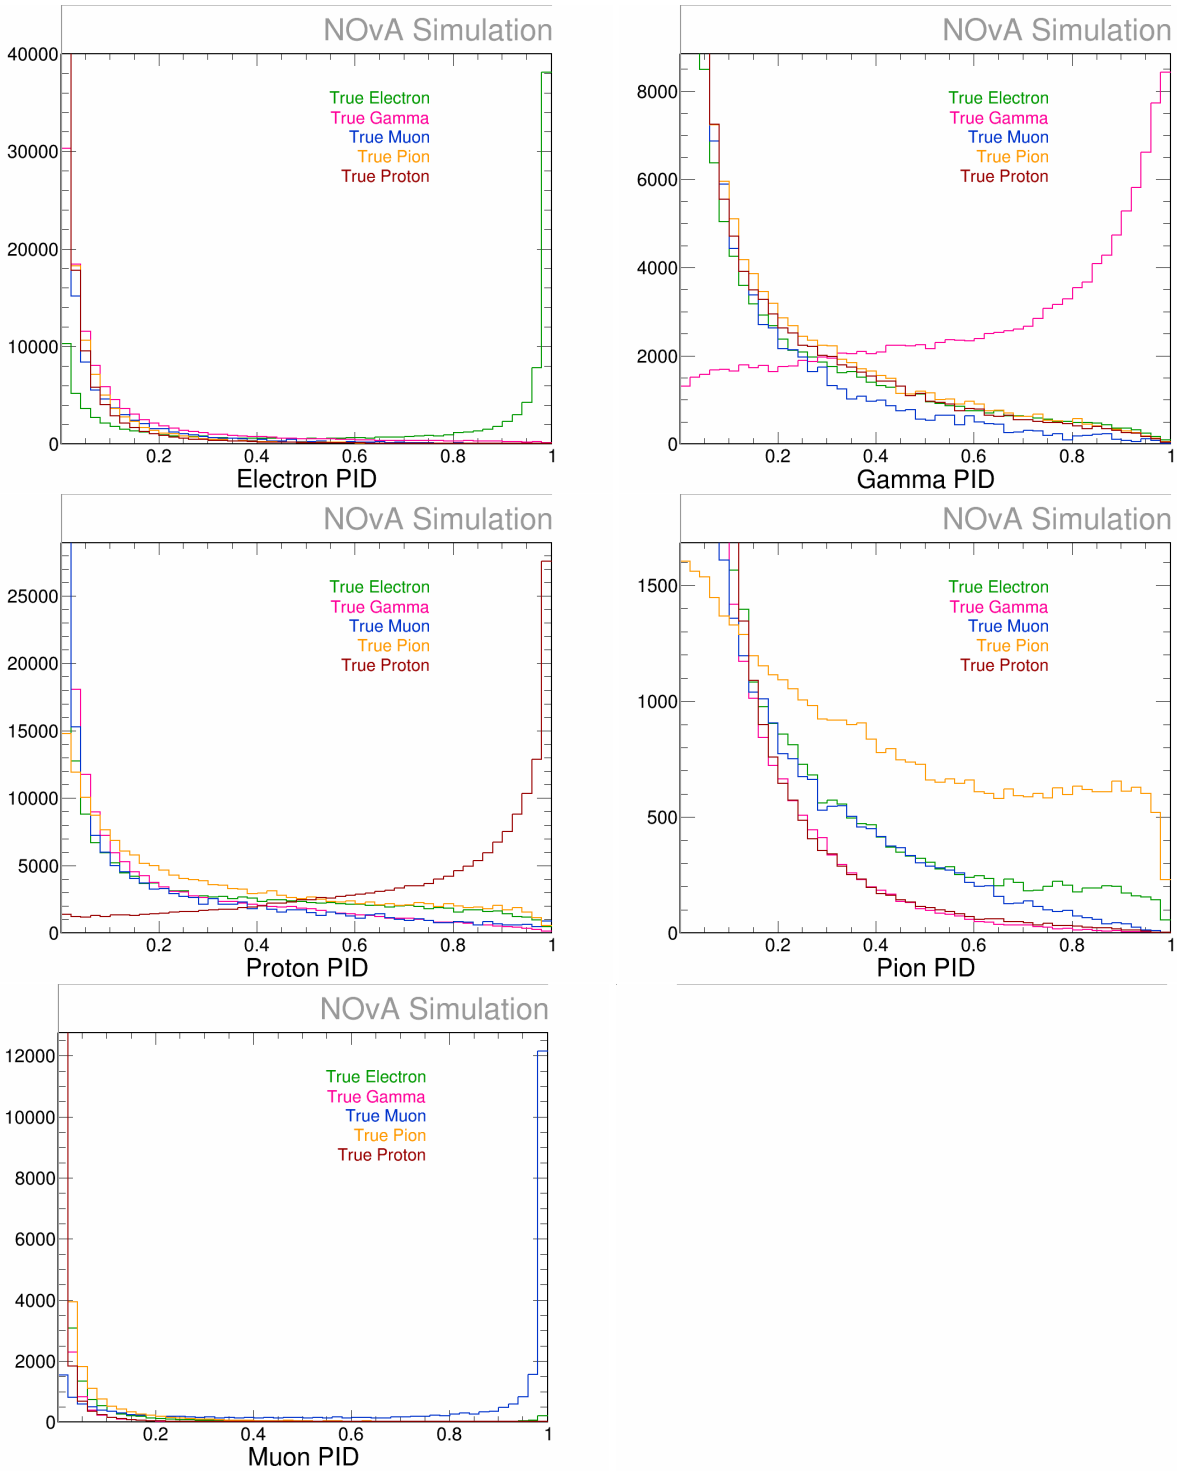

Figure 5.35: Prong CVN PID values for all categories. The sample is FD MC neutrino events.

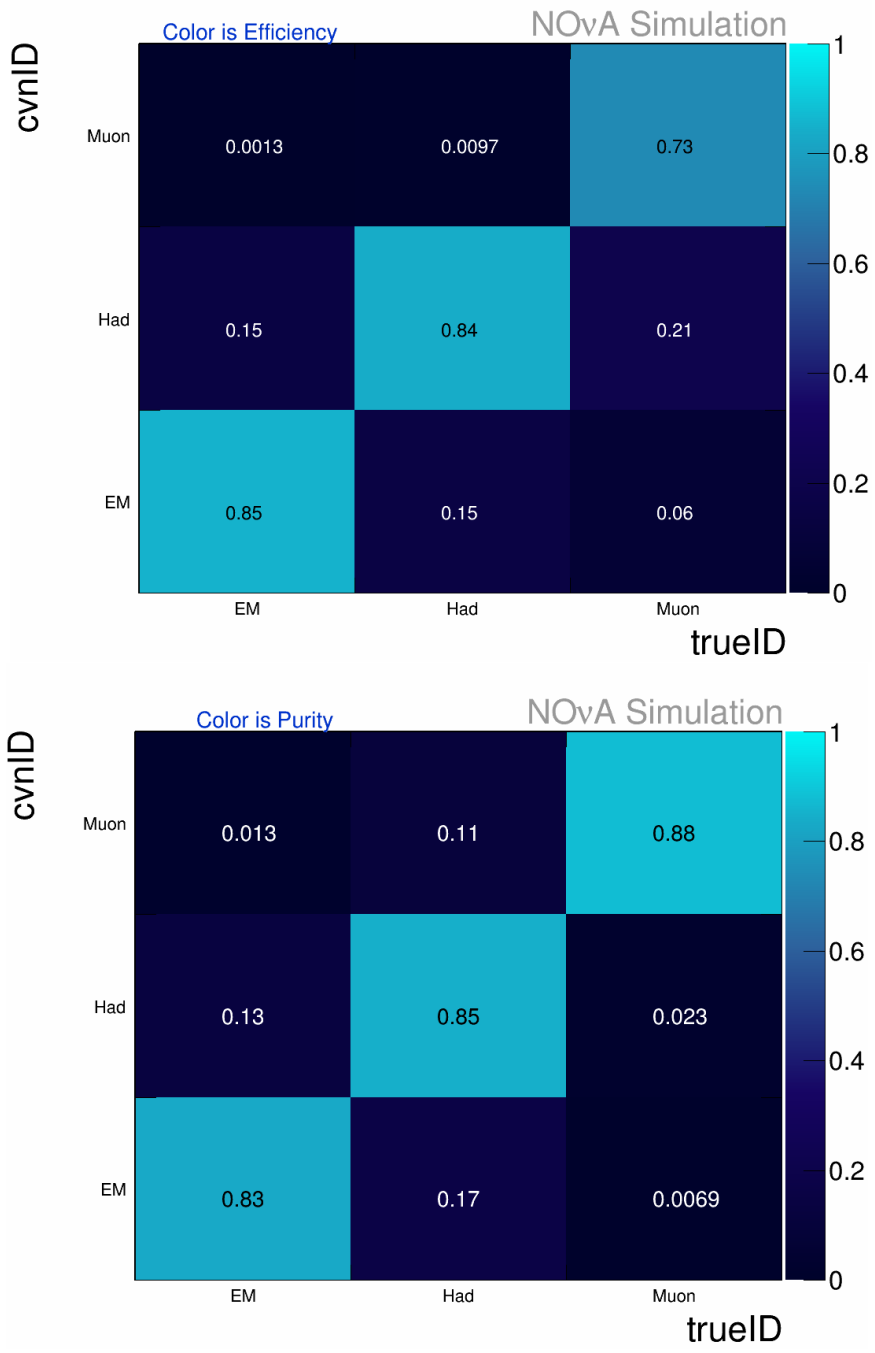

Figure 5.36: Classification Matrix for reduced labels. Top: The diagonal elements are efficiency. Bottom: The diagonal elements are purity. The sample is FD MC neutrino events.

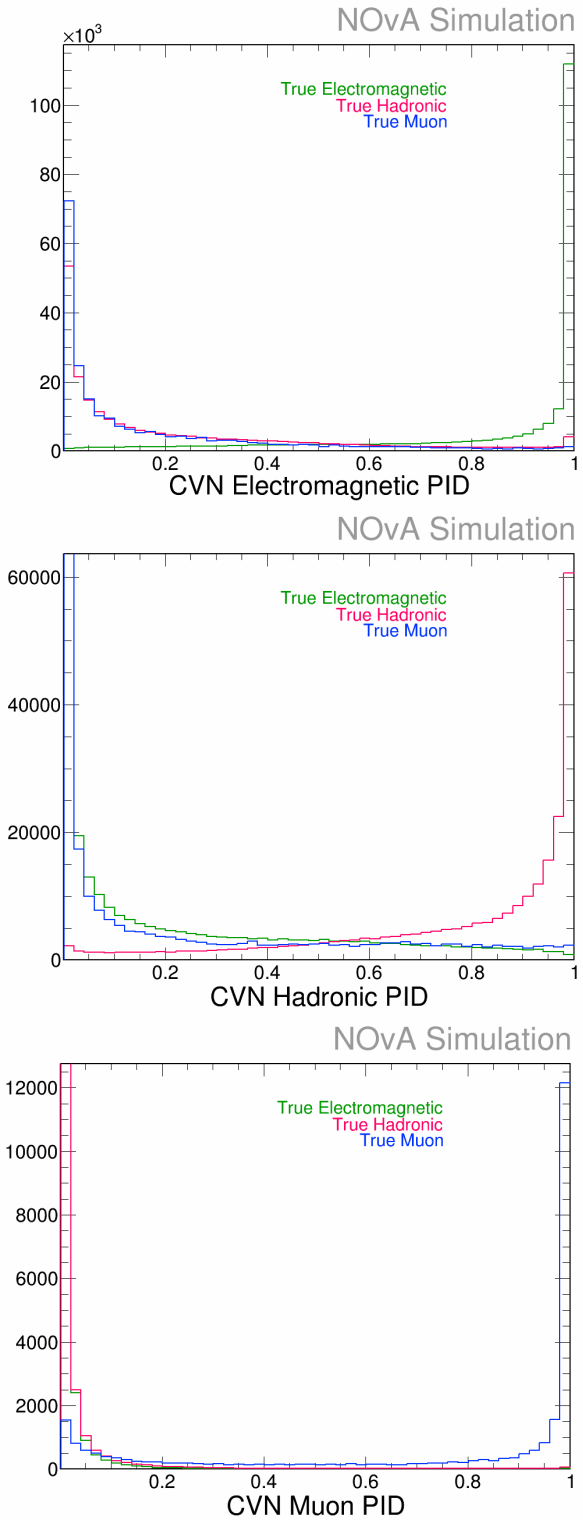

Figure 5.37: Prong CVN PID values for reduced categories. The sample is FD MC neutrino events.

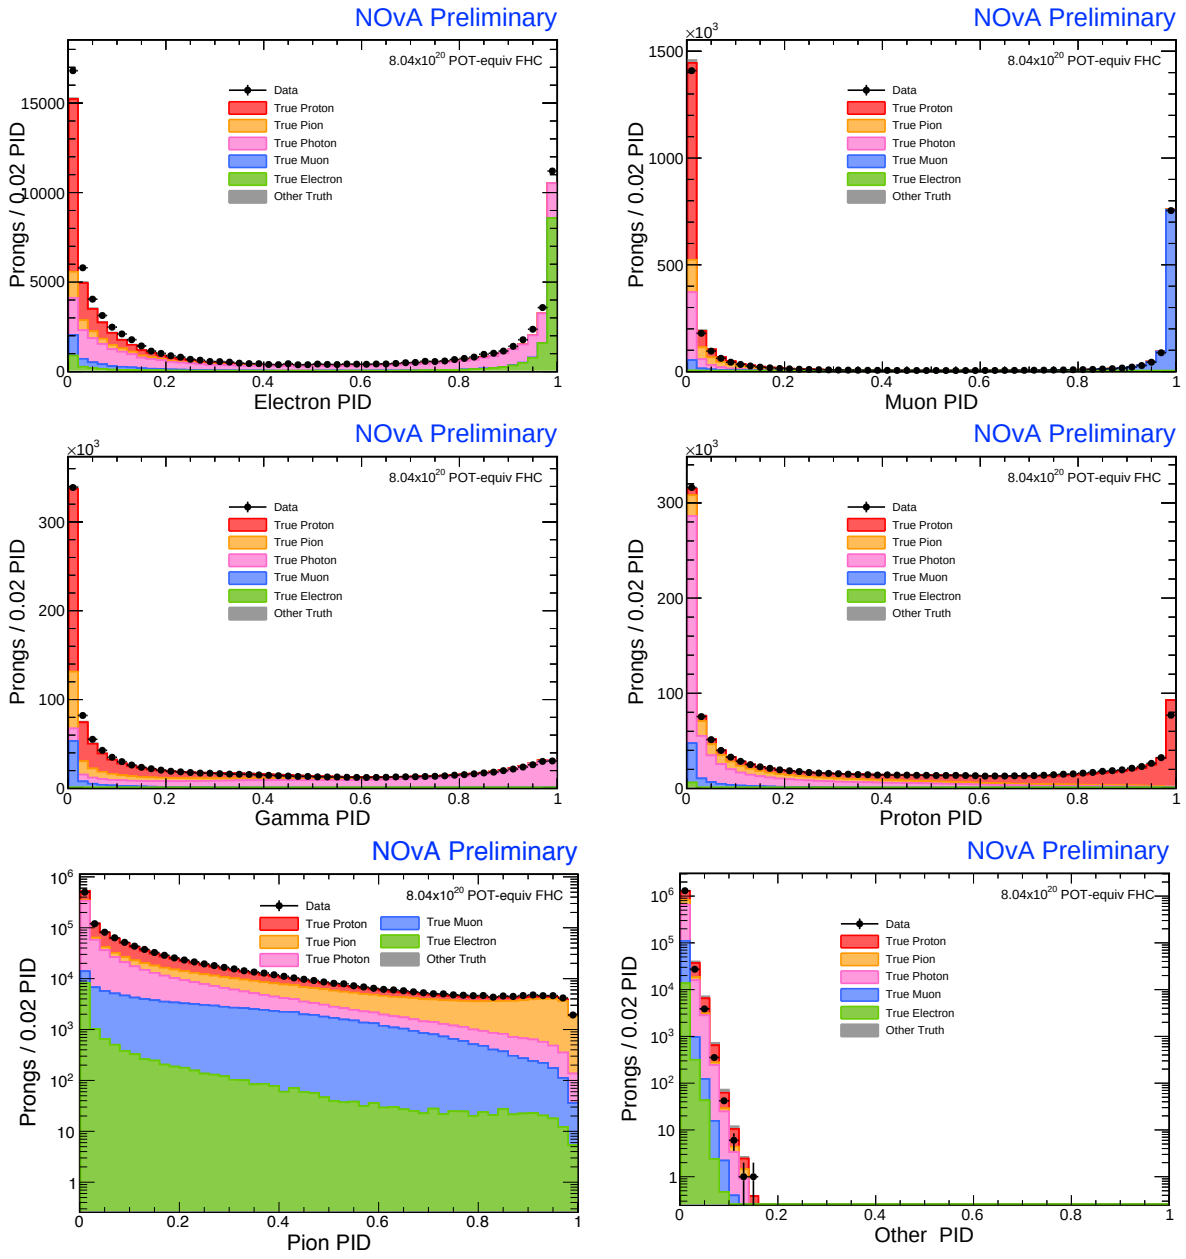

Figure 5.38: Particle CNN PID values for selected events in the ND. Top: $\nu_e CC$ selection (left) and $\nu_\mu CC$ selection (right). Middle and bottom: Neutral current event selection.

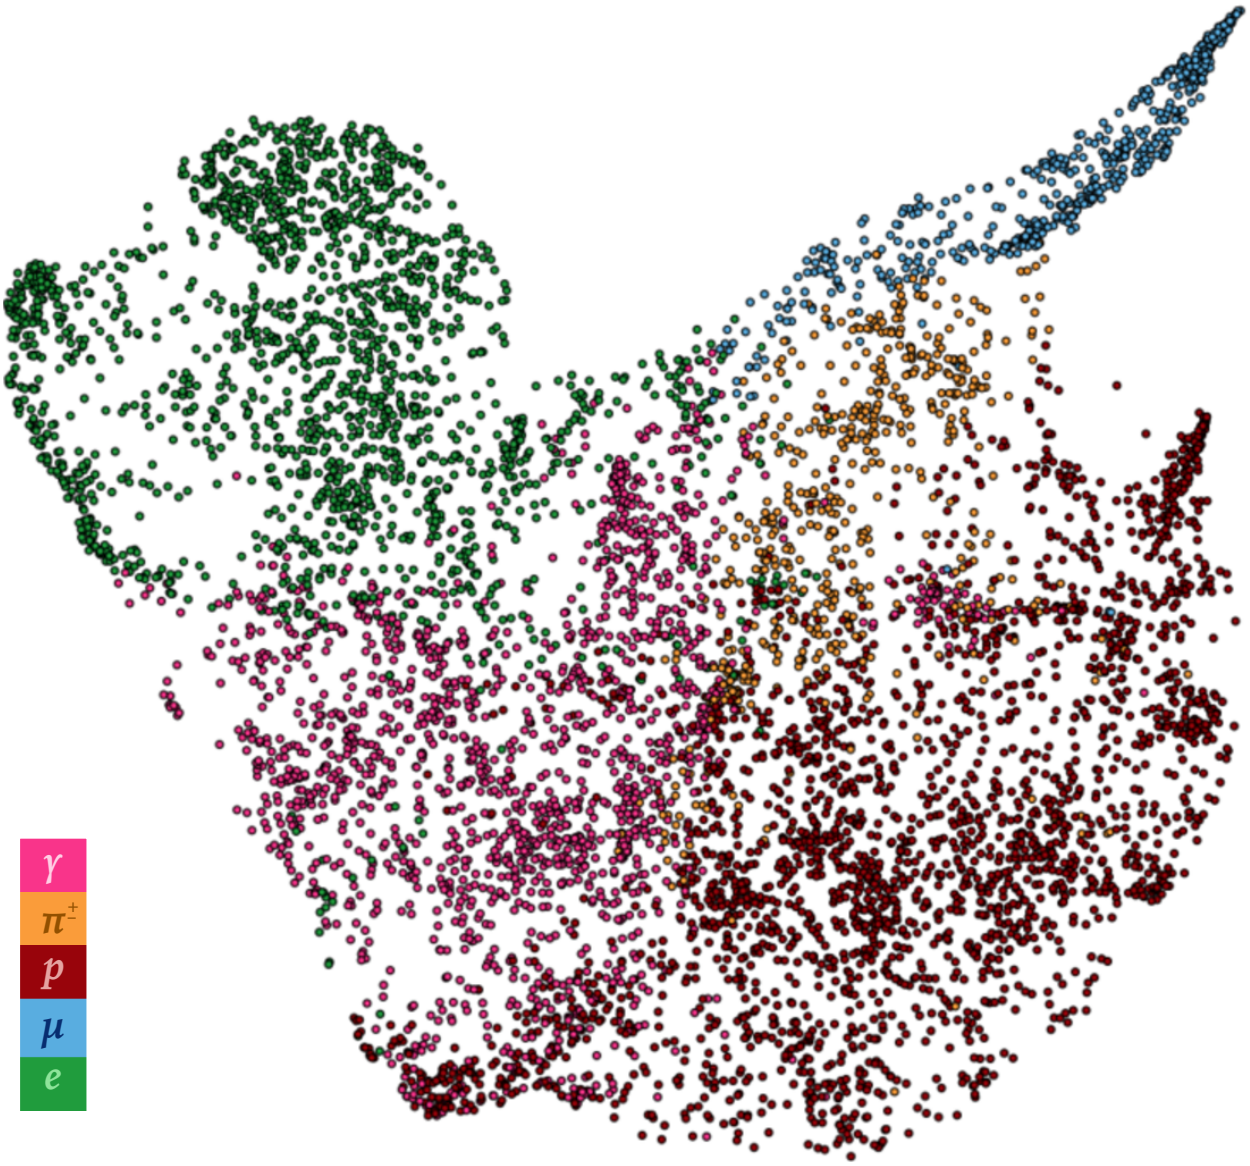

Figure 5.39: t-sne for cluster classifier. A clear separation can be seen between muons (blue) and electrons (green). The expected overlap is visible between pions (yellow) and protons (red). [83]

pion identification indicated by the classification matrices, the PID distribbutions shown in Fig. 5.35 suggest that a higher purity pion selection can be achieved with further optimization.

The PID values for both the full categories shown in Fig.5.35 also reflects the inefficiency for pions in the corresponding distribution. The efficiencies discussed can also be matched to the observed features in the t-SNE space in Fig.5.40, where it is visible that good separation is achieved for the leptons, protons and photons, also shown in the data/MC comparisons seen in Fig.5.35. Fig.5.40 also shows the expected overlap in this "separation power" space of the pion category, especially on the regions populated by protons and photons.

π^0 MASS WITH PRONGCVN SELECTION

One handle on both the network's performance in data and our detector energy response is the reconstruction of the π^0 invariant mass. This quantity is a common standard candle in many particle physics experiments given the abundance of π^0 production in particle interactions, its main decay channel to $\gamma\gamma$, and the well known mass of the π^0. The electromagnetic showers left by the photons from this decay are characteristically displaced from the interaction point (see Fig. 5.40) and, thus, expected to be some distance away from the reconstructed vertex in the event.

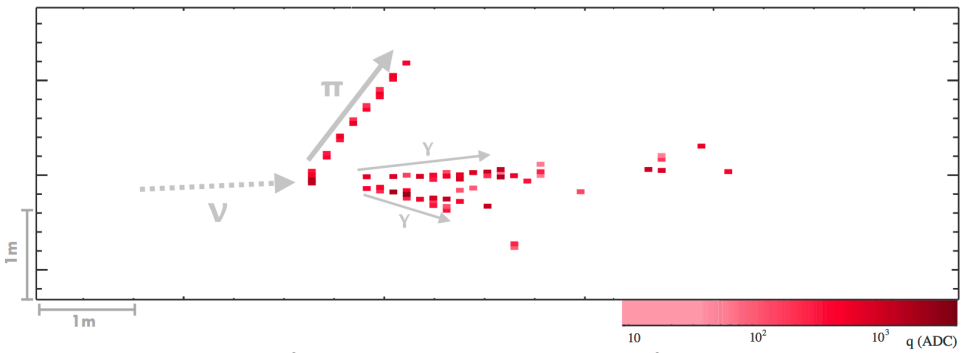

Figure 5.40: The event topology for a π^0 decay in a ND data event. The π^0 originates from the neutrino interaction and then decays promptly into two photons, which travel for one radiation length on average before pair producing.

The procedure consists of selecting events with two FuzzyK prongs consistent with two photon showers at high efficiency. Under the assumption that these two photons are produced in the same decay, the invariant mass of the mother particle is given by:

$$M_{inv.} = \sqrt{2E_{\gamma 1}E_{\gamma 2}(1 - cos\theta)} \tag{5.9}$$

where $E_{\gamma 1}$ and $E_{\gamma 2}$ are the energy of the first and second photons, respectively and θ is the opening angle between the two photons. For event reconstruction we take E_i as the sum of the calorimetric energy from the hits in prong, and the reconstructed opening angle ebtween the prongs as θ.

For a selection of $\pi^0 \rightarrow \gamma\gamma$ events with suficiently high efficiency a peak corresponding to $M_{\pi^0} = 134.9$ MeV, the π^0 invariant mass, will emerge in the distribution. This peak, however, will be initially located bellow M_{π^0}, as some of the energy will be lost as particles cross the inactive PVC regions of the detector. The necessary correction on the calorimetric energy to shift the peak to the correct value of M_{π^0} yields a measure of the fraction of the energy which is lost in "dead" material.

The original selection used for identification of events with two γs was used kNN based cuts to remove ν_μCC-like events, as well as a cut on $\frac{dE}{dx}$ to remove events with predominantly hadronic-like depositions.

The following selection was used to identify these events in ND data and MC:

Containment: Reconstructed vertex and prong activity at least one radiation length (40 cm) away from the front and back of the active region of the ND and half a radiation length from the sides of the detector.

Prong Quality : No gaps (planes without activity) within either prongs and no prongs crossing more

Selection	Dataset	Data	MC	Background	Purity	Efficiency	Dead Material
Old Selection	FHC	19560	20108	4072.5	0.797	0.244	0.357
ProngCVN Selection	FHC	17409	18268.4	1461.2	0.920	0.256	0.355
ProngCVN Selection	RHC	3701	3915.1	235.1	0.940	0.316	0.337

Table 5.3: The purity, efficiency, and dead material correction for the old selection, the CVN selection for FHC (neutrino mode), and the CVN selection for RHC (antineutrino mode).

than four planes per view.

ProngCVN: Photon PID value > 0.75.

The purity and efficiency of the two selections can be seen in table 5.3. The result is an increase in purity of about 12% over the old selection at the same efficiency. The cuts presented have not yet been optimized for efficiency, but were chosen to perform a comparison of prongCVN performance on data and MC with respect to the known selection. Hand-scanning of events revealed that the prongCVN selection at removing backgrounds was primarily from rejection of events with hadronic activity and events with showers starting near the vertex.

Fig.5.41 shows the corrected M_{π^0} peak, reconstructed with events selected with traditional selection as well as a selection based on the ParticleCVN output for neutrino and antineutrino running modes.

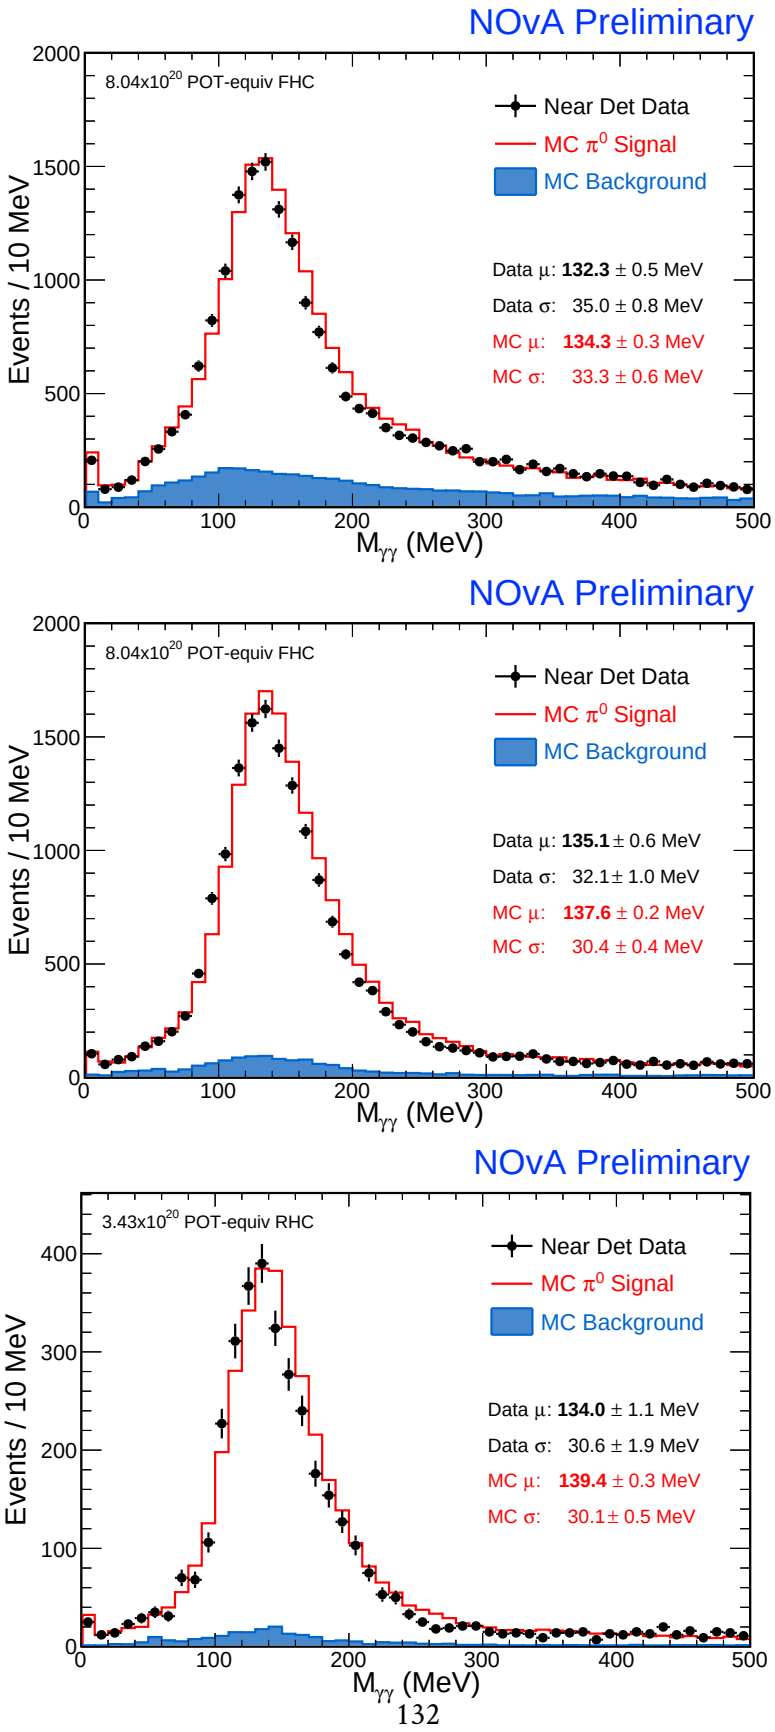

Figure 5.41: Corrected π^0 mass peak for the old (top) and prongCVN (middle and bottom) selection. The bottom plot shows events from reverse horn current (RHC) mode, thus, contains predominantly antineutrinos events.

5.3 Energy Reconstruction

The key measurable for neutrino oscillations is the energy of the incoming neutrino, E_ν, which interacts in the detector material. Since the only handle on E_ν are the energy and momenta of the neutrino daughters (see Sec. 3.1), a reconstructed energy is obtained from the depositions of these outgoing particles and our knowledge of how they interact in the detector described in Sec. 3.2.

In NOvA, the process of energy calibration makes corrections in order to standardize the light signal conversion into deposited energy. After calibration, we develop specific methods to reconstruct the energy of signal events by combining the reconstructed objects discussed in Sec. 5.1 and the identification methods discussed in Sec. 5.2. This section details the methods used to calibrate the energy response of the NOvA detectors in Subsections 5.3.1 and 5.3.2, as well as the methods to reconstruct the energy of signal events in 5.4 and 5.5.

5.3.1 Calibration

The response of the detectors to energy depositions depends on many factors such as the type of particle and its energy, the distance from the hits to the readout electronics, the light yield of the liquid scintillator, the quantum efficiency of the APDs, and the amount of dead material the particles traverse in the form of PVC. Thus, corrections must be made to account for these losses or inefficiencies introduced by all of these factors, in order to reconstruct the energy of the events of interest.

As discussed in Sec. 3.2, the energy deposition of muons at the energies seen in our detectors is well-described by the Bethe-Bloch curve shown in Fig.3.6. The high rate of cosmic ray muons in our detectors (approximately 150 kHz in the FD) allows for the use of this well known signature in order to get a handle on the corrections necessary for energy calibration. The two methods in which cosmic

ray muons are used for calibration are the *attenuation calibration*, a correction for the attenuation of the signals as a function of distance to the readout, and the *absolute energy calibration*, a determination of the energy scale to convert the signal in PE (photo-electrons) into deposited energy.

The attenuation calibration is a cell-by-cell correction using through-going muon tracks to avoid the stopping muon behavior explained in Fig.5.42. The correction function for each cell is calculated from the distribution of the signal in PE/cm vs distance from the readout. This is depicted in Fig.5.42 whose shape is given by:

$$y = C + A\left(e^{x/X} + e^{-(3L/2+x)/X}\right), \tag{5.10}$$

where L is the length of the cell, x is the distance away from the cell center, X is the attenuation length. Here, the factor $3L/2$ accounts for the fact that the light propagates in both directions along the fiber.

The behavior at the extrema of the function corresponds to light reflection near the ends of the cells. This behavior is corrected by the factor $1–\alpha(x–x_{end})$ where α is a factor found empirically for each end.

The absolute energy calibration is done by examining the dE/dx at the end of muon tracks across the detector. This is another well known distribution shown in Fig.5.42. By selecting the end of muon tracks and comparing with the expected behavior an absolute conversion from PECorr (attenuation corrected PE) to GeV of total deposited energy in the cell.

The selection of muon hits used for the absolute energy calibration is called *tri-cell*. These are triplets of adjacent hits along the muon's path. Tri-cell hits are also a handle on detection efficiency. The middle of three cells where hits are visible is expected to see a signal of magnitude correlated with the other two. Thus, a measure of efficiency comes from the comparison between the energy

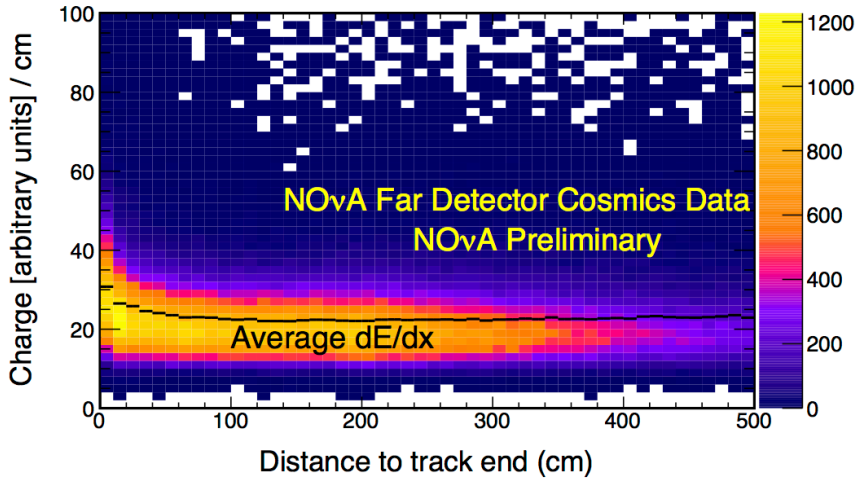

Figure 5.42: The charge deposited at the end of muon tracks as a function of the distance to the track ends. This well-known behavior of muon energy depositions is used for calibration.

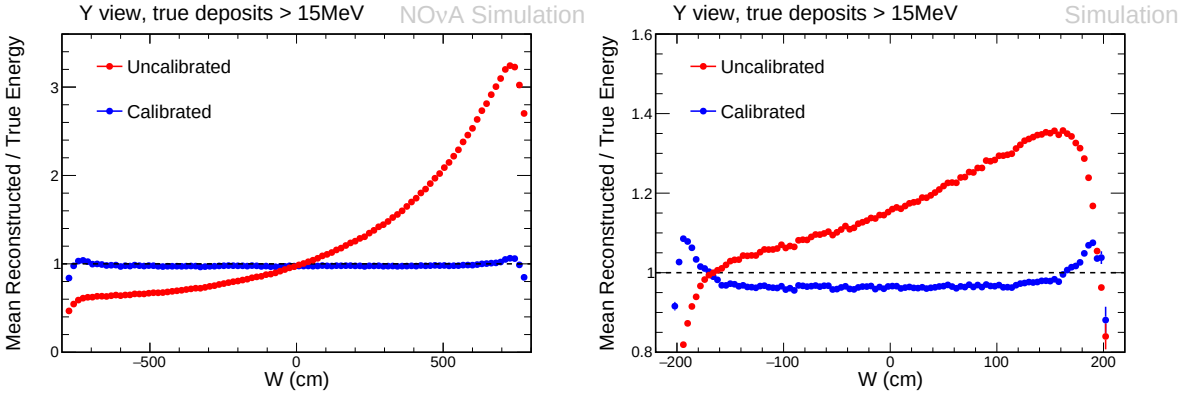

Figure 5.43: Calibration performance in terms of mean reconstructed/true energy as a function of position on the cell. The largest x value corresponds to the side closest to the readout for the FD (left) and ND (right). Note that the binning on right and left is different, with 5 bins-per-meter on the left and 20 bins per-meter-on the left.

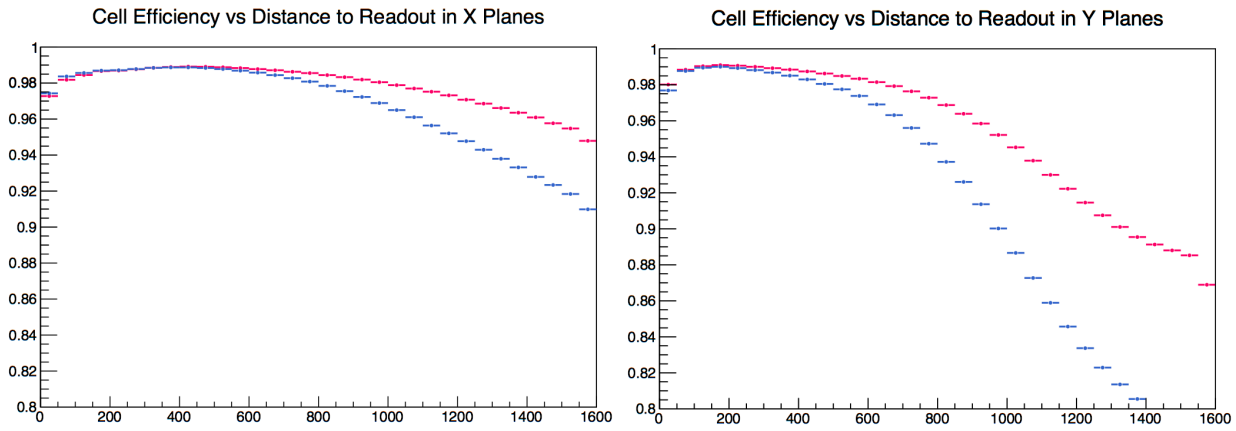

Figure 5.44: Efficiency as a function of distance to the readout electronics for horizontal planes (left) and vertical planes (right). The two distributions correspond to two gain settings for the APDs Blue: Gain 100. Pink: Gain 150.

deposition on the middle cell with the expectation given the other two hits. Fig.5.44 shows the change in this efficiency as a function of distance to the readout as well as the gain set on the APDs.

5.3.2 Energy Response

After the calibration conversion of raw signal into deposited energy, the task of estimating the event energy starts with an approximation of the energy losses as a fraction of the total deposited energy in the active regions. Given that the losses are small and well understood for electrons and photons as explained in Sec. 3.2, techniques like the π^0 invariant mass reconstruction can be employed to estimate the fraction of losses to dead material in the detector. Specifically, for the distributions shown in Fig.5.41, a dead material correction of 35.7% was found for the traditional (ν_μ anti-PID) selection. The correction for the CVN based selection was 35.5%.

A typical signal event has both electromagnetic and hadronic depositions, as seen in Fig.5.45. Thus, the conversion from the π^0 invariant mass reconstruction is appropriate for mostly electromagnetic depositions in the near detector, but it is not sufficient to make a correction for the whole event. When reconstructing the energy of our signal events and in particular in the case of charge

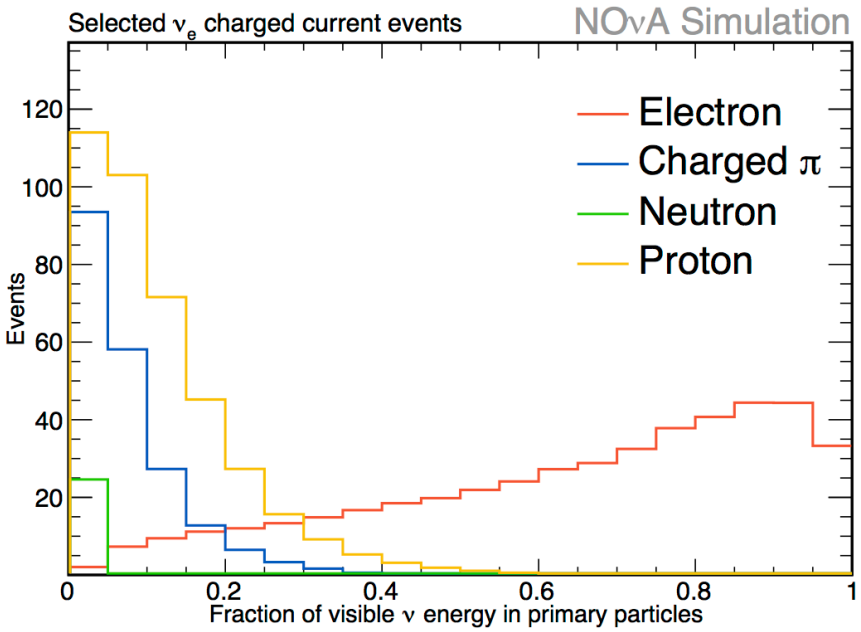

Figure 5.45: Left: The composition of simulated ν_e events in fraction of energy from different particles. Electrons (red), dominate the energy content, but protons (yellow) and charged pions (orange) can contribute up to 40% of the event's energy as well.

current interactions a different treatment is necessary for the leptonic and the hadronic components of the event due to their differences in energy losses.

The response of calorimeters, the fraction of energy deposited which is detected, is different for electromagnetic and hadronic depositions. The fractional difference between these responses is measured in calorimetry by the factor e/h, the ratio of electromagnetic to hadronic detector response. If e/h is known it can be used for energy reconstruction, provided that the fraction of electromagnetic depositions in the event f_{em} is also known.

Given that our reconstruction approaches must deal with both the hadronic and electromagnetic components of the neutrino interactions in the same slice a study was done to focus on determining the detector response as a function of f_{em}.

In this study, f_{em} was defined for each interaction as:

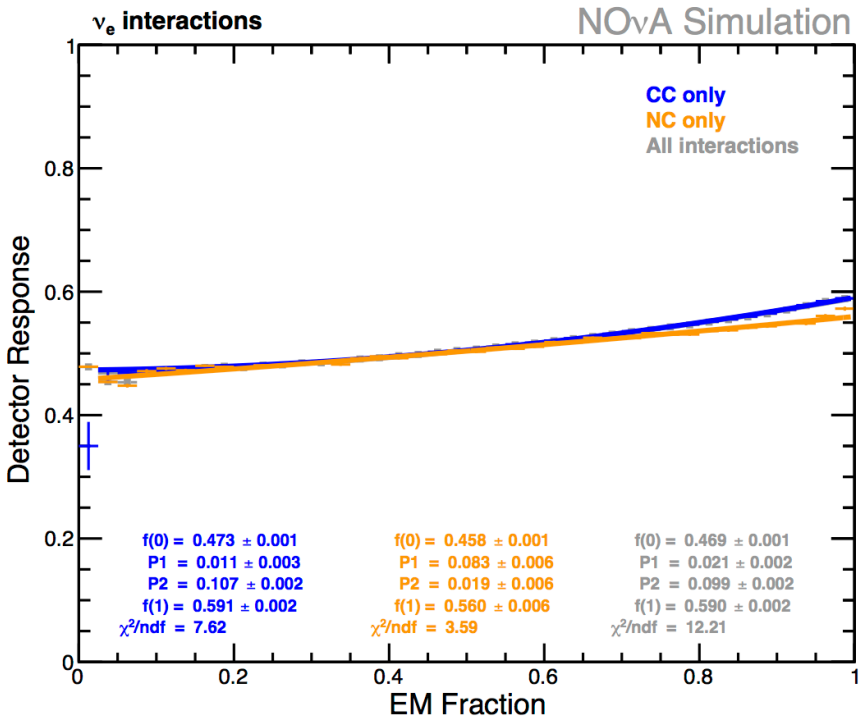

Figure 5.46: Detector response as a function of f_{em}. The e/h fraction is calculated from the value of this distribution at 0 and 1.

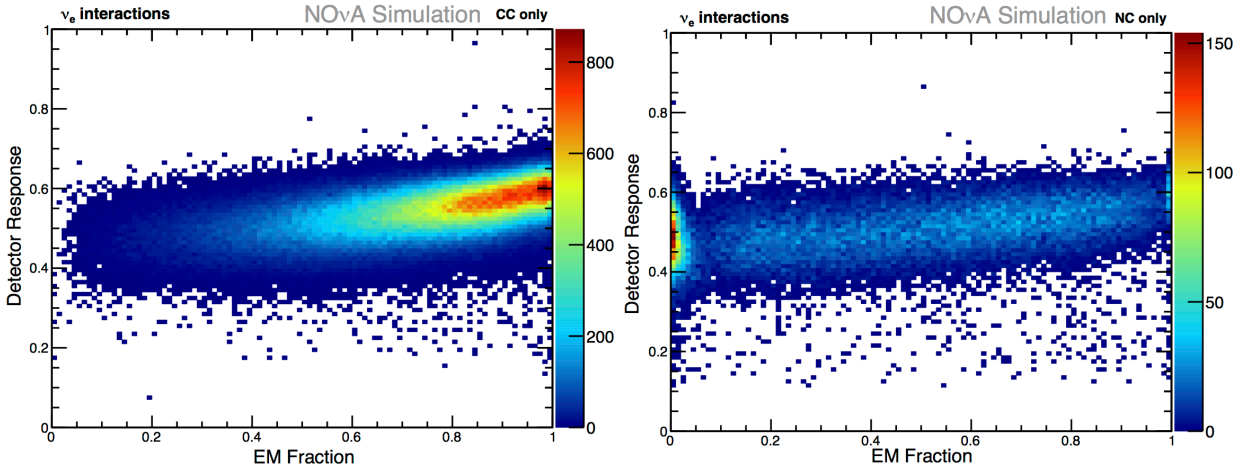

Figure 5.47: Contributions to detector response as a function of f_{em} from CC and NC interactions.

$$f_{em} = \frac{\text{Energy deposited by electrons or photons}}{\text{Total energy deposited in the interaction.}} \tag{5.11}$$

e/h was calculated from the ratio of the detector response, or fraction of true visible energy seen in deposited energy for f_{em} = 1 and f_{em} = 0. Where events with f_{em} = 1 are mostly electromagnetic and f_{em} = 0 are mostly hadronic, shown in Fig.5.46.

The study was done for simulated electron neutrino interactions in the far detector after containment and basic quality pre-selection as well as slice efficiency (see Sec. 5.1.1) above 90%.

The value obtained from this study was e/h = 1.26 $\pm$ 0.02.

5.4 Electron Neutrino Energy Reconstruction

The approach at reconstructing the energy of signal events for the ν_e appearance analysis is based on the knowledge that, as discussed in the previous section, the response of the detectors is inherently different for electromagnetic and hadronic depositions. Both of these are present in our signal events. While the reconstruction algorithms detailed in Sec. 5.1 aim to separate the contribution of each particle into individual clusters, the inefficiencies in such process and the inefficiencies in identification prevent us from simply applying e/h as a correction for each component of the event energy.

The energy reconstruction developed for the ν_e signal events aims to account for these inefficiencies while also accounting for both types of energy depositions separately.

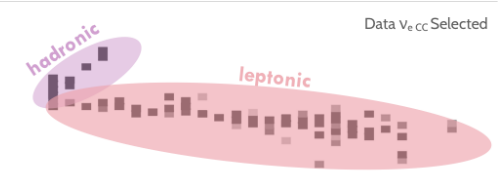

Figure 5.48: ν_e **Energy Diagram**

The ParticleCVN classifier is used to identify prongs as hadronic or electromagnetic. This is done by taking the reduced categories shown in Fig.5.34 and assigning a label (hadronic or electromagnetic) corresponding to the largest CVN score for the prong. Clusters are classified into electromagnetic or

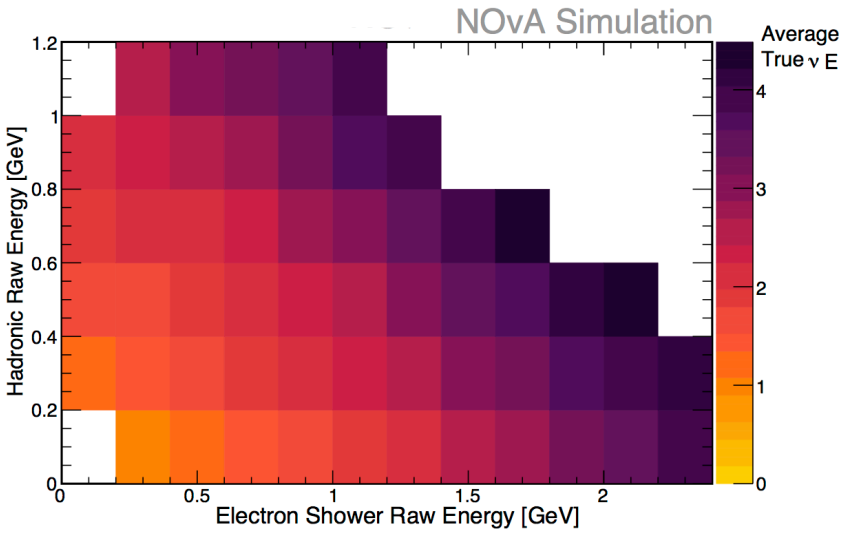

Figure 5.49: ν_e **energy reconstruction is done with the function obtained from fitting a quadratic function to this distribution of hadronic energy (**y **axis) vs electromagnetic energy (**x **axis). The color is the average true energy per bin.**

hadronic and the rest of the energy in the event (coming from un-matched prongs or unclustered hits) is identified as hadronic.

Rather than reconstructing the hadronic and electromagnetic components independently, in order to reconstruct the energy of signal events we make a fit to true neutrino energy as a function of the reconstructed electromagnetic and hadronic energy shown in Fig.5.48. In this case this is the sum of the calibrated energy deposited in the prongs or hits identified for each kind.

A quadratic function is fit to the distribution on Fig.5.49 and the resulting function is:

$$E_{true} \approx \frac{1}{1+0.057}(0.996E_{EM} + 0.869E_{HAD} + 0.025E_{EM}^2 + 0.504E_{HAD}^2)\text{GeV}$$

To ameliorate the shaping of the resulting distribution to that of the training sample, the function is also corrected by weighting the contributions of events according to the energy spectrum of the sample used for training. This yields a worse resolution overall, but yields a more even resolution as a function of true energy, across the 1-4 GeV range of interest.

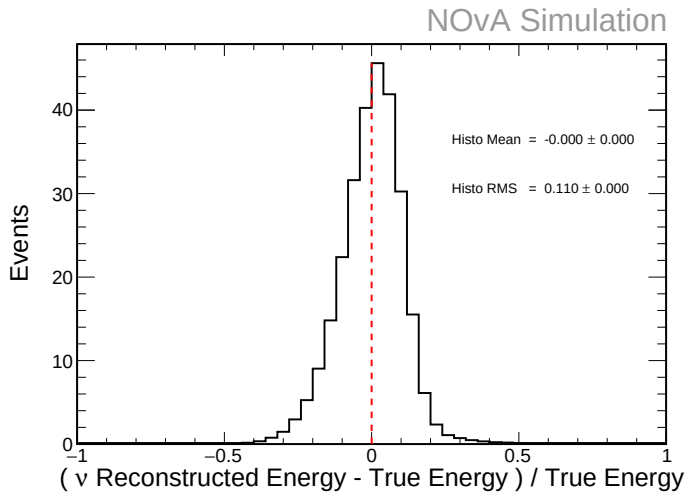

Figure 5.50: ν_e **Energy Resolution (reconstructed - true / true energy).**

This function is then used to reconstruct the energy of data and MC events. The performance of this fit can be seen in Fig.5.50.

5.5 Muon Neutrino Energy Reconstruction

For the ν_μ disappearance channel, the signal events are characterized by the presence of a muon. The presence of muons in these events represent an advantage from the standpoint of energy reconstruction given the well understood losses of muons at these energies, as described in Sec. 3.2. In this section two approaches at energy reconstruction are described, both separating the contribution of the muon for the rest of the event, but whose treatment of the remaining energy is different for each.

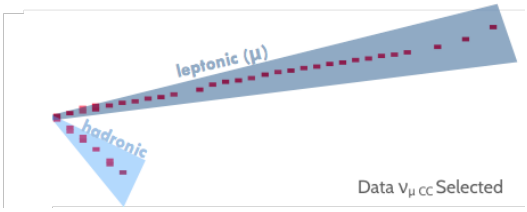

Figure 5.51: ν_μ **Energy Diagram**

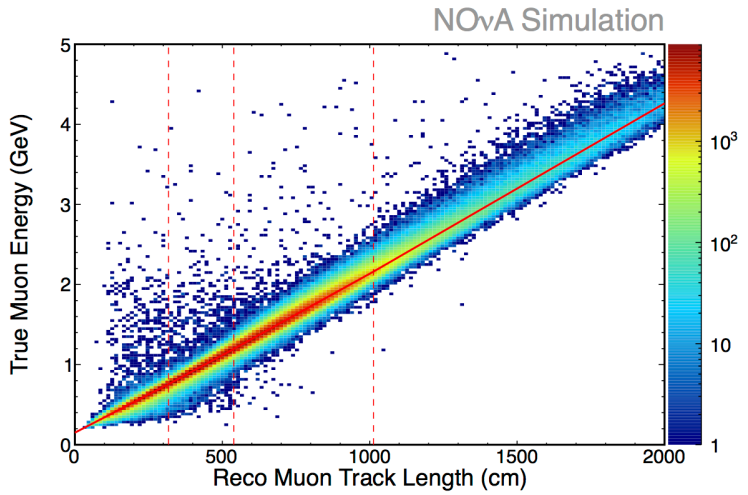

Figure 5.52: **Muon energy reconstruction is done with the relation from this MC distribution of true muon energy vs reconstructed muon track length.**

5.5.1 Muon Energy Reconstruction

The 0.5 to 5 GeV energy range of neutrino events on NOvA is within the minimum ionization region of muons on the Bethe-Bloch curve shown in Fig.3.6. This MIP behavior makes for well defined muon trajectories in the detectors. The length of muon tracks is related the true energy of the muons and is expected to be mostly linear. Fig.5.52 shows this relation in the MonteCarlo, which we use to reconstruct the energy of muons. The resolution of the muon energy is 3.5% with this simple method, which uses no information about the amount of energy deposited along the track for energy reconstruction. Two factors contribute to the small resolution of muon energies: the well understood deposition profile of MIPs, and the large rates of cosmic ray muons in the detectors, which are used for hit calibration prior to this reconstruction as described in Sec. 5.3.1. While the energy calibration does not contribute to this method for energy estimation directly, it does allow for better track reconstruction and identification of the muon hits.

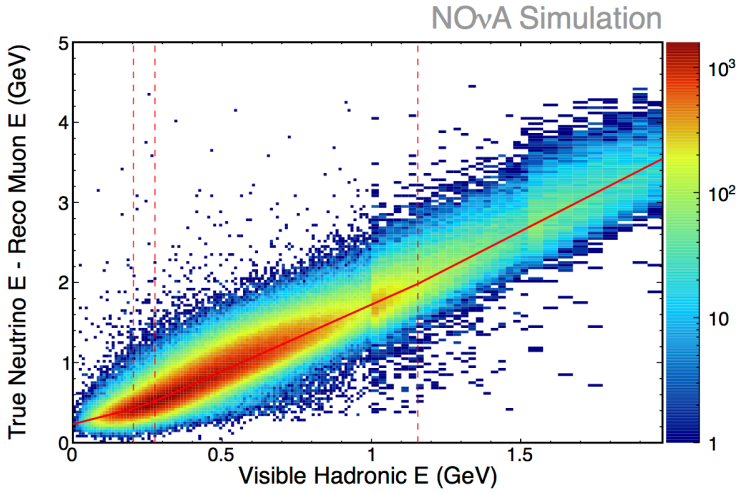

Figure 5.53: ν_μ **energy reconstruction is done with the relation from this MC distribution of true neutrino - reconstructed muon energy versus visible hadronic energy.**

5.5.2 Hadronic Energy Reconstruction

The simplest approach to reconstructing the energy of the neutrino is by fitting the visible deposited energy in the rest of the event to the missing energy in the neutrino interaction, once the muon has been accounted for and its hits removed. A linear fit is done to this distribution, as depicted in Fig.5.53. It is visible from the comparison of the distributions from Fig.5.52 and Fig.5.53, that the resolution in hadronic energy is substantially worse than it is for the muon track.

The energy resolution of the hadronic component is 30%, which yields a total neutrino energy of 11% overall.

The substantial difference in energy resolution between the two components can be attributed to both the well understood behavior of muons as MIPs, and the poorly reconstructed energy of protons and charged pions under the global reconstruction of all the hadronic energy. The energy deposition profiles for protons and pions are distinct, as is their scattering behavior. An improved reconstruction approach using ProngCVN and BPF is currently being developed to separate the contributions from

protons an pions, and employ the BPF energy estimate for the corresponding particle assumption.

Chapter 6

Analysis of Neutrino Oscillations

The analysis presented in the rest of this document is a measurement of oscillations in the $\nu_\mu \rightarrow \nu_e$ and $\nu_\mu \rightarrow \nu_\mu$ channels.

In the following sections, I describe the methodology for the ν_e appearance and ν_μ disappearance channels. The methodology which is common to both is described in Sec. 6.1, followed by the details of event selection and expected signal prediction for each channel. The treatment of systematic uncertainties is detailed in Sec. 6.6.

6.1 Methodology

Neutrino oscillation analyses require information of the signal spectrum in two locations. In the case of NOvA, this is a measurement of the ν_μ spectrum at the near detector, and a measurement of the spectrum of signal events at the far detector. Signal events for the $\nu_\mu \rightarrow \nu_e$ channel (or ν_e appearance) are ν_e CC (charged current) interactions, and for the $\nu_\mu \rightarrow \nu_\mu$ channel (or ν_μ disappearance) are ν_μ CC interactions, as described in Sec. 3.1. Given that the measurable is the neutrino energy, the main steps of the methodology are common to both analyses, as are the processing of the data, calibration, and basic reconstruction.

The common methodology for both channels is as follows:

i **Simulation** The flux delivered by the NuMI beam, the neutrino interactions and the detector response are all modeled by simulations.

ii **Measurement of ν_μ spectrum at the near detector.** This yields an energy spectrum of selected events, which is used as the initial flux in the analysis or $|\nu_\alpha >= e^{-iE_k t}$ at t = 0.

iii **Produce extrapolated prediction for the far detector.** Start from the measurement made from step *i*, make an assumption for its breakdown in different background components and extrapolate it to the far detector location to make a prediction.

iv **Uncertainty Calculation.** Systematic uncertainties are studied to account for detector response and other variables. See Sec. 6.6.

v **Measure the signal events ν_{signal} at the far detector.** For details of the selection for signal events see Sec. 6.3. The result for this analysis is shown in Chapter 7.

vi **Fit to model.** Fit the observed spectra to the PMNS model using known constraints to measure the parameters of interest. Details of the fit can be found on Sec. 7.2.

6.2 Simulation

Simulations are necessary to qualify our understanding of the underlying physics of our experiment. On NOvA, we use a suite of software and techniques to produce simulations of mainly the flux from the beam, the neutrino interactions, the propagation of particles through different materials in our detectors, and the detector response to these interactions. The following are the main pieces of our simulation:

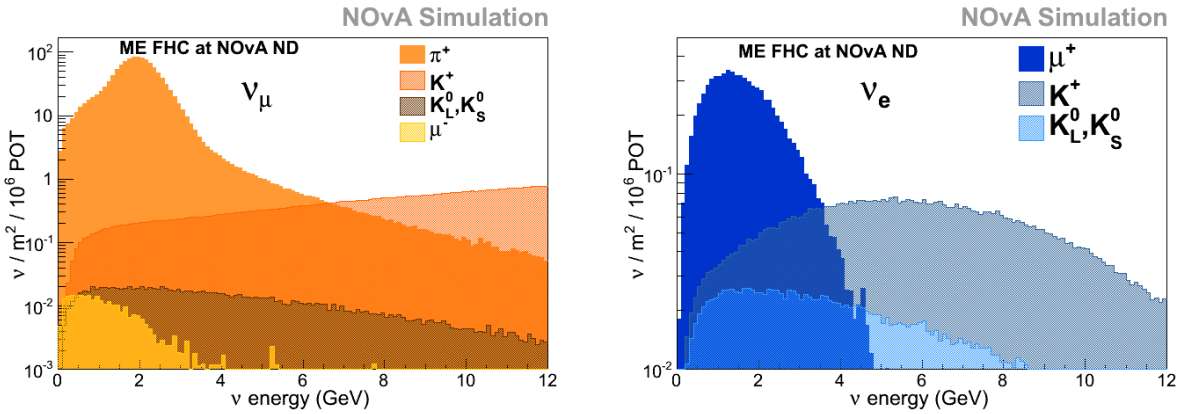

Figure 6.1: Parents of muon neutrinos (left) and electron neutrinos (right) in the NOvA near detector. Neutrinos up to 7 GeV come mostly from pions and higher energy neutrinos come mostly from kaons.

NUMI FLUX SIMULATION

We model the flux expected from the NuMI beam, as described in Chapter 4, at both the near and far detectors. This is done using G4NuMI [84], which is a Geant4 [85] based Monte Carlo simulation. G4NuMI incorporates the knowledge of the proton beam, the geometry and composition of the target, as well as the geometry and physics of the NuMI beamline. These simulations start from the generation of the hadrons produced in the collision and simulate their decay and interaction within the NuMI beamline. Package to Predict the FluX, or PPFX [86], is used to weight the hadron production. This package combines external data from multiple sources to constrain and quantify the uncertainties on hadron production from the NuMI target. PPFX provides a set of weights, which are then applied to the Geant4 simulation.

A central value for the NuMI flux is used, derived from the combination of Geant4 simulation and PPFX weights. The output contains information about the neutrino flux, including parentage, as shown in Fig.6.1.

INTERACTION SIMULATION

We use the GENIE [87] event generator to simulate neutrino interactions. GENIE uses the G4NuMI flux, as well as models of neutrino interaction cross sections to generate neutrino interactions in a volume larger than and centered on each detector. The simulated interactions retain the parentage information, the generated vertex (interaction location) and type of interaction, as well as information of direction and momenta of the daughter particles produced in the interaction.

Geant4 is also used to propagate the generated secondaries, this time through the rock/building surrounding the detector as well as the detector itself, whose detailed geometry is an input to Geant4 at this step. The output of this step are simulated hits, which are true energy depositions throughout the detector geometry.

DETECTOR SIMULATION

The detector response is simulated with NOvA internal software [88,89], which takes the true hits as input and simulates scintillation and Cherenkov light, APD response, and signal processing. It uses the measured response of our liquid scintillator and the measured reflectivity of the PVC (see Chapter 4) to simulate the scintillation light and then transport it to the electronics. This software also uses our knowledge of the APD response to simulate noise and other effects of the electronics. The output of the detector simulation, apart from the MC truth information from the simulation, are timed raw hits in the same format as the raw data taken from the detector.

6.3 Event Selection

This section details the data quality selection for NOvA data as well as the event selection for each channel. The selection for both channels includes cuts to ensure data quality, background rejection and signal selection. Each analysis has a different optimization as well as different challenges for background rejection.

6.3.1 Data Quality

The data quality selection is applied on a file-by-file basis, as it relates to data-taking conditions at the time the data was recorded, such as detector operating conditions and beam quality.

The following data quality requirements are shown in Fig.6.2:

Number of active diblocks > 4: Ensures that a region of the detector large enough to contain events was active for data taking.

Timestamp: Correctly timed data synchronized with the NuMI beam spills.

Live Time: Stable data-taking for more than a minimum continuous time.

Hit Rate: Overall hit rate stable with respect to the averages shown in Sec. 4.2.4 for the maintenance system. This ensures a noise rate within expected values and a data-taking rate consistent with the average for the trigger used.

Tracking: A high fraction of 3D tracks shows stability in the DCM to DCM timing synchronization. Given that tracks are reconstructed in 2D and then matched, a region out of time with the rest of the detector would impact view matching.

Slicing: Average number of slices (see Sec. 5.1.1) ensures complete readouts were taken as well as stable timing resolution.

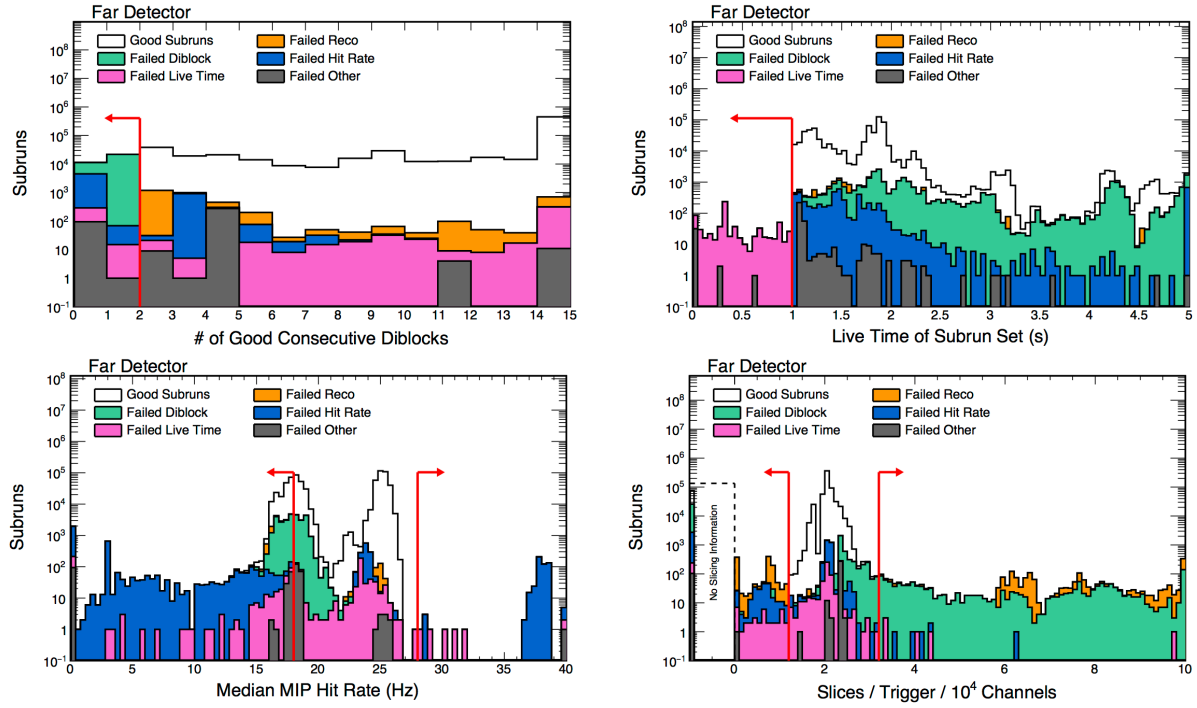

Figure 6.2: Data quality metrics with subruns (files) failing cuts highlighted in the corresponding color in all distributions. Top Left: Number of contiguous diblocks. Top Right: Total accumulated livetime calculated from the exposure of each event stored. Bottom left: Hit rate in the MIP region of ADC, as seen in Sec. 5.3.1. Bottom right: Slices per trigger, whose value is roughly constant for triggers of the same length.

6.3.2 ν_e Event Selection

The characteristic topology of the electron shower is the main feature of these signal events. The main sources of background are misidentified neutral current events, especially those containing electromagnetic showers, as well as neutrons from cosmic ray particle interactions in the barite overburden of the far detector, which may resemble the ν_e CC topology in a similar way. Another source of background is the original ν_e contamination in the NuMI beam, for which there is no handle in the selection from topology. Fig.6.3 shows an example of these backgrounds compared to a signal event.

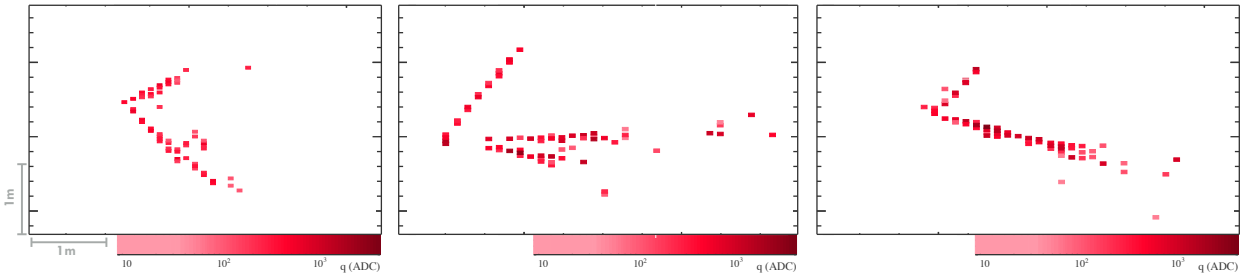

Figure 6.3: ν_e **CC signal (right) and background topologies from neutral current events (center) and cosmic ray neutrons (left).**

The selection for these events is done using the score for the ν_e CC category from the CVN event classifier, described in Sec. 5.2.2. A cut on this value is made after quality cuts and cuts for cosmic rejection are applied to the data collected in coincidence with the 10 μsec spills from the NuMI beam. The cuts, in the order they are applied, are the following:

BASIC CUTS

A set of basic cuts are applied to the analysis sample. Veto is applied before the reconstruction phase to eliminate obvious cosmic ray muon events and reconstruction failures.

Veto: Basic cosmic ray rejection is done by removing vertical tracks and through-going tracks.

Analysis Mask: Only periods with four or more contiguous diblocks are included in the analysis.

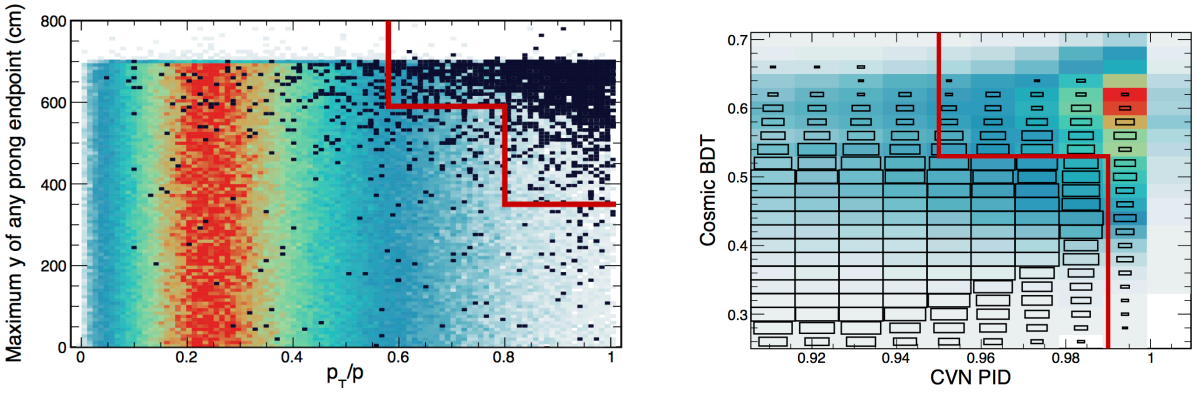

Figure 6.4: Left: Slice Max Y vs Pt/p for signal (color) and cosmic ray (black) events. Right: BDT score vs CVN score for signal (color) and cosmic ray (black) events. Red line delimits cut values.

Event Quality: A set of cuts to remove events with reconstruction failures or APD saturation effects. It checks for the existence of a reconstructed vertex and prongs, as well as no more than 8 hits per plane.

COSMIC REJECTION

Containment: Removes events for which any hit in reconstructed prongs is near one of the 6 detector walls. This removes events that pass through the walls of the detector as well as uncontained events.

- Distance to the Top > 63 cm and Bottom > 12 cm
- Distance to the Front > 18 cm and Back > 18 cm
- Distance to the East > 12 cm and West > 12 cm

Pt/p A cut on the transverse momentum, p_T/p, removes events that are not aligned with the NuMI beam axis. p_T/p < 0.58, or 0.58 < p_T/p < 0.8 and all hits within 590 cm from the top, or p_T/p > 0.8 and all hits within 350 cm from the top.

Backward photon: Rejects cosmic photon events that occur in the back 200 cm of the detector with a positive value for sparseness asymmetry, which is defined as the difference divided by sum of planes without any hits in the first 8 and the last 8 planes of the prong.

PRESELECTION

These cuts are applied to select events in the most signal-like region.

- Deposited energy in the range 1-4 GeV.
- Number of hits within 30-150.
- Length of longest prong between 100 and 500 cm.

ν_e CC - LIKE TOPOLOGY

CVN: Events are selected using the ν_e CC score from the event classifier discussed in Sec. 5.2.2 greater than 0.75.

NEAREST SLICE

Cut on the nearest slice to the selected event. This slice must be at least 100 ns or 500 cm from the event or 50 cm from the top of the detector if CVN > 0.87 or 400 cm from the top of the detector.

Events that pass all of these cuts make up the core sample of the analysis. The expected distribution of core sample events with this selection and assuming the best fit values from our 2016 analysis is shown in Fig.6.5. Table 6.1 and Fig.6.7 show the surviving signal and background for each cut.

FAILED CORE CUTS

In addition to the core sample, events which fail either the cosmic rejection or preselection cuts are included in the analysis. This peripheral sample consists of events that fail some core cuts, but have high CVN event classifier scores. The cuts used for the peripheral differ from the core sample after the basic cuts as follows:

Events that pass basic core sample cuts but fail any of the cosmic rejection cuts or preselection.

PRESELECTION

For the peripheral sample the preselection is a cut on the reconstructed ν_e energy. Higher energy events are allowed in the peripheral than in the core sample, up to 4.5 GeV.

ν_e CC - LIKE TOPOLOGY

CVN and Cosmic BDT: Two types of events are allowed in the peripheral sample. Events with ν_e event scores greater than 0.99 or 0.95 passing a cut given by a boosted decision tree trained for cosmic ray discrimination. The cosmic BDT uses sparseness asymmetry, CVN, number of hits, and distance to the top as discriminators.

Events that pass these cuts make up the peripheral sample of the analysis. The expected distribution of peripheral events with this selection and assuming the best fit values from our 2016 analysis is shown in Fig.6.6. Table 6.2 and Fig.6.7 show the signal and background efficiencies for each cut.

Cut	ν_e CC	beam ν_e	$\bar{\nu}_e$ CC	NC	ν_μ CC	ν_τ CC	cosmic
No Cuts	55.44	1.11	15.87	424.95	242.50	2.98	3.26e+06
Basic Cuts	53.50	1.08	15.12	304.79	222.14	2.79	2.88e+06
Cosmic Rejection	40.02	0.82	9.98	215.97	101.49	1.69	71734.5
Preselection	38.28	0.71	7.47	108.43	43.75	1.07	15734.7
CVN	31.09	0.58	5.76	5.08	0.93	0.31	2.34
Nearest Slice	30.98	0.58	5.74	5.05	0.92	0.31	2.02

Table 6.1: ν_e **CC core signal selection. This table shows the surviving events after each cut as shown in Fig.6.7.**

Cut	ν_e CC	$\bar{\nu}_e$ CC	beam ν_e	NC	ν_μ CC	ν_τ CC	cosmic
Peripheral before cuts	15.22	0.37	7.65	196.37	178.39	1.72	2.87e+06
Peripheral Preselection	14.98	0.33	6.30	188.62	156.32	1.42	2.76e+06
BDT+CVN	4.51	0.07	0.94	0.15	0.06	0.03	2.25
Nearest Slice	4.49	0.07	0.94	0.15	0.06	0.03	2.23

Table 6.2: ν_e **CC peripheral sample signal selection. This table shows the surviving events after each cut as shown in Fig.6.8.**

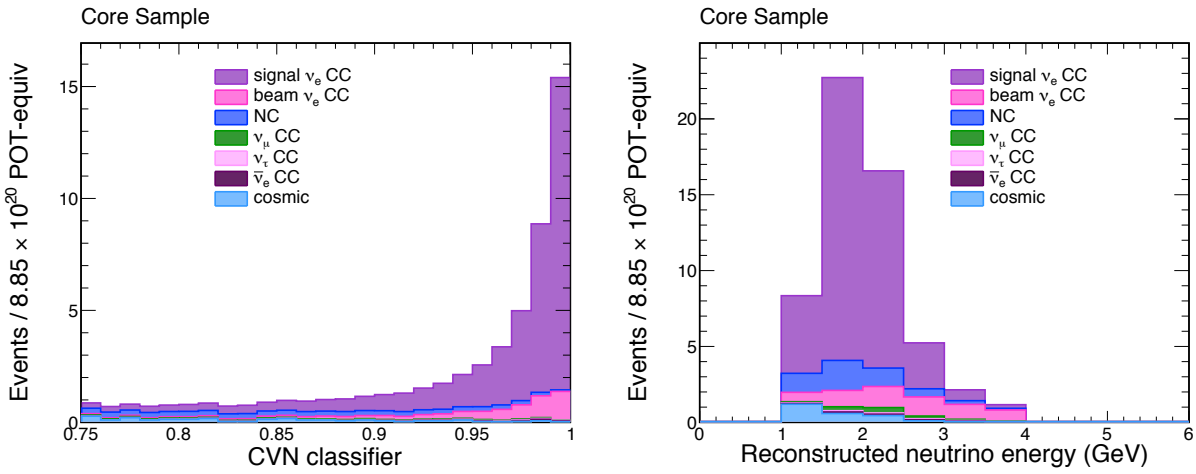

Figure 6.5: Expected CVN PID (left) and reconstructed energy (right) for ν_eCC Core sample. Background has been separated by component.

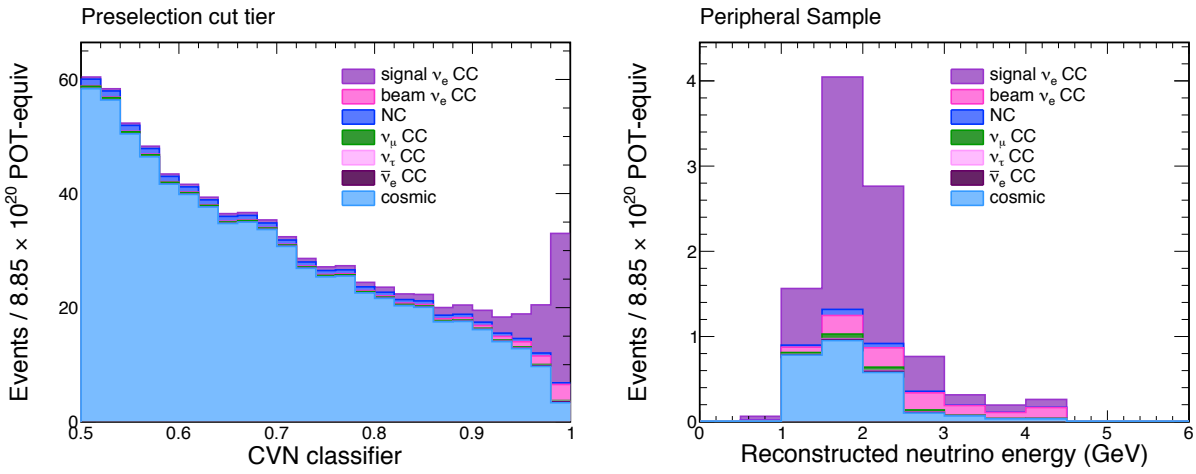

Figure 6.6: Expected CVN PID (left) and reconstructed energy (right) for ν_eCC Core sample. Background has been separated by component.

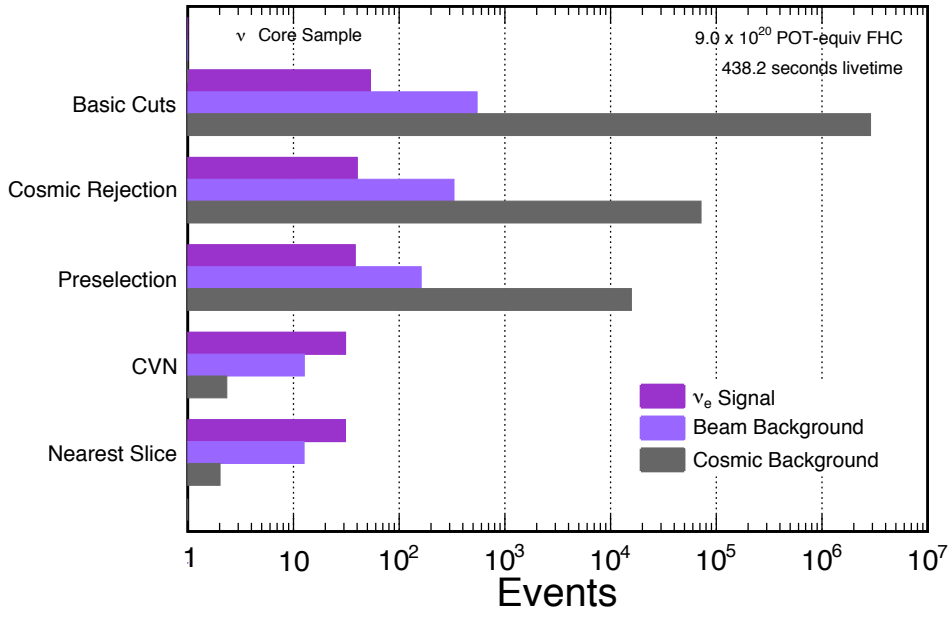

Figure 6.7: Selection survival counts for ν_eCC signal, beam backgrounds and cosmic ray backgrounds in the core sample.

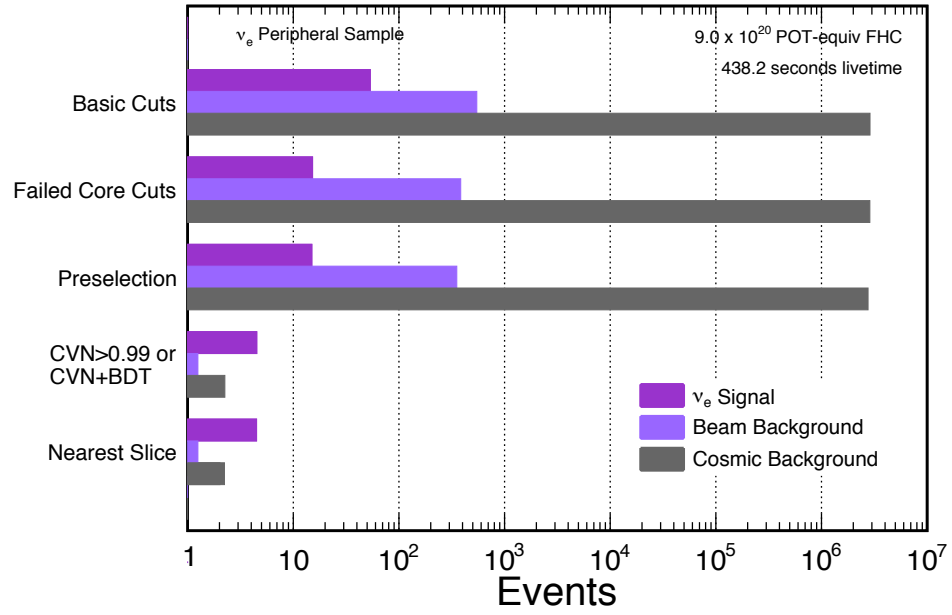

Figure 6.8: Selection survival counts for ν_eCC signal, beam backgrounds and cosmic ray backgrounds in the peripheral sample.

6.3.3 ν_μ Event Selection

The characteristic topology of ν_μ CC signal events is the muon track, characterized by the length of the tracks, and minimum ionizing depositions. The main sources of background in this case come from muons from cosmic rays as well as neutral current events, especially those containing long tracks for pions which mimic low energy muon tracks. Fig.6.9 shows an example of these backgrounds compared to a signal event.

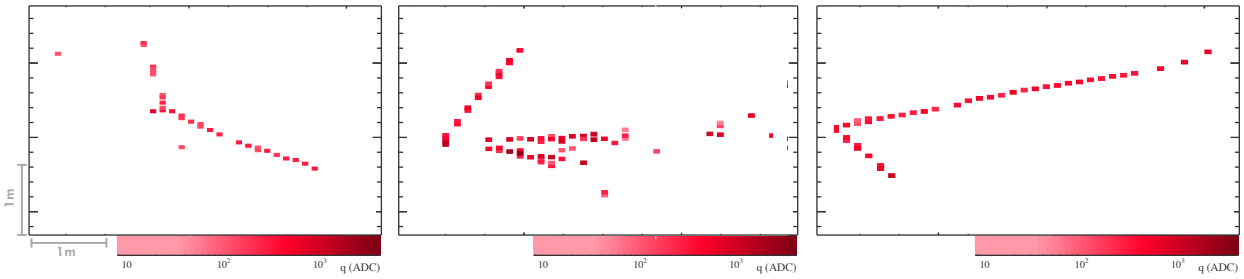

Figure 6.9: ν_μ CC signal (right) and background topologies from neutral current event (center) and from selected cosmic ray event (left).

Two track builder algorithms are used for some of the cuts applied to these events. One is a Kalman filter [90], the other is a simple straight line fitter. Additionally, a Nearest Neighbors algorithm called ReMID is used for this selection. This uses characteristics of the Kalman tracks to select events with tracks which mostly resemble muons in length, angle, and other variables [91].

Similar to the ν_e selection, the ν_μ CC event selection is done using the ν_μ CC score from the CVN event classifier, shown in Sec. 5.2.2 in combination with ReMID. The cuts are applied in the following order:

QUALITY

A set of cuts are applied to ensure complete events with reconstruction applied.

Event Quality: Events with problems in the DCMs are cut. This includes events with blank DCMs, such as those that are out of sync, or with hit pattern discontinuities at the DCM boundaries.

Reconstruction Quality: Require that a minimum number of hits in the slice, and that hits exist in at least four continuous planes. Also requires that at least one reconstructed track have non-zero PID scores.

CONTAINMENT

Removes events for which any hit in the event, prongs or tracks is near one of the 6 detector walls. This rejects cosmic ray events and ensures that all the energy from the ν_μ interaction is contained within the detector volume. The values for the containment cuts are:

- Kalman Track Forward Cell > 6 (cells)
- Kalman Track Backward Cell > 6 (cells)
- Simple Track Forward Cell > 0 (cells)
- Simple Track Backward Cell > 7 (cells)
- Planes to Front > 1 (planes)
- Planes to Back > 1 (planes)
- Min Prong Distance to FD Top > 60 (cm)
- Min Prong Distance to FD Bottom > 12 (cm)
- Min Prong Distance to FD East > 16 (cm)
- Min Prong Distance to FD West > 12 (cm)
- Min Prong Distance to FD Front > 18 (cm)
- Min Prong Distance to FD Back > 18 (cm)

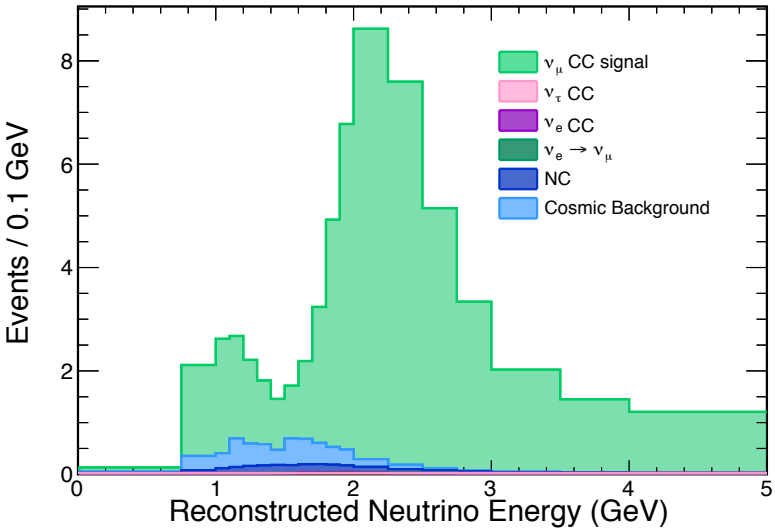

Figure 6.10: Expected reconstructed energy for ν_μCC. Background has been separated by component.

ν_μ CC-LIKE

Number of hits in the slice less than 400 plus two cuts to select ν_μ CC topologies.

CVN: Events are selected using the ν_μ CC score from the event classifier with scores greater than 0.5.

ReMID: Events are also required to have a ReMID score greater than 0.5. This cut helps to reduce the cosmic background.

COSMIC REJECTION

Cosmic BDT: The same style cosmic BDT used in selecting peripheral events for the ν_e appearance analysis is used to reject cosmic rays.

Kalman Angle: A cut is placed on the angle of the Kalman track to remove events that are not aligned with the beam axis.

Events that pass these cuts make up the ν_μ disappearance analysis sample. The analysis binning is discussed in Sec. 6.4. Fig.6.11 shows the surviving events for each cut.

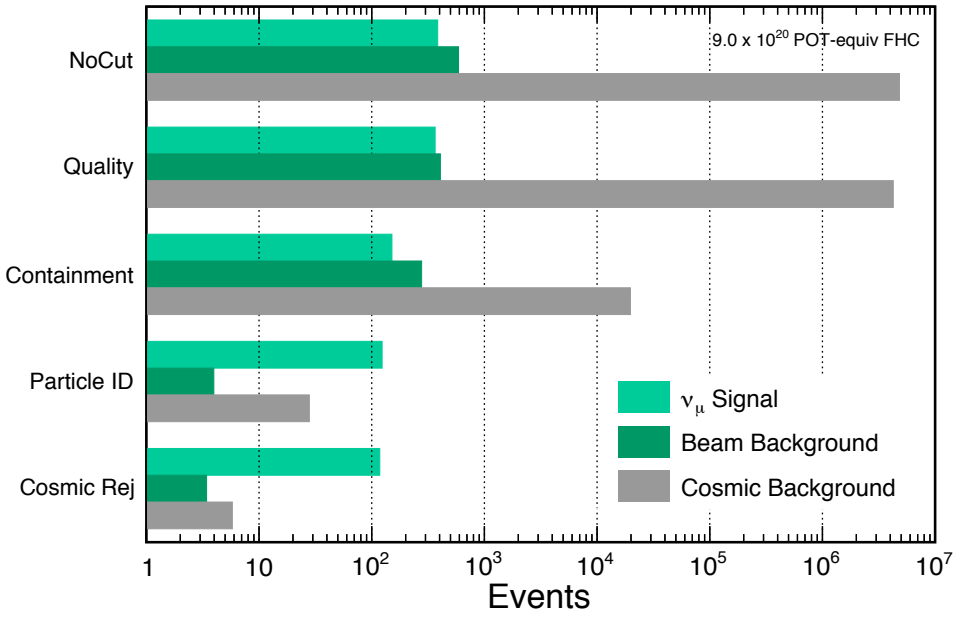

Figure 6.11: Selection survival counts for ν_μCC signal, beam backgrounds and cosmic ray backgrounds.

Cut	ν_μ CC	ν_e to ν_μ CC	ν_e	NC	ν_τ CC	cosmic
NoCut	385.01	0.56	83.25	500.57	4.55	4.82e+06
Quality	365.34	0.54	78.93	323.03	4.32	4.25e+06
Containment	151.13	0.30	54.66	219.73	2.47	19760.3
Particle ID	123.63	0.26	0.30	2.96	0.44	27.95
Cosmic Rej	117.91	0.25	0.24	2.49	0.42	5.79

Table 6.3: ν_μ CC signal selection. Surviving events after each cut.

6.4 Binning Optimization

Further optimization to the projected sensitivity of the results through studies of binning effects was performed for both channels.

ν_e BINNING

For ν_e signal, the optimization used treats the signal in a two dimensional spectrum of reconstructed energy and CVN score cut, where the three bins of selection correspond to different regions of the PID (CVN) value .These bins are 0.75–0.87, 0.87–0.95, and 0.95–1.0 for the core sample. The fourth bin corresponds to a peripheral sample, for which cuts have been relaxed, as described in Sec. 6.3.2. Fig.6.12 shows a breakdown of the energy by component in the expected far detector spectrum, presented in this binning. The peripheral sample is shown in one bin in the figure and in the results section. Given that some events in the peripheral sample might not be fully contained, this sample is included in the fits as a counting experiment, rather than an energy distribution.

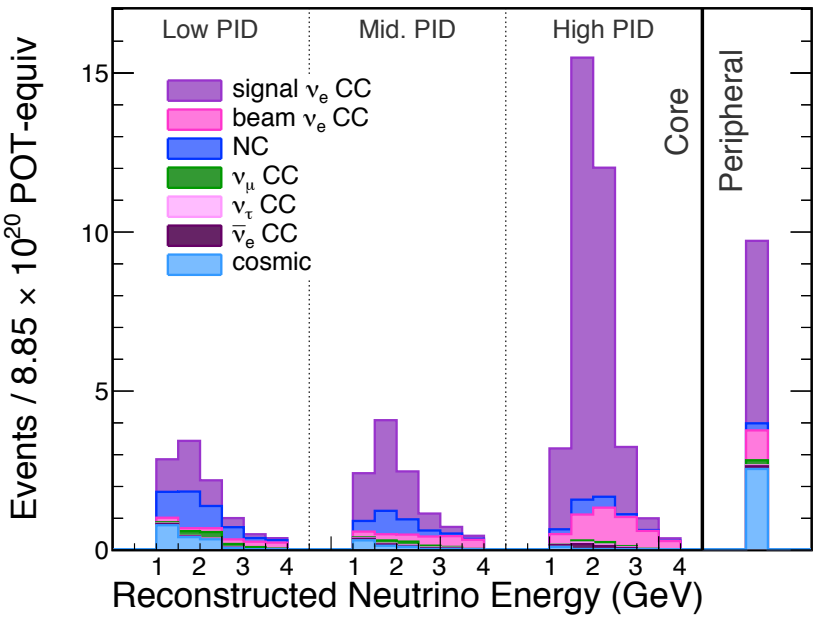

Figure 6.12: Expected ν_eCC in optimized binning.

ν_μ ENERGY BINNING

For ν_μ signal events, there are two optimizations on the binning: bin size in the energy region of 1-2 GeV and energy resolution binning. Given that a maximal value of the amplitude of the $\nu_\mu \rightarrow \nu_\mu$ oscillation probability ($\sin^2\theta_{23}$) is still allowed under our current experimental constraints, smaller bins were used in the region were the largest disappearance is expected, for an increase in sensitivity to maximal mixing.

As discussed in Sec. 5.5, the overall energy resolution (percent difference between reconstructed and true energy) for ν_μ CC selected events is 9% overall. However, the fraction of each component (muonic or hadronic) which contributes to the total neutrino energy varies largely event by event, mapping to different resolution for total neutrino energy. Splitting the events into regions of energy resolution increases the sensitivity and splits the sample roughly into regions of fraction of hadronic energy in the event. Resulting from studies of the sensitivity of the appearance analysis for different regions of the $E_{Had.}/E_\nu$ phase space [92], a separation into four energy quantiles was chosen for the analysis for it's optimal performance in sensitivity. As seen in Fig.6.13, the energy resolution is 6%, 8%, 10%, and 12%, for events in quantiles 1-4, respectively.

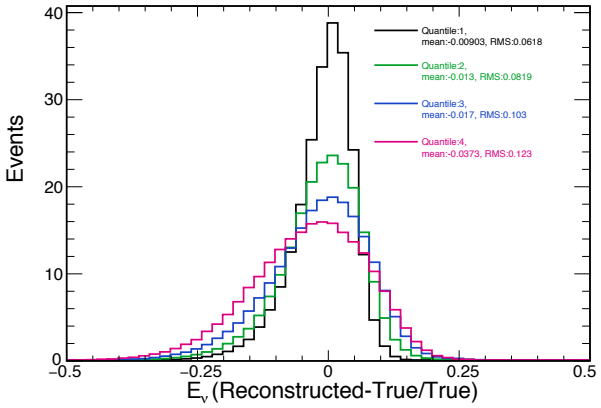

Figure 6.13: Overall ν_μ Energy Resolution. The difference in resolution in each quantile correlates with the fraction of hadronic energy in each one.

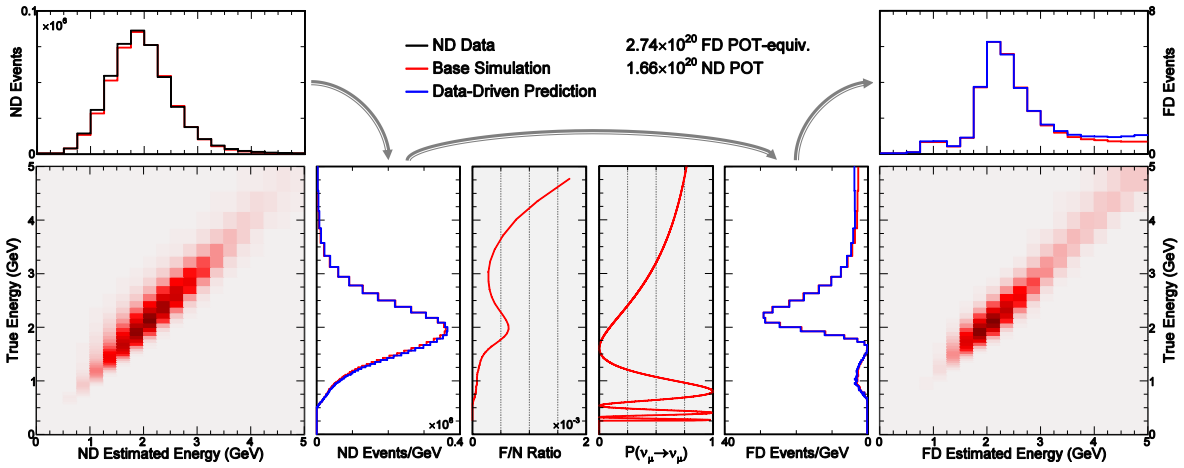

Figure 6.14: Extrapolation diagram

6.5 Prediction of the Energy Spectrum

The methodology followed to make predictions uses the near detector data and Monte Carlo (MC) to produce a prediction of the energy spectrum at the far detector (FD). FD predictions are reconstructed energy spectra from FD MC, which have been corrected using extrapolated ND energy spectra and an assumption for the oscillation probability.

6.5.1 Extrapolation

Fig.6.14 depicts the methodology of the extrapolation. Using a matrix conversion which relates true neutrino energy to reconstructed energy (M_{ND} in the figure), the reconstructed energy spectra in ND data and MC are converted into true energy spectra bin-by-bin (A in the figure).

The converted energy spectra are then extrapolated using a FD/ND ratios from the simulation. This yields binned true energy spectra, which are then converted to reconstructed energy spectra using corresponding conversion matrices (M_{FD} in the figure).

The extrapolation is done component by component for each signal and background on both anal-

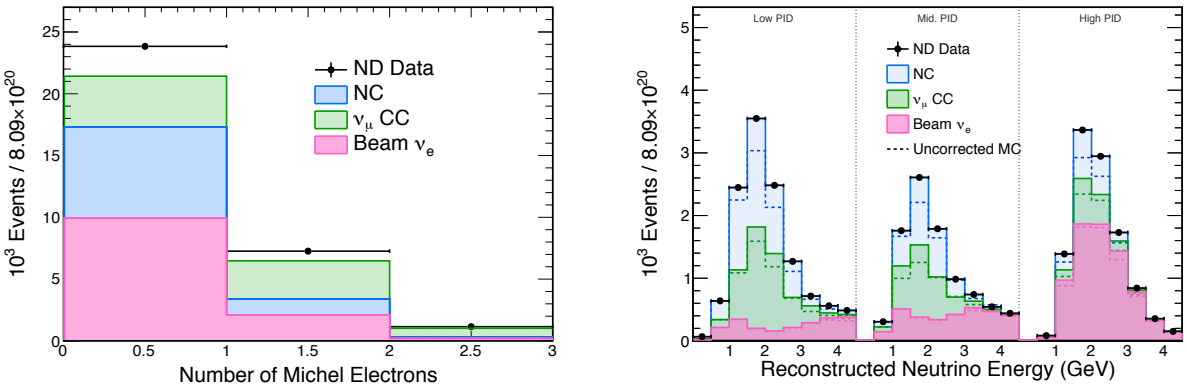

Figure 6.15: Left: Number of Michel electrons for each background component. Right: Corrected MC for backgrounds in the near detector, after corrections from Michel decomposition.

yses. As such, it depends on appropriately decomposing the spectra from the ND before extrapolating. Different decomposition methods are used for separating the selected events in data.

PROPORTIONAL DECOMPOSITION

This method, used for the ν_μ sample, assumes the composition of the MC for each type of background is the same in the data.

BEN DECOMPOSITION

This method uses truth information from ND selected ν_μ CC events, specifically, whether they come from kaon or pion decay, in order to estimate the ν_e backgrounds coming from the NuMI beam [93].

MICHEL DECOMPOSITION

This method measures the number of Michel electrons in selected ν_μ events in the near detector. This metric gives a handle on the fraction of events which are NC backgrounds, as shown in Fig.6.15.

6.6 Treatment of Systematics

The approach for calculating systematic uncertainties utilizes both far and near detector simulations, with the advantage (by design) that the functionally equivalent technology in the near and far detectors reduces, and in some cases cancels, these uncertainties. The uncertainty calculations are made using far detector predictions made with nominal versus "shifted" Monte Carlo samples produced for this purpose.

The effect of each systematic shift is evaluated through bin-by-bin comparisons of spectra of selected events from the shifted predictions, with respect to the nominal. For ν_μ disappearance, these comparisons are made on the reconstructed energy spectrum, separated by quantiles, and for ν_e appearance the comparisons are made on the CVN x reconstructed energy spectra. Here, "shifted" Monte Carlo refers to a sample made by systematically shifting the quantity of interest, corresponding to the uncertainty which is to be calculated. These shifts are made by either reweighting events according to the change in value of the parameter of interest, or by direct changes in its simulated value event by event. Both the generation of shifted samples and the comparisons to the nominal are performed independently for each systematic.

The shifted predictions are produced with the same method as the nominal, as described in Sec. 6.5. In the case of the shifted predictions, however, the near detector data is extrapolated to the far detector using the systematically modified samples, except for the evaluation of relative near-to-far effects, where only one of the samples is shifted.

All far detector predictions (nominal and shifted) were made assuming the following oscillation parameters: $L = 810$km, $\rho = 2.84\text{g/cm}^3$, $\Delta m^2_{32} = 2.67e-03$ eV^2, $\Delta m^2_{12} = 7.53e-05$ eV^2, $\sin^2 2\theta_{12} = 0.851$, $\sin^2 2\theta_{13} = 0.082$, $\sin^2 \theta_{23} = 0.404$, $\delta_{CP} = 1.48\pi$.

The effect of each systematic uncertainty is quantified and reported as the percentage change in the far detector prediction of the number of background events and that of signal events independently.

The uncertainties presented here were evaluated for a set of effects which can be divided into the following categories:

- Both ν_e and ν_μ
  - Calibration
  - Light levels and Cerenkov light
  - Beam transport
  - Hadron production
  - Cross section uncertainties
- ν_e only
  - Extrapolation systematics
  - Rock uncertainty
  - Normalization
- ν_μ only
  - Energy scale shifts for both muons and hadrons.

The following sections describe the dominant systematic uncertainties, Appendix B contains tables for the remaining small uncertainties.

6.7 Calibration

In order to cover data-MC discrepancies post-calibration, deliberate mis-calibrations are introduced to the events in the nominal MC. This method preserves the number of hits and other simulation information in order to assess the effects of the mis-calibration in isolation from other effects. These

artificial mis-calibrations take the form of an absolute calibration shift and a calibration "shape" uncertainty. These shifts are motivated by data-driven studies, like the one shown in Fig.5.41, which have shown both flat and spatially dependent disagreement in the detector response in data and MC after calibration corrections have been applied [94].

The effect of a 5% absolute calibration uncertainty was studied, as well as the shape effects seen in Fig.5.43.

6.8 Cherenkov and Light Levels

The light yield of the scintillator as charged particles travel through it is modeled by two effects: scintillation, the main detection mode, and Cherenkov radiation. NOvA's light yield simulation includes a correction to the dE/dx dependence of the scintillation yield to account for quenching at high energies. This quenching factor is incorporated though Birks attenuation [95], an empirical relationship that describes the light yield

$$LY \propto \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx} + k_C \left(\frac{dE}{dx}\right)^2}, \tag{6.1}$$

where k_B and k_C are material dependent parameters.

We account for uncertainties related to light levels and thresholds by altering the light level by $\pm$10% with a compensating change made to the absolute calibration constants [96]. In addition, an adjustment to the Cherenkov model that shifts proton response down by 2.6% [96].

6.9 Flux Uncertainties

Two main categories of flux uncertainties were considered: *beam transport*, which refer to possible differences between the simulation and the actual working conditions of the NuMI beam (e.g. small shifts in the positions of some components), and *hadron production*, an uncertainty concerning the production of pions and kaons in the simulation.

6.9.1 Beam transport uncertainties

These uncertainties include the horn and target position, the horn current, the position of the beam on the target and the beam spot size. The effect of each uncertainty is below 5% on the flux spectrum at the near or far detector for most cases. Given that the effect of each individual beam systematic uncertainty was evaluated separately and found to be small, their total effect is summed in quadrature.

6.9.2 Hadron Production Uncertainties

These uncertainties are assessed by generating an number of alternative weightings where the constraints from the fixed target data and theoretical assumptions are allowed to float within their uncertainty. Uncertainties are treated by choosing a statistical ensemble of randomly generated universes, each of which has interaction parameters allowed by the data uncertainty. These universes are then propagated through the simulation chain in order to estimate the variance in each neutrino energy bin as well as the correlations between different bins and build the covariance matrix shown in Fig.6.16.

Diagonalizing this covariance matrix from Fig.6.16 yields a set of uncorrelated weights or principal components for these uncertainties. Four of these principal components are used in the fit as uncertainties, the rest are taken as one uncertainty, summed in quadrature with other small effects.

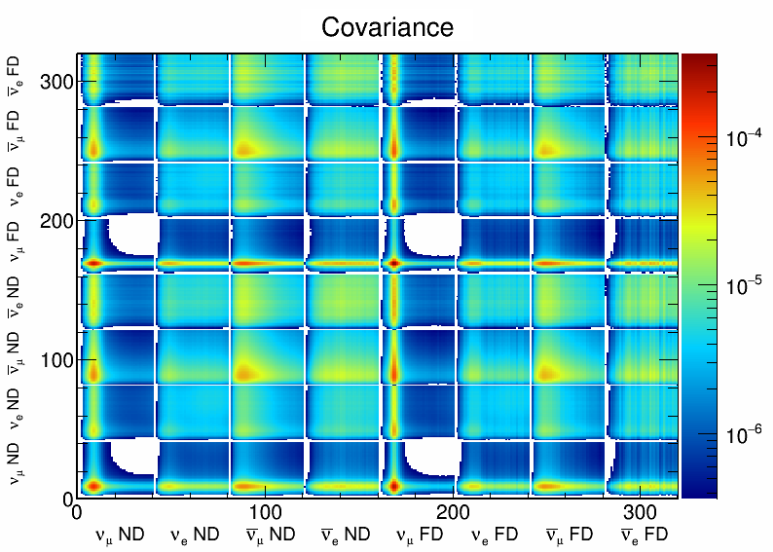

Figure 6.16: PPFX Covariance Matrix.

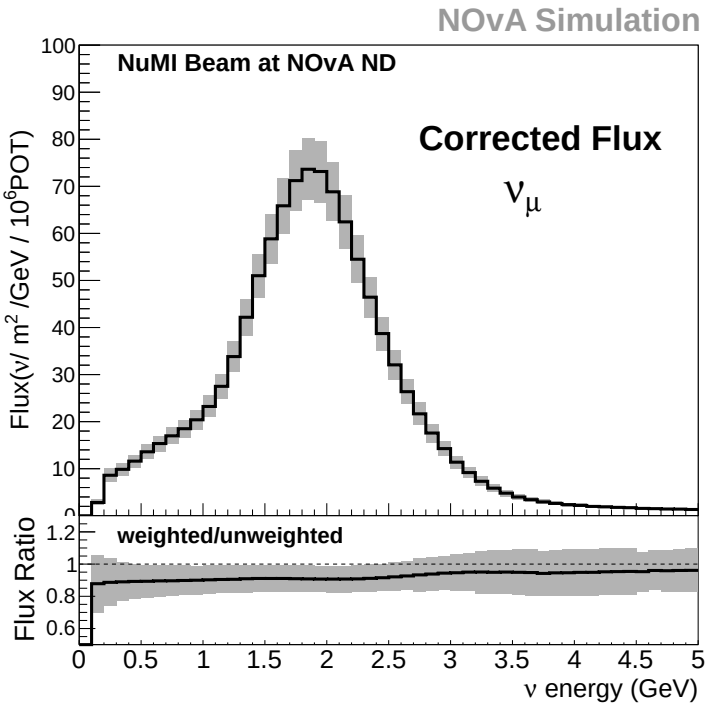

Figure 6.17: Hadron Production Uncertainty. The muon neutrino flux integrated in the NOvA near detector (-176 < X < 177 cm, -172 < Y < 179 cm and 25 < Z < 1150 cm). The band corresponds to the total beam uncertainty (hadron production plus focusing added in quadrature). The ratio with respect to the G4NuMI FTFP output is also shown on the bottom.

6.10 GENIE Systematics

Several neutrino interaction systematics are evaluated by using the reweighting tools built into GENIE. These tools compute a weighting factor that can be applied to simulated events to enhance or suppress a particular type of interaction. The size of the systematic modification of each parameter is the recommendation of the GENIE authors, based on a careful survey of the interaction model. The effects considered fall into three categories: cross-section uncertainties, hadronization model uncertainties, and uncertainties due to final state interactions. The GENIE systematics being considered for this analysis include:

- An uncertainty in the shape of the MEC (meson exchange current) cross-section by allowing its q_0 ($\sim$hadronic energy) shape to vary between quasielastic-like and resonance-like.
- Systematic uncertainties on the shape of the RPA suppression, one when applied to quasielastic events, and another one when the suppression, calculated for QE, is extrapolated to resonance events.
- An uncorrelated 2% uncertainty on the ν_e/ν_μ cross section ratio to account for radiative corrections [97].
- A 2% uncertainty anticorrelated between ν_e and $\bar{\nu}_e$ to allow for second class currents [97].
- Adjust M_A in the Rein-Sehgal resonance cross section model by $\pm 20\%$ for both CC and NC interactions.
- Adjust M_V in the Rein-Sehgal resonance cross section model by $\pm 10\%$ for CC interactions.
- Various components of the DIS cross section (see Appendix B) .
- Adjust the M_A in the Llewellyn-Smith quasielastic cross section by $\pm 5\%$.
- The additional cross section systematics can be found Appendix B are summed together in quadrature and added as an additional systematic.

6.10.1 Rock Contamination (ν_e only)

Neutrinos coming from the beam represent additional backgrounds when they interact in the building or rock surrounding the detector hall. These, called rock events, are not included in any of the systematics samples, instead it is added to the final result when fitting, applying 100% uncertainty to it.

6.10.2 Normalization

The ν_e normalization systematic accounts for uncertainties in POT accounting, detector mass, and near detector reconstruction efficiency.

6.10.3 Extrapolation systematics

While the near and far detectors are very similar, they do differ; most obviously in their sizes. For instance, kinematic distributions are sculpted given the smaller size of the ND in a way which does not happen in the large FD volume. Two types of such uncertainties are considered for this analysis:

Background uncertainty: The flux difference in different regions of the ND (east/west, top/bottom, front/back, inside/outside) can cause the background predictions in the FD to be different. The maximum 1.3% deviation of any detector region for the nominal prediction is taken as a systematic uncertainty.

Signal uncertainty: The FD ν_e signal distribution is estimated from selected ND ν_μ events. These two samples have somewhat different kinematics (particularly in Q^2), which could affect the extrapolation. This effect is estimated by reweighting the ND data and MC to match the FD Monte Carlo signal kinematics.

6.10.4 Energy scale (ν_μ only)

There are three energy scale shifts which are considered for the ν_μ signal.

- A 5% absolute hadronic energy scale shift.
- A 1% absolute muon energy scale shift in both detectors.
- A roughly 0.27% relative muon energy scale shift in the FD only.

The muon energy scale shifts are applied to the reconstructed track length of the muon, and were determined after a detailed literature review and GEANT4 study which is detailed in [98, 99].

The hadronic energy scale shift was determined by studying the calibration calorimetric responses of the detectors, it is motivated by a discrepancy for protons in the ND in both MC and data.

6.10.5 Systematic Uncertainties

6.10.6 ν_e Systematics Tables

The following figures and table show a summary of the predominant systematic effects on the ν_e signal and background:

Extrap. Bkg	Bin1		Bin2		Bin3		Bin4		Tot		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
SmallXsecsJoint	+4.92	-9.55	+6.98	-9.41	+4.32	-6.63	+3.89	-4.56	+5.14	-8.08	0.046	0.134
Lightlevel	+2.34	+0.08	-1.41	+6.02	-0.75	-3.92	+0.54	-4.89	+0.15	-0.38	0.096	0.124
Calibration	-1.65	-1.17	-0.30	+1.52	-1.32	+3.72	-4.83	+4.68	-1.44	+1.72	0.074	0.100
Cherenkov	+4.49	–	+2.57	–	-1.60	–	-1.17	–	+1.31	–	0.073	–
CalibShape	+4.13	–	+1.02	–	+1.22	–	-4.23	–	+1.65	–	0.047	–
MECq0Shape	+1.03	+0.94	+1.50	+1.07	+3.23	+0.01	+2.96	-1.04	+2.09	+0.48	0.015	0.005
extrapsignalkin	-1.16	+1.38	+1.86	-1.65	+1.60	-1.66	+1.26	-1.28	+0.77	-0.68	0.010	0.011
MaCCRES	-0.18	+0.36	+1.44	-1.32	+2.91	-2.84	+3.58	-3.12	+1.63	-1.48	0.009	0.008
MvCCRES	-0.10	+0.16	+0.86	-0.79	+1.66	-1.49	+1.98	-1.68	+0.94	-0.82	0.003	0.002
extrapbkg	+1.21	-1.19	+1.32	-1.09	+1.11	-1.28	+1.24	-1.36	+1.21	-1.21	0.002	0.002

Extrap. Sig	Bin1		Bin2		Bin3		Bin4		Tot		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
Calibration	-0.95	+1.20	+0.85	-0.50	+1.83	-3.56	+2.08	-3.27	+1.48	-2.67	1.326	1.375
SmallXsecsJoint	+2.23	-6.23	+0.26	-1.26	-3.94	+11.54	-4.61	+13.62	-2.88	+8.41	0.057	0.491
MECq0Shape	+4.23	+1.44	+3.27	+1.94	-4.19	+8.94	-3.43	+9.37	-2.29	+7.33	0.072	0.274
CalibShape	-0.40	–	-0.68	–	-1.12	–	-3.05	–	-1.23	–	0.113	–
Lightlevel	+1.52	+1.25	+0.71	+1.02	+0.82	+0.20	+2.56	-2.27	+1.07	+0.11	0.036	0.108
RPAShapeRES	+4.23	+0.00	+3.28	+0.00	+4.01	+0.01	+4.55	+0.01	+3.99	+0.01	0.069	0.000
MaCCRES	+5.52	-4.84	+4.69	-4.37	+0.82	-0.24	+1.25	-0.62	+1.84	-1.28	0.029	0.022
radcorrnue	+1.95	-1.95	+1.94	-1.94	+1.97	-1.97	+1.97	-1.97	+1.96	-1.96	0.016	0.016
2ndclasscurr	+1.90	-1.90	+1.89	-1.89	+1.94	-1.94	+1.94	-1.94	+1.93	-1.93	0.016	0.016
MvCCRES	+3.48	-3.03	+2.96	-2.62	+0.62	-0.42	+0.89	-0.66	+1.24	-1.00	0.012	0.009
RPAShapeenh	+0.21	+1.04	+0.59	+0.42	+1.68	-1.02	+1.90	-1.14	+1.42	-0.65	0.010	0.004
Cherenkov	+0.03	–	+0.50	–	+0.03	–	+0.11	–	+0.11	–	0.009	–
DISvnCC1pi	+3.09	-3.12	+2.16	-2.19	+0.52	-0.53	+0.63	-0.63	+0.99	-1.00	0.008	0.008
MaCCQEreduced	+0.08	+0.17	+0.25	-0.03	+1.21	-1.05	+1.44	-1.22	+1.00	-0.82	0.005	0.004

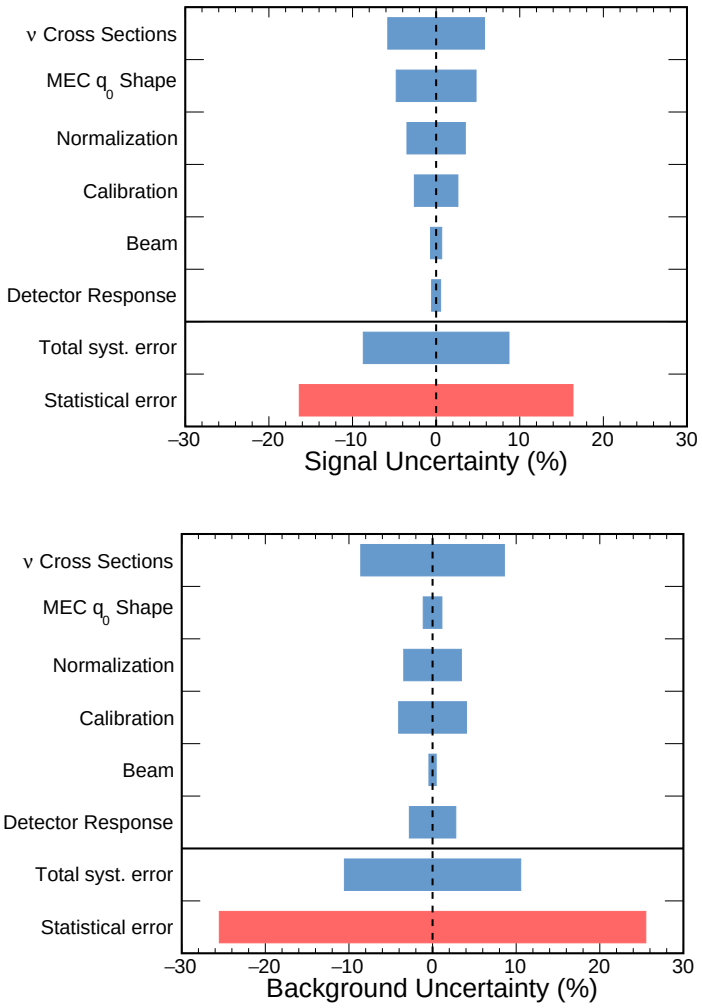

Figure 6.18: Sizes of ν_e systematic uncertainties.

6.10.7 ν_μ Systematics Tables

Extrap. Sig	Quant1		Quant2		Quant3		Quant4		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
MECq0Shape	-1.92	+3.45	-5.26	+12.00	-3.63	+7.94	-0.41	-0.61	0.182	0.734
Calibration	-4.59	+1.63	-4.47	+3.13	-5.65	+4.98	-3.99	+4.08	0.528	0.442
RelHadEScale	+5.36	-5.06	+0.47	-0.30	-3.26	+3.22	-5.70	+6.13	0.413	0.441
CalibrationEnergy	-4.03	+2.61	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.426	0.321
AbsHadEScale	-1.62	+1.65	-2.71	+3.23	-4.43	+4.78	-4.69	+5.79	0.321	0.410
SmallXsecsNumu	-0.88	+3.15	-1.23	+5.29	-0.87	+5.55	-1.73	+3.67	0.026	0.312
RPAShapeRES	+1.67	+0.00	+4.61	+0.00	+6.06	+0.00	+2.21	+0.00	0.197	0.000
RelHadEScaleE	-0.78	+0.78	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.178	0.178
RPAShapeRESE	+3.46	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.151	0.000
CalibShape	-3.30	–	-0.80	–	+0.03	–	+0.70	–	0.129	–
MaCCRES	-0.05	+0.21	+1.31	-0.76	+3.81	-2.44	+2.21	-1.34	0.103	0.057
MaCCRESEnergy	+2.23	-1.49	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.098	0.058
AbsMuEScale	+2.19	-1.99	+1.50	-1.40	+1.09	-1.06	+0.98	-0.95	0.081	0.069

Extrap. Bkg	Quant1		Quant2		Quant3		Quant4		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
SmallXsecsNumu	+46.56	-45.83	+51.03	-49.60	+58.67	-57.21	+69.29	-65.81	1.448	1.322
Calibration	-13.29	+13.77	-11.82	+11.06	-13.41	+11.22	-19.38	+21.87	0.124	0.154
MECq0Shape	-8.76	+14.18	+10.55	+8.99	+12.90	+2.83	+1.91	+0.35	0.044	0.014
MaNCRES	+6.58	-6.11	+8.44	-7.36	+8.22	-6.67	+8.44	-6.05	0.038	0.025
MaNCRESEnergy	+8.29	-6.26	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.037	0.025
Lightlevel	+6.81	+6.09	+1.86	+1.64	+2.88	+2.76	+7.09	+3.30	0.026	0.016
Cherenkov	+4.25	–	-0.47	–	+1.17	–	+3.33	–	0.018	–
ppfxpc01	-7.39	+7.39	-7.34	+7.34	-7.18	+7.18	-7.08	+7.08	0.017	0.017
AbsHadEScale	-7.32	+7.65	-5.17	+5.16	-3.95	+4.66	+0.88	-1.14	0.012	0.014
MaCCRES	+1.81	-2.03	+6.15	-5.61	+6.72	-5.13	+2.90	-2.15	0.010	0.006
AllBeamTransport	+4.59	-4.71	+4.72	-4.75	+5.00	-4.92	+5.35	-5.22	0.009	0.009
RelHadEScale	+4.12	-3.66	+2.27	-2.73	+2.51	-2.34	-0.64	+0.55	0.004	0.004
MvCCRES	+0.89	-0.83	+3.28	-2.76	+3.48	-2.71	+1.52	-1.19	0.003	0.002

It is clear that statistical uncertainties dominate with the amount of data considered for the present analysis (see next chapter). The leading systematic uncertainties on both channels are the calibration and cross-setions uncertainties. While these summary tables are useful indications of the relative size of the systematics, they are applied as a shape to the fits described in Sec. 7.2.

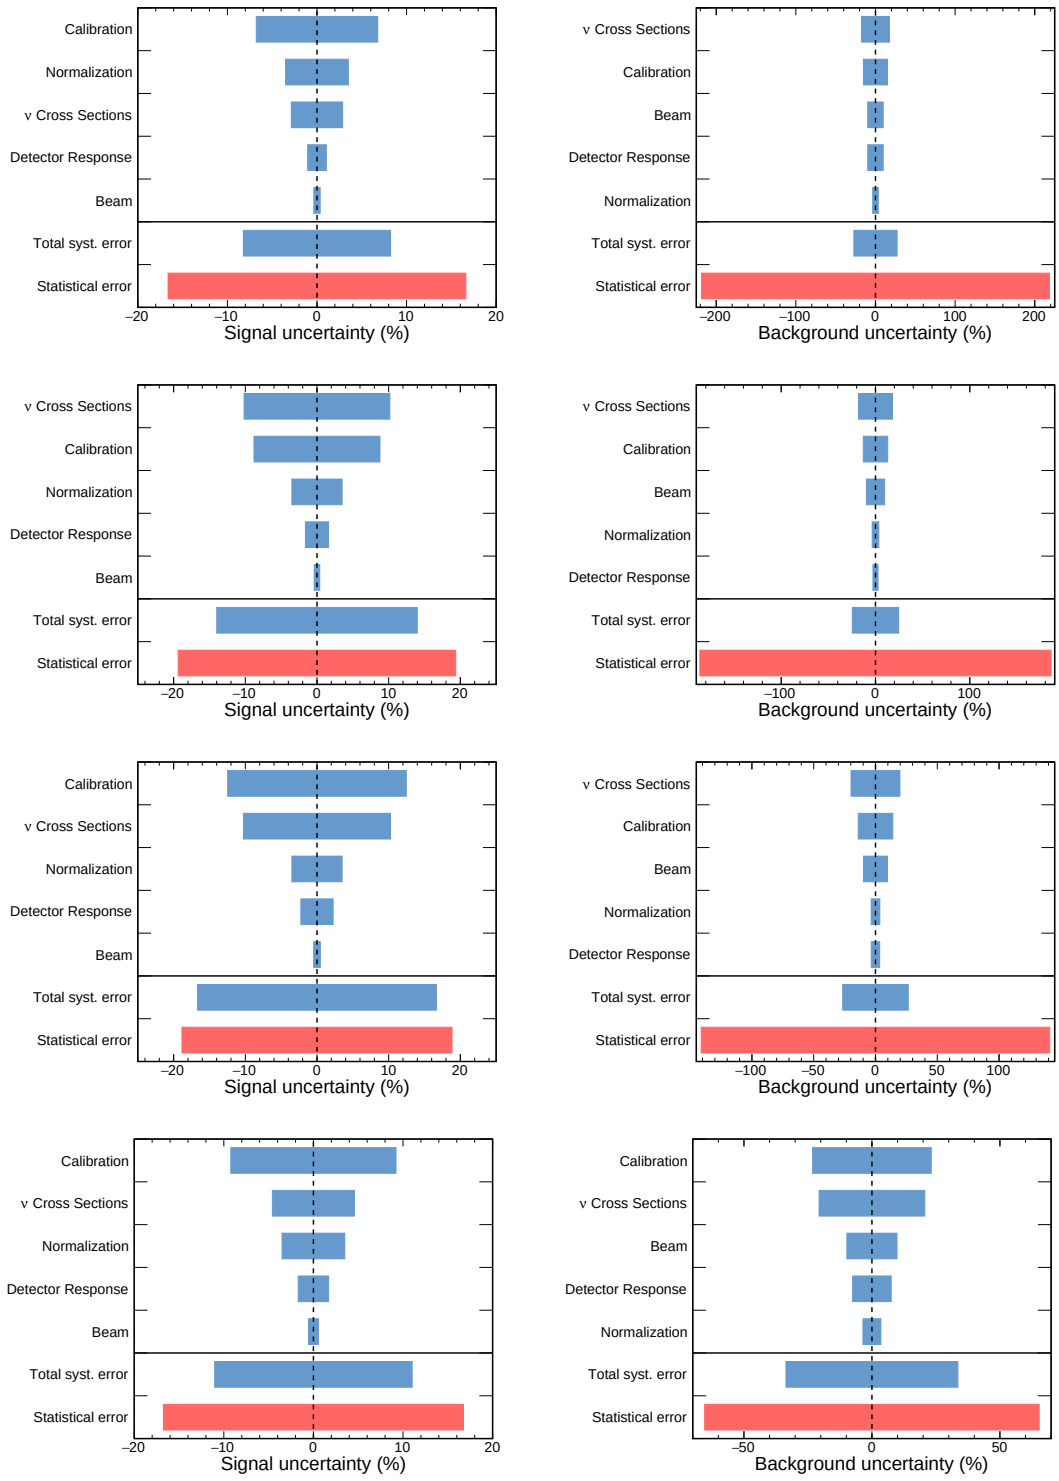

Figure 6.19: Sizes ν_μ systematic uncertainties. Quantiles 1-4 shown from top to bottom.

Chapter 7

Results and Discussion

This analysis, presented on January 2018 at [100], implemented the contributions detailed in this document and is NOvA's official analysis for 2017. This chapter presents the resulting spectra of observed appeared ν_e events and surviving ν_μ events at the far detector, with the selection detailed in Secs. 6.3.2 and 6.3.3, with the binning described in Secs. 6.4 and 6.4, respectively.

The dataset which was analyzed constitutes 8.85×10^{20} POT equivalent, which corresponds to the data collected in the period starting on February 2014 and up to March 2017, as shown in Fig.7.1.

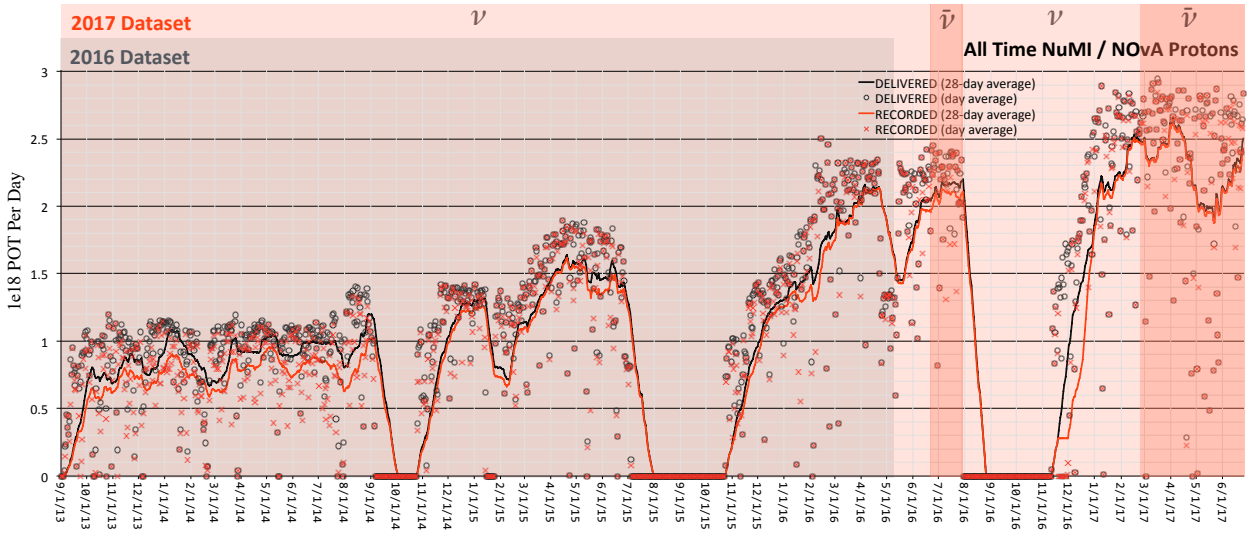

Figure 7.1: NuMI POT recorded for this dataset is shown in light orange. Dark orange: Anti-neutrino (rhc) data, to be analyzed later this year. Gray: Data taken while the detector was being commissioned.

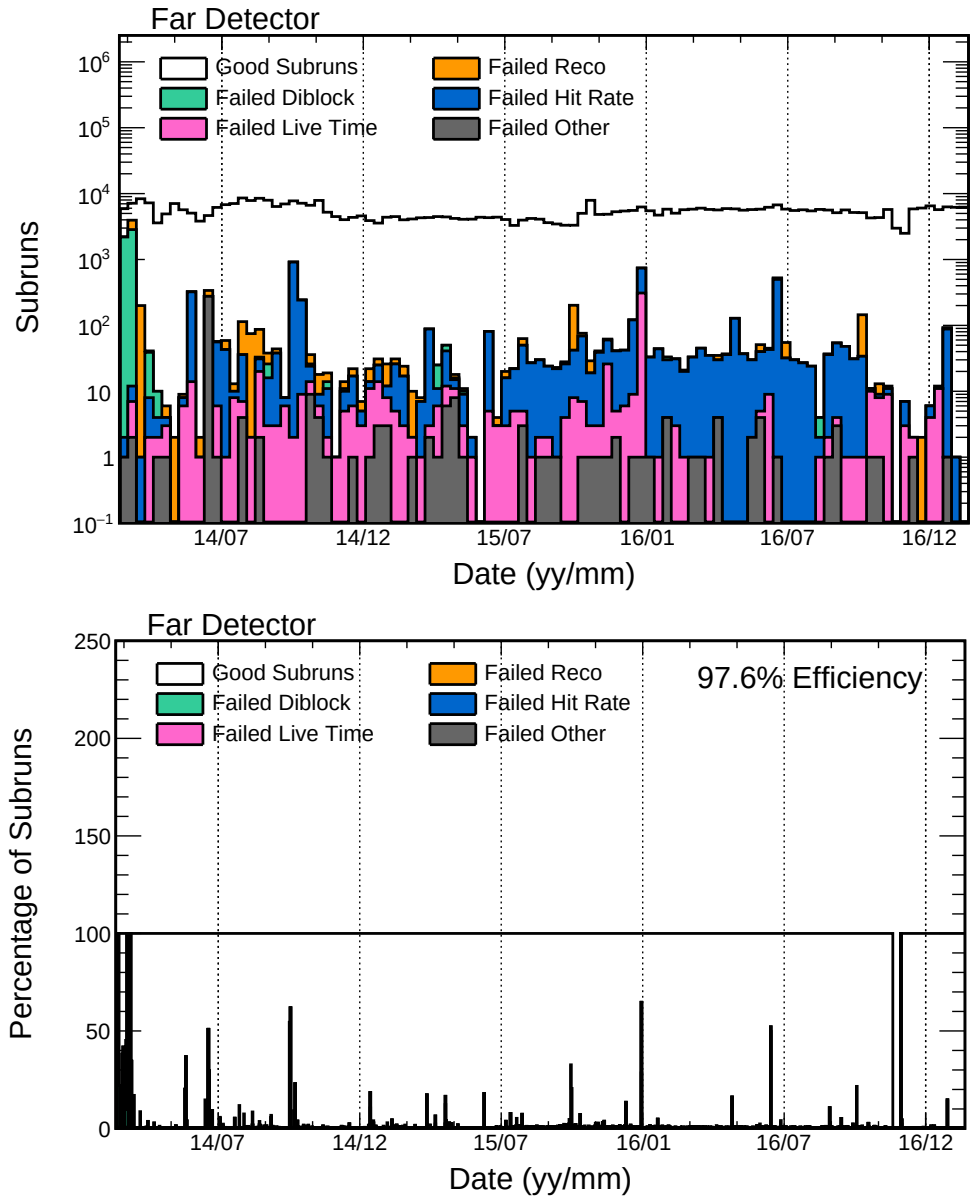

Figure 7.2: Data quality for this dataset. Stats on February 2014 and is 97.6% of the data collected for that period. The plots show the number of Subruns (corresponding to files of data collected in periods of time of about 1-2 min.) which do not pass basic data quality cuts. The plot on the top shows most common failure modes in log scale, and shows that the most common failure is the Hit Rate and LiveTime, both of which occur commonly on short Subruns, which happen when a run is stopped or started incorrectly, but are still fairly rare (less than 3%).

The anti-neutrino data shown in the figure is not included in this result, but will be included in the analysis to be presented by the collaboration in the summer of 2018.

In the period corresponding to this dataset, the fraction of collected POT was 95%, including the commissioning era (gray Fig.7.1) and the fraction of data used in the analysis was 97.6% after data quality, as shown in Fig.7.2.

7.1 Results

ν_e APPEARANCE RESULTS

The energy spectra at the far detector are shown in Figs.7.4 and 7.3 for appeared ν_e events selected as detailed in Sec. 6.3.2. The energy of events in the core sample has been reconstructed with the method outlined in Sec. 5.4, and binned according to the optimization described in Sec. 6.4. The energy spectrum for the peripheral sample is shown in Fig.7.3, but the sample is only taken as an event count into the fit, since the selection allows for uncontained events.

In this channel we see 66 events over 20.5 expected background. There are distributed as seen in the Table 7.1. The following figures and tables show the data compared to the best fit prediction.

NOvA has seen unambiguous appearance of electron neutrinos in a muon neutrino beam since our first result in 2014, but the increase of statistics and the improvements in the analysis have an impact in our sensitivity to the mass hierarchy, δ_{CP} and the octant of θ_{23}. Details of the fit to this observation are described in Sec. 7.2.

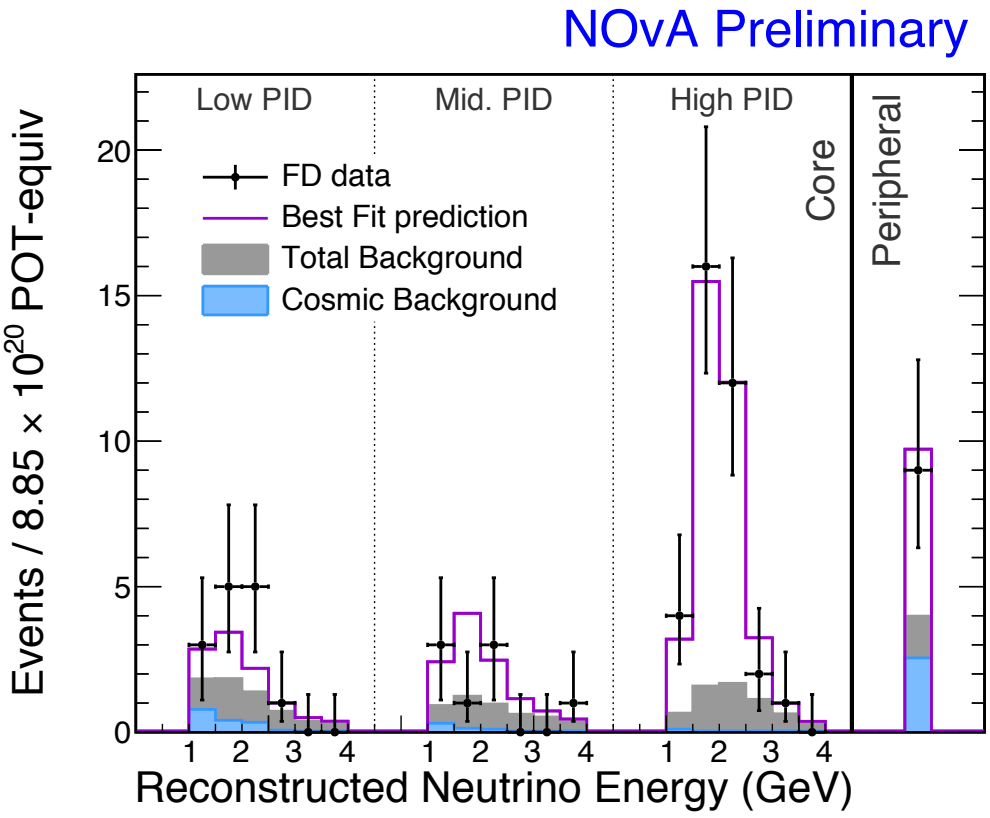

Figure 7.3: Selected data and MC reconstructed energy for ν_eCC Core sample in bins of CVN, and event count for peripheral sample.

	Total	Min PID	Mid PID	High PID	Periph.
Observed Events	66	14	8	35	14
Predicted Total	66.8	10	11	36	9.8
Background Exp	20.5	6	4.5	5	4

Table 7.1: Observed appeared ν_e events at the far detector, compared to predictions.

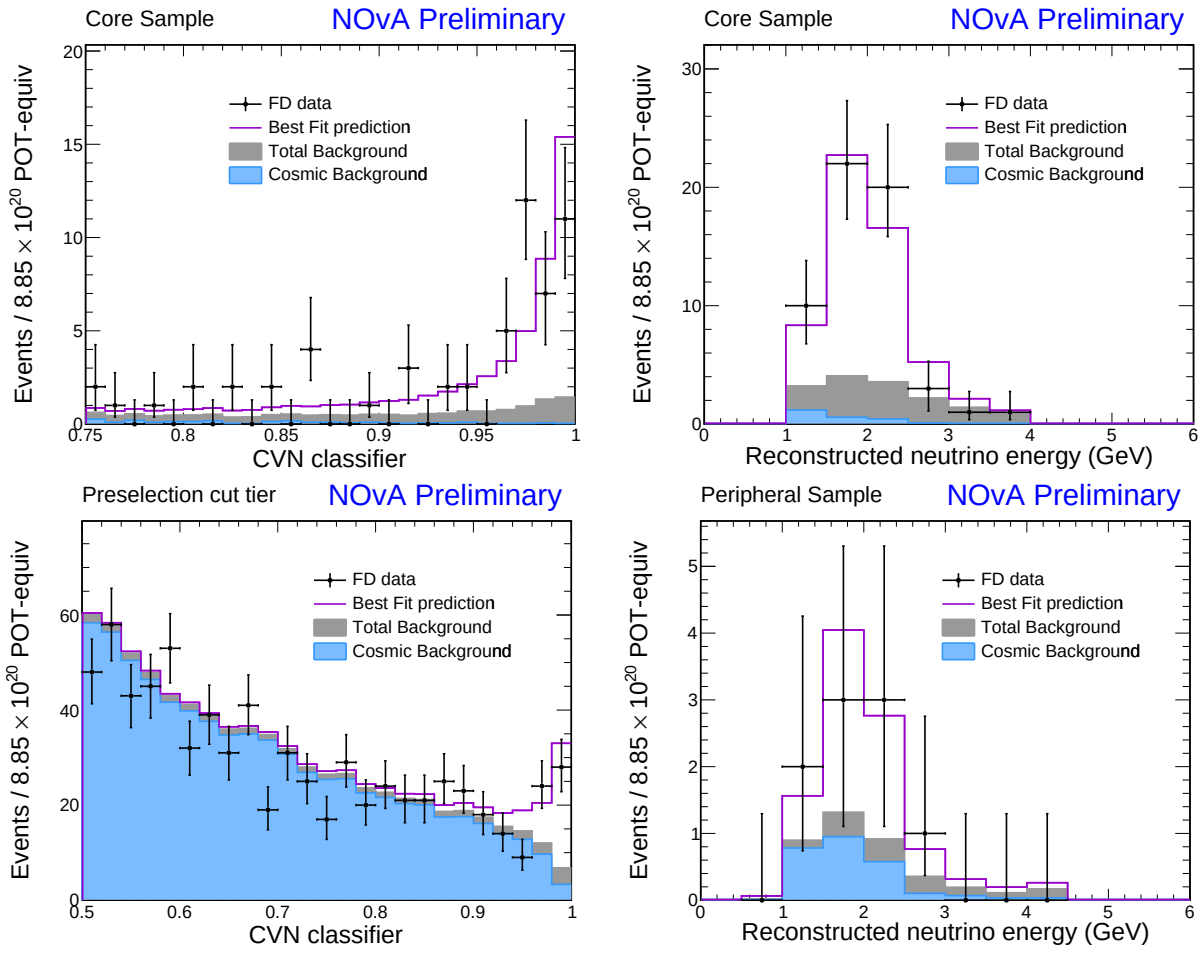

Figure 7.4: Signal CVN PID (left) and reconstructed energy (right) for ν_eCC Core sample. Signal CVN PID (left) and reconstructed energy (right) for ν_eCC Core sample. Signal Spectra for ν_eCC Sample. Signal CVN PID x reconstructed energy for ν_eCC sample.

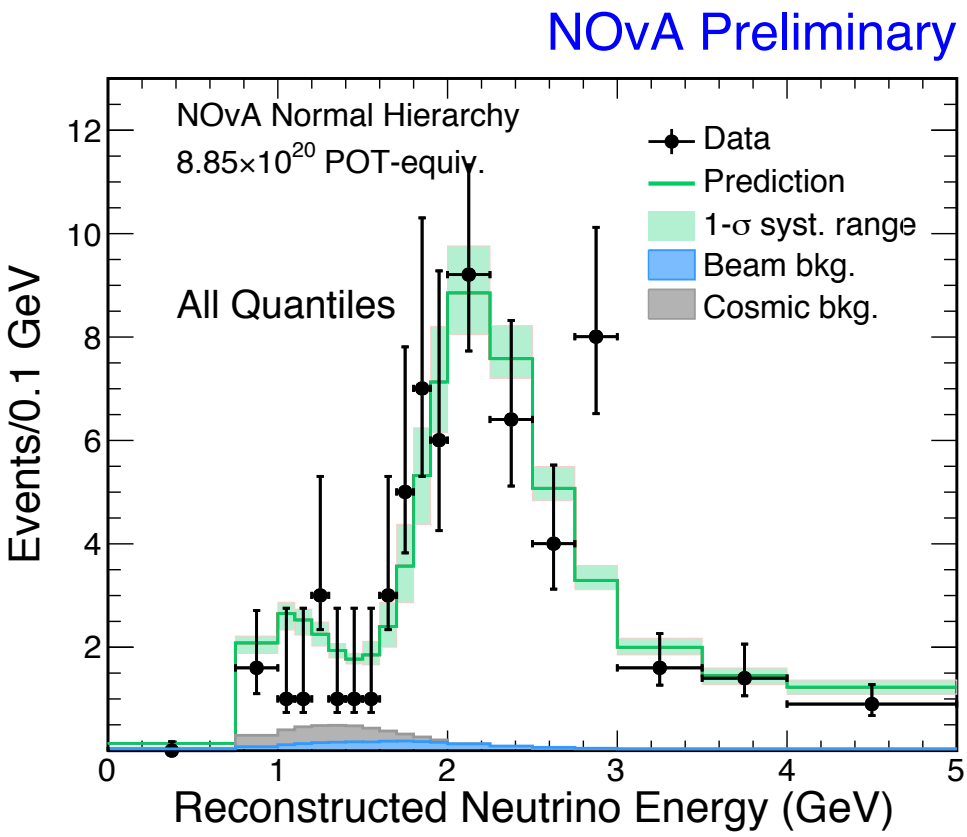

Figure 7.5: Selected data and MC reconstructed energy for ν_μCC.

	Total	Quantile 1	Quantile 2	Quantile 3	Quantile 4
Observed Events	**126**	36	24	26	40
Predicted Total	**120.2**	34.35	25.35	26.8	33.7
Background Exp	**9.2**	0.35	0.45	0.95	7.4

Table 7.2: Observed surviving ν_μ events at the far detector, compared to predictions.

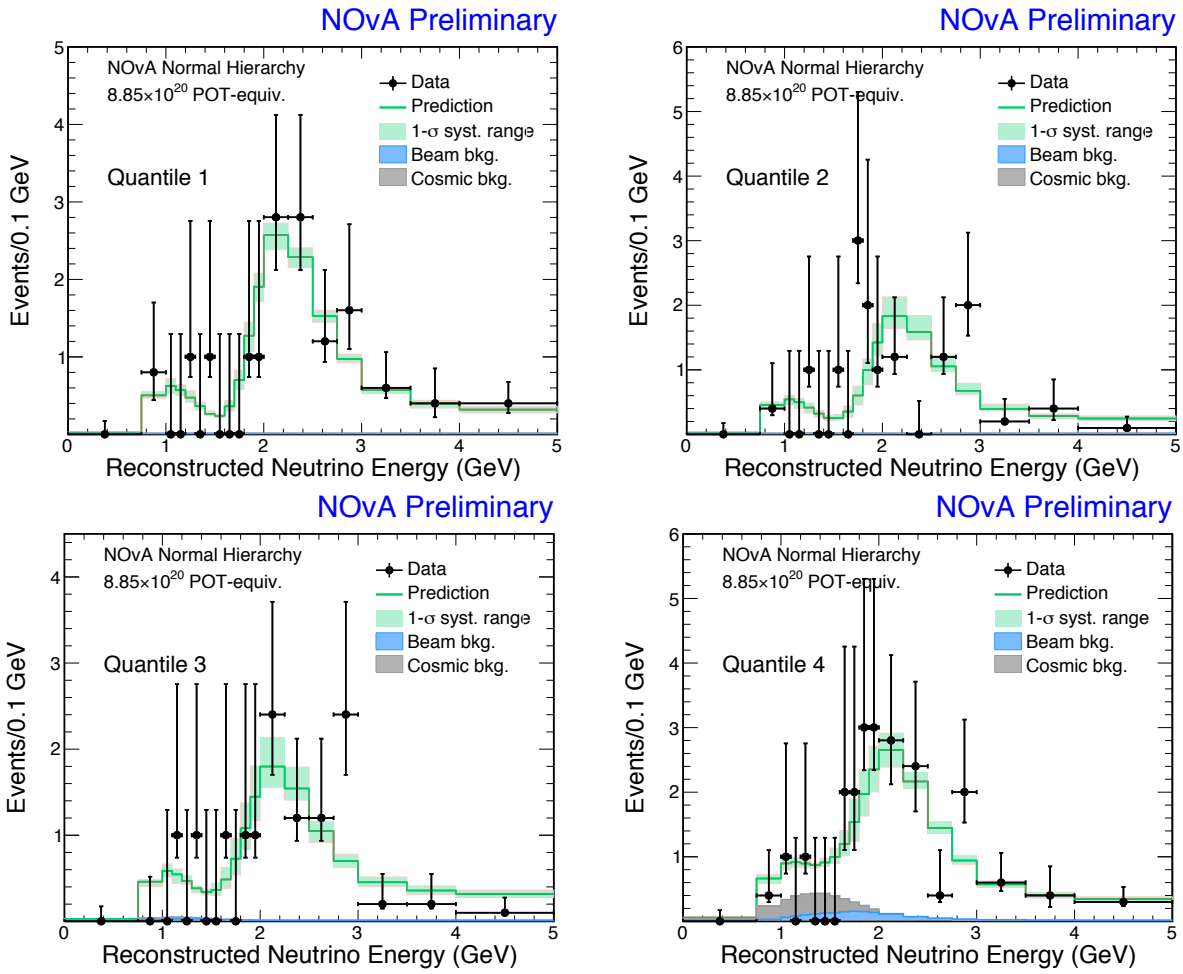

Figure 7.6: Signal and reconstructed energy for ν_μCC for each quantile.

ν_μ DISAPPEARANCE RESULTS

The event spectra at the far detector are shown in Figs.7.5 and 7.6 for surviving ν_μ events. In this channel we see 126 events over 9.2 expected background. Events are distributed in the four quantiles of fraction of hadronic energy as seen in the Table 7.2. The following figures and tables show the data compared to the best fit prediction.

One interesting characteristic to observe in the quantile distributions is that the background is mostly expected in the fourth quantile, and that the events in the region of maximal mixing (around 1.4 GeV) are distributed amongst all samples. Naturally, the closest this mixing is to maximal, the fewer events will be expected in that energy region and the higher precision will be needed to resolve this question.

Disappearance of muon neutrinos from NuMI had been observed unambiguously by NOvA since our first result in 2014, but the increase of statistics and the improvements in the analysis have an impact in our sensitivity to Δm^2_{32} , $\sin^2\theta_{23}$ and whether this mixing is maximal.

Details of the fit to this observation are described in Sec. 7.2, where results are shown for a fit including the results from both the appearance and the disappearance channels. Even without a joint fit, the disappearance channel has sensitivity to the parameters Δm^2_{32} and $\sin^2\theta_{23}$, which is shown in Fig.7.7.

7.2 Fit to oscillations

"It is not unscientific to make a guess, although many people who are not in science think it is."
– Richard Feynman

The resulting constraints on the parameters of the mixing matrix were computed from the results from both measurements in combination with the existing limits from other experimental results for the well known parameters.

This section presents the results of fits made to the appearance and disappearance results shown in the previous section, in order to constrain the values of the oscillation parameters Δm^2_{32} , $\sin^2\theta_{23}$, and δ_{CP} . Fig.7.7 shows a comparison of the limits to Δm^2_{32} and $\sin^2\theta_{23}$ from a fit to the disappearance data only, in comparison with results from other experiments.

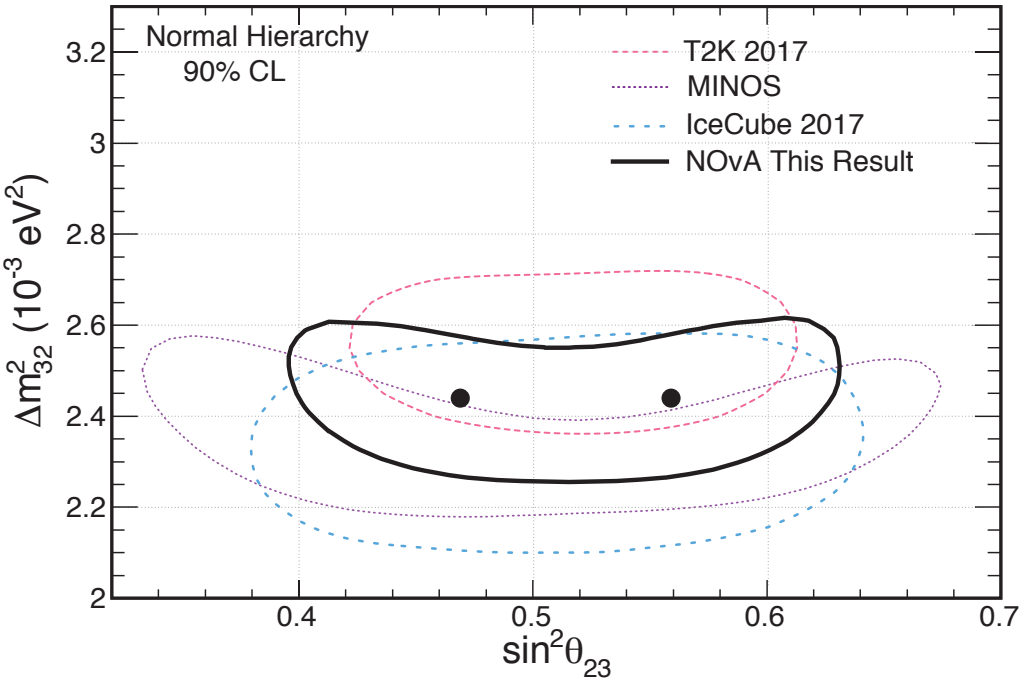

Figure 7.7: Results: 90% CL constraints for $\sin^2\theta_{23} \times \Delta m^2_{32}$ from a fit to the ν_μ disappearance only (black solid line), compared with results from T2K [101] in pink dashed, MINOS [17] in purple dotted and IceCube [102] in blue dashed. All contours are 90% confidence levels.

Figures 7.8, 7.9, and 7.9 show constraints obtained from joint fits of the dissappearance and appearance results, and Fig.7.11 shows the exclusion significance over the full range of δ_{CP} and for best fit values. The values used for the remaining oscillation parameters are the best fit values given by [43], including the world average for $\sin^2 2\theta_{13} = 0.082 \pm 0.004$ as a constraint, currently driven by the results of reactor experiments.

The best fits were obtained by profiling over $\sin^2 2\theta_{13}$, δ_{CP}, Δm^2_{32}, and the systematic uncertainties shown in Sec. 6.10.6, and include corrections made with the Feldman-Cousins procedure [103]. The best fit points are reported in the following table:

Ordering	θ_{23} Octant	δ_{CP}	$\sin^2\theta_{23}$	Δm^2_{32}	$\sin^2 2\theta_{13}$	LL
Normal	Upper	1.21 π	0.558	2.445	0.082	84.57
Normal	Lower	1.46 π	0.474	2.435	0.082	84.70
Inverted	Upper	1.47 π	0.558	-2.510	0.083	87.11
Inverted	Upper	-	-	-	-	-

Table 7.3: Best fit values of the oscillation parameters.

This result achieves an rejection of the inverted ordering with $\delta_{CP} = \pi/2$ above 3σ, and is approaching a 2σ rejection of the inverted ordering over the whole region of δ_{CP} values. The values obtained for Δm^2_{32} are competitive with existing limits, and the best fit value of $\sin^2\theta_{23}$ is consistent with maximal mixing. NOvA will continue to take data until 2024, alternating between neutrino and anti-neutrino beam running. For the assumed best fits on the normal hierarchy and the upper octant, this yields a sensitivity larger than 4σ for the determination of the hierarchy and larger than 2σ for CP violation.

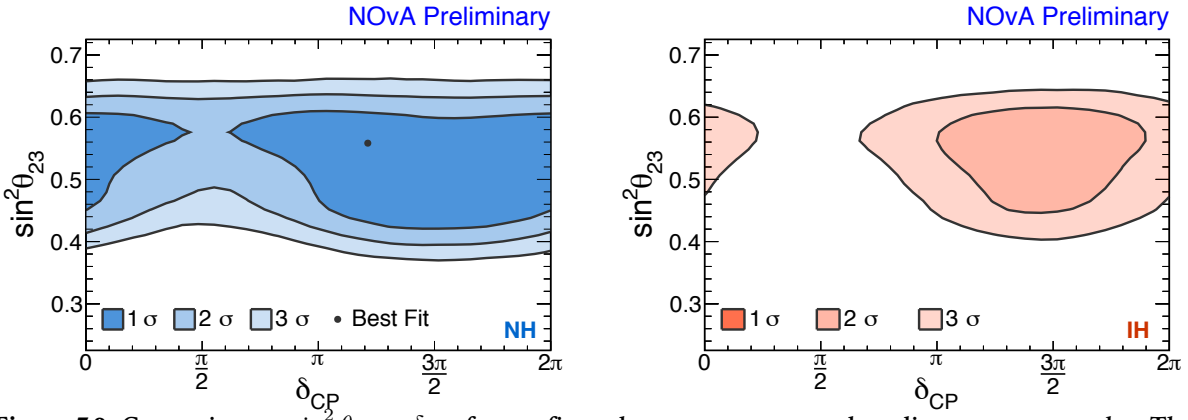

Figure 7.8: Constraints on $\sin^2\theta_{23} \times \delta_{CP}$ from a fit to the ν_e appearance and ν_μ disappearance results. The 1, 2, and 3σ allowed regions are shown from darkest to lightest, for the normal (blue (left)) and inverted (orange (right)) ordering.

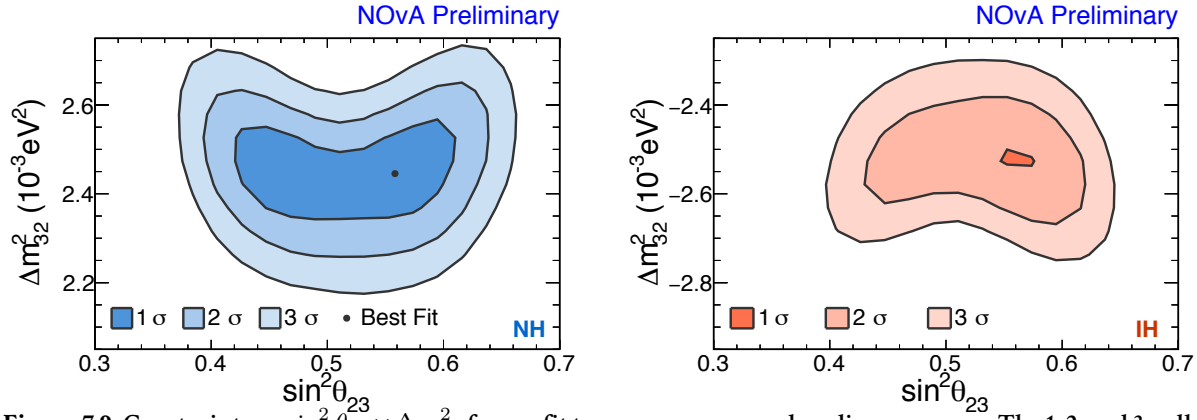

Figure 7.9: Constraints on $\sin^2\theta_{23} \times \Delta m^2_{32}$ from a fit to ν_e appearance and ν_μ disappearance. The 1, 2, and 3σ allowed regions are shown from darkest to lightest, for the normal (blue (left)) and inverted (orange (right)) ordering.

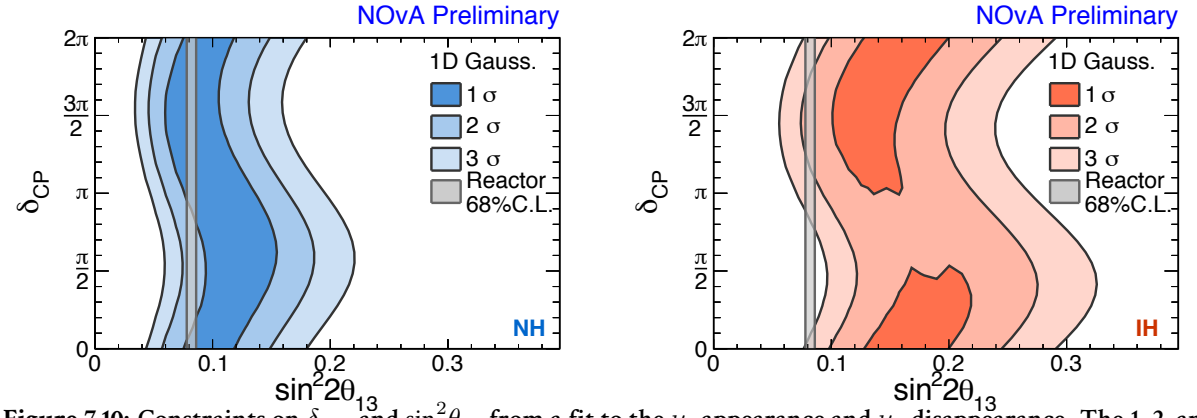

Figure 7.10: Constraints on δ_{CP} and $\sin^2\theta_{23}$ from a fit to the ν_e appearance and ν_μ disappearance. The 1, 2, and 3σ allowed regions are shown from darkest to lightest, for the normal (blue (left)) and inverted (orange (right)) ordering. The reactor constraint (gray) is not used in this fit. Feldman-Cousins corrections are not applied.

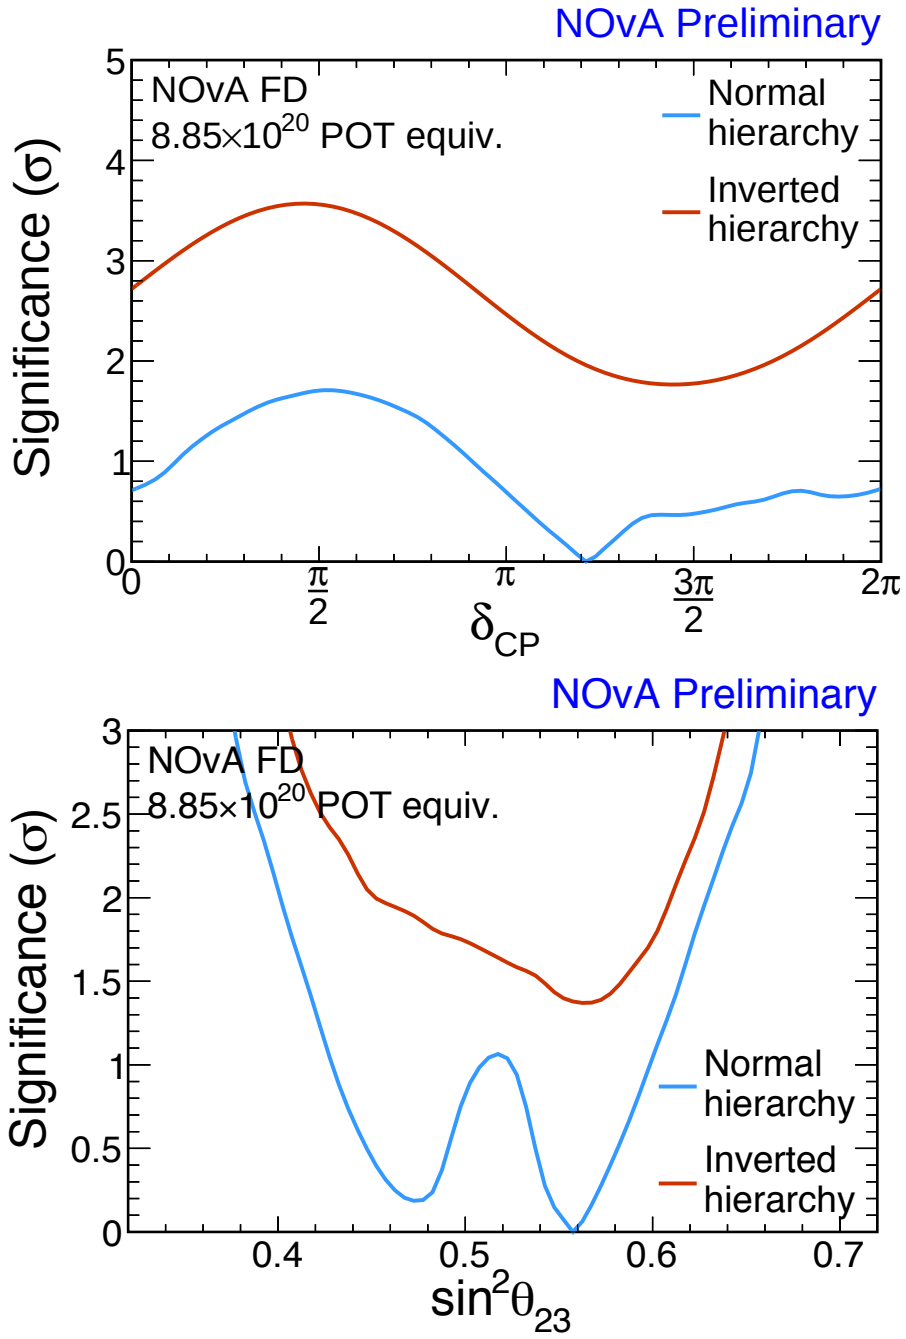

Figure 7.11: Significance of the excluded values of δ_{CP} (top) and $\sin^2\theta_{23}$ (bottom), according to the joint fit described in this section.

7.3 Conclusions

In this dissertation I have shown NOvA's current results for three-flavor oscillations, which are narrowing down the allowed values of the remaining oscillation parameters available at long baselines. While the hierarchy of the neutrino masses may be determined by NOvA in the coming years, the question of CP violation may remain for experiments like Hyper-K and DUNE to resolve at high significance.

NOvA has not only contributed these results to the understanding of neutrinos, we have also contributed the new deep learning based analysis techniques which I have shown here. While I have shown the improvements in terms of signal selection and single particle identification provided by these techniques, their development was always intended to benefit the reconstruction capabilities of NOvA as a whole, beyond the analyses here presented.

The techniques presented in this document are used on NOvA's sterile neutrino searches [104], and are being implemented in signal selection of specific final states for cross sections measurements. Since our first application of CVN, similar techniques have also been employed in other neutrino experiments [78, 105–107].

APPENDIX A

CNN Concepts, Glossary and Architectures

This appendix includes a description of elements common to artificial neural networks NNs and convolutional neural networks CNNs. It starts from a brief description of the functions performed by nodes in simple NN architectures as well as the basic concepts of learning methods, loss functions, etc.

A.0.1 Basic Structure of Neural Networks

Neural Networks are computing algorithms that progressively **improve their own performance** at—or in other words, *learn*—a given task. There are many tasks for which NNs can be designed but in practice all tasks are to approximate an unknown function F of multiple input parameters $x_0, x_1, ...x_n$ given a collection of data to learn from. Data to learn from is usually a collection of examples for which x_i and $F(x_0 i)$ are known or have been simulated. Regression tasks involve approximating a function whose output is a single value (like estimating the total energy of a particle) while *classification* tasks approximate functions whose output is a vector of probabilities to belong to each of multiple classes.[1]

[1] Note that knowing the form of the target function is not necessarily important. What NNs try to do in practice is produce the same output given the same set of inputs. For example, a NN trained to reproduce the invariant mass of an initial state does not need to know decay kinematics but it learns to reproduce the same principles by training on final state

A good way to illustrate the training process is to explain it in the context of a binary classification task. The behavior of NNs as they *train* to improve at binary (signal or background, for example) classification tasks can be understood in three steps which they will iterate on: compute scores or probabilities i.e. propose an estimated function $f(x_i)$, evaluate its performance, and readjust the estimate.

COMPUTING PROBABILITIES

The basic structure of a fully connected network. The output (also called *score*) depicted is a probability of belonging to one of the two categories, signal and background in this example.

$$f(x_i; w_j, b) = p_s ig \tag{A.1}$$

In this case let's assume that $f(x_i; w_j, b) = 1$ is 100% probability of being signal, rather than background, and vice-versa. The way in which each node computes it's contribution to this score is through an *activation function* $f_n(x_i; w_{n,i}, b_n)$. At each iteration, the output from a neuron is computed for a set of values a_i from the training data.

Activation functions which compute scores as inner products of vectors of weights w and variables x are common for classification tasks. An example, which we use in the classification tasks described in Sec. 5.2.1 is the *softmax* activation function. This functions computes the inner product describes and includes an additional normalization function to guarantee score values between 0 and 1. This normalization factor is given by

$$\sigma(z)_j = \frac{e^{z_j}}{\sum_{k=1}^{K} e^{z_k}} \tag{A.2}$$

vectors with corresponding invariant mass values.

which is what is referred to as the *softmax fuction* in the literature. Here z is a vector of the inputs to the output layer (if you have 10 output units, then there are 10 elements in z). And again, j indexes the output units, so j = 1, 2, ..., K.

A.0.2 Calculating Losses

Given the above description, the goal of training a network is to learn how to perform the task (in this case, classify as signal or background) by modifying the weights w according to the performance of $f(x_i; w_j, b)$ or how closely it maps the behavior $p_sig(x_i, c)$, which is the function we are trying to approximate where c is the true class for an example with parameters x_i. In other words, the goal is to minimize a function $L(p_sig(x_i, c), f(x_i; w_j, b))$ called the *loss function* by varying w_i and b. A loss function is designed to give a measure of the performance or goodness of fit of f with respect to p_sig in general, but which incorporates other considerations:

- The output value of L must be smaller for better agreement between the true function $p_sig(x_i, c)$ and the modeled estimate $f(x_i; w_j, b)$. Usually a positive valued function is used.
- L must not be degenerate for different values of w_i and b.
- L is to be calculated over a number N of examples or test data.

A common form for a loss function which satisfies these considerations is

$$L = \frac{1}{N}\sum_{n}^{N} L_n + \lambda R(w_i) \tag{A.3}$$

where the first term is the loss calculation over the data and the second term is a penalty called *regularization*, which ensures L is singled valued for a given set of weights and λ controls the relative variation of the weight regularization with respect to the data loss.

In practice for the case of a *softmax* classifier typically use]it cross-entropy loss functions. This function, in terms of the real probability for each class indexed k, p_k and the approximation made by the NN f_k is given by:

$$L_n = -log\left(\frac{e^{p_{k=c}(x_j,c)}}{\sum_k^K e^{f_k(x_i,w_j)}}\right). \tag{A.4}$$

A.0.3 Network Architectures

The following pages contain the caffe network architectures described in Sec. 5.2.2 and 5.2.3. The input data layers for these networks consist of pixel maps of dimensions 80×200 pixels, in the scale shown in Fig.5.18. The following architectures are shown in caffe prototxt format:

Siamese 2-tower Architecture. This is the original network from [81], used for neutrino event classification. The input data to the towers are event views, as shown in Fig.5.23.

Siamese 2-tower Architecture v2. This is the optimized network used for neutrino event classification on neutrino (fhc) and antineutrino (rhc) events. The number of inception modules is lower and the depth of the network is reduced with respect to the v1 network. The input data to the towers are event views, as shown in Fig.5.23.

Siamese-Siamese (4-tower) Architecture. This network is used for single particle classification. The input data to the four towers are prong views and full event views, as shown in Fig.5.28.

```
name: "Siamese-Siamese (4-tower) Architecture
      for Prong Classification"

layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TRAIN
  }
  transform_param {
    mirror: true
  }
  data_param {
    source: "CVN_TRAINSAMPLE_PATH"
    batch_size: 16
    prefetch: 10
    backend: LEVELDB
  }
}
layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TEST
  }
  data_param {
    source: "CVN_TESTSAMPLE_PATH"
    batch_size: 64
    prefetch: 10
    backend: LEVELDB
  }
}
layer {
  name: "jitter"
  type: "DummyData"
  top: "jitter"
  include {
    phase: TRAIN
  }
  dummy_data_param {
    data_filler {
        type: "gaussian"
        mean: 1.0
        std:  0.01
      }
    shape {
       dim: 16
       dim: 4
       dim: 100
       dim: 80
      }
  }
}

layer {
  name: "jitter"
  type: "DummyData"
  top: "jitter"
  include {
    phase: TEST
  }
  dummy_data_param {
    data_filler {
        type: "constant"
        value:  1.0
      }
    shape {
       dim: 64
       dim: 4
       dim: 100
       dim: 80
      }
  }
}
layer {
  name: "jitteredData"
  type: "Eltwise"
  bottom: "data"
  bottom: "jitter"
  top: "jitteredData"
  eltwise_param {
   operation:PROD
  }
}

layer {
  name: "slice"
  type: "Slice"
  bottom: "jitteredData"
  top: "data_x"
  top: "data_y"
  top: "data_px"
  top: "data_py"
  slice_param {
     slice_dim: 1
     slice_point: 1
     slice_point: 2
     slice_point: 3
  }
}

########X View######

layer {
  name: "conv1/7x7_s2_x"
  type: "Convolution"
  bottom: "data_x"
  top: "conv1/7x7_s2_x"
  param {
    name: "conv1/7x7_s2_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv1/7x7_s2_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    pad: 3
    kernel_size: 7
    stride: 2
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv1/relu_7x7_x"
  type: "ReLU"
  bottom: "conv1/7x7_s2_x"
  top: "conv1/7x7_s2_x"
}


layer {
  name: "pool1/3x3_s2_x"
  type: "Pooling"
  bottom: "conv1/7x7_s2_x"
  top: "pool1/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "pool1/norm1_x"
  type: "LRN"
  bottom: "pool1/3x3_s2_x"
  top: "pool1/norm1_x"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

layer {
  name: "conv2/3x3_reduce_x"
  type: "Convolution"
  bottom: "pool1/norm1_x"
  top: "conv2/3x3_reduce_x"
  param {
    name: "conv2/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "conv2/3x3_reduce_x"
  top: "conv2/3x3_reduce_x"
}

layer {
  name: "conv2/3x3a_x"
  type: "Convolution"
  bottom: "conv2/3x3_reduce_x"
  top: "conv2/3x3a_x"
  param {
    name: "conv2/3x3a_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3a_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3a_x"
  type: "ReLU"
  bottom: "conv2/3x3a_x"
  top: "conv2/3x3a_x"
}

layer {
  name: "conv2/3x3_x"
  type: "Convolution"
  bottom: "conv2/3x3a_x"
  top: "conv2/3x3_x"
  param {
    name: "conv2/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    stride: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_x"
  type: "ReLU"
  bottom: "conv2/3x3_x"
  top: "conv2/3x3_x"
}

layer {
  name: "conv2/norm2_x"
  type: "LRN"
  bottom: "conv2/3x3_x"
  top: "conv2/norm2_x"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "pool2/3x3_s2_x"
  type: "Pooling"
  bottom: "conv2/norm2_x"
  top: "pool2/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "inception_3a/1x1_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/1x1_x"
  param {
    name: "inception_3a/1x1_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/1x1_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_1x1_x"
  type: "ReLU"
  bottom: "inception_3a/1x1_x"
  top: "inception_3a/1x1_x"
}
layer {
  name: "inception_3a/3x3_reduce_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/3x3_reduce_x"
  param {
    name: "inception_3a/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 48
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "inception_3a/3x3_reduce_x"
  top: "inception_3a/3x3_reduce_x"
}
layer {
  name: "inception_3a/3x3_x"
  type: "Convolution"
  bottom: "inception_3a/3x3_reduce_x"
  top: "inception_3a/3x3_x"
  param {
    name: "inception_3a/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_x"
  type: "ReLU"
  bottom: "inception_3a/3x3_x"
  top: "inception_3a/3x3_x"
}

layer {
  name: "inception_3a/5x5_reduce_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/5x5_reduce_x"
  param {
    name: "inception_3a/5x5_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 8
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_reduce_x"
  type: "ReLU"
  bottom: "inception_3a/5x5_reduce_x"
  top: "inception_3a/5x5_reduce_x"
}

layer {
  name: "inception_3a/5x5_x"
  type: "Convolution"
  bottom: "inception_3a/5x5_reduce_x"
  top: "inception_3a/5x5_x"
  param {
    name: "inception_3a/5x5_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    pad: 2
    kernel_size: 5
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_x"
  type: "ReLU"
  bottom: "inception_3a/5x5_x"
  top: "inception_3a/5x5_x"
}

layer {
  name: "inception_3a/pool_x"
  type: "Pooling"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/pool_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 1
    pad: 1
  }
}
layer {
  name: "inception_3a/pool_proj_x"
  type: "Convolution"
  bottom: "inception_3a/pool_x"
  top: "inception_3a/pool_proj_x"
  param {
    name: "inception_3a/pool_proj_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/pool_proj_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_pool_proj_x"
  type: "ReLU"
  bottom: "inception_3a/pool_proj_x"
  top: "inception_3a/pool_proj_x"
}
layer {
  name: "inception_3a/output_x"
  type: "Concat"
  bottom: "inception_3a/1x1_x"
  bottom: "inception_3a/3x3_x"
  bottom: "inception_3a/5x5_x"
  bottom: "inception_3a/pool_proj_x"
  top: "inception_3a/output_x"
}

layer {
  name: "pool3a/3x3_s2_x"
  type: "Pooling"
  bottom: "inception_3a/output_x"
  top: "pool3a/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}

########Y View######

layer {
  name: "conv1/7x7_s2_y"
  type: "Convolution"
  bottom: "data_y"
  top: "conv1/7x7_s2_y"
  param {
    name: "conv1/7x7_s2_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv1/7x7_s2_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    pad: 3
    kernel_size: 7
    stride: 2
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv1/relu_7x7_y"
  type: "ReLU"
  bottom: "conv1/7x7_s2_y"
  top: "conv1/7x7_s2_y"
}


layer {
  name: "pool1/3x3_s2_y"
  type: "Pooling"
  bottom: "conv1/7x7_s2_y"
  top: "pool1/3x3_s2_y"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "pool1/norm1_y"
  type: "LRN"
  bottom: "pool1/3x3_s2_y"
  top: "pool1/norm1_y"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

layer {
  name: "conv2/3x3_reduce_y"
  type: "Convolution"
  bottom: "pool1/norm1_y"
  top: "conv2/3x3_reduce_y"
  param {
    name: "conv2/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_reduce_y"
  type: "ReLU"
  bottom: "conv2/3x3_reduce_y"
  top: "conv2/3x3_reduce_y"
}

layer {
  name: "conv2/3x3a_y"
  type: "Convolution"
  bottom: "conv2/3x3_reduce_y"
  top: "conv2/3x3a_y"
  param {
    name: "conv2/3x3a_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3a_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3a_y"
  type: "ReLU"
  bottom: "conv2/3x3a_y"
  top: "conv2/3x3a_y"
}

layer {
  name: "conv2/3x3_y"
  type: "Convolution"
  bottom: "conv2/3x3a_y"
  top: "conv2/3x3_y"
  param {
    name: "conv2/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    stride: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_y"
  type: "ReLU"
  bottom: "conv2/3x3_y"
  top: "conv2/3x3_y"
}
```

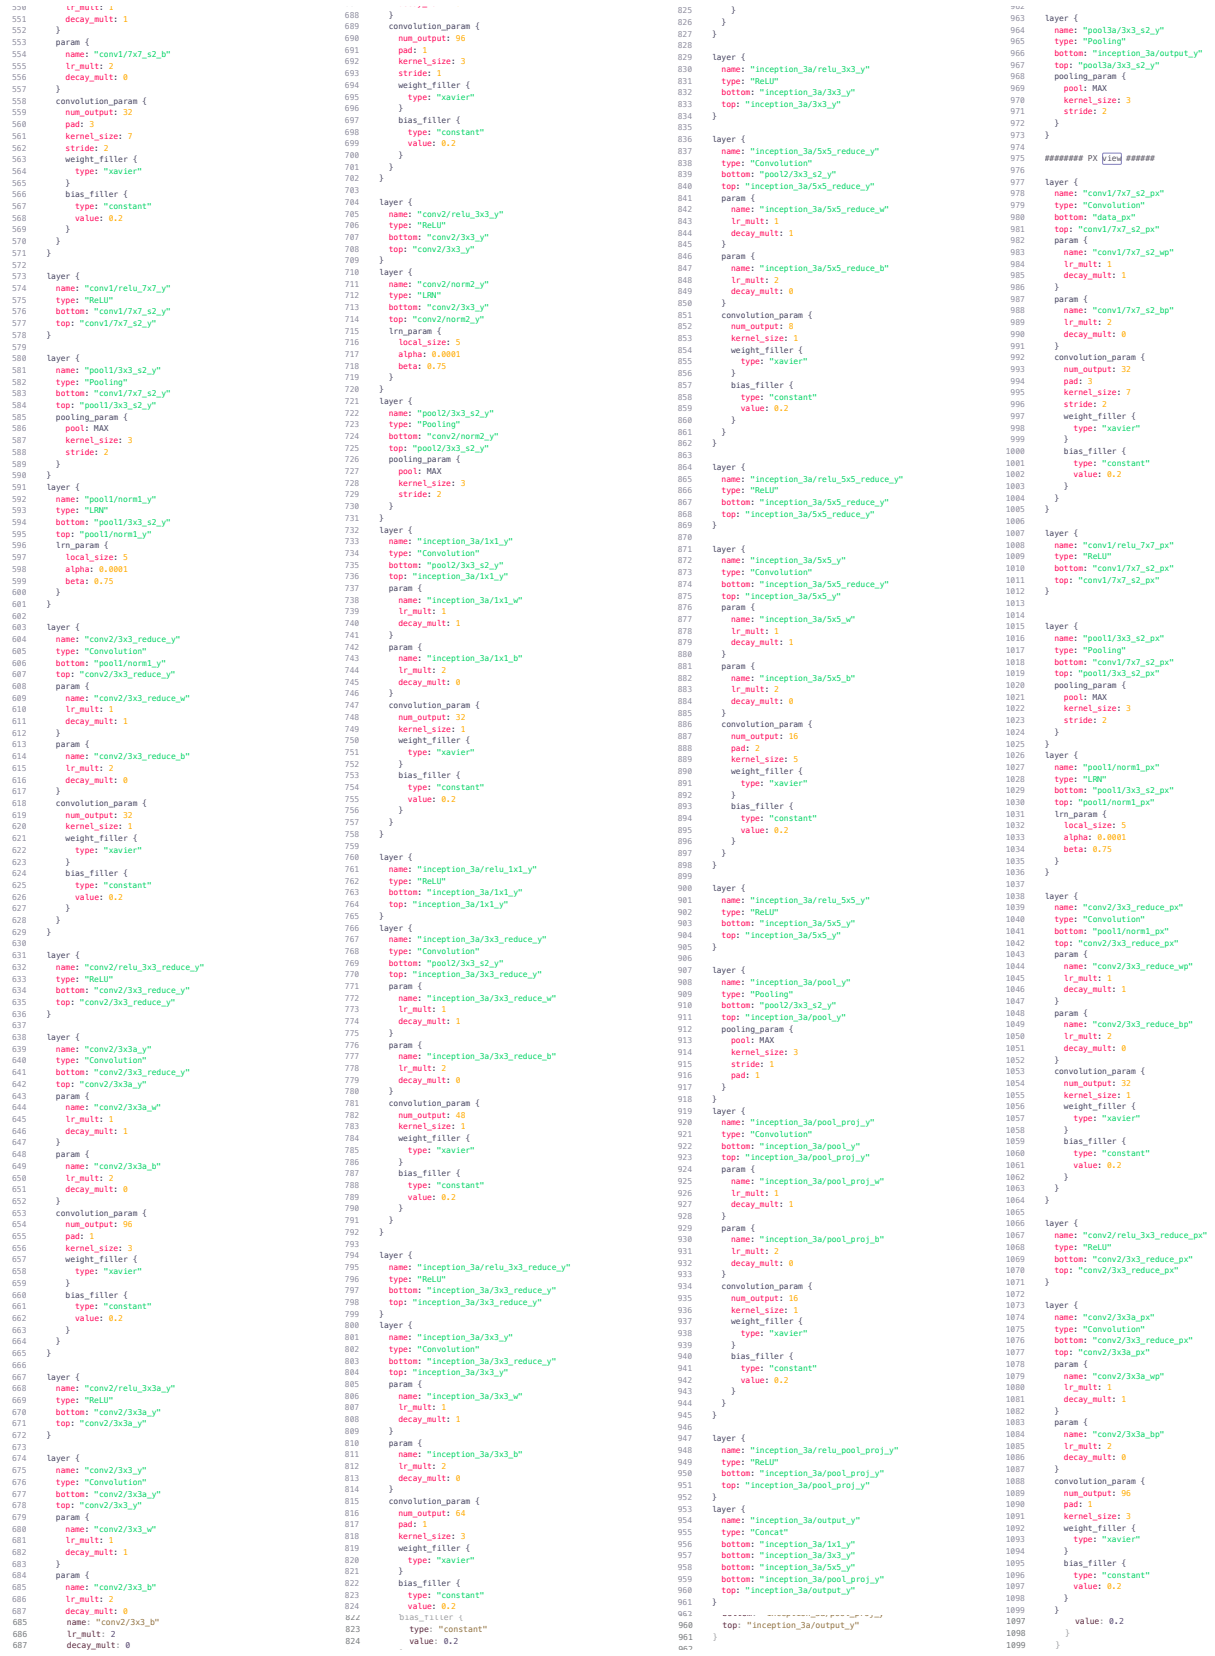

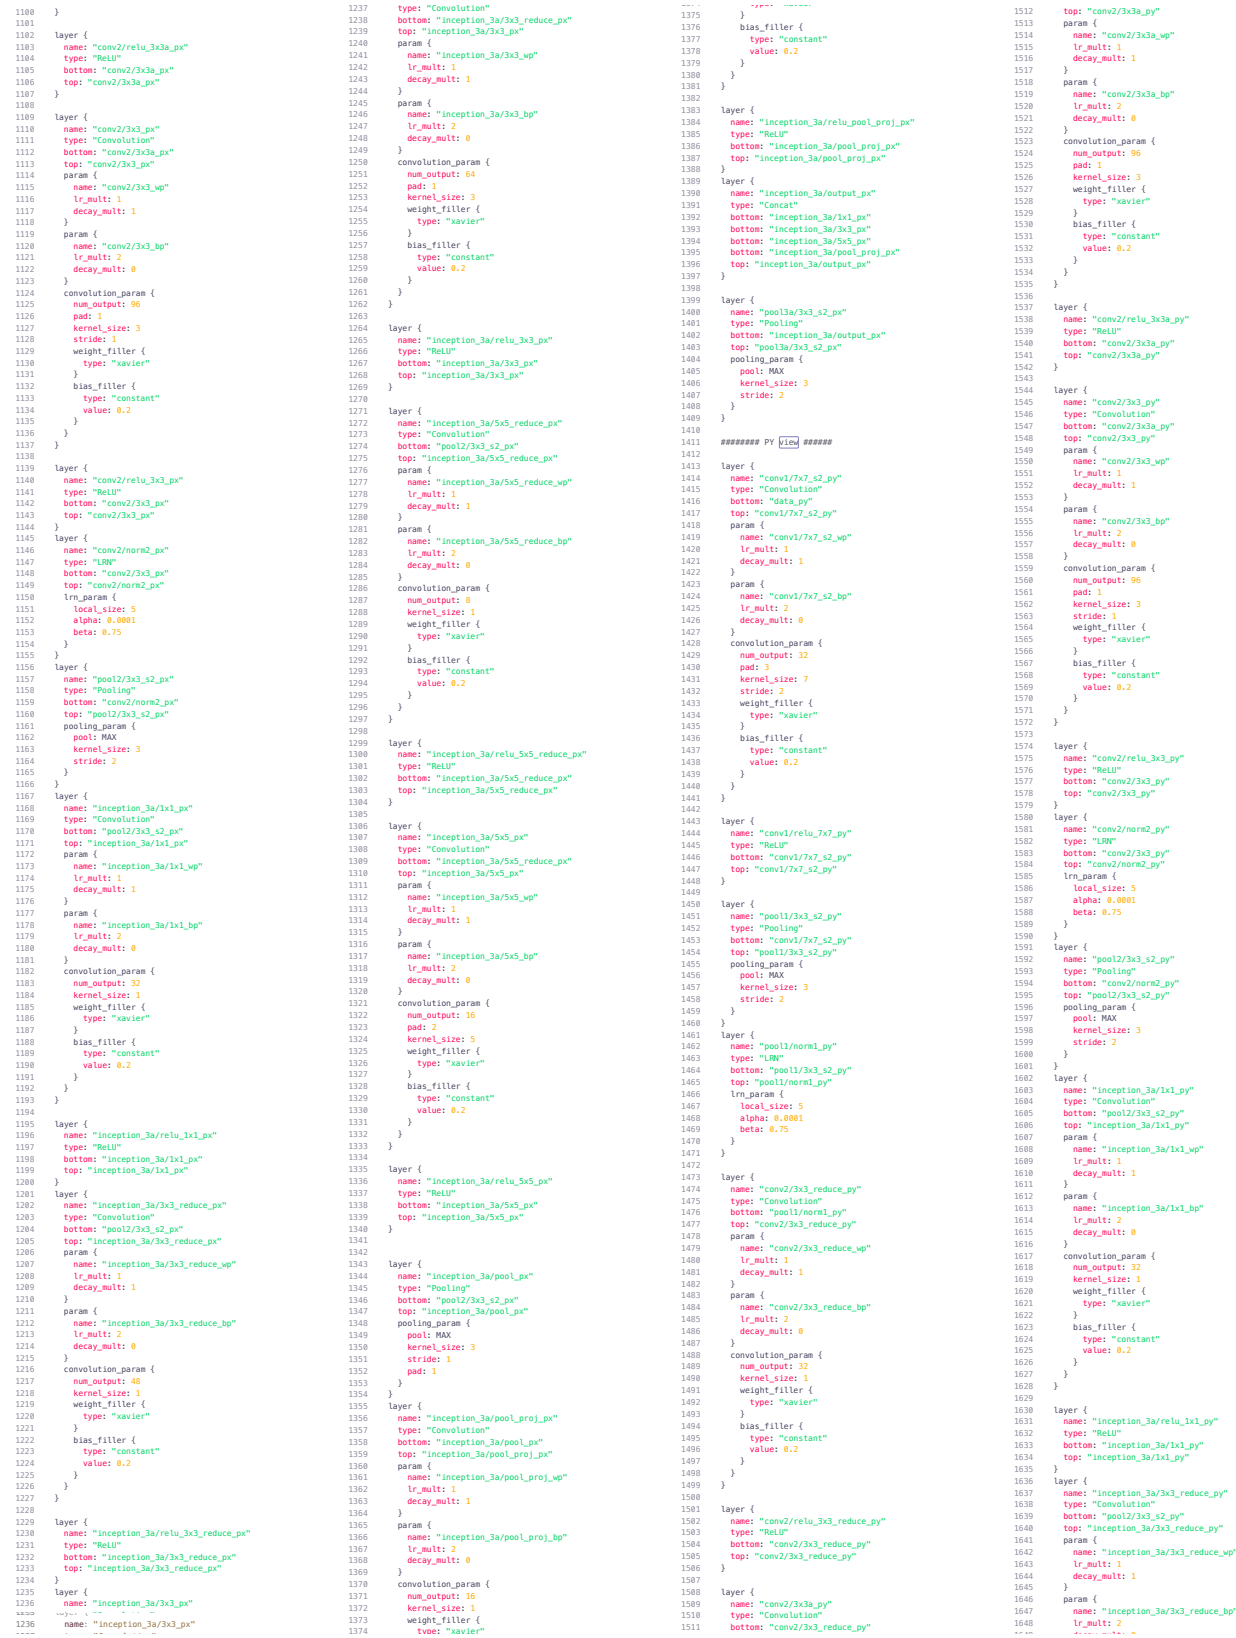

```
}

layer {
  name: "conv2/relu_3x3a_px"
  type: "ReLU"
  bottom: "conv2/3x3a_px"
  top: "conv2/3x3a_px"
}

layer {
  name: "conv2/3x3_px"
  type: "Convolution"
  bottom: "conv2/3x3a_px"
  top: "conv2/3x3_px"
  param {
    name: "conv2/3x3_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    stride: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_px"
  type: "ReLU"
  bottom: "conv2/3x3_px"
  top: "conv2/3x3_px"
}
layer {
  name: "conv2/norm2_px"
  type: "LRN"
  bottom: "conv2/3x3_px"
  top: "conv2/norm2_px"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "pool2/3x3_s2_px"
  type: "Pooling"
  bottom: "conv2/norm2_px"
  top: "pool2/3x3_s2_px"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "inception_3a/1x1_px"
  type: "Convolution"
  bottom: "pool2/3x3_s2_px"
  top: "inception_3a/1x1_px"
  param {
    name: "inception_3a/1x1_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/1x1_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_1x1_px"
  type: "ReLU"
  bottom: "inception_3a/1x1_px"
  top: "inception_3a/1x1_px"
}
layer {
  name: "inception_3a/3x3_reduce_px"
  type: "Convolution"
  bottom: "pool2/3x3_s2_px"
  top: "inception_3a/3x3_reduce_px"
  param {
    name: "inception_3a/3x3_reduce_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_reduce_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 48
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_reduce_px"
  type: "ReLU"
  bottom: "inception_3a/3x3_reduce_px"
  top: "inception_3a/3x3_reduce_px"
}
layer {
  name: "inception_3a/3x3_px"
  name: "inception_3a/3x3_px"
  type: "Convolution"
  bottom: "inception_3a/3x3_reduce_px"
  top: "inception_3a/3x3_px"
  param {
    name: "inception_3a/3x3_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_px"
  type: "ReLU"
  bottom: "inception_3a/3x3_px"
  top: "inception_3a/3x3_px"
}

layer {
  name: "inception_3a/5x5_reduce_px"
  type: "Convolution"
  bottom: "pool2/3x3_s2_px"
  top: "inception_3a/5x5_reduce_px"
  param {
    name: "inception_3a/5x5_reduce_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_reduce_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 8
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_reduce_px"
  type: "ReLU"
  bottom: "inception_3a/5x5_reduce_px"
  top: "inception_3a/5x5_reduce_px"
}

layer {
  name: "inception_3a/5x5_px"
  type: "Convolution"
  bottom: "inception_3a/5x5_reduce_px"
  top: "inception_3a/5x5_px"
  param {
    name: "inception_3a/5x5_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    pad: 2
    kernel_size: 5
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_px"
  type: "ReLU"
  bottom: "inception_3a/5x5_px"
  top: "inception_3a/5x5_px"
}


layer {
  name: "inception_3a/pool_px"
  type: "Pooling"
  bottom: "pool2/3x3_s2_px"
  top: "inception_3a/pool_px"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 1
    pad: 1
  }
}
layer {
  name: "inception_3a/pool_proj_px"
  type: "Convolution"
  bottom: "inception_3a/pool_px"
  top: "inception_3a/pool_proj_px"
  param {
    name: "inception_3a/pool_proj_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/pool_proj_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_pool_proj_px"
  type: "ReLU"
  bottom: "inception_3a/pool_proj_px"
  top: "inception_3a/pool_proj_px"
}
layer {
  name: "inception_3a/output_px"
  type: "Concat"
  bottom: "inception_3a/1x1_px"
  bottom: "inception_3a/3x3_px"
  bottom: "inception_3a/5x5_px"
  bottom: "inception_3a/pool_proj_px"
  top: "inception_3a/output_px"
}

layer {
  name: "pool3a/3x3_s2_px"
  type: "Pooling"
  bottom: "inception_3a/output_px"
  top: "pool3a/3x3_s2_px"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}

######## PY view ######

layer {
  name: "conv1/7x7_s2_py"
  type: "Convolution"
  bottom: "data_py"
  top: "conv1/7x7_s2_py"
  param {
    name: "conv1/7x7_s2_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv1/7x7_s2_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    pad: 3
    kernel_size: 7
    stride: 2
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv1/relu_7x7_py"
  type: "ReLU"
  bottom: "conv1/7x7_s2_py"
  top: "conv1/7x7_s2_py"
}

layer {
  name: "pool1/3x3_s2_py"
  type: "Pooling"
  bottom: "conv1/7x7_s2_py"
  top: "pool1/3x3_s2_py"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "pool1/norm1_py"
  type: "LRN"
  bottom: "pool1/3x3_s2_py"
  top: "pool1/norm1_py"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

layer {
  name: "conv2/3x3_reduce_py"
  type: "Convolution"
  bottom: "pool1/norm1_py"
  top: "conv2/3x3_reduce_py"
  param {
    name: "conv2/3x3_reduce_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_reduce_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_reduce_py"
  type: "ReLU"
  bottom: "conv2/3x3_reduce_py"
  top: "conv2/3x3_reduce_py"
}

layer {
  name: "conv2/3x3a_py"
  type: "Convolution"
  bottom: "conv2/3x3_reduce_py"
  top: "conv2/3x3a_py"
  param {
    name: "conv2/3x3a_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3a_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3a_py"
  type: "ReLU"
  bottom: "conv2/3x3a_py"
  top: "conv2/3x3a_py"
}

layer {
  name: "conv2/3x3_py"
  type: "Convolution"
  bottom: "conv2/3x3a_py"
  top: "conv2/3x3_py"
  param {
    name: "conv2/3x3_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    pad: 1
    kernel_size: 3
    stride: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_py"
  type: "ReLU"
  bottom: "conv2/3x3_py"
  top: "conv2/3x3_py"
}
layer {
  name: "conv2/norm2_py"
  type: "LRN"
  bottom: "conv2/3x3_py"
  top: "conv2/norm2_py"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "pool2/3x3_s2_py"
  type: "Pooling"
  bottom: "conv2/norm2_py"
  top: "pool2/3x3_s2_py"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "inception_3a/1x1_py"
  type: "Convolution"
  bottom: "pool2/3x3_s2_py"
  top: "inception_3a/1x1_py"
  param {
    name: "inception_3a/1x1_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/1x1_bp"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_1x1_py"
  type: "ReLU"
  bottom: "inception_3a/1x1_py"
  top: "inception_3a/1x1_py"
}
layer {
  name: "inception_3a/3x3_reduce_py"
  type: "Convolution"
  bottom: "pool2/3x3_s2_py"
  top: "inception_3a/3x3_reduce_py"
  param {
    name: "inception_3a/3x3_reduce_wp"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_reduce_bp"
    lr_mult: 2
    decay_mult: 0
  }
```

```
      decay_mult: 0
    }
    convolution_param {
      num_output: 48
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_3a/relu_3x3_reduce_py"
    type: "ReLU"
    bottom: "inception_3a/3x3_reduce_py"
    top: "inception_3a/3x3_reduce_py"
  }
  layer {
    name: "inception_3a/3x3_py"
    type: "Convolution"
    bottom: "inception_3a/3x3_reduce_py"
    top: "inception_3a/3x3_py"
    param {
      name: "inception_3a/3x3_wp"
      lr_mult: 1
      decay_mult: 1
    }
    param {
      name: "inception_3a/3x3_bp"
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      pad: 1
      kernel_size: 3
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_3a/relu_3x3_py"
    type: "ReLU"
    bottom: "inception_3a/3x3_py"
    top: "inception_3a/3x3_py"
  }

  layer {
    name: "inception_3a/5x5_reduce_py"
    type: "Convolution"
    bottom: "pool2/3x3_s2_py"
    top: "inception_3a/5x5_reduce_py"
    param {
      name: "inception_3a/5x5_reduce_wp"
      lr_mult: 1
      decay_mult: 1
    }
    param {
      name: "inception_3a/5x5_reduce_bp"
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 8
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_3a/relu_5x5_reduce_py"
    type: "ReLU"
    bottom: "inception_3a/5x5_reduce_py"
    top: "inception_3a/5x5_reduce_py"
  }


  layer {
    name: "inception_3a/5x5_py"
    type: "Convolution"
    bottom: "inception_3a/5x5_reduce_py"
    top: "inception_3a/5x5_py"
    param {
      name: "inception_3a/5x5_wp"
      lr_mult: 1
      decay_mult: 1
    }
    param {
      name: "inception_3a/5x5_bp"
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 16
      pad: 2
      kernel_size: 5
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_3a/relu_5x5_py"
    type: "ReLU"
    bottom: "inception_3a/5x5_py"
    top: "inception_3a/5x5_py"
  }

  layer {
    name: "inception_3a/pool_py"
    type: "Pooling"
    bottom: "pool2/3x3_s2_py"
    top: "inception_3a/pool_py"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 1
      pad: 1
    }
  }
  layer {
    name: "inception_3a/pool_proj_py"
    type: "Convolution"
    bottom: "inception_3a/pool_py"
    top: "inception_3a/pool_proj_py"
    param {
      name: "inception_3a/pool_proj_wp"
      lr_mult: 1
      decay_mult: 1
    }
    param {
      name: "inception_3a/pool_proj_bp"
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 16
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_3a/relu_pool_proj_py"
    type: "ReLU"
    bottom: "inception_3a/pool_proj_py"
    top: "inception_3a/pool_proj_py"
  }
  layer {
    name: "inception_3a/output_py"
    type: "Concat"
    bottom: "inception_3a/1x1_py"
    bottom: "inception_3a/3x3_py"
    bottom: "inception_3a/5x5_py"
    bottom: "inception_3a/pool_proj_py"
    top: "inception_3a/output_py"
  }

  layer {
    name: "pool3a/3x3_s2_py"
    type: "Pooling"
    bottom: "inception_3a/output_py"
    top: "pool3a/3x3_s2_py"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 2
    }
  }

  ######## Combine Towers ######

  layer {
    name: "merge_x_y"
    type: "Concat"
    bottom: "pool3a/3x3_s2_x"
    bottom: "pool3a/3x3_s2_y"
    bottom: "pool3a/3x3_s2_px"
    bottom: "pool3a/3x3_s2_py"
    top: "merge_x_y"
    concat_param {
      axis: 1
    }
  }
  layer {
    name: "inception_5b/1x1"
    type: "Convolution"
    bottom: "merge_x_y"
    top: "inception_5b/1x1"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 192
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_1x1"
    type: "ReLU"
    bottom: "inception_5b/1x1"
    top: "inception_5b/1x1"
  }
  layer {
    name: "inception_5b/3x3_reduce"
    type: "Convolution"
    bottom: "merge_x_y"
    top: "inception_5b/3x3_reduce"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 96
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_3x3_reduce"
    type: "ReLU"
    bottom: "inception_5b/3x3_reduce"
    top: "inception_5b/3x3_reduce"
  }
  layer {
    name: "inception_5b/3x3"
    type: "Convolution"
    bottom: "inception_5b/3x3_reduce"
    top: "inception_5b/3x3"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 192
      pad: 1
      kernel_size: 3
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_3x3"
    type: "ReLU"
    bottom: "inception_5b/3x3"
    top: "inception_5b/3x3"
  }
  layer {
    name: "inception_5b/5x5_reduce"
    type: "Convolution"
    bottom: "merge_x_y"
    top: "inception_5b/5x5_reduce"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 24
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_5x5_reduce"
    type: "ReLU"
    bottom: "inception_5b/5x5_reduce"
    top: "inception_5b/5x5_reduce"
  }

  layer {
    name: "inception_5b/5x5"
    type: "Convolution"
    bottom: "inception_5b/5x5_reduce"
    top: "inception_5b/5x5"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      pad: 2
      kernel_size: 5
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_5x5"
    type: "ReLU"
    bottom: "inception_5b/5x5"
    top: "inception_5b/5x5"
  }

  layer {
    name: "inception_5b/pool"
    type: "Pooling"
    bottom: "merge_x_y"
    top: "inception_5b/pool"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 1
      pad: 1
    }
  }
  layer {
    name: "inception_5b/pool_proj"
    type: "Convolution"
    bottom: "inception_5b/pool"
    top: "inception_5b/pool_proj"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_pool_proj"
    type: "ReLU"
    bottom: "inception_5b/pool_proj"
    top: "inception_5b/pool_proj"
  }
  layer {
    name: "inception_5b/output"
    type: "Concat"
    bottom: "inception_5b/1x1"
    bottom: "inception_5b/3x3"
    bottom: "inception_5b/5x5"
    bottom: "inception_5b/pool_proj"
    top: "inception_5b/output"
  }
  layer {
    name: "pool5/7x7_s1"
    type: "Pooling"
    bottom: "inception_5b/output"
    top: "pool5/7x7_s1"
    pooling_param {
      pool: AVE
      ## Adapted!!! Originally kernel_size: 7
      kernel_h: 6
      kernel_w: 5
      stride: 1
    }
  }
  layer {
    name: "pool5/drop_7x7_s1"
    type: "Dropout"
    bottom: "pool5/7x7_s1"
    top: "pool5/7x7_s1"
    dropout_param {
      dropout_ratio: 0.4
    }
  }
  layer {
    name: "loss3/classifier15"
    type: "InnerProduct"
    bottom: "pool5/7x7_s1"
    top: "loss3/classifier15"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    inner_product_param {
      num_output: 8
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0
      }
    }
  }
  layer {
    name: "loss3/loss3"
    type: "SoftmaxWithLoss"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/loss3"
    loss_weight: 1
  }
  layer {
    name: "loss3/top-1"
    type: "Accuracy"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/top-1"
    include {
      phase: TEST
    }
  }
  layer {
    name: "loss3/top-5"
    type: "Accuracy"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/top-5"
    include {
      phase: TEST
    }
    accuracy_param {
      top_k: 5
    }
  }
```

```
name: "Siamese Architecture
       for Event Classification"

layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TRAIN
  }
  transform_param {
    mirror: true
    #scale: 0.00390625
  }
  data_param {
    source: "CVN_TRAINSAMPLE_PATH"
    batch_size: 16
    prefetch: 40
    backend: LEVELDB
  }
}
layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TEST
  }
  data_param {
    source: "CVN_TESTSAMPLE_PATH"
    batch_size: 64
    prefetch: 400
    backend: LEVELDB
  }
}
layer {
  name: "jitter"
  type: "DummyData"
  top: "jitter"
  include {
    phase: TRAIN
  }
  dummy_data_param {
    data_filler {
       type: "gaussian"
       mean: 1.0
       std:  0.01
       }
    shape {
       dim: 16
       dim: 2
       dim: 100
       dim: 80
       }
    }
}

layer {
  name: "jitter"
  type: "DummyData"
  top: "jitter"
  include {
    phase: TEST
  }
  dummy_data_param {
    data_filler {
       type: "constant"
       value:  1.0
       }
    shape {
       dim: 64
       dim: 2
       dim: 100
       dim: 80
       }
    }
}
layer {
  name: "jitteredData"
  type: "Eltwise"
  bottom: "data"
  bottom: "jitter"
  top: "jitteredData"
  eltwise_param {
   operation:PROD
  }
}

layer {
  name: "slice"
  type: "Slice"
  bottom: "jitteredData"
  top: "data_x"
  top: "data_y"
  slice_param {
     slice_dim: 1
     slice_point: 1
  }
}

########################################
# X view ##############################

layer {
  name: "conv1/7x7_s2_x"
  type: "Convolution"
  bottom: "data_x"
  top: "conv1/7x7_s2_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    pad: 3
    kernel_size: 7
    stride: 2
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "conv1/relu_7x7_x"
  type: "ReLU"
  bottom: "conv1/7x7_s2_x"
  top: "conv1/7x7_s2_x"
}
layer {
  name: "pool1/3x3_s2_x"
  type: "Pooling"
  bottom: "conv1/7x7_s2_x"
  top: "pool1/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "pool1/norm1_x"
  type: "LRN"
  bottom: "pool1/3x3_s2_x"
  top: "pool1/norm1_x"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "conv2/3x3_reduce_x"
  type: "Convolution"
  bottom: "pool1/norm1_x"
  top: "conv2/3x3_reduce_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "conv2/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "conv2/3x3_reduce_x"
  top: "conv2/3x3_reduce_x"
}
layer {
  name: "conv2/3x3_x"
  type: "Convolution"
  bottom: "conv2/3x3_reduce_x"
  top: "conv2/3x3_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "conv2/relu_3x3_x"
  type: "ReLU"
  bottom: "conv2/3x3_x"
  top: "conv2/3x3_x"
}
layer {
  name: "conv2/norm2_x"
  type: "LRN"
  bottom: "conv2/3x3_x"
  top: "conv2/norm2_x"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "pool2/3x3_s2_x"
  type: "Pooling"
  bottom: "conv2/norm2_x"
  top: "pool2/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "inception_3a/1x1_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/1x1_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3a/relu_1x1_x"
  type: "ReLU"
  bottom: "inception_3a/1x1_x"
  top: "inception_3a/1x1_x"
}
layer {
  name: "inception_3a/3x3_reduce_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/3x3_reduce_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3a/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "inception_3a/3x3_reduce_x"
  top: "inception_3a/3x3_reduce_x"
}
layer {
  name: "inception_3a/3x3_x"
  type: "Convolution"
  bottom: "inception_3a/3x3_reduce_x"
  top: "inception_3a/3x3_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3a/relu_3x3_x"
  type: "ReLU"
  bottom: "inception_3a/3x3_x"
  top: "inception_3a/3x3_x"
}
layer {
  name: "inception_3a/5x5_reduce_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/5x5_reduce_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3a/relu_5x5_reduce_x"
  type: "ReLU"
  bottom: "inception_3a/5x5_reduce_x"
  top: "inception_3a/5x5_reduce_x"
}
layer {
  name: "inception_3a/5x5_x"
  type: "Convolution"
  bottom: "inception_3a/5x5_reduce_x"
  top: "inception_3a/5x5_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    pad: 2
    kernel_size: 5
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3a/relu_5x5_x"
  type: "ReLU"
  bottom: "inception_3a/5x5_x"
  top: "inception_3a/5x5_x"
}
layer {
  name: "inception_3a/pool_x"
  type: "Pooling"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/pool_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 1
    pad: 1
  }
}
layer {
  name: "inception_3a/pool_proj_x"
  type: "Convolution"
  bottom: "inception_3a/pool_x"
  top: "inception_3a/pool_proj_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3a/relu_pool_proj_x"
  type: "ReLU"
  bottom: "inception_3a/pool_proj_x"
  top: "inception_3a/pool_proj_x"
}
layer {
  name: "inception_3a/output_x"
  type: "Concat"
  bottom: "inception_3a/1x1_x"
  bottom: "inception_3a/3x3_x"
  bottom: "inception_3a/5x5_x"
  bottom: "inception_3a/pool_proj_x"
  top: "inception_3a/output_x"
}
layer {
  name: "inception_3b/1x1_x"
  type: "Convolution"
  bottom: "inception_3a/output_x"
  top: "inception_3b/1x1_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3b/relu_1x1_x"
  type: "ReLU"
  bottom: "inception_3b/1x1_x"
  top: "inception_3b/1x1_x"
}
layer {
  name: "inception_3b/3x3_reduce_x"
  type: "Convolution"
  bottom: "inception_3a/output_x"
  top: "inception_3b/3x3_reduce_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3b/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "inception_3b/3x3_reduce_x"
  top: "inception_3b/3x3_reduce_x"
}
layer {
  name: "inception_3b/3x3_x"
  type: "Convolution"
  bottom: "inception_3b/3x3_reduce_x"
  top: "inception_3b/3x3_x"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_3b/relu_3x3_x"
  type: "ReLU"
  bottom: "inception_3b/3x3_x"
  top: "inception_3b/3x3_x"
}
layer {
```

```
    name: "inception_3b/3x3_reduce_x"
    type: "Convolution"
    bottom: "inception_3a/output_x"
    top: "inception_3b/5x5_reduce_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 32
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_3b/relu_5x5_reduce_x"
    type: "ReLU"
    bottom: "inception_3b/5x5_reduce_x"
    top: "inception_3b/5x5_reduce_x"
  }
  layer {
    name: "inception_3b/5x5_x"
    type: "Convolution"
    bottom: "inception_3b/5x5_reduce_x"
    top: "inception_3b/5x5_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 96
      pad: 2
      kernel_size: 5
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_3b/relu_5x5_x"
    type: "ReLU"
    bottom: "inception_3b/5x5_x"
    top: "inception_3b/5x5_x"
  }
  layer {
    name: "inception_3b/pool_x"
    type: "Pooling"
    bottom: "inception_3a/output_x"
    top: "inception_3b/pool_x"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 1
      pad: 1
    }
  }
  layer {
    name: "inception_3b/pool_proj_x"
    type: "Convolution"
    bottom: "inception_3b/pool_x"
    top: "inception_3b/pool_proj_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_3b/relu_pool_proj_x"
    type: "ReLU"
    bottom: "inception_3b/pool_proj_x"
    top: "inception_3b/pool_proj_x"
  }
  layer {
    name: "inception_3b/output_x"
    type: "Concat"
    bottom: "inception_3b/1x1_x"
    bottom: "inception_3b/3x3_x"
    bottom: "inception_3b/5x5_x"
    bottom: "inception_3b/pool_proj_x"
    top: "inception_3b/output_x"
  }
  layer {
    name: "pool3/3x3_s2_x"
    type: "Pooling"
    bottom: "inception_3b/output_x"
    top: "pool3/3x3_s2_x"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 2
    }
  }
  layer {
    name: "inception_4a/1x1_x"
    type: "Convolution"
    bottom: "pool3/3x3_s2_x"
    top: "inception_4a/1x1_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 192
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_4a/relu_1x1_x"
    type: "ReLU"
    bottom: "inception_4a/1x1_x"
    top: "inception_4a/1x1_x"
  }
  layer {
    name: "inception_4a/3x3_reduce_x"
    type: "Convolution"
    bottom: "pool3/3x3_s2_x"
    top: "inception_4a/3x3_reduce_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 96
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_4a/relu_3x3_reduce_x"
    type: "ReLU"
    bottom: "inception_4a/3x3_reduce_x"
    top: "inception_4a/3x3_reduce_x"
  }
  layer {
    name: "inception_4a/3x3_x"
    type: "Convolution"
    bottom: "inception_4a/3x3_reduce_x"
    top: "inception_4a/3x3_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 208
      pad: 1
      kernel_size: 3
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_4a/relu_3x3_x"
    type: "ReLU"
    bottom: "inception_4a/3x3_x"
    top: "inception_4a/3x3_x"
  }
  layer {
    name: "inception_4a/5x5_reduce_x"
    type: "Convolution"
    bottom: "pool3/3x3_s2_x"
    top: "inception_4a/5x5_reduce_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 16
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_4a/relu_5x5_reduce_x"
    type: "ReLU"
    bottom: "inception_4a/5x5_reduce_x"
    top: "inception_4a/5x5_reduce_x"
  }
  layer {
    name: "inception_4a/5x5_x"
    type: "Convolution"
    bottom: "inception_4a/5x5_reduce_x"
    top: "inception_4a/5x5_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 48
      pad: 2
      kernel_size: 5
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_4a/relu_5x5_x"
    type: "ReLU"
    bottom: "inception_4a/5x5_x"
    top: "inception_4a/5x5_x"
  }
  layer {
    name: "inception_4a/pool_x"
    type: "Pooling"
    bottom: "pool3/3x3_s2_x"
    top: "inception_4a/pool_x"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 1
      pad: 1
    }
  }
  layer {
    name: "inception_4a/pool_proj_x"
    type: "Convolution"
    bottom: "inception_4a/pool_x"
    top: "inception_4a/pool_proj_x"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_4a/relu_pool_proj_x"
    type: "ReLU"
    bottom: "inception_4a/pool_proj_x"
    top: "inception_4a/pool_proj_x"
  }
  layer {
    name: "inception_4a/output_x"
    type: "Concat"
    bottom: "inception_4a/1x1_x"
    bottom: "inception_4a/3x3_x"
    bottom: "inception_4a/5x5_x"
    bottom: "inception_4a/pool_proj_x"
    top: "inception_4a/output_x"
  }

  ######################################
  # Y view ############################

  layer {
    name: "conv1/7x7_s2_y"
    type: "Convolution"
    bottom: "data_y"
    top: "conv1/7x7_s2_y"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      pad: 3
      kernel_size: 7
      stride: 2
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "conv1/relu_7x7_y"
    type: "ReLU"
    bottom: "conv1/7x7_s2_y"
    top: "conv1/7x7_s2_y"
  }
  layer {
    name: "pool1/3x3_s2_y"
    type: "Pooling"
    bottom: "conv1/7x7_s2_y"
    top: "pool1/3x3_s2_y"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 2
    }
  }
  layer {
    name: "pool1/norm1_y"
    type: "LRN"
    bottom: "pool1/3x3_s2_y"
    top: "pool1/norm1_y"
    lrn_param {
      local_size: 5
      alpha: 0.0001
      beta: 0.75
    }
  }
  layer {
    name: "conv2/3x3_reduce_y"
    type: "Convolution"
    bottom: "pool1/norm1_y"
    top: "conv2/3x3_reduce_y"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "conv2/relu_3x3_reduce_y"
    type: "ReLU"
    bottom: "conv2/3x3_reduce_y"
    top: "conv2/3x3_reduce_y"
  }
  layer {
    name: "conv2/3x3_y"
    type: "Convolution"
    bottom: "conv2/3x3_reduce_y"
    top: "conv2/3x3_y"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 192
      pad: 1
      kernel_size: 3
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "conv2/relu_3x3_y"
    type: "ReLU"
    bottom: "conv2/3x3_y"
    top: "conv2/3x3_y"
  }
  layer {
    name: "conv2/norm2_y"
    type: "LRN"
    bottom: "conv2/3x3_y"
    top: "conv2/norm2_y"
    lrn_param {
      local_size: 5
      alpha: 0.0001
      beta: 0.75
    }
  }
  layer {
    name: "pool2/3x3_s2_y"
    type: "Pooling"
    bottom: "conv2/norm2_y"
    top: "pool2/3x3_s2_y"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 2
    }
  }
  layer {
    name: "inception_3a/1x1_y"
    type: "Convolution"
    bottom: "pool2/3x3_s2_y"
    top: "inception_3a/1x1_y"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 64
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_3a/relu_1x1_y"
    type: "ReLU"
    bottom: "inception_3a/1x1_y"
    top: "inception_3a/1x1_y"
  }
  layer {
    name: "inception_3a/3x3_reduce_y"
    type: "Convolution"
    bottom: "pool2/3x3_s2_y"
    top: "inception_3a/3x3_reduce_y"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 96
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }
  layer {
    name: "inception_3a/relu_3x3_reduce_y"
    type: "ReLU"
    bottom: "inception_3a/3x3_reduce_y"
    top: "inception_3a/3x3_reduce_y"
  }
  layer {
    name: "inception_3a/3x3_y"
    type: "Convolution"
    bottom: "inception_3a/3x3_reduce_y"
    top: "inception_3a/3x3_y"
    param {
```

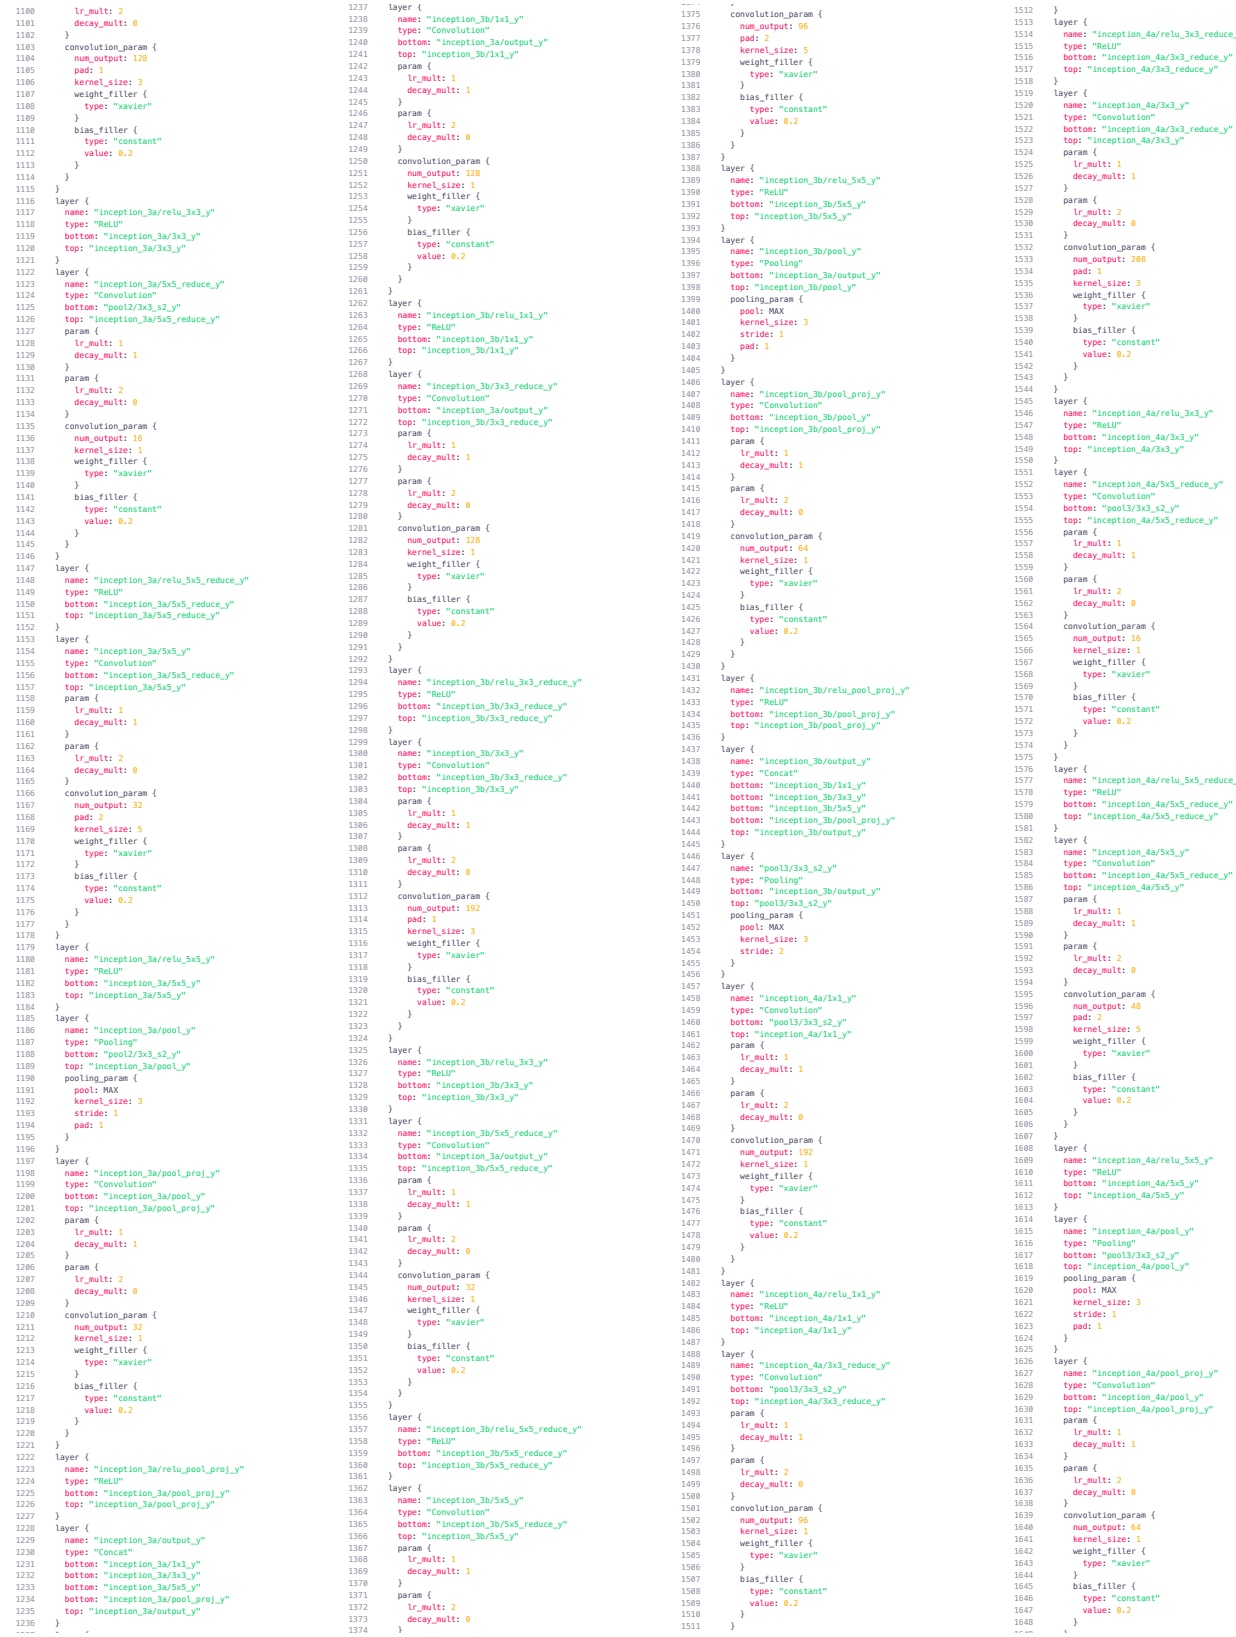

```
  }
}
layer {
  name: "inception_4a/relu_pool_proj_y"
  type: "ReLU"
  bottom: "inception_4a/pool_proj_y"
  top: "inception_4a/pool_proj_y"
}
layer {
  name: "inception_4a/output_y"
  type: "Concat"
  bottom: "inception_4a/1x1_y"
  bottom: "inception_4a/3x3_y"
  bottom: "inception_4a/5x5_y"
  bottom: "inception_4a/pool_proj_y"
  top: "inception_4a/output_y"
}

##################################
# Combo ##########################

layer {
  name: "merge_x_y"
  type: "Concat"
  bottom: "inception_4a/output_x"
  bottom: "inception_4a/output_y"
  top: "merge_x_y"
  concat_param {
    axis: 1
  }
}
layer {
  name: "inception_5b/1x1"
  type: "Convolution"
  bottom: "merge_x_y"
  top: "inception_5b/1x1"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 384
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_5b/relu_1x1"
  type: "ReLU"
  bottom: "inception_5b/1x1"
  top: "inception_5b/1x1"
}
layer {
  name: "inception_5b/3x3_reduce"
  type: "Convolution"
  bottom: "merge_x_y"
  top: "inception_5b/3x3_reduce"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_5b/relu_3x3_reduce"
  type: "ReLU"
  bottom: "inception_5b/3x3_reduce"
  top: "inception_5b/3x3_reduce"
}
layer {
  name: "inception_5b/3x3"
  type: "Convolution"
  bottom: "inception_5b/3x3_reduce"
  top: "inception_5b/3x3"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 384
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_5b/relu_3x3"
  type: "ReLU"
  bottom: "inception_5b/3x3"
  top: "inception_5b/3x3"
}
layer {
  name: "inception_5b/5x5_reduce"
  type: "Convolution"
  bottom: "merge_x_y"
  top: "inception_5b/5x5_reduce"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
```

```
  convolution_param {
    num_output: 48
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_5b/relu_5x5_reduce"
  type: "ReLU"
  bottom: "inception_5b/5x5_reduce"
  top: "inception_5b/5x5_reduce"
}
layer {
  name: "inception_5b/5x5"
  type: "Convolution"
  bottom: "inception_5b/5x5_reduce"
  top: "inception_5b/5x5"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    pad: 2
    kernel_size: 5
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_5b/relu_5x5"
  type: "ReLU"
  bottom: "inception_5b/5x5"
  top: "inception_5b/5x5"
}
layer {
  name: "inception_5b/pool"
  type: "Pooling"
  bottom: "merge_x_y"
  top: "inception_5b/pool"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 1
    pad: 1
  }
}
layer {
  name: "inception_5b/pool_proj"
  type: "Convolution"
  bottom: "inception_5b/pool"
  top: "inception_5b/pool_proj"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}
layer {
  name: "inception_5b/relu_pool_proj"
  type: "ReLU"
  bottom: "inception_5b/pool_proj"
  top: "inception_5b/pool_proj"
}
layer {
  name: "inception_5b/output"
  type: "Concat"
  bottom: "inception_5b/1x1"
  bottom: "inception_5b/3x3"
  bottom: "inception_5b/5x5"
  bottom: "inception_5b/pool_proj"
  top: "inception_5b/output"
}
layer {
  name: "pool5/7x7_s1"
  type: "Pooling"
  bottom: "inception_5b/output"
  top: "pool5/7x7_s1"
  pooling_param {
    pool: AVE
    ## Adapted!!! Originally kernel_size: 7
    kernel_h: 6
    kernel_w: 5
    stride: 1
  }
}
layer {
  name: "pool5/drop_7x7_s1"
  type: "Dropout"
  bottom: "pool5/7x7_s1"
  top: "pool5/7x7_s1"
  dropout_param {
    dropout_ratio: 0.4
  }
}
layer {
  name: "loss3/classifier15"
  type: "InnerProduct"
  bottom: "pool5/7x7_s1"
  top: "loss3/classifier15"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
```

```
  inner_product_param {
    num_output: 392
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0
    }
  }
}
layer {
  name: "loss3/loss3"
  type: "SoftmaxWithLoss"
  bottom: "loss3/classifier15"
  bottom: "label"
  top: "loss3/loss3"
  loss_weight: 1
}
layer {
  name: "loss3/top-1"
  type: "Accuracy"
  bottom: "loss3/classifier15"
  bottom: "label"
  top: "loss3/top-1"
  include {
    phase: TEST
  }
}
layer {
  name: "loss3/top-5"
  type: "Accuracy"
  bottom: "loss3/classifier15"
  bottom: "label"
  top: "loss3/top-5"
  include {
    phase: TEST
  }
  accuracy_param {
    top_k: 5
  }
}
```

```
name: "name: "Siamese Architecture
       for Event Classification v2"

layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TRAIN
  }
  transform_param {
    mirror: true
    #scale: 0.00390625
  }
  data_param {
    source: "CVN_TRAINSAMPLE_PATH"
    batch_size: 16
    prefetch: 10
    backend: LEVELDB
  }
}
layer {
  name: "data"
  type: "Data"
  top: "data"
  top: "label"
  include {
    phase: TEST
  }
  data_param {
    source: "CVN_TESTSAMPLE_PATH"
    batch_size: 64
    prefetch: 10
    backend: LEVELDB
  }
}
layer {
  name: "jitter"
  type: "DummyData"
  top: "jitter"
  include {
    phase: TRAIN
  }
  dummy_data_param {
    data_filler {
        type: "gaussian"
        mean: 1.0
        std:  0.01
        }
    shape {
        dim: 16
        dim: 2
        dim: 100
        dim: 80
        }
  }
}

layer {
  name: "jitter"
  type: "DummyData"
  top: "jitter"
  include {
    phase: TEST
  }
  dummy_data_param {
    data_filler {
        type: "constant"
        value:  1.0
        }
    shape {
        dim: 64
        dim: 2
        dim: 100
        dim: 80
        }
  }
}
layer {
  name: "jitteredData"
  type: "Eltwise"
  bottom: "data"
  bottom: "jitter"
  top: "jitteredData"
  eltwise_param {
   operation:PROD
  }
}

layer {
  name: "slice"
  type: "Slice"
  bottom: "jitteredData"
  top: "data_x"
  top: "data_y"
  slice_param {
     slice_dim: 1
     slice_point: 1
  }
}

########################################
# X view ##############################

layer {
  name: "conv1/7x7_s2_x"
  type: "Convolution"
  bottom: "data_x"
  top: "conv1/7x7_s2_x"
  param {
    name: "conv1/7x7_s2_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv1/7x7_s2_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    pad: 3
    kernel_size: 7
    stride: 2
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv1/relu_7x7_x"
  type: "ReLU"
  bottom: "conv1/7x7_s2_x"
  top: "conv1/7x7_s2_x"
}

layer {
  name: "pool1/3x3_s2_x"
  type: "Pooling"
  bottom: "conv1/7x7_s2_x"
  top: "pool1/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "pool1/norm1_x"
  type: "LRN"
  bottom: "pool1/3x3_s2_x"
  top: "pool1/norm1_x"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

layer {
  name: "conv2/3x3_reduce_x"
  type: "Convolution"
  bottom: "pool1/norm1_x"
  top: "conv2/3x3_reduce_x"
  param {
    name: "conv2/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "conv2/3x3_reduce_x"
  top: "conv2/3x3_reduce_x"
}

layer {
  name: "conv2/3x3a_x"
  type: "Convolution"
  bottom: "conv2/3x3_reduce_x"
  top: "conv2/3x3a_x"
  param {
    name: "conv2/3x3a_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3a_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3a_x"
  type: "ReLU"
  bottom: "conv2/3x3a_x"
  top: "conv2/3x3a_x"
}

layer {
  name: "conv2/3x3_x"
  type: "Convolution"
  bottom: "conv2/3x3a_x"
  top: "conv2/3x3_x"
  param {
    name: "conv2/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    pad: 1
    kernel_size: 3
    stride: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_x"
  type: "ReLU"
  bottom: "conv2/3x3_x"
  top: "conv2/3x3_x"
}
layer {
  name: "conv2/norm2_x"
  type: "LRN"
  bottom: "conv2/3x3_x"
  top: "conv2/norm2_x"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "pool2/3x3_s2_x"
  type: "Pooling"
  bottom: "conv2/norm2_x"
  top: "pool2/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "inception_3a/1x1_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/1x1_x"
  param {
    name: "inception_3a/1x1_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/1x1_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_1x1_x"
  type: "ReLU"
  bottom: "inception_3a/1x1_x"
  top: "inception_3a/1x1_x"
}
layer {
  name: "inception_3a/3x3_reduce_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/3x3_reduce_x"
  param {
    name: "inception_3a/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_reduce_x"
  type: "ReLU"
  bottom: "inception_3a/3x3_reduce_x"
  top: "inception_3a/3x3_reduce_x"
}
layer {
  name: "inception_3a/3x3_x"
  type: "Convolution"
  bottom: "inception_3a/3x3_reduce_x"
  top: "inception_3a/3x3_x"
  param {
    name: "inception_3a/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_x"
  type: "ReLU"
  bottom: "inception_3a/3x3_x"
  top: "inception_3a/3x3_x"
}

layer {
  name: "inception_3a/5x5_reduce_x"
  type: "Convolution"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/5x5_reduce_x"
  param {
    name: "inception_3a/5x5_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_reduce_x"
  type: "ReLU"
  bottom: "inception_3a/5x5_reduce_x"
  top: "inception_3a/5x5_reduce_x"
}

layer {
  name: "inception_3a/5x5_x"
  type: "Convolution"
  bottom: "inception_3a/5x5_reduce_x"
  top: "inception_3a/5x5_x"
  param {
    name: "inception_3a/5x5_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    pad: 2
    kernel_size: 5
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_x"
  type: "ReLU"
  bottom: "inception_3a/5x5_x"
  top: "inception_3a/5x5_x"
}

layer {
  name: "inception_3a/pool_x"
  type: "Pooling"
  bottom: "pool2/3x3_s2_x"
  top: "inception_3a/pool_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 1
    pad: 1
  }
}
layer {
  name: "inception_3a/pool_proj_x"
  type: "Convolution"
  bottom: "inception_3a/pool_x"
  top: "inception_3a/pool_proj_x"
  param {
    name: "inception_3a/pool_proj_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/pool_proj_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_pool_proj_x"
  type: "ReLU"
  bottom: "inception_3a/pool_proj_x"
  top: "inception_3a/pool_proj_x"
}
layer {
  name: "inception_3a/output_x"
  type: "Concat"
  bottom: "inception_3a/1x1_x"
  bottom: "inception_3a/3x3_x"
  bottom: "inception_3a/5x5_x"
  bottom: "inception_3a/pool_proj_x"
  top: "inception_3a/output_x"
}

layer {
  name: "pool3a/3x3_s2_x"
  type: "Pooling"
  bottom: "inception_3a/output_x"
  top: "pool3a/3x3_s2_x"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}

################################
# Y view #######################

layer {
  name: "conv1/7x7_s2_y"
  type: "Convolution"
  bottom: "data_y"
  top: "conv1/7x7_s2_y"
  param {
    name: "conv1/7x7_s2_w"
    lr_mult: 1
    decay_mult: 1
```

```
    decay_mult: 1
  }
  param {
    name: "conv1/7x7_s2_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    pad: 3
    kernel_size: 7
    stride: 2
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv1/relu_7x7_y"
  type: "ReLU"
  bottom: "conv1/7x7_s2_y"
  top: "conv1/7x7_s2_y"
}

layer {
  name: "pool1/3x3_s2_y"
  type: "Pooling"
  bottom: "conv1/7x7_s2_y"
  top: "pool1/3x3_s2_y"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "pool1/norm1_y"
  type: "LRN"
  bottom: "pool1/3x3_s2_y"
  top: "pool1/norm1_y"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

layer {
  name: "conv2/3x3_reduce_y"
  type: "Convolution"
  bottom: "pool1/norm1_y"
  top: "conv2/3x3_reduce_y"
  param {
    name: "conv2/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_reduce_y"
  type: "ReLU"
  bottom: "conv2/3x3_reduce_y"
  top: "conv2/3x3_reduce_y"
}

layer {
  name: "conv2/3x3a_y"
  type: "Convolution"
  bottom: "conv2/3x3_reduce_y"
  top: "conv2/3x3a_y"
  param {
    name: "conv2/3x3a_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3a_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3a_y"
  type: "ReLU"
  bottom: "conv2/3x3a_y"
  top: "conv2/3x3a_y"
}

layer {
  name: "conv2/3x3_y"
  type: "Convolution"
  bottom: "conv2/3x3a_y"
  top: "conv2/3x3_y"
  param {
    name: "conv2/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "conv2/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "conv2/relu_3x3_y"
  type: "ReLU"
  bottom: "conv2/3x3_y"
  top: "conv2/3x3_y"
}
layer {
  name: "conv2/norm2_y"
  type: "LRN"
  bottom: "conv2/3x3_y"
  top: "conv2/norm2_y"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}
layer {
  name: "pool2/3x3_s2_y"
  type: "Pooling"
  bottom: "conv2/norm2_y"
  top: "pool2/3x3_s2_y"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}
layer {
  name: "inception_3a/1x1_y"
  type: "Convolution"
  bottom: "pool2/3x3_s2_y"
  top: "inception_3a/1x1_y"
  param {
    name: "inception_3a/1x1_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/1x1_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 64
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_1x1_y"
  type: "ReLU"
  bottom: "inception_3a/1x1_y"
  top: "inception_3a/1x1_y"
}
layer {
  name: "inception_3a/3x3_reduce_y"
  type: "Convolution"
  bottom: "pool2/3x3_s2_y"
  top: "inception_3a/3x3_reduce_y"
  param {
    name: "inception_3a/3x3_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 96
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_reduce_y"
  type: "ReLU"
  bottom: "inception_3a/3x3_reduce_y"
  top: "inception_3a/3x3_reduce_y"
}
layer {
  name: "inception_3a/3x3_y"
  type: "Convolution"
  bottom: "inception_3a/3x3_reduce_y"
  top: "inception_3a/3x3_y"
  param {
    name: "inception_3a/3x3_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/3x3_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 128
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_3x3_y"
  type: "ReLU"
  bottom: "inception_3a/3x3_y"
  top: "inception_3a/3x3_y"
}

layer {
  name: "inception_3a/5x5_reduce_y"
  type: "Convolution"
  bottom: "pool2/3x3_s2_y"
  top: "inception_3a/5x5_reduce_y"
  param {
    name: "inception_3a/5x5_reduce_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_reduce_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 16
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_reduce_y"
  type: "ReLU"
  bottom: "inception_3a/5x5_reduce_y"
  top: "inception_3a/5x5_reduce_y"
}

layer {
  name: "inception_3a/5x5_y"
  type: "Convolution"
  bottom: "inception_3a/5x5_reduce_y"
  top: "inception_3a/5x5_y"
  param {
    name: "inception_3a/5x5_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/5x5_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    pad: 2
    kernel_size: 5
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_5x5_y"
  type: "ReLU"
  bottom: "inception_3a/5x5_y"
  top: "inception_3a/5x5_y"
}

layer {
  name: "inception_3a/pool_y"
  type: "Pooling"
  bottom: "pool2/3x3_s2_y"
  top: "inception_3a/pool_y"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 1
    pad: 1
  }
}
layer {
  name: "inception_3a/pool_proj_y"
  type: "Convolution"
  bottom: "inception_3a/pool_y"
  top: "inception_3a/pool_proj_y"
  param {
    name: "inception_3a/pool_proj_w"
    lr_mult: 1
    decay_mult: 1
  }
  param {
    name: "inception_3a/pool_proj_b"
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 32
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_3a/relu_pool_proj_y"
  type: "ReLU"
  bottom: "inception_3a/pool_proj_y"
  top: "inception_3a/pool_proj_y"
}
layer {
  name: "inception_3a/output_y"
  type: "Concat"
  bottom: "inception_3a/1x1_y"
  bottom: "inception_3a/3x3_y"
  bottom: "inception_3a/5x5_y"
  bottom: "inception_3a/pool_proj_y"
  top: "inception_3a/output_y"
}

layer {
  name: "pool3a/3x3_s2_y"
  type: "Pooling"
  bottom: "inception_3a/output_y"
  top: "pool3a/3x3_s2_y"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}

##############################
# Combo ######################

layer {
  name: "merge_x_y"
  type: "Concat"
  bottom: "pool3a/3x3_s2_x"
  bottom: "pool3a/3x3_s2_y"
  top: "merge_x_y"
  concat_param {
    axis: 1
  }
}
layer {
  name: "inception_5b/1x1"
  type: "Convolution"
  bottom: "merge_x_y"
  top: "inception_5b/1x1"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 384
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_5b/relu_1x1"
  type: "ReLU"
  bottom: "inception_5b/1x1"
  top: "inception_5b/1x1"
}
layer {
  name: "inception_5b/3x3_reduce"
  type: "Convolution"
  bottom: "merge_x_y"
  top: "inception_5b/3x3_reduce"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 192
    kernel_size: 1
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_5b/relu_3x3_reduce"
  type: "ReLU"
  bottom: "inception_5b/3x3_reduce"
  top: "inception_5b/3x3_reduce"
}
layer {
  name: "inception_5b/3x3"
  type: "Convolution"
  bottom: "inception_5b/3x3_reduce"
  top: "inception_5b/3x3"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 384
    pad: 1
    kernel_size: 3
    weight_filler {
      type: "xavier"
    }
    bias_filler {
      type: "constant"
      value: 0.2
    }
  }
}

layer {
  name: "inception_5b/relu_3x3"
  type: "ReLU"
  bottom: "inception_5b/3x3"
  top: "inception_5b/3x3"
}
layer {
  name: "inception_5b/5x5_reduce"
  type: "Convolution"
  bottom: "merge_x_y"
  top: "inception_5b/5x5_reduce"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
  convolution_param {
    num_output: 48
    kernel_size: 1
```

```
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_5x5_reduce"
    type: "ReLU"
    bottom: "inception_5b/5x5_reduce"
    top: "inception_5b/5x5_reduce"
  }

  layer {
    name: "inception_5b/5x5"
    type: "Convolution"
    bottom: "inception_5b/5x5_reduce"
    top: "inception_5b/5x5"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 128
      pad: 2
      kernel_size: 5
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_5x5"
    type: "ReLU"
    bottom: "inception_5b/5x5"
    top: "inception_5b/5x5"
  }

  layer {
    name: "inception_5b/pool"
    type: "Pooling"
    bottom: "merge_x_y"
    top: "inception_5b/pool"
    pooling_param {
      pool: MAX
      kernel_size: 3
      stride: 1
      pad: 1
    }
  }
  layer {
    name: "inception_5b/pool_proj"
    type: "Convolution"
    bottom: "inception_5b/pool"
    top: "inception_5b/pool_proj"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    convolution_param {
      num_output: 128
      kernel_size: 1
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0.2
      }
    }
  }

  layer {
    name: "inception_5b/relu_pool_proj"
    type: "ReLU"
    bottom: "inception_5b/pool_proj"
    top: "inception_5b/pool_proj"
  }
  layer {
    name: "inception_5b/output"
    type: "Concat"
    bottom: "inception_5b/1x1"
    bottom: "inception_5b/3x3"
    bottom: "inception_5b/5x5"
    bottom: "inception_5b/pool_proj"
    top: "inception_5b/output"
  }
  layer {
    name: "pool5/7x7_s1"
    type: "Pooling"
    bottom: "inception_5b/output"
    top: "pool5/7x7_s1"
    pooling_param {
      pool: AVE
      ## Adapted!!! Originally kernel_size: 7
      kernel_h: 6
      kernel_w: 5
      stride: 1
    }
  }
  layer {
    name: "pool5/drop_7x7_s1"
    type: "Dropout"
    bottom: "pool5/7x7_s1"
    top: "pool5/7x7_s1"
    dropout_param {
      dropout_ratio: 0.4
    }
  }
  layer {
    name: "loss3/classifier15"
    type: "InnerProduct"
    bottom: "pool5/7x7_s1"
    top: "loss3/classifier15"
    param {
      lr_mult: 1
      decay_mult: 1
    }
    param {
      lr_mult: 2
      decay_mult: 0
    }
    inner_product_param {
      num_output: 392
      weight_filler {
        type: "xavier"
      }
      bias_filler {
        type: "constant"
        value: 0
      }
    }
  }
  layer {
    name: "loss3/loss3"
    type: "SoftmaxWithLoss"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/loss3"
    loss_weight: 1
  }
  layer {
    name: "loss3/top-1"
    type: "Accuracy"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/top-1"
    include {
      phase: TEST
    }
  }
  layer {
    name: "loss3/top-5"
    type: "Accuracy"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/top-5"
    include {
      phase: TEST
    }
    accuracy_param {
      top_k: 5
    }
  }
  layer {
    name: "loss3/top-10"
    type: "Accuracy"
    bottom: "loss3/classifier15"
    bottom: "label"
    top: "loss3/top-10"
    include {
      phase: TEST
    }
    accuracy_param {
      top_k: 10
    }
  }
```

APPENDIX B

Systematics Details

"There are sadistic scientists who hurry

to hunt down errors instead of establishing the truth"

– Marie Curie

This appendix contains the full details of the systematic uncertainties. For each uncertainty considered there are plots showing the shape difference applied to the fits in Chapter 7, followed by a table, detailing the average effect on each bin (or quantile) of the analyses. The following pages show the full detailed information for the largest contributing systematics and contain summary tables for all the systematics studied, the smallest of which are combined into larger categories and applied in combination.

ALL ν_e SYSTEMATICS

Extrap. Sig	Bin1		Bin2		Bin3		Bin4		Tot		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
Calibration	-0.95	+1.20	+0.85	-0.50	+1.83	-3.56	+2.08	-3.27	+1.48	-2.67	1.326	1.375
SmallXsecsJoint	+2.23	-6.23	+0.26	-1.26	-3.94	+11.54	-4.61	+13.62	-2.88	+8.41	0.057	0.491
MECq0Shape	+4.23	+1.44	+3.27	+1.94	-4.19	+8.94	-3.43	+9.37	-2.29	+7.33	0.072	0.274
CalibShape	-0.40	–	-0.68	–	-1.12	–	-3.05	–	-1.23	–	0.113	–
Lightlevel	+1.52	+1.25	+0.71	+1.02	+0.82	+0.20	+2.56	-2.27	+1.07	+0.11	0.036	0.108
RPAShapeRES	+4.23	+0.00	+3.28	+0.00	+4.01	+0.01	+4.55	+0.01	+3.99	+0.01	0.069	0.000
MaCCRES	+5.52	-4.84	+4.69	-4.37	+0.82	-0.24	+1.25	-0.62	+1.84	-1.28	0.029	0.022
radcorrnue	+1.95	-1.95	+1.94	-1.94	+1.97	-1.97	+1.97	-1.97	+1.96	-1.96	0.016	0.016
2ndclasscurr	+1.90	-1.90	+1.89	-1.89	+1.94	-1.94	+1.94	-1.94	+1.93	-1.93	0.016	0.016
MvCCRES	+3.48	-3.03	+2.96	-2.62	+0.62	-0.42	+0.89	-0.66	+1.24	-1.00	0.012	0.009
RPAShapeenh	+0.21	+1.04	+0.59	+0.42	+1.68	-1.02	+1.90	-1.14	+1.42	-0.65	0.010	0.004
Cherenkov	+0.03	–	+0.50	–	+0.03	–	+0.11	–	+0.11	–	0.009	–
DISvnCC1pi	+3.09	-3.12	+2.16	-2.19	+0.52	-0.53	+0.63	-0.63	+0.99	-1.00	0.008	0.008
MaCCQEreduced	+0.08	+0.17	+0.25	-0.03	+1.21	-1.05	+1.44	-1.22	+1.00	-0.82	0.005	0.004
extrapsignalkin	+0.42	-0.37	+0.37	-0.32	+0.32	-0.28	+0.35	-0.31	+0.34	-0.30	0.003	0.003

AllBeamTransport	+0.45	-0.40	+0.47	-0.42	+0.63	-0.53	+0.73	-0.61	+0.60	-0.52	0.002	0.001
ppfxpc01	+0.23	-0.20	+0.22	-0.19	+0.26	-0.23	+0.30	-0.26	+0.25	-0.23	0.000	0.000
ppfxpc04	+0.15	-0.14	+0.15	-0.14	+0.22	-0.21	+0.24	-0.23	+0.21	-0.19	0.000	0.000
ppfxpc00	-0.05	+0.05	-0.05	+0.05	-0.15	+0.15	-0.18	+0.18	-0.13	+0.13	0.000	0.000
ppfxpc03	+0.13	-0.13	+0.13	-0.14	+0.11	-0.11	+0.10	-0.10	+0.11	-0.12	0.000	0.000
ppfxpc02	+0.03	-0.04	+0.02	-0.03	-0.03	+0.03	-0.06	+0.07	-0.02	+0.02	0.000	0.000
MaNCRES	-0.06	+0.04	-0.06	+0.04	-0.05	+0.04	-0.06	+0.04	-0.06	+0.04	0.000	0.000
radcorrnuebar	+0.05	-0.05	+0.06	-0.06	+0.03	-0.03	+0.03	-0.03	+0.04	-0.04	0.000	0.000
extrapbkg	+0.00	+0.00	+0.00	+0.00	+0.01	+0.01	+0.01	+0.01	+0.01	+0.01	0.000	0.000

Extrap. Bkg	Bin1		Bin2		Bin3		Bin4		Tot		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
SmallXsecsJoint	+4.92	-9.55	+6.98	-9.41	+4.32	-6.63	+3.89	-4.56	+5.14	-8.08	0.046	0.134
Lightlevel	+2.34	+0.08	-1.41	+6.02	-0.75	-3.92	+0.54	-4.89	+0.15	-0.38	0.096	0.124
Calibration	-1.65	-1.17	-0.30	+1.52	-1.32	+3.72	-4.83	+4.68	-1.44	+1.72	0.074	0.100
Cherenkov	+4.49	–	+2.57	–	-1.60	–	-1.17	–	+1.31	–	0.073	–
CalibShape	+4.13	–	+1.02	–	+1.22	–	-4.23	–	+1.65	–	0.047	–
MECq0Shape	+1.03	+0.94	+1.50	+1.07	+3.23	+0.01	+2.96	-1.04	+2.09	+0.48	0.015	0.005
extrapsignalkin	-1.16	+1.38	+1.86	-1.65	+1.60	-1.66	+1.26	-1.28	+0.77	-0.68	0.010	0.011
MaCCRES	-0.18	+0.36	+1.44	-1.32	+2.91	-2.84	+3.58	-3.12	+1.63	-1.48	0.009	0.008
MvCCRES	-0.10	+0.16	+0.86	-0.79	+1.66	-1.49	+1.98	-1.68	+0.94	-0.82	0.003	0.002
extrapbkg	+1.21	-1.19	+1.32	-1.09	+1.11	-1.28	+1.24	-1.36	+1.21	-1.21	0.002	0.002
DISvnCC1pi	-0.57	+0.79	-0.46	+0.60	+0.06	-0.03	+0.26	-0.21	-0.25	+0.37	0.001	0.001
RPAShapeRES	+0.13	+0.01	+0.29	+0.11	-0.10	-0.08	+0.67	-0.06	+0.13	-0.00	0.001	0.000
MaNCRES	-0.74	+0.79	-0.19	+0.16	+0.41	-0.36	-0.14	+0.13	-0.14	+0.17	0.001	0.001
AllBeamTransport	+0.28	-0.07	+0.41	-0.02	+0.35	-0.50	+0.23	-0.38	+0.33	-0.24	0.000	0.001
RPAShapeenh	+0.08	-0.01	+0.39	+0.23	+0.29	+0.10	+0.71	+0.05	+0.28	+0.10	0.000	0.000
MaCCQEreduced	-0.01	+0.04	+0.12	+0.12	-0.03	-0.01	+0.35	-0.25	+0.04	+0.02	0.000	0.000
ppfxpc00	-0.01	-0.02	-0.10	+0.08	-0.37	+0.39	-0.08	+0.09	-0.17	+0.16	0.000	0.000
ppfxpc03	+0.24	-0.24	+0.23	-0.24	+0.17	-0.17	+0.07	-0.07	+0.20	-0.21	0.000	0.000
ppfxpc02	+0.01	+0.01	+0.11	-0.12	+0.16	-0.17	-0.03	+0.03	+0.09	-0.09	0.000	0.000
ppfxpc01	+0.10	-0.04	+0.04	+0.03	+0.23	-0.12	+0.19	-0.09	+0.14	-0.06	0.000	0.000
2ndclasscurr	+0.03	-0.03	+0.10	-0.12	+0.09	-0.09	+0.19	-0.19	+0.08	-0.09	0.000	0.000
radcorrnue	+0.01	-0.00	+0.09	-0.11	+0.06	-0.06	+0.17	-0.17	+0.06	-0.06	0.000	0.000
ppfxpc04	+0.10	-0.09	+0.06	-0.04	+0.11	-0.08	+0.15	-0.12	+0.10	-0.08	0.000	0.000
radcorrnuebar	-0.00	+0.00	-0.05	+0.03	-0.04	+0.04	-0.02	+0.02	-0.03	+0.02	0.000	0.000

Extrap. Sig	Quant1		Quant2		Quant3		Quant4		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
MECq0Shape	-1.92	+3.45	-5.26	+12.00	-3.63	+7.94	-0.41	-0.61	0.182	0.734
Calibration	-4.59	+1.63	-4.47	+3.13	-5.65	+4.98	-3.99	+4.08	0.528	0.442
RelHadEScale	+5.36	-5.06	+0.47	-0.30	-3.26	+3.22	-5.70	+6.13	0.413	0.441
CalibrationEnergy	-4.03	+2.61	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.426	0.321
AbsHadEScale	-1.62	+1.65	-2.71	+3.23	-4.43	+4.78	-4.69	+5.79	0.321	0.410
SmallXsecsNumu	-0.88	+3.15	-1.23	+5.29	-0.87	+5.55	-1.73	+3.67	0.026	0.312
SmallXsecsJoint	-0.88	+3.15	-1.23	+5.29	-0.87	+5.55	-1.73	+3.67	0.026	0.312
RPAShapeRES	+1.67	+0.00	+4.61	+0.00	+6.06	+0.00	+2.21	+0.00	0.197	0.000
RelHadEScaleE	-0.78	+0.78	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.178	0.178
RPAShapeRESEner	+3.46	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.151	0.000
CalibShape	-3.30	–	-0.80	–	+0.03	–	+0.70	–	0.129	–
MaCCRES	-0.05	+0.21	+1.31	-0.76	+3.81	-2.44	+2.21	-1.34	0.103	0.057
MaCCRESEnergy	+2.23	-1.49	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.098	0.058
Lightlevel	+0.59	+0.09	+1.42	+0.71	+1.95	+0.55	+2.38	-0.77	0.068	0.028

Cherenkov	-1.00	–	+0.32	–	+0.91	–	+0.13	–	0.052	–
MvCCRES	+0.02	+0.01	+0.91	-0.70	+2.25	-1.71	+1.12	-0.80	0.034	0.022
RelMuEScale	+0.95	-0.93	+0.45	-0.46	+0.06	-0.05	-0.27	+0.30	0.013	0.013
beamTranspComb	+0.03	+0.05	+0.03	+0.06	+0.04	+0.07	+0.14	-0.03	0.013	0.013
RPAShapeenh	+0.49	-0.34	+1.30	-0.86	+1.16	+0.15	+0.27	+0.55	0.011	0.010
DISvnCC1pi	-0.07	+0.07	+0.07	-0.07	+0.66	-0.70	+0.70	-0.78	0.009	0.011
ppfxpc00	+0.29	-0.29	+0.31	-0.31	+0.34	-0.35	+0.24	-0.25	0.006	0.006

Extrap. Bkg	Quant1		Quant2		Quant3		Quant4		Chi2Test	
Syst	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+(%)	- (%)	+	-
SmallXsecsNumu	+46.56	-45.83	+51.03	-49.60	+58.67	-57.21	+69.29	-65.81	1.448	1.322
SmallXsecsNumuE	+64.68	-61.86	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	1.428	1.306
Calibration	-13.29	+13.77	-11.82	+11.06	-13.41	+11.22	-19.38	+21.87	0.124	0.154
CalibrationEnergy	-17.45	+18.81	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.113	0.132
MECq0Shape	-8.76	+14.18	+10.55	+8.99	+12.90	+2.83	+1.91	+0.35	0.044	0.014
MaNCRES	+6.58	-6.11	+8.44	-7.36	+8.22	-6.67	+8.44	-6.05	0.038	0.025
MaNCRESEnergy	+8.29	-6.26	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.037	0.025
Lightlevel	+6.81	+6.09	+1.86	+1.64	+2.88	+2.76	+7.09	+3.30	0.026	0.016
Cherenkov	+4.25	–	-0.47	–	+1.17	–	+3.33	–	0.018	–
ppfxpc01	-7.39	+7.39	-7.34	+7.34	-7.18	+7.18	-7.08	+7.08	0.017	0.017
ppfxpc01Energy	-7.14	+7.14	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.017	0.017
]LightlevelEnergy	+5.97	+3.25	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.016	0.008
AbsHadEScale	-7.32	+7.65	-5.17	+5.16	-3.95	+4.66	+0.88	-1.14	0.012	0.014
MaCCRES	+1.81	-2.03	+6.15	-5.61	+6.72	-5.13	+2.90	-2.15	0.010	0.006
AllBeamTransport	+4.59	-4.71	+4.72	-4.75	+5.00	-4.92	+5.35	-5.22	0.009	0.009
AllBeamTransportE	+5.19	-5.10	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.009	0.009
AbsHadEScaleE	-0.88	+0.83	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.008	0.009
MECq0ShapeEnerg	+3.66	+2.34	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.008	0.002
CherenkovEnergy	+2.73	–	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.008	–
MaCCRESEnergy	+3.69	-2.89	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.007	0.004
CalibShape	-0.19	–	+1.26	–	-1.63	–	+1.66	–	0.006	–
RelHadEScale2017	+4.12	-3.66	+2.27	-2.73	+2.51	-2.34	-0.64	+0.55	0.004	0.004
ppfxpc04	-3.58	+3.58	-3.53	+3.53	-3.41	+3.41	-3.33	+3.33	0.004	0.004
MvCCRES	+0.89	-0.83	+3.28	-2.76	+3.48	-2.71	+1.52	-1.19	0.003	0.002
CalibShapeEnergy	+1.01	–	+0.00	+0.00	+0.00	+0.00	+0.00	+0.00	0.002	–
MvCCRESEnergy	+1.93	-1.53	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.002	0.001
RPAShapeenh	+3.84	-3.19	+4.01	-2.88	+2.25	-0.84	+0.75	-0.18	0.002	0.001
AbsMuEScale	-0.88	+1.32	-1.11	+0.33	-0.83	+1.16	+0.54	-0.58	0.001	0.001
ppfxpc00	+1.85	-1.85	+1.78	-1.78	+1.58	-1.58	+1.47	-1.47	0.001	0.001
ppfxpc00Energy	+1.54	-1.54	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.001	0.001
ppfxpc02	+1.64	-1.64	+1.64	-1.64	+1.65	-1.65	+1.61	-1.61	0.001	0.001
RPAShapeenhE	+1.45	-0.70	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.001	0.000
ppfxpc03	+1.31	-1.31	+1.42	-1.42	+1.48	-1.48	+1.44	-1.44	0.001	0.001
ppfxpc03Energy	+1.43	-1.43	+0.00	+0.00	+0.00	+0.00	+450.00	+0.00	0.001	0.001
DISvnCC1pi	+0.28	-0.28	+0.54	-0.54	+1.02	-1.02	+1.09	-1.09	0.000	0.000

B.1 Systematic Uncertainties

Table B.5: Calibration shape - ν_e

	Bin 1			Bin 2			Bin 3			Bin 4		
	Nom.	Shift	% Diff.	Nom.	Shift	% Diff.	Nom.	Shift	% Diff.	Nom.	Shift	% Diff.
ν_e signal	3.6	3.6	-0.4	6.2	6.2	-0.7	27	27	-1.1	5.2	5.1	-3.0
Tot beam bkg	4.7	4.9	+4.1	3.8	3.8	+1.0	5.4	5.4	+1.2	1.2	1.1	-4.2
Beam ν_e CC	0.85	0.86	+1.0	1.5	1.5	+0.0	3.9	4	+0.9	0.89	0.83	-6.7
NC	3.2	3.4	+4.0	1.8	1.9	+0.5	1.1	1.1	+3.8	0.2	0.21	+7.6
ν_μ CC	0.49	0.55	+12.3	0.28	0.31	+8.5	0.16	0.15	-6.2	0.055	0.055	+0.4
ν_τ CC	0.11	0.1	-4.1	0.13	0.13	+3.0	0.17	0.17	-0.2	0.033	0.027	-16.0

Table B.6: Calibration Shape - ν_μ

(a) Quantile 1

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	34.053	32.928	–	-3.3	–
Tot beam bkg	0.211	0.210	–	-0.2	–
ν_e CC	0.000	0.000	–	-11.8	–
NC	0.078	0.077	–	-0.9	–
ν_μApp	0.041	0.041	–	+0.8	–
ν_τ CC	0.092	0.092	–	+0.0	–

(b) Quantile 2

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	25.435	25.231	–	-0.8	–
Tot beam bkg	0.295	0.298	–	+1.3	–
ν_e CC	0.002	0.002	–	-3.6	–
NC	0.139	0.144	–	+3.3	–
ν_μApp	0.040	0.040	–	+0.0	–
ν_τ CC	0.113	0.113	–	-0.7	–

(c) Quantile 3

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	26.638	26.646	–	+0.0	–
Tot beam bkg	0.515	0.507	–	-1.6	–
ν_e CC	0.015	0.015	–	+0.9	–
NC	0.330	0.324	–	-2.0	–
ν_μApp	0.039	0.038	–	-2.1	–
ν_τ CC	0.131	0.130	–	-0.8	–

(d) Quantile 4

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	32.445	32.671	–	+0.7	–
Tot beam bkg	2.377	2.417	–	+1.7	–
ν_e CC	0.234	0.247	–	+5.5	–
NC	1.952	1.977	–	+1.3	–
ν_μApp	0.043	0.042	–	-1.9	–
ν_τ CC	0.149	0.151	–	+1.5	–

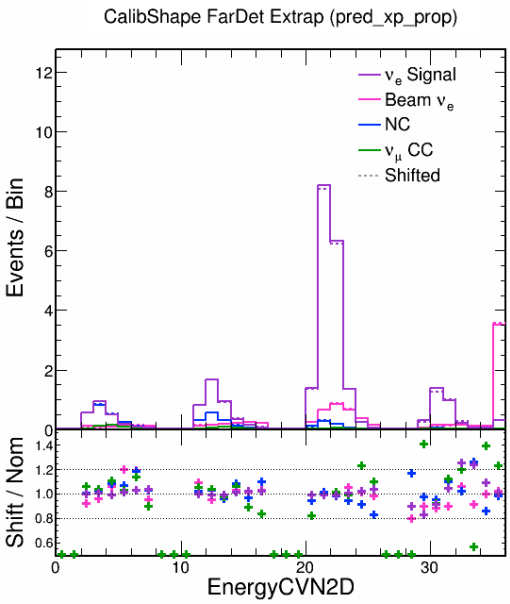
CalibShape FarDet Extrap (pred_xp_prop)
ν_e Signal
Beam ν_e
NC
ν_μ CC
Shifted
Events / Bin
Shift / Nom
EnergyCVN2D

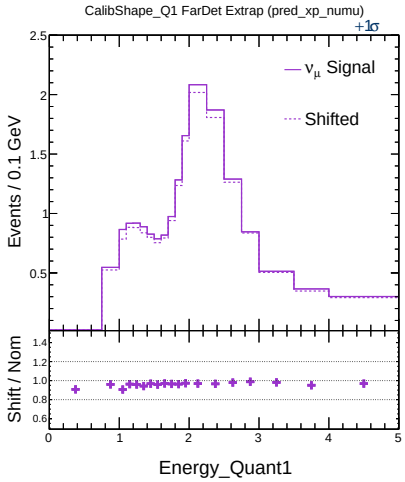
CalibShape_Q1 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant1

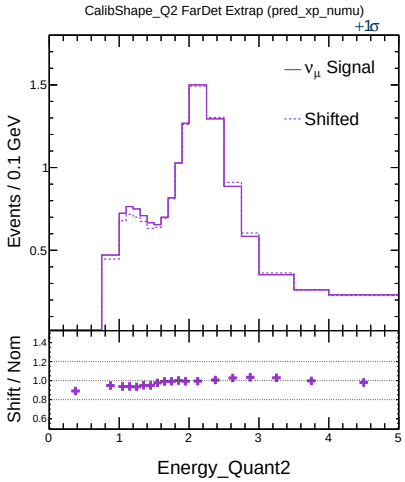
CalibShape_Q2 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant2

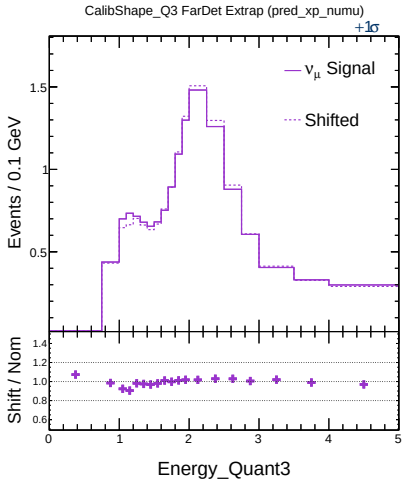
CalibShape_Q3 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant3

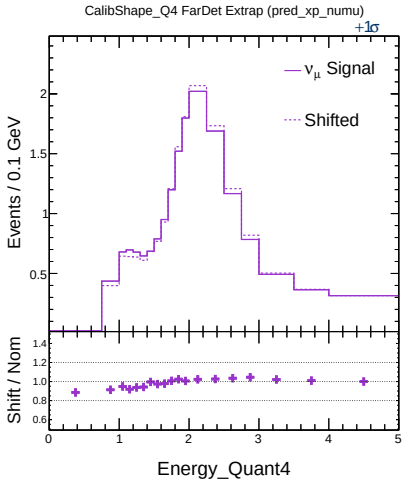
CalibShape_Q4 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

Table B.7: Absolute calibration - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	-0.9	+1.2	6.2	6.3	6.2	+0.8	-0.5	27	28	26	+1.8	-3.6	5.2	5.3	5.1	+2.1	-3.3
Tot beam bkg	4.7	4.6	4.6	-1.7	-1.2	3.8	3.8	3.8	-0.3	+1.5	5.4	5.3	5.6	-1.3	+3.7	1.2	1.1	1.2	-4.8	+4.7
Beam ν_e CC	0.85	0.81	0.91	-5.0	+6.5	1.5	1.4	1.6	-4.7	+2.8	3.9	3.7	4.2	-4.9	+7.4	0.89	0.83	0.95	-7.3	+6.7
NC	3.2	3.3	3.1	+0.4	-5.8	1.8	1.9	1.8	+4.5	+0.2	1.1	1.2	1	+12.6	-7.8	0.2	0.21	0.2	+5.6	-0.5
ν_μ CC	0.49	0.45	0.57	-9.0	+16.6	0.28	0.27	0.29	-5.3	+0.1	0.16	0.15	0.15	-3.9	-8.0	0.055	0.055	0.05	+0.5	-8.1
ν_τ CC	0.11	0.1	0.1	-1.9	-2.1	0.13	0.12	0.14	-5.8	+8.0	0.17	0.16	0.18	-4.0	+3.3	0.033	0.029	0.033	-10.2	+2.4

Table B.8: Absolute calibration - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	32	35	-4.6	+1.6
Tot beam bkg	0.21	0.18	0.24	-13.3	+13.8
NC	0.078	0.063	0.093	-18.6	+19.7
ν_μ App	0.041	0.037	0.045	-10.1	+8.6
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.082	0.1	-10.1	+11.1

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	24	26	-4.5	+3.1
Tot beam bkg	0.29	0.26	0.33	-11.8	+11.1
NC	0.14	0.11	0.17	-20.8	+20.0
ν_μ App	0.04	0.04	0.041	+0.4	+2.6
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.11	0.12	-4.7	+2.5

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	25	28	-5.6	+5.0
Tot beam bkg	0.52	0.45	0.57	-13.4	+11.2
NC	0.33	0.27	0.39	-19.5	+17.6
ν_μ App	0.039	0.039	0.036	+1.7	-8.3
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.13	0.13	-1.4	-0.5

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	31	34	-4.0	+4.1
Tot beam bkg	2.4	1.9	2.9	-19.4	+21.9
NC	2	1.6	2.4	-20.6	+22.7
ν_μ App	0.043	0.041	0.041	-4.0	-2.9
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.14	0.16	-7.9	+7.6

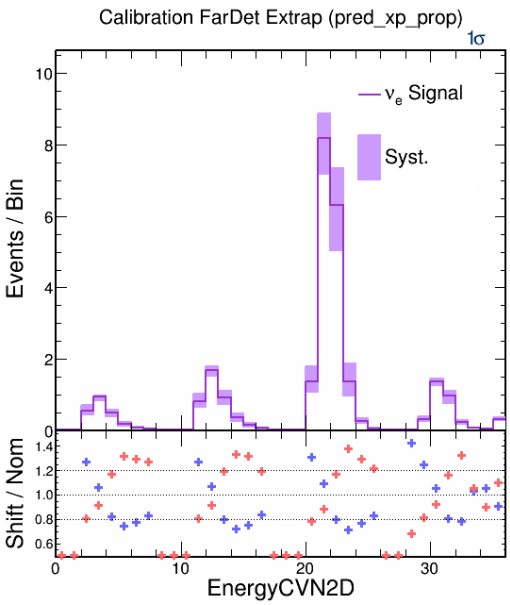
Calibration FarDet Extrap (pred_xp_prop)
1σ
ν_e Signal
Syst.
Events / Bin
Shift / Nom
EnergyCVN2D

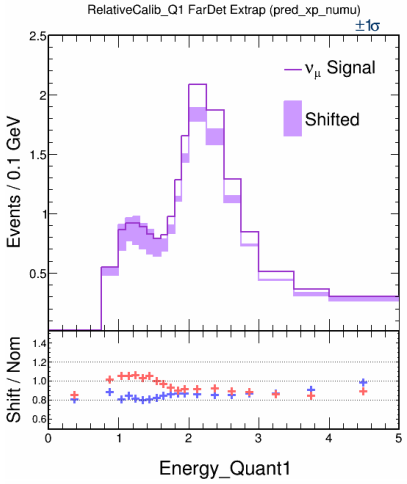
RelativeCalib_Q1 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant1

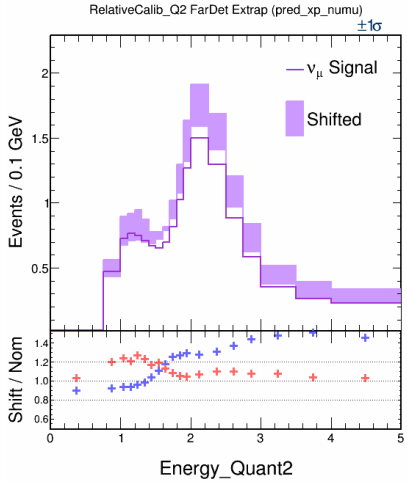
RelativeCalib_Q2 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant2

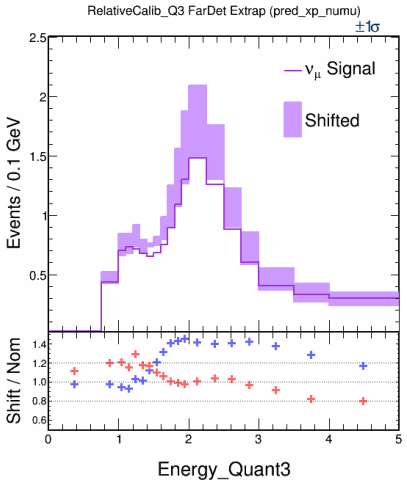
RelativeCalib_Q3 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant3

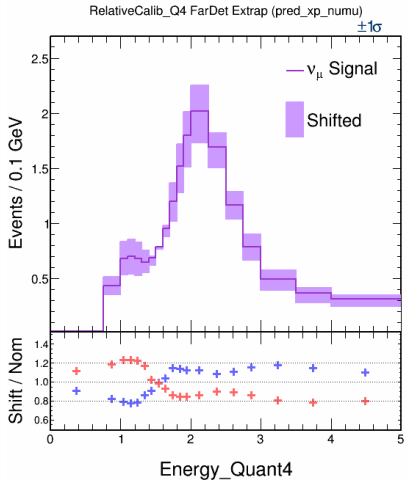
RelativeCalib_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

Table B.9: Cherenkov - ν_e

	Bin 1			Bin 2			Bin 3			Bin 4		
	Nom.	Shift	% Diff.	Nom.	Shift	% Diff.	Nom.	Shift	% Diff.	Nom.	Shift	% Diff.
ν_e signal	3.5	3.5	+0.0	6.1	6.1	+0.5	27	27	+0.0	5.1	5.1	+0.1
Tot beam bkg	4.5	4.7	+4.5	3.5	3.6	+2.6	5.4	5.4	-1.6	1.2	1.2	-1.2
Beam ν_e CC	0.87	0.9	+3.8	1.5	1.5	+1.5	4	3.9	-1.7	0.91	0.89	-2.1
NC	3.1	3.2	+4.5	1.7	1.7	+3.3	1.2	1.2	-1.3	0.21	0.22	+4.1
ν_μ CC	0.49	0.53	+6.4	0.28	0.29	+4.1	0.15	0.14	-4.1	0.046	0.043	-7.0
ν_τ CC	0.11	0.11	+1.5	0.12	0.13	+2.0	0.17	0.17	+0.6	0.032	0.032	+0.3

Table B.10: Cherenkov - ν_μ

(a) Quantile 1

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	34.692	34.344	–	-1.0	–
Tot beam bkg	0.216	0.225	–	+4.2	–
ν_e CC	0.000	0.001	–	+149.5	–
NC	0.082	0.091	–	+10.7	–
ν_μApp	0.042	0.043	–	+1.4	–
ν_τ CC	0.091	0.091	–	-0.7	–

(b) Quantile 2

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	24.898	24.976	–	+0.3	–
Tot beam bkg	0.297	0.295	–	-0.5	–
ν_e CC	0.002	0.002	–	-2.8	–
NC	0.144	0.142	–	-1.2	–
ν_μApp	0.039	0.039	–	+0.5	–
ν_τ CC	0.112	0.112	–	+0.2	–

(c) Quantile 3

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	26.178	26.418	–	+0.9	–
Tot beam bkg	0.536	0.542	–	+1.2	–
ν_e CC	0.017	0.016	–	-3.3	–
NC	0.348	0.355	–	+2.0	–
ν_μApp	0.038	0.037	–	-2.1	–
ν_τ CC	0.134	0.134	–	+0.4	–

(d) Quantile 4

	Nominal	Sh(+)	Sh(-)	%(+)	%(-)
ν_μ signal	31.915	31.957	–	+0.1	–
Tot beam bkg	2.410	2.490	–	+3.3	–
ν_e CC	0.227	0.244	–	+7.5	–
NC	1.994	2.061	–	+3.4	–
ν_μApp	0.042	0.041	–	-1.8	–
ν_τ CC	0.147	0.145	–	-2.0	–

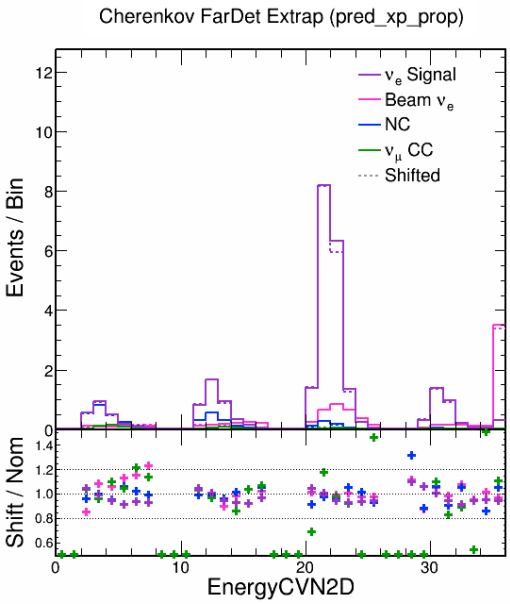
Cherenkov FarDet Extrap (pred_xp_prop)
ν_e Signal
Beam ν_e
NC
ν_μ CC
Shifted
Events / Bin
Shift / Nom
EnergyCVN2D

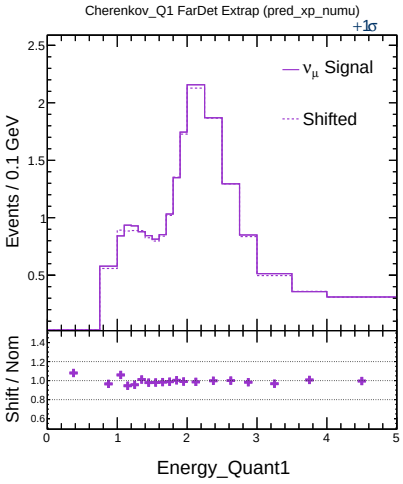
Cherenkov_Q1 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant1

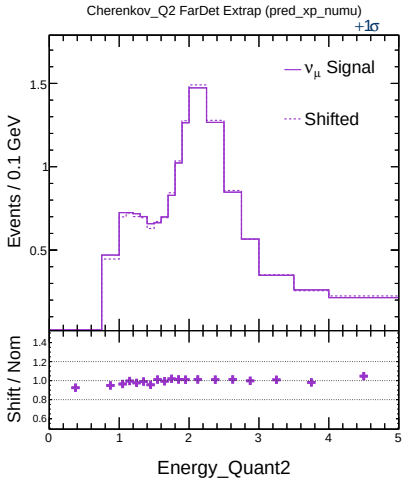
Cherenkov_Q2 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant2

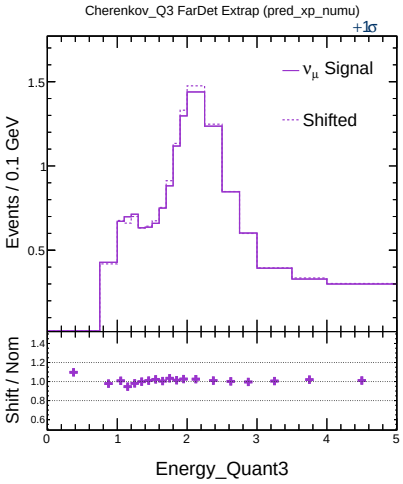
Cherenkov_Q3 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant3

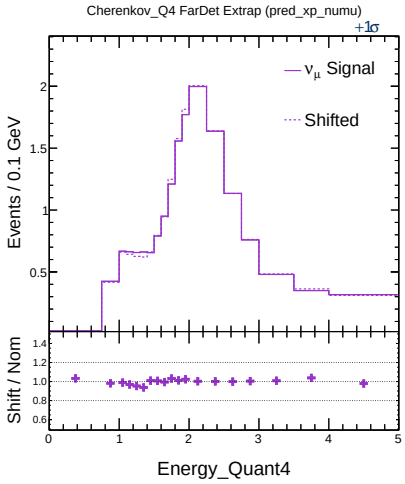
Cherenkov_Q4 FarDet Extrap (pred_xp_numu)
+1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

Table B.11: Light level - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.5	3.6	3.5	+1.5	+1.3	6.1	6.1	6.1	+0.7	+1.0	27	27	27	+0.8	+0.2	5.1	5.2	5	+2.6	-2.3
Tot beam bkg	4.5	4.6	4.5	+2.3	+0.1	3.5	3.5	3.7	-1.4	+6.0	5.4	5.4	5.2	-0.7	-3.9	1.2	1.2	1.1	+0.5	-4.9
Beam ν_e CC	0.87	0.88	0.94	+1.1	+7.8	1.5	1.5	1.5	-0.0	+3.0	4	4	3.9	+1.5	-1.2	0.91	0.91	0.88	-0.5	-3.6
NC	3.1	3.1	2.9	+1.8	-3.7	1.7	1.6	1.8	-3.0	+9.8	1.2	1.1	0.99	-8.4	-16.1	0.21	0.21	0.18	+1.6	-13.5
ν_μ CC	0.49	0.53	0.54	+7.8	+9.5	0.28	0.28	0.28	+1.8	+0.6	0.15	0.14	0.17	-2.4	+15.6	0.046	0.051	0.047	+11.1	+2.1
ν_τ CC	0.11	0.11	0.11	+2.1	+2.6	0.12	0.12	0.13	-3.7	+3.1	0.17	0.17	0.17	+2.1	+1.7	0.032	0.034	0.033	+8.1	+3.6

Table B.12: Light level - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	35	35	35	+0.6	+0.1
Tot beam bkg	0.22	0.23	0.23	+6.8	+6.1
NC	0.082	0.094	0.093	+14.5	+13.4
ν_μ App	0.042	0.043	0.043	+1.9	+1.4
ν_e CC	0.00026	0.00026	0.00026	+0.0	+0.0
ν_τ CC	0.091	0.093	0.093	+1.9	+1.4

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	25	25	+1.4	+0.7
Tot beam bkg	0.3	0.3	0.3	+1.9	+1.6
NC	0.14	0.15	0.15	+3.8	+2.2
ν_μ App	0.039	0.04	0.04	+1.8	+3.4
ν_e CC	0.002	0.002	0.002	+0.0	+0.0
ν_τ CC	0.11	0.11	0.11	-0.8	+0.4

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	26	27	26	+2.0	+0.5
Tot beam bkg	0.54	0.55	0.55	+2.9	+2.8
NC	0.35	0.36	0.36	+3.9	+3.2
ν_μ App	0.038	0.039	0.039	+2.3	+3.0
ν_e CC	0.017	0.017	0.017	+0.0	+0.0
ν_τ CC	0.13	0.13	0.13	-0.4	+0.5

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	33	32	+2.4	-0.8
Tot beam bkg	2.4	2.6	2.5	+7.1	+3.3
NC	2	2.1	2.1	+7.4	+2.9
ν_μ App	0.042	0.042	0.042	+1.0	+0.3
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.15	0.15	-0.5	-0.7

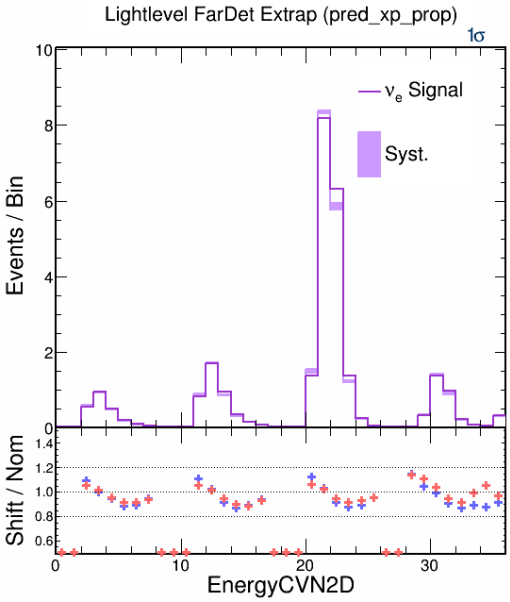
Lightlevel FarDet Extrap (pred_xp_prop)
1σ
ν_e Signal
Syst.
Events / Bin
Shift / Nom
EnergyCVN2D

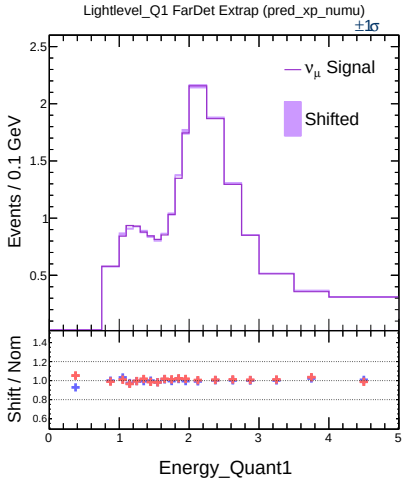
Lightlevel_Q1 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant1

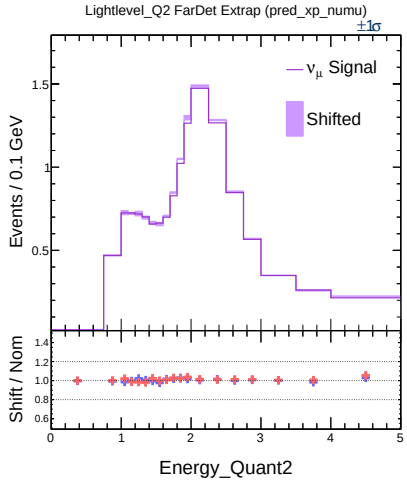
Lightlevel_Q2 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant2

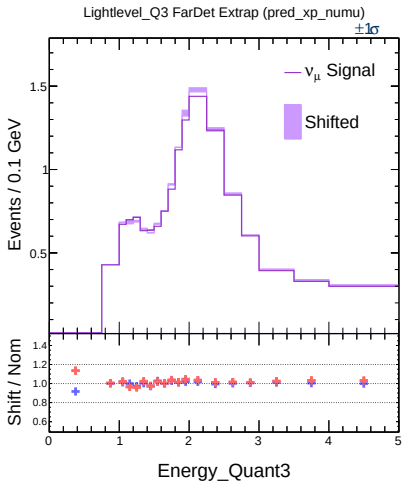
Lightlevel_Q3 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant3

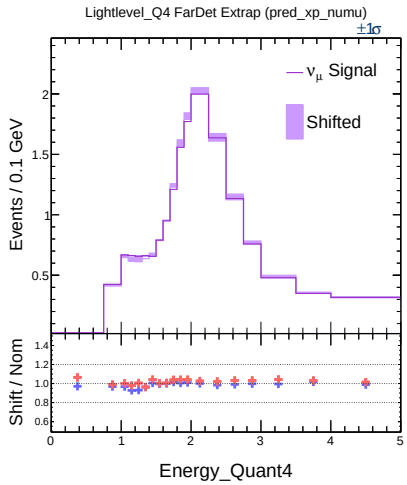
Lightlevel_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

Table B.13: MEC q0 Shape - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.7	3.6	+4.2	+1.4	6.2	6.4	6.3	+3.3	+1.9	27	26	30	-4.2	+8.9	5.2	5.1	5.7	-3.4	+9.4
Tot beam bkg	4.7	4.7	4.7	+1.0	+0.9	3.8	3.8	3.8	+1.5	+1.1	5.4	5.5	5.4	+3.2	+0.0	1.2	1.2	1.2	+3.0	-1.0
Beam ν_e CC	0.85	0.89	0.86	+4.2	+0.8	1.5	1.5	1.5	+0.7	-0.1	3.9	3.9	4.1	-0.9	+3.7	0.89	0.88	0.91	-1.7	+2.1
NC	3.2	3.2	3.3	-0.7	+0.5	1.8	1.9	1.9	+0.6	+0.7	1.1	1.3	0.89	+17.3	-18.2	0.2	0.23	0.16	+17.4	-18.6
ν_μ CC	0.49	0.49	0.49	+0.4	-0.6	0.28	0.28	0.28	-1.2	-0.5	0.16	0.16	0.17	-2.5	+3.2	0.055	0.062	0.054	+12.9	-1.3
ν_τ CC	0.11	0.14	0.13	+31.3	+22.9	0.13	0.16	0.16	+30.6	+24.4	0.17	0.2	0.22	+14.6	+26.1	0.033	0.041	0.039	+25.6	+20.8

Table B.14: MEC q0 Shape - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	33	35	-1.9	+3.4
Tot beam bkg	0.21	0.19	0.24	-8.8	+14.2
NC	0.078	0.077	0.077	-0.2	-0.2
ν_μ App	0.041	0.036	0.046	-13.5	+10.6
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.079	0.12	-13.9	+28.1

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	24	28	-5.3	+12.0
Tot beam bkg	0.29	0.33	0.32	+10.5	+9.0
NC	0.14	0.14	0.14	-0.3	-0.3
ν_μ App	0.04	0.041	0.038	+2.4	-4.2
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.14	0.14	+26.9	+25.2

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	26	29	-3.6	+7.9
Tot beam bkg	0.52	0.58	0.53	+12.9	+2.8
NC	0.33	0.33	0.33	-0.0	-0.0
ν_μ App	0.039	0.041	0.039	+5.3	-0.4
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.2	0.15	+49.3	+11.3

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	32	32	-0.4	-0.6
Tot beam bkg	2.4	2.4	2.4	+1.9	+0.4
NC	2	1.9	1.9	-0.1	-0.1
ν_μ App	0.043	0.043	0.043	+0.7	+0.2
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.2	0.16	+32.2	+7.5

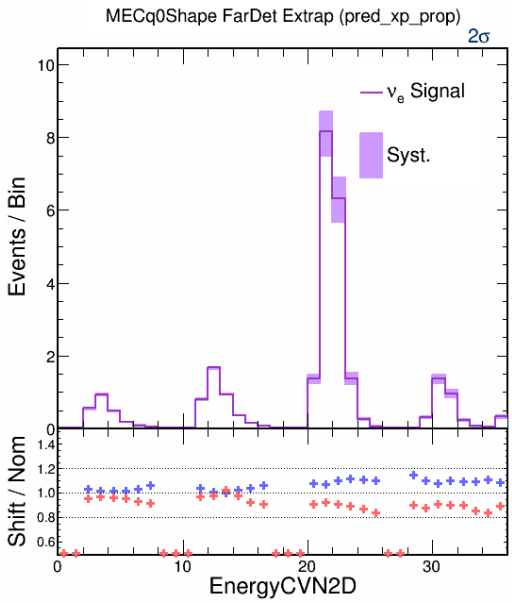
MECq0Shape FarDet Extrap (pred_xp_prop)
2σ
ν_e Signal
Syst.
Events / Bin
Shift / Nom
EnergyCVN2D

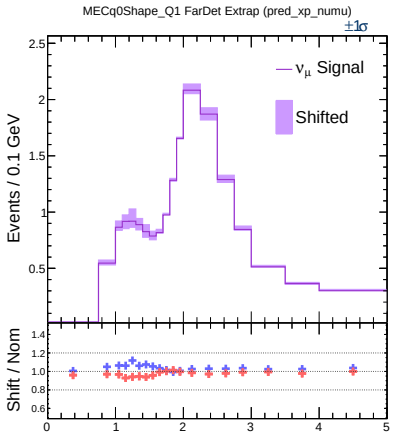
MECq0Shape_Q1 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant1

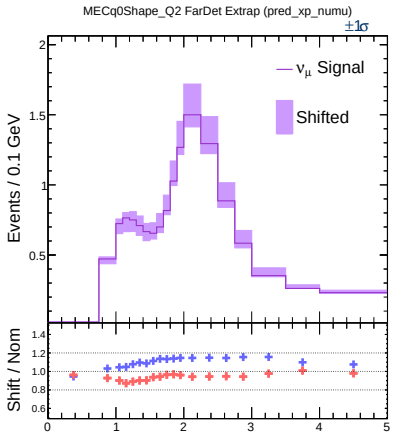
MECq0Shape_Q2 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant2

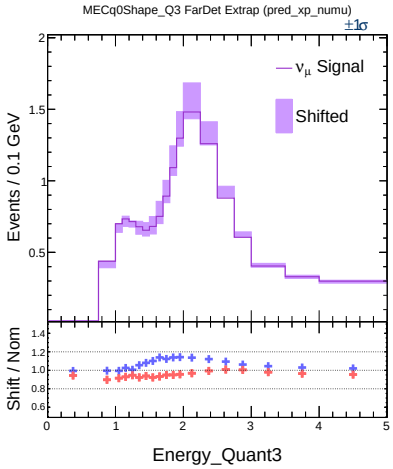
MECq0Shape_Q3 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant3

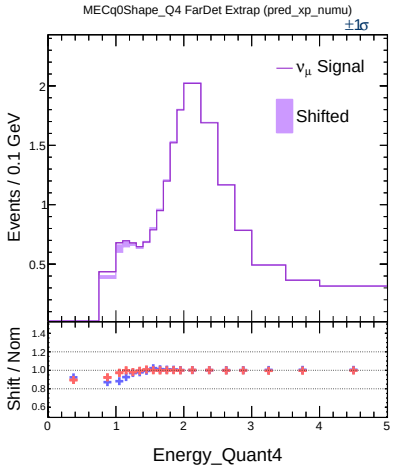
MECq0Shape_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

Table B.15: RPA Shape Resonance - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.7	3.6	+4.2	+0.0	6.2	6.4	6.2	+3.3	+0.0	27	28	27	+4.0	+0.0	5.2	5.5	5.2	+4.6	+0.0
Tot beam bkg	4.7	4.7	4.7	+0.1	+0.0	3.8	3.8	3.8	+0.3	+0.1	5.4	5.3	5.3	-0.1	-0.1	1.2	1.2	1.2	+0.7	-0.1
Beam ν_e CC	0.85	0.85	0.85	+0.5	+0.1	1.5	1.5	1.5	+0.0	-0.0	3.9	4	3.9	+0.7	+0.0	0.89	0.91	0.89	+1.5	+0.0
NC	3.2	3.2	3.2	-0.1	+0.0	1.8	1.9	1.8	+0.3	+0.3	1.1	1	1.1	-4.1	-0.5	0.2	0.19	0.2	-4.2	-0.5
ν_μ CC	0.49	0.49	0.49	+0.3	-0.1	0.28	0.29	0.28	+0.8	-0.0	0.16	0.17	0.16	+5.3	+0.2	0.055	0.057	0.055	+4.2	+0.3
ν_τ CC	0.11	0.11	0.11	+2.1	+0.0	0.13	0.13	0.13	+1.6	+0.0	0.17	0.17	0.17	+1.0	+0.0	0.033	0.033	0.033	+1.0	+0.0

Table B.16: RPA Shape Resonance - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	35	34	+1.7	+0.0
Tot beam bkg	0.21	0.21	0.21	-1.4	+0.0
NC	0.078	0.078	0.078	+0.0	+0.0
ν_μ App	0.041	0.039	0.041	-5.1	+0.0
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.091	0.092	-0.9	+0.0

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	27	25	+4.6	+0.0
Tot beam bkg	0.29	0.29	0.29	-0.4	+0.0
NC	0.14	0.14	0.14	+0.0	+0.0
ν_μ App	0.04	0.037	0.04	-6.3	+0.0
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.11	0.11	+1.2	+0.0

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	28	27	+6.1	+0.0
Tot beam bkg	0.52	0.52	0.52	+0.6	+0.0
NC	0.33	0.33	0.33	+0.0	+0.0
ν_μ App	0.039	0.038	0.039	-2.9	+0.0
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.14	0.13	+3.4	+0.0

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	33	32	+2.2	+0.0
Tot beam bkg	2.4	2.4	2.4	+0.2	+0.0
NC	2	2	2	+0.0	+0.0
ν_μ App	0.043	0.043	0.043	-0.0	+0.0
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.15	0.15	+3.4	+0.0

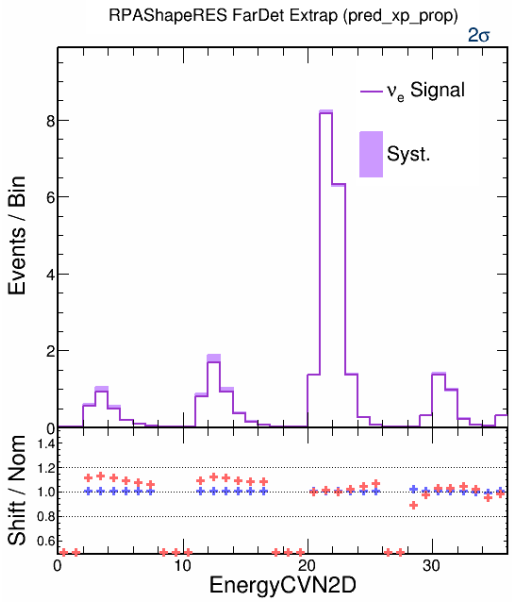
RPAShapeRES FarDet Extrap (pred_xp_prop)
2σ
ν_e Signal
Syst.
Events / Bin
Shift / Nom
EnergyCVN2D

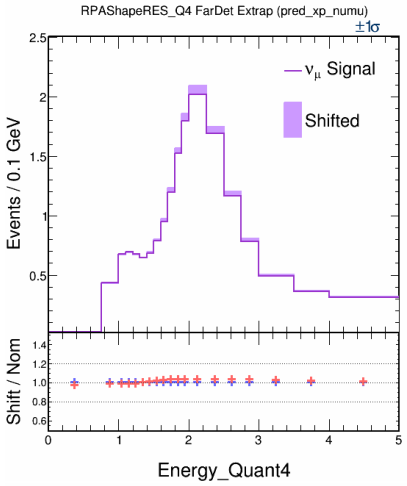
RPAShapeRES_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

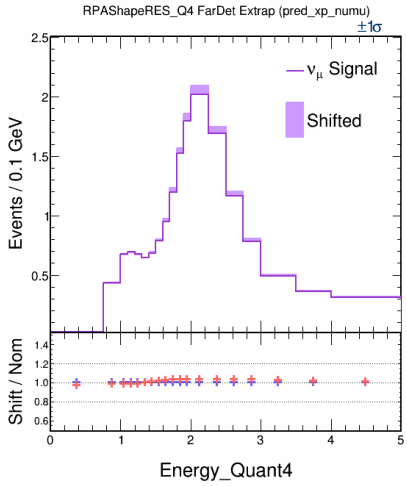
RPAShapeRES_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

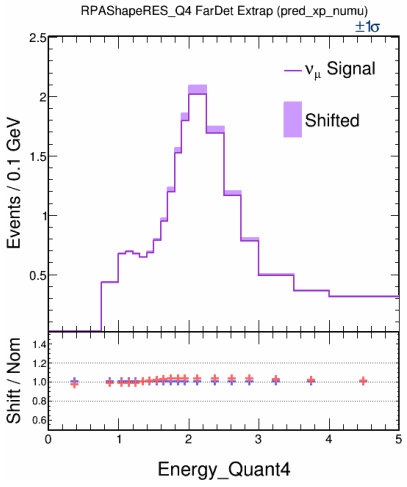
RPAShapeRES_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

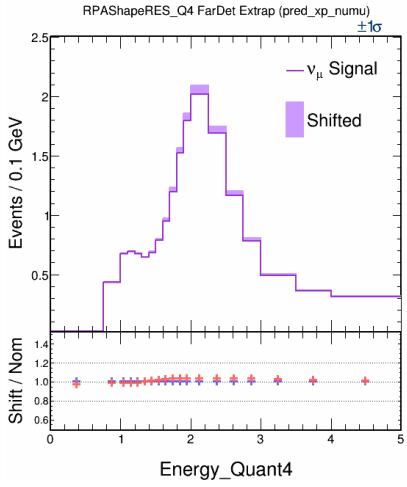
RPAShapeRES_Q4 FarDet Extrap (pred_xp_numu)
±1σ
ν_μ Signal
Shifted
Events / 0.1 GeV
Shift / Nom
Energy_Quant4

Table B.17: Ma CC Resonance - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.8	3.4	+5.5	-4.8	6.2	6.5	5.9	+4.7	-4.4	27	27	27	+0.8	-0.2	5.2	5.3	5.2	+1.2	-0.6
Tot beam bkg	4.7	4.7	4.7	-0.2	+0.4	3.8	3.8	3.7	+1.4	-1.3	5.4	5.5	5.2	+2.9	-2.8	1.2	1.2	1.1	+3.6	-3.1
Beam ν_e CC	0.85	0.89	0.81	+5.1	-4.3	1.5	1.6	1.4	+5.0	-4.6	3.9	4.1	3.8	+3.6	-3.4	0.89	0.93	0.86	+4.4	-3.8
NC	3.2	3.2	3.3	-2.3	+2.1	1.8	1.8	1.9	-2.0	+1.7	1.1	1.1	1.1	+0.9	-1.8	0.2	0.2	0.19	+1.0	-1.9
ν_μ CC	0.49	0.5	0.48	+1.4	-0.9	0.28	0.29	0.29	+0.5	+0.3	0.16	0.15	0.17	-6.5	+8.2	0.055	0.051	0.059	-5.9	+7.3
ν_τ CC	0.11	0.12	0.095	+14.7	-10.9	0.13	0.14	0.11	+12.3	-9.5	0.17	0.19	0.16	+9.1	-7.3	0.033	0.036	0.03	+11.6	-8.4

Table B.18: Ma CC Resonance - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	-0.1	+0.2
Tot beam bkg	0.21	0.21	0.21	+1.8	-2.0
NC	0.078	0.078	0.078	+0.0	+0.0
ν_μ App	0.041	0.042	0.04	+2.5	-3.2
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.094	0.089	+3.0	-3.2

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	26	25	+1.3	-0.8
Tot beam bkg	0.29	0.31	0.28	+6.1	-5.6
NC	0.14	0.14	0.14	+0.0	+0.0
ν_μ App	0.04	0.043	0.036	+8.4	-9.2
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.13	0.1	+12.9	-11.2

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	28	26	+3.8	-2.4
Tot beam bkg	0.52	0.55	0.49	+6.7	-5.1
NC	0.33	0.33	0.33	+0.0	+0.0
ν_μ App	0.039	0.044	0.034	+14.0	-13.3
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.16	0.11	+21.3	-15.4

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	36	28	+12.1	-10.8
Tot beam bkg	2.4	2.4	2.3	+2.9	-2.2
NC	2	2	2	+0.0	+0.0
ν_μ App	0.043	0.048	0.038	+12.5	-10.6
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.19	0.12	+25.0	-16.0

Table B.19: Ma NC Resonance - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	-0.1	+0.0	6.2	6.2	6.2	-0.1	+0.0	27	27	27	-0.1	+0.0	5.2	5.2	5.2	-0.1	+0.0
Tot beam bkg	4.7	4.7	4.7	-0.7	+0.8	3.8	3.8	3.8	-0.2	+0.2	5.4	5.4	5.3	+0.4	-0.4	1.2	1.2	1.2	-0.1	+0.1
Beam ν_e CC	0.85	0.85	0.85	-0.5	+0.3	1.5	1.5	1.5	-0.4	+0.3	3.9	3.9	4	-0.5	+0.4	0.89	0.89	0.9	-0.5	+0.4
NC	3.2	3.2	3.3	-1.2	+1.3	1.8	1.8	1.8	-0.2	+0.2	1.1	1.1	1.1	+3.6	-3.1	0.2	0.2	0.2	+1.1	-0.8
ν_μ CC	0.49	0.5	0.48	+1.4	-1.5	0.28	0.29	0.28	+0.9	-0.9	0.16	0.16	0.16	+0.6	-0.4	0.055	0.055	0.054	+0.6	-0.4
ν_τ CC	0.11	0.11	0.11	+0.0	+0.0	0.13	0.13	0.13	+0.0	+0.0	0.17	0.17	0.17	+0.0	+0.0	0.033	0.033	0.033	+0.0	+0.0

Table B.20: Ma NC Resonance - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	-0.0	+0.0
Tot beam bkg	0.21	0.22	0.2	+6.6	-6.1
NC	0.078	0.091	0.065	+17.9	-16.6
ν_μ App	0.041	0.041	0.041	+0.0	+0.0
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.092	0.092	+0.0	+0.0

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	25	25	-0.0	+0.0
Tot beam bkg	0.29	0.32	0.27	+8.4	-7.4
NC	0.14	0.16	0.12	+17.8	-15.5
ν_μ App	0.04	0.04	0.04	+0.0	+0.0
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.11	0.11	+0.0	+0.0

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	27	-0.1	+0.0
Tot beam bkg	0.52	0.56	0.48	+8.2	-6.7
NC	0.33	0.37	0.3	+12.8	-10.4
ν_μ App	0.039	0.039	0.039	+0.0	+0.0
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.13	0.13	+0.0	+0.0

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	32	32	-0.2	+0.1
Tot beam bkg	2.4	2.6	2.2	+8.4	-6.0
NC	2	2.2	1.8	+10.3	-7.4
ν_μ App	0.043	0.043	0.043	+0.0	+0.0
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.15	0.15	+0.0	+0.0

Table B.21: Mv CC Resonance - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.7	3.5	+3.5	-3.0	6.2	6.4	6	+3.0	-2.6	27	27	27	+0.6	-0.4	5.2	5.3	5.2	+0.9	-0.7
Tot beam bkg	4.7	4.7	4.7	-0.1	+0.2	3.8	3.8	3.7	+0.9	-0.8	5.4	5.4	5.3	+1.7	-1.5	1.2	1.2	1.2	+2.0	-1.7
Beam ν_e CC	0.85	0.87	0.83	+2.7	-2.2	1.5	1.5	1.5	+2.7	-2.4	3.9	4	3.9	+2.1	-1.8	0.89	0.92	0.87	+2.5	-2.1
NC	3.2	3.2	3.3	-1.2	+1.1	1.8	1.8	1.9	-0.9	+0.6	1.1	1.1	1.1	+0.5	-0.8	0.2	0.2	0.2	+0.6	-0.8
ν_μ CC	0.49	0.49	0.49	+0.7	-0.5	0.28	0.28	0.29	-0.1	+0.2	0.16	0.15	0.17	-3.8	+3.8	0.055	0.053	0.057	-3.3	+3.4
ν_τ CC	0.11	0.11	0.1	+7.6	-5.8	0.13	0.13	0.12	+6.5	-5.0	0.17	0.18	0.17	+4.8	-3.8	0.033	0.035	0.031	+6.1	-4.5

Table B.22: Mv CC Resonance - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	+0.0	+0.0
Tot beam bkg	0.21	0.21	0.21	+0.9	-0.8
NC	0.078	0.078	0.078	+0.0	+0.0
ν_μ App	0.041	0.042	0.041	+1.1	-1.2
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.093	0.09	+1.5	-1.4

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	26	25	+0.9	-0.7
Tot beam bkg	0.29	0.3	0.29	+3.3	-2.8
NC	0.14	0.14	0.14	+0.0	+0.0
ν_μ App	0.04	0.041	0.038	+4.4	-4.1
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.12	0.11	+6.9	-5.6

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	26	+2.2	-1.7
Tot beam bkg	0.52	0.53	0.5	+3.5	-2.7
NC	0.33	0.33	0.33	+0.0	+0.0
ν_μ App	0.039	0.042	0.036	+7.8	-6.8
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.15	0.12	+10.8	-8.2

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	33	32	+1.1	-0.8
Tot beam bkg	2.4	2.4	2.3	+1.5	-1.2
NC	2	2	2	+0.0	+0.0
ν_μ App	0.043	0.046	0.04	+6.8	-5.7
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.17	0.14	+12.2	-8.8

Table B.23: RPA Shape Enh - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	+0.2	+1.0	6.2	6.2	6.2	+0.6	+0.4	27	28	27	+1.7	-1.0	5.2	5.3	5.2	+1.9	-1.1
Tot beam bkg	4.7	4.7	4.7	+0.1	-0.0	3.8	3.8	3.8	+0.4	+0.2	5.4	5.4	5.4	+0.3	+0.1	1.2	1.2	1.2	+0.7	+0.1
Beam ν_e CC	0.85	0.86	0.85	+1.2	+0.2	1.5	1.5	1.5	+1.1	-0.1	3.9	4	3.9	+1.4	-0.6	0.89	0.91	0.89	+1.7	-0.5
NC	3.2	3.2	3.2	-0.4	+0.0	1.8	1.8	1.9	-0.6	+0.7	1.1	1	1.1	-4.8	+3.3	0.2	0.19	0.2	-4.9	+3.4
ν_μ CC	0.49	0.49	0.49	+0.1	+0.2	0.28	0.29	0.29	+0.1	+0.4	0.16	0.16	0.16	+0.4	-0.3	0.055	0.055	0.055	+1.2	-0.3
ν_τ CC	0.11	0.11	0.1	+7.3	-3.7	0.13	0.13	0.12	+7.2	-4.2	0.17	0.18	0.16	+7.0	-4.5	0.033	0.035	0.031	+6.8	-3.9

Table B.24: RPA Shape Enh - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	+0.5	-0.3
Tot beam bkg	0.21	0.22	0.2	+3.8	-3.2
NC	0.078	0.078	0.078	+0.0	+0.0
ν_μ App	0.041	0.043	0.039	+4.9	-4.3
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.098	0.087	+6.6	-5.4

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	26	25	+1.3	-0.9
Tot beam bkg	0.29	0.31	0.29	+4.0	-2.9
NC	0.14	0.14	0.14	+0.0	+0.0
ν_μ App	0.04	0.042	0.038	+6.3	-5.0
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.12	0.11	+8.2	-5.7

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	27	+1.2	+0.2
Tot beam bkg	0.52	0.53	0.51	+2.2	-0.8
NC	0.33	0.33	0.33	+0.0	+0.0
ν_μ App	0.039	0.041	0.038	+5.3	-3.2
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.14	0.13	+6.9	-2.2

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	33	33	+0.3	+0.5
Tot beam bkg	2.4	2.4	2.4	+0.8	-0.2
NC	2	2	2	+0.0	+0.0
ν_μ App	0.043	0.044	0.042	+3.6	-1.7
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.16	0.15	+5.8	-0.6

Table B.25: DIS vn CC 1π - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.7	3.5	+3.1	-3.1	6.2	6.3	6.1	+2.2	-2.2	27	27	27	+0.5	-0.5	5.2	5.3	5.2	+0.6	-0.6
Tot beam bkg	4.7	4.7	4.7	-0.6	+0.8	3.8	3.7	3.8	-0.5	+0.6	5.4	5.4	5.4	+0.1	-0.0	1.2	1.2	1.2	+0.3	-0.2
Beam ν_e CC	0.85	0.88	0.82	+3.5	-3.5	1.5	1.5	1.5	+1.6	-1.5	3.9	3.9	3.9	-0.1	+0.2	0.89	0.9	0.89	+0.2	-0.2
NC	3.2	3.2	3.3	-2.2	+2.6	1.8	1.8	1.9	-2.9	+3.3	1.1	1.1	1.1	-0.1	+0.1	0.2	0.2	0.2	-0.0	-0.0
ν_μ CC	0.49	0.5	0.47	+2.6	-3.2	0.28	0.3	0.27	+4.4	-5.2	0.16	0.17	0.15	+4.3	-5.1	0.055	0.055	0.054	+1.3	-1.2
ν_τ CC	0.11	0.11	0.11	+0.9	-0.9	0.13	0.13	0.12	+0.7	-0.7	0.17	0.17	0.17	+0.4	-0.4	0.033	0.033	0.032	+0.6	-0.6

Table B.26: DIS vn CC 1π - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	-0.1	+0.1
Tot beam bkg	0.21	0.21	0.21	+0.3	-0.3
NC	0.078	0.078	0.078	+0.0	+0.0
ν_μ App	0.041	0.041	0.041	+0.5	-0.5
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.092	0.091	+0.3	-0.3

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	25	25	+0.1	-0.1
Tot beam bkg	0.29	0.3	0.29	+0.5	-0.5
NC	0.14	0.14	0.14	+0.0	+0.0
ν_μ App	0.04	0.04	0.039	+1.1	-1.1
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.11	0.11	+0.6	-0.6

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	26	+0.7	-0.7
Tot beam bkg	0.52	0.52	0.51	+1.0	-1.0
NC	0.33	0.33	0.33	+0.0	+0.0
ν_μ App	0.039	0.04	0.037	+3.3	-3.3
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.13	0.13	+1.7	-1.7

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	33	32	+0.7	-0.8
Tot beam bkg	2.4	2.4	2.4	+1.1	-1.1
NC	2	2	2	+0.0	+0.0
ν_μ App	0.043	0.045	0.04	+6.1	-6.1
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.15	0.14	+3.0	-3.0

Table B.27: Ma CC Quasi-elastic - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	+0.1	+0.2	6.2	6.2	6.2	+0.3	-0.0	27	28	27	+1.2	-1.1	5.2	5.3	5.2	+1.4	-1.2
Tot beam bkg	4.7	4.7	4.7	-0.0	+0.0	3.8	3.8	3.8	+0.1	+0.1	5.4	5.4	5.4	-0.0	-0.0	1.2	1.2	1.2	+0.3	-0.2
Beam ν_e CC	0.85	0.85	0.85	+0.5	-0.1	1.5	1.5	1.5	+0.3	-0.3	3.9	4	3.9	+0.8	-0.6	0.89	0.9	0.89	+1.1	-0.8
NC	3.2	3.2	3.2	-0.2	+0.2	1.8	1.8	1.9	-0.2	+0.6	1.1	1	1.1	-3.5	+2.6	0.2	0.19	0.2	-3.5	+2.6
ν_μ CC	0.49	0.49	0.49	+0.1	-0.1	0.28	0.29	0.28	+0.2	-0.1	0.16	0.16	0.16	+1.0	-0.5	0.055	0.055	0.054	+1.2	-0.7
ν_τ CC	0.11	0.11	0.1	+2.7	-2.0	0.13	0.13	0.12	+2.7	-2.1	0.17	0.18	0.17	+2.7	-2.2	0.033	0.034	0.032	+2.8	-2.2

Table B.28: Ma CC Quasi-elastic - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	+0.6	-0.5
Tot beam bkg	0.21	0.21	0.21	+1.8	-1.8
NC	0.078	0.078	0.078	+0.0	+0.0
ν_μ App	0.041	0.042	0.04	+1.9	-2.1
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.095	0.089	+3.2	-3.1

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	26	25	+0.9	-0.7
Tot beam bkg	0.29	0.3	0.29	+1.2	-1.0
NC	0.14	0.14	0.14	+0.0	+0.0
ν_μ App	0.04	0.04	0.039	+1.5	-1.3
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.12	0.11	+2.7	-2.1

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	27	+0.7	-0.5
Tot beam bkg	0.52	0.52	0.51	+0.7	-0.5
NC	0.33	0.33	0.33	+0.0	+0.0
ν_μ App	0.039	0.039	0.038	+1.0	-0.8
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.13	0.13	+2.2	-1.5

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	32	32	+0.1	-0.0
Tot beam bkg	2.4	2.4	2.4	+0.1	-0.1
NC	2	2	2	+0.0	+0.0
ν_μ App	0.043	0.043	0.043	+0.3	-0.2
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.15	0.15	+1.9	-1.2

Table B.29: Other cross section systematics - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.7	3.4	+2.2	-6.2	6.2	6.2	6.1	+0.3	-1.3	27	26	30	-3.9	+11.5	5.2	5	6	-4.6	+13.6
Tot beam bkg	4.7	4.9	4.2	+4.9	-9.6	3.8	4	3.4	+7.0	-9.4	5.4	5.6	5	+4.3	-6.6	1.2	1.2	1.1	+3.9	-4.6
Beam ν_e CC	0.85	0.87	0.88	+2.4	+3.0	1.5	1.6	1.6	+2.8	+3.1	3.9	3.9	4.3	-0.8	+8.7	0.89	0.89	0.97	-0.5	+9.1
NC	3.2	3.4	2.8	+5.7	-14.4	1.8	2	1.5	+9.6	-19.5	1.1	1.3	0.43	+19.1	-60.4	0.2	0.24	0.072	+19.5	-63.8
ν_μ CC	0.49	0.47	0.52	-3.1	+7.1	0.28	0.28	0.29	-1.9	+2.0	0.16	0.15	0.18	-3.9	+10.5	0.055	0.055	0.058	-0.3	+6.4
ν_τ CC	0.11	0.15	0.065	+39.6	-38.7	0.13	0.17	0.079	+38.1	-37.3	0.17	0.23	0.11	+35.1	-34.3	0.033	0.045	0.021	+37.7	-37.0

Table B.30: Other cross section systematics - ν_e

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	31	44	17	+44.1	-44.3
Tot beam bkg	0.21	0.31	0.11	+46.6	-45.8
NC	0.078	0.12	0.035	+56.2	-54.4
ν_μ App	0.041	0.059	0.023	+43.4	-43.9
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.13	0.056	+39.8	-39.4

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	30	43	16	+44.9	-45.2
Tot beam bkg	0.29	0.44	0.15	+51.0	-49.6
NC	0.14	0.22	0.059	+60.5	-58.0
ν_μ App	0.04	0.057	0.022	+45.1	-45.3
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.16	0.067	+41.7	-40.8

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	30	45	15	+49.7	-49.6
Tot beam bkg	0.52	0.82	0.22	+58.7	-57.2
NC	0.33	0.55	0.12	+65.5	-63.3
ν_μ App	0.039	0.059	0.019	+51.1	-50.3
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.19	0.074	+44.7	-43.7

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	51	13	+59.6	-58.2
Tot beam bkg	2.4	4	0.81	+69.3	-65.8
NC	2	3.4	0.61	+72.7	-68.5
ν_μ App	0.043	0.069	0.017	+61.6	-59.8
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.22	0.08	+47.5	-46.4

Table B.31: Summed beam transport systematics - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	+0.5	-0.4	6.2	6.2	6.2	+0.5	-0.4	27	27	27	+0.6	-0.5	5.2	5.3	5.2	+0.7	-0.6
Tot beam bkg	4.7	4.7	4.7	+0.3	-0.1	3.8	3.8	3.8	+0.4	-0.0	5.4	5.4	5.3	+0.4	-0.5	1.2	1.2	1.2	+0.2	-0.4
Beam ν_e CC	0.85	0.84	0.86	-1.2	+0.8	1.5	1.5	1.5	-0.9	+0.5	3.9	3.9	4	-0.5	+0.5	0.89	0.89	0.9	-0.6	+0.5
NC	3.2	3.3	3.2	+0.4	-0.0	1.8	1.9	1.8	+1.0	+0.0	1.1	1.1	1	+2.4	-3.4	0.2	0.2	0.19	+2.4	-3.2
ν_μ CC	0.49	0.49	0.49	+1.1	-0.8	0.28	0.29	0.28	+1.6	-1.1	0.16	0.16	0.16	+2.6	-1.2	0.055	0.056	0.054	+2.7	-1.3
ν_τ CC	0.11	0.11	0.1	+4.9	-4.9	0.13	0.13	0.12	+4.9	-4.9	0.17	0.18	0.16	+4.9	-4.8	0.033	0.034	0.031	+4.8	-4.7

Table B.32: Summed beam transport systematics - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	+0.0	+0.1
Tot beam bkg	0.21	0.22	0.2	+4.6	-4.7
NC	0.078	0.082	0.073	+5.4	-5.3
ν_μ App	0.041	0.042	0.04	+2.8	-3.6
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.096	0.087	+4.8	-4.7

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	25	25	+0.0	+0.1
Tot beam bkg	0.29	0.31	0.28	+4.7	-4.8
NC	0.14	0.15	0.13	+5.3	-5.2
ν_μ App	0.04	0.041	0.038	+2.8	-3.6
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.12	0.11	+4.7	-4.6

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	27	+0.0	+0.1
Tot beam bkg	0.52	0.54	0.49	+5.0	-4.9
NC	0.33	0.35	0.31	+5.4	-5.3
ν_μ App	0.039	0.04	0.037	+2.7	-3.6
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.14	0.13	+4.6	-4.4

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	32	32	+0.1	-0.0
Tot beam bkg	2.4	2.5	2.3	+5.3	-5.2
NC	2	2.1	1.8	+5.5	-5.3
ν_μ App	0.043	0.044	0.041	+2.7	-3.7
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.16	0.14	+4.5	-4.4

Table B.33: PPFX Systematics (Principal component 00) - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	-0.1	+0.1	6.2	6.2	6.2	-0.1	+0.0	27	27	27	-0.2	+0.1	5.2	5.2	5.2	-0.2	+0.2
Tot beam bkg	4.7	4.7	4.7	-0.0	-0.0	3.8	3.8	3.8	-0.1	+0.1	5.4	5.3	5.4	-0.4	+0.4	1.2	1.2	1.2	-0.1	+0.1
Beam ν_e CC	0.85	0.85	0.85	+0.4	-0.3	1.5	1.5	1.5	+0.2	-0.2	3.9	3.9	3.9	-0.0	+0.1	0.89	0.89	0.89	+0.2	-0.1
NC	3.2	3.2	3.2	-0.2	+0.1	1.8	1.8	1.9	-0.5	+0.4	1.1	1.1	1.1	-2.1	+2.0	0.2	0.19	0.2	-1.8	+1.7
ν_μ CC	0.49	0.49	0.49	-0.3	+0.4	0.28	0.28	0.29	-0.6	+0.6	0.16	0.16	0.16	-0.6	+0.6	0.055	0.055	0.055	+0.3	-0.3
ν_τ CC	0.11	0.11	0.1	+2.9	-2.9	0.13	0.13	0.12	+2.8	-2.8	0.17	0.18	0.17	+2.7	-2.7	0.033	0.034	0.032	+3.0	-3.0

Table B.34: PPFX Systematics (Principal component 00) - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	+0.3	-0.3
Tot beam bkg	0.21	0.21	0.21	+1.8	-1.8
NC	0.078	0.078	0.078	+0.1	-0.1
ν_μ App	0.041	0.042	0.04	+1.8	-1.8
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.095	0.088	+3.4	-3.4

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	26	25	+0.3	-0.3
Tot beam bkg	0.29	0.3	0.29	+1.8	-1.8
NC	0.14	0.14	0.14	+0.5	-0.5
ν_μ App	0.04	0.04	0.039	+1.7	-1.7
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.12	0.11	+3.4	-3.4

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	27	+0.3	-0.3
Tot beam bkg	0.52	0.52	0.51	+1.6	-1.6
NC	0.33	0.33	0.33	+0.9	-0.9
ν_μ App	0.039	0.039	0.038	+1.7	-1.7
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.14	0.13	+3.5	-3.5

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	33	32	+0.2	-0.2
Tot beam bkg	2.4	2.4	2.3	+1.5	-1.5
NC	2	2	1.9	+1.5	-1.5
ν_μ App	0.043	0.043	0.042	+1.7	-1.7
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.15	0.14	+3.5	-3.5

Table B.35: PPFX Systematics (Principal component 01) - ν_e

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	+0.2	-0.2	6.2	6.2	6.2	+0.2	-0.2	27	27	27	+0.3	-0.2	5.2	5.3	5.2	+0.3	-0.3
Tot beam bkg	4.7	4.7	4.7	+0.1	-0.0	3.8	3.8	3.8	+0.0	+0.0	5.4	5.4	5.3	+0.2	-0.1	1.2	1.2	1.2	+0.2	-0.1
Beam ν_e CC	0.85	0.85	0.85	+0.2	-0.1	1.5	1.5	1.5	+0.2	-0.1	3.9	4	3.9	+0.4	-0.3	0.89	0.9	0.89	+0.3	-0.2
NC	3.2	3.3	3.2	+0.4	-0.3	1.8	1.9	1.8	+0.5	-0.4	1.1	1.1	1.1	+1.0	-0.8	0.2	0.2	0.2	+1.0	-0.8
ν_μ CC	0.49	0.49	0.49	-0.1	+0.1	0.28	0.28	0.28	-0.1	+0.0	0.16	0.16	0.16	-0.0	+0.1	0.055	0.055	0.055	-0.3	+0.3
ν_τ CC	0.11	0.098	0.12	-8.1	+8.1	0.13	0.12	0.14	-8.1	+8.1	0.17	0.16	0.19	-8.0	+8.0	0.033	0.03	0.035	-7.9	+7.9

Table B.36: PPFX Systematics (Principal component 01) - ν_μ

(a) Quantile 1

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	34	34	34	+0.1	-0.1
Tot beam bkg	0.21	0.2	0.23	-7.4	+7.4
NC	0.078	0.072	0.083	-7.0	+7.0
ν_μ App	0.041	0.038	0.044	-6.6	+6.6
ν_e CC	0.00047	0.00047	0.00047	+0.0	+0.0
ν_τ CC	0.092	0.084	0.099	-8.1	+8.1

(b) Quantile 2

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	25	25	25	+0.1	-0.1
Tot beam bkg	0.29	0.27	0.32	-7.3	+7.3
NC	0.14	0.13	0.15	-7.1	+7.1
ν_μ App	0.04	0.037	0.042	-6.5	+6.5
ν_e CC	0.0022	0.0022	0.0022	+0.0	+0.0
ν_τ CC	0.11	0.1	0.12	-8.0	+8.0

(c) Quantile 3

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	27	27	27	+0.1	-0.1
Tot beam bkg	0.52	0.48	0.55	-7.2	+7.2
NC	0.33	0.31	0.35	-7.0	+7.0
ν_μ App	0.039	0.036	0.041	-6.5	+6.5
ν_e CC	0.015	0.015	0.015	+0.0	+0.0
ν_τ CC	0.13	0.12	0.14	-7.8	+7.8

(d) Quantile 4

Integral shift	Nominal	Shift (+)	Shift (-)	% Diff. (+)	% Diff. (-)
ν_μ signal	32	32	32	+0.2	-0.2
Tot beam bkg	2.4	2.2	2.5	-7.1	+7.1
NC	2	1.8	2.1	-7.1	+7.1
ν_μ App	0.043	0.04	0.045	-6.5	+6.5
ν_e CC	0.23	0.23	0.23	+0.0	+0.0
ν_τ CC	0.15	0.14	0.16	-7.8	+7.8

Table B.37: Extrapolation systematics - ν_e signal

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	+0.0	+0.0	6.2	6.2	6.2	+0.0	+0.0	28	28	28	+0.0	+0.0	5.4	5.4	5.4	+0.0	+0.0
Tot beam bkg	4.1	4.1	4.1	+0.0	+0.0	3.6	3.6	3.6	+0.0	+0.0	4.8	4.8	4.8	+0.0	+0.0	1.2	1.2	1.2	+0.0	+0.0
Beam ν_e CC	0.82	0.82	0.82	+0.0	+0.0	1.5	1.5	1.5	+0.0	+0.0	3.7	3.7	3.7	+0.0	+0.0	0.98	0.98	0.98	+0.0	+0.0
NC	2.8	2.8	2.8	+0.0	+0.0	1.7	1.7	1.7	+0.0	+0.0	0.82	0.82	0.82	+0.0	+0.0	0.15	0.15	0.15	+0.0	+0.0
ν_μ CC	0.42	0.42	0.42	+0.0	+0.0	0.26	0.26	0.26	+0.0	+0.0	0.14	0.14	0.14	+0.0	+0.0	0.05	0.05	0.05	+0.0	+0.0
ν_τ CC	0.11	0.11	0.11	+0.0	+0.0	0.13	0.13	0.13	+0.0	+0.0	0.17	0.17	0.17	+0.0	+0.0	0.035	0.035	0.035	+0.0	+0.0

Table B.38: Extrapolation systematics - ν_e background

	Bin 1					Bin 2					Bin 3					Bin 4				
	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)	Nom.	(+)	(-)	%(+)	%(-)
ν_e signal	3.6	3.6	3.6	+0.0	+0.0	6.2	6.2	6.2	+0.0	+0.0	28	28	28	+0.0	+0.0	5.4	5.4	5.4	+0.0	+0.0
Tot beam bkg	4.1	4.2	4.1	+1.2	-1.2	3.6	3.6	3.5	+1.2	-1.2	4.8	4.9	4.8	+1.2	-1.2	1.2	1.2	1.2	+1.3	-1.3
Beam ν_e CC	0.82	0.83	0.81	+1.2	-1.2	1.5	1.5	1.5	+1.2	-1.2	3.7	3.8	3.7	+1.2	-1.2	0.98	0.99	0.97	+1.3	-1.3
NC	2.8	2.8	2.8	+1.2	-1.2	1.7	1.7	1.7	+1.2	-1.2	0.82	0.83	0.81	+1.2	-1.2	0.15	0.16	0.15	+1.3	-1.3
ν_μ CC	0.42	0.43	0.42	+1.2	-1.2	0.26	0.26	0.25	+1.2	-1.2	0.14	0.14	0.13	+1.2	-1.2	0.05	0.05	0.049	+1.3	-1.3
ν_τ CC	0.11	0.11	0.11	+1.2	-1.2	0.13	0.13	0.12	+1.2	-1.2	0.17	0.17	0.17	+1.2	-1.2	0.035	0.036	0.035	+1.3	-1.3

APPENDIX C

Event Display Gallery

C.1 ν_e Core Sample Selected Events

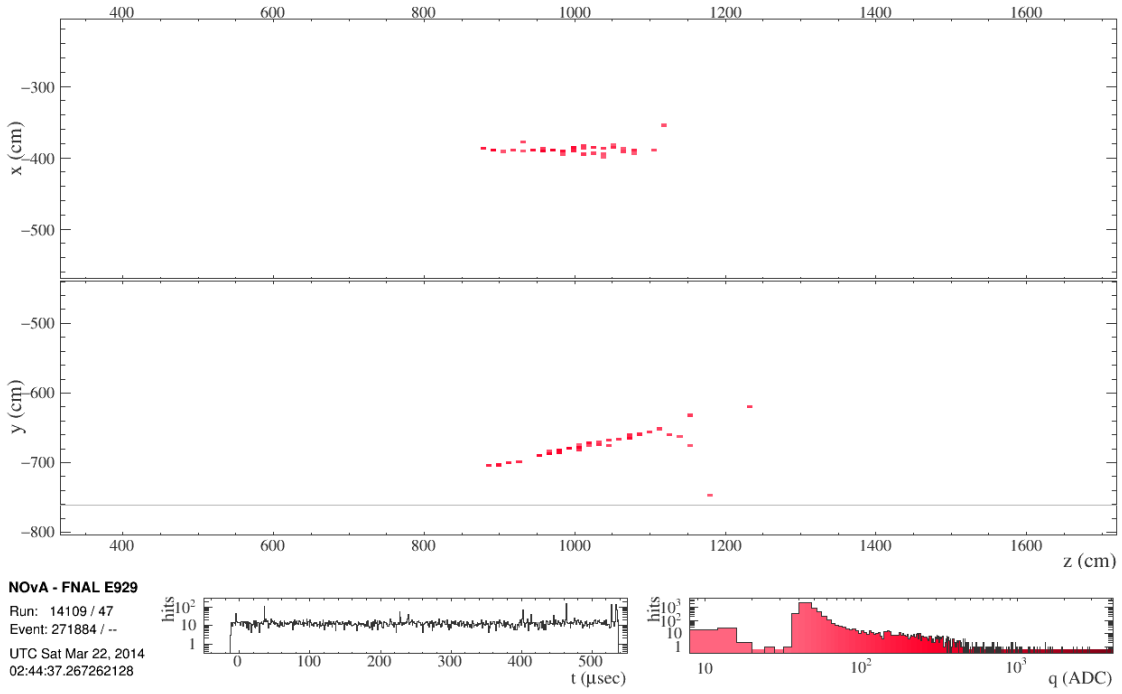

Figure C.1: Selected ν_e core event.

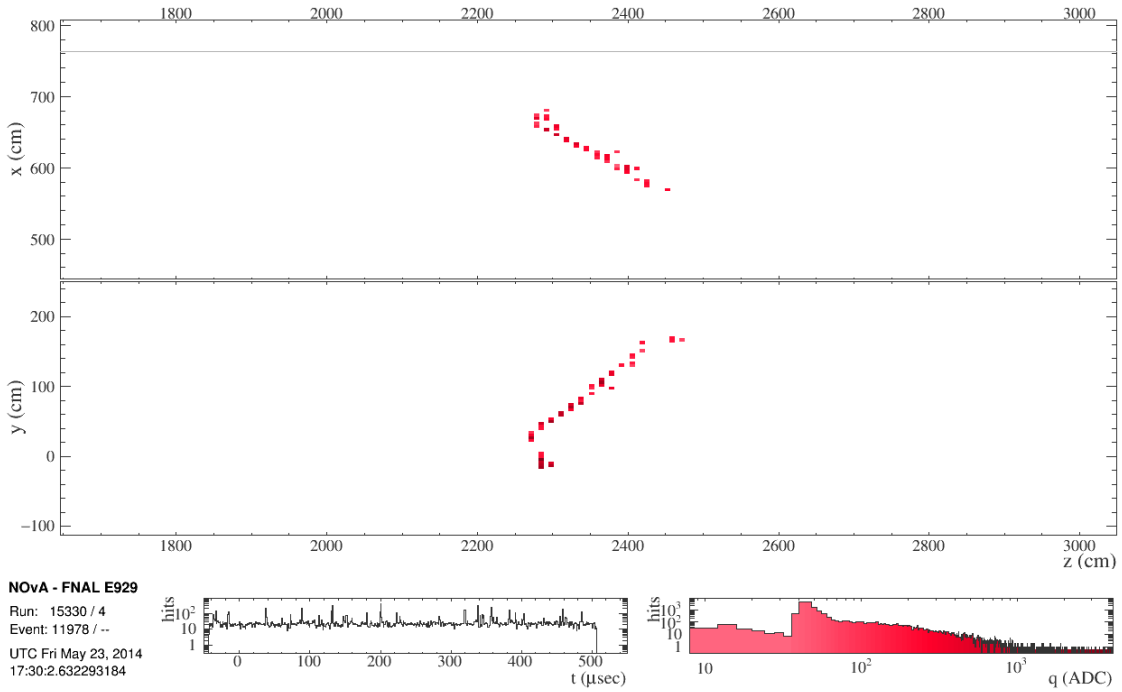

Figure C.2: Selected ν_e core event.

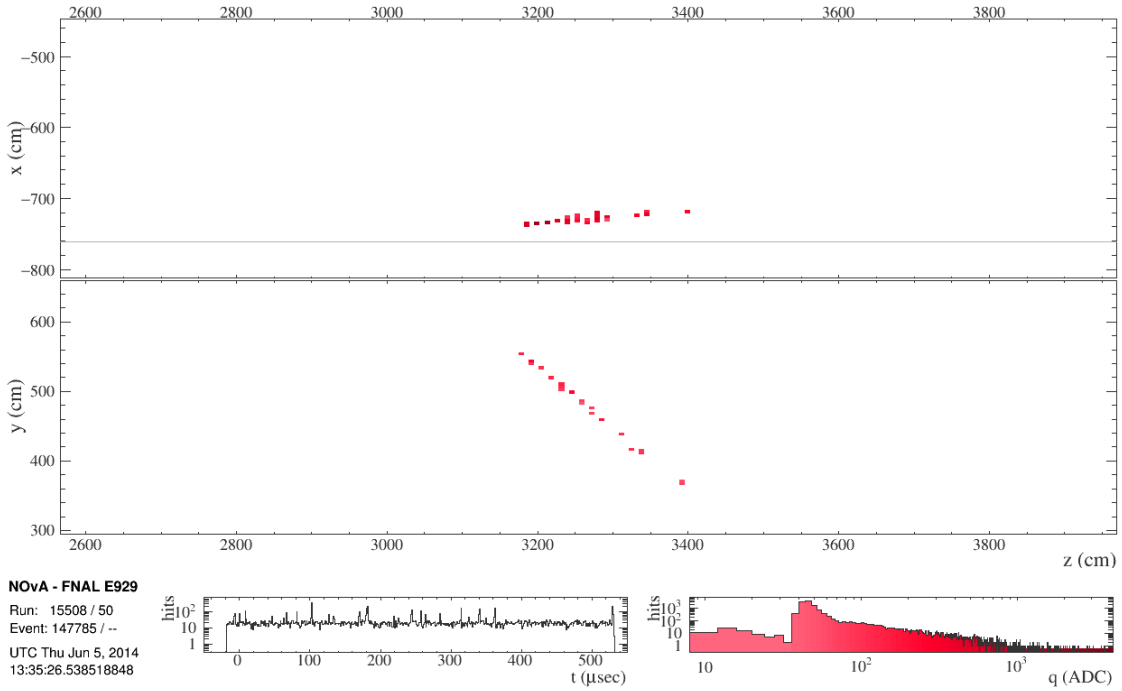

Figure C.3: Selected ν_e core event.

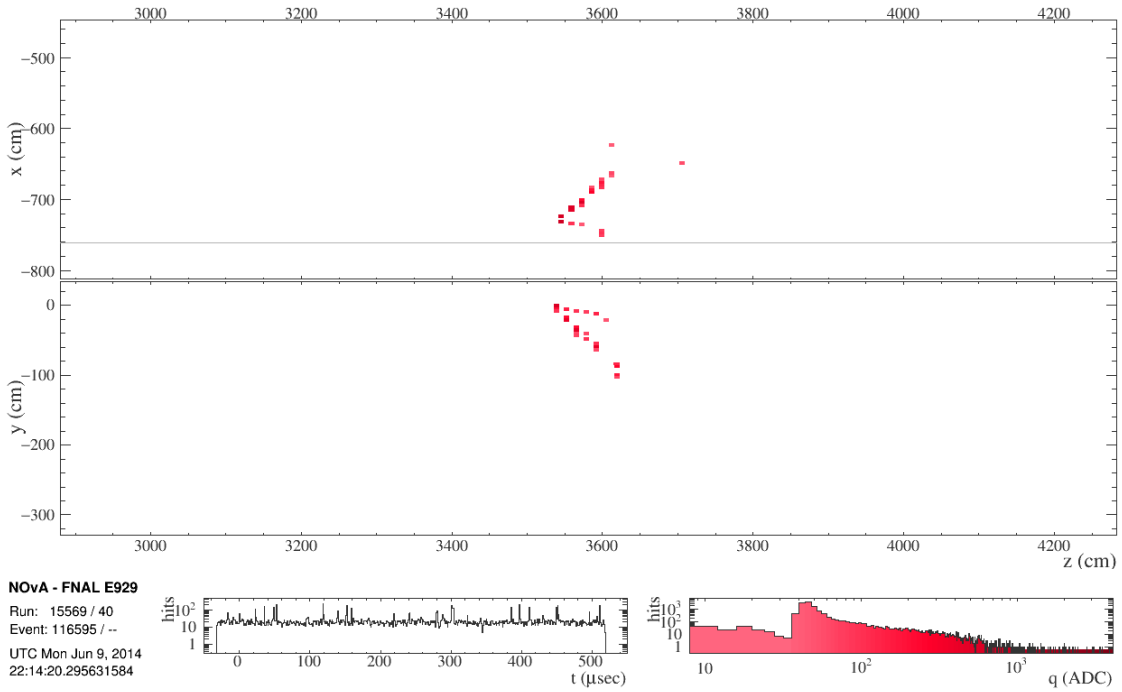

Figure C.4: Selected ν_e core event.

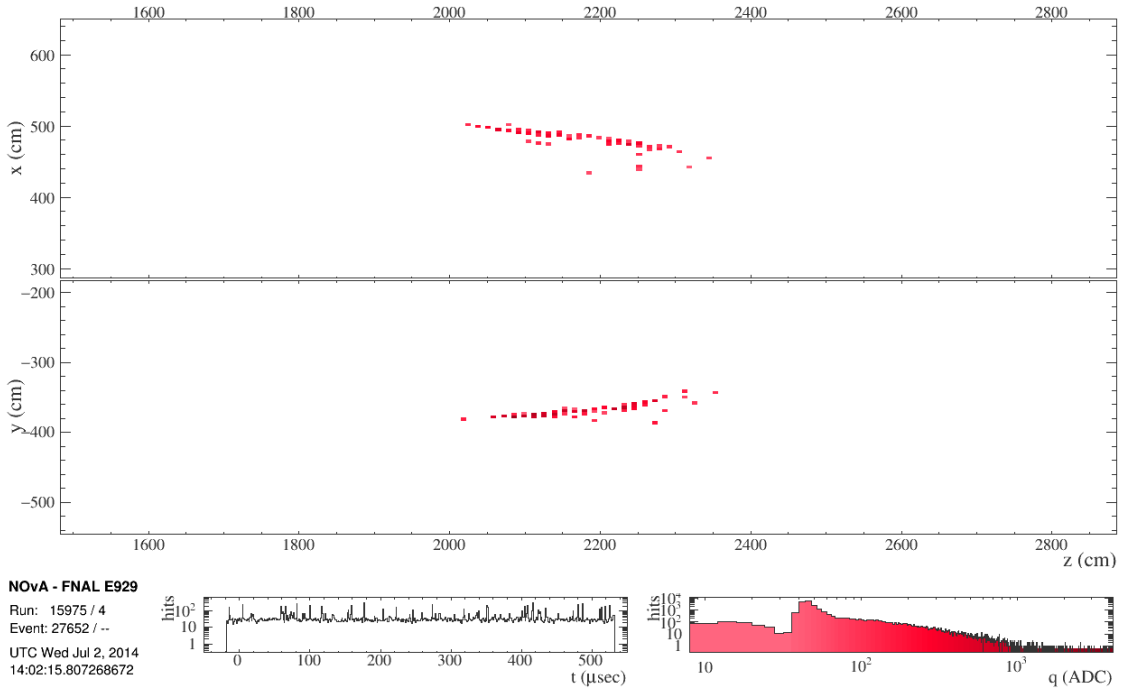

Figure C.5: Selected ν_e core event.

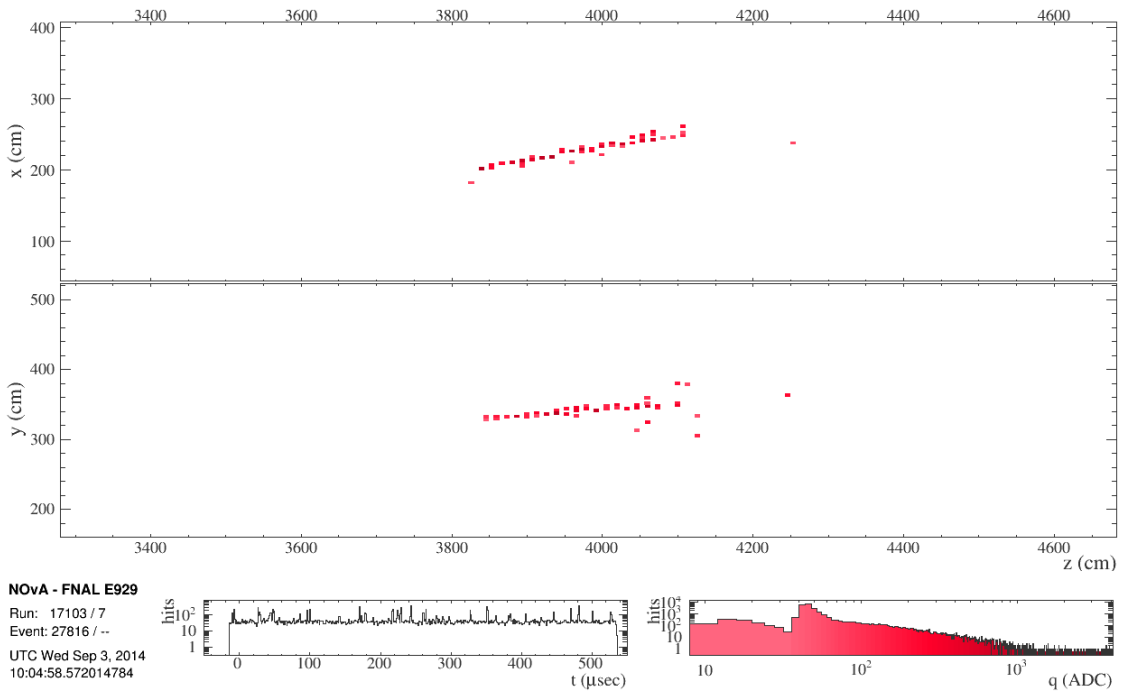

Figure C.6: Selected ν_e core event.

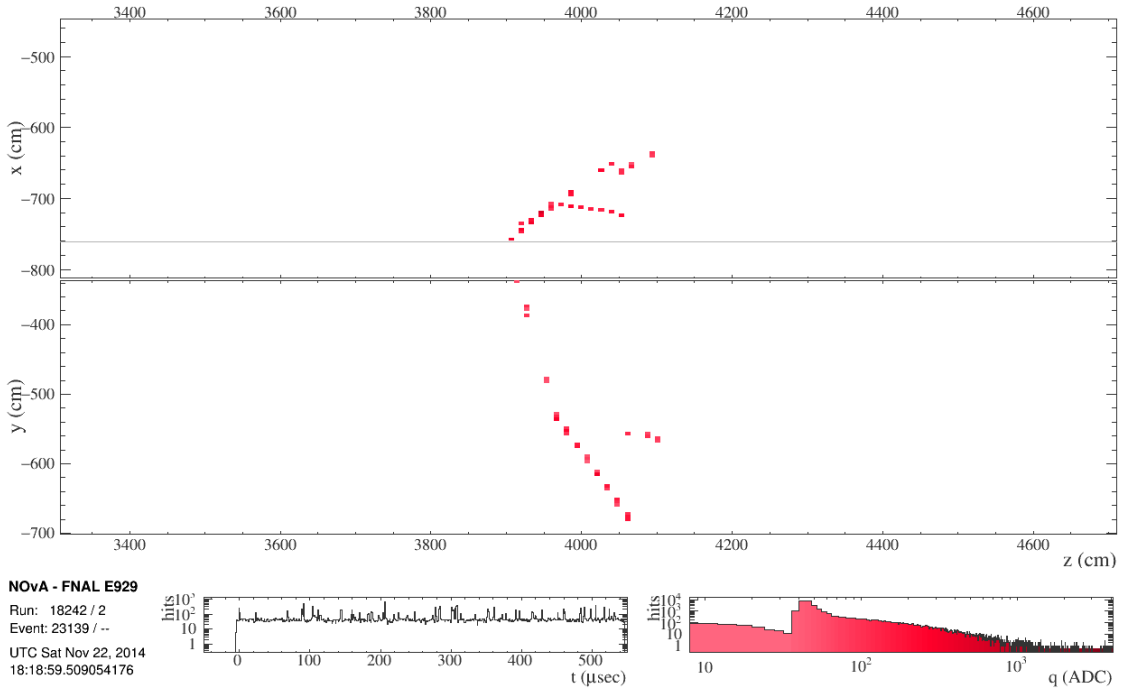

Figure C.7: Selected ν_e core event.

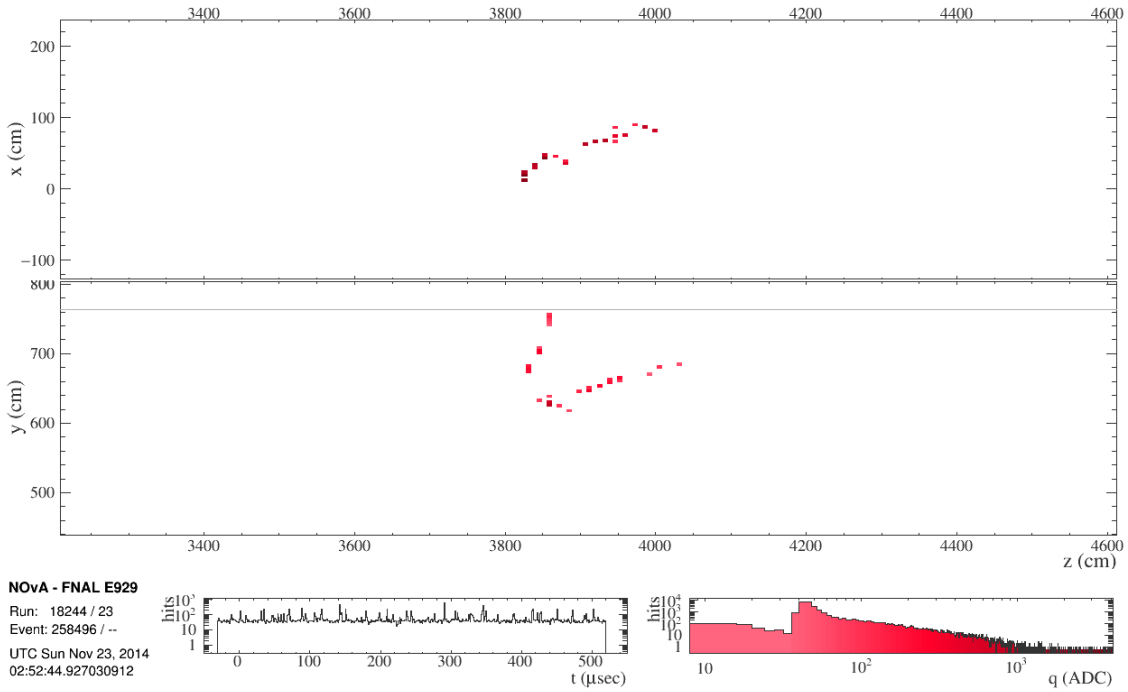

Figure C.8: Selected ν_e core event.

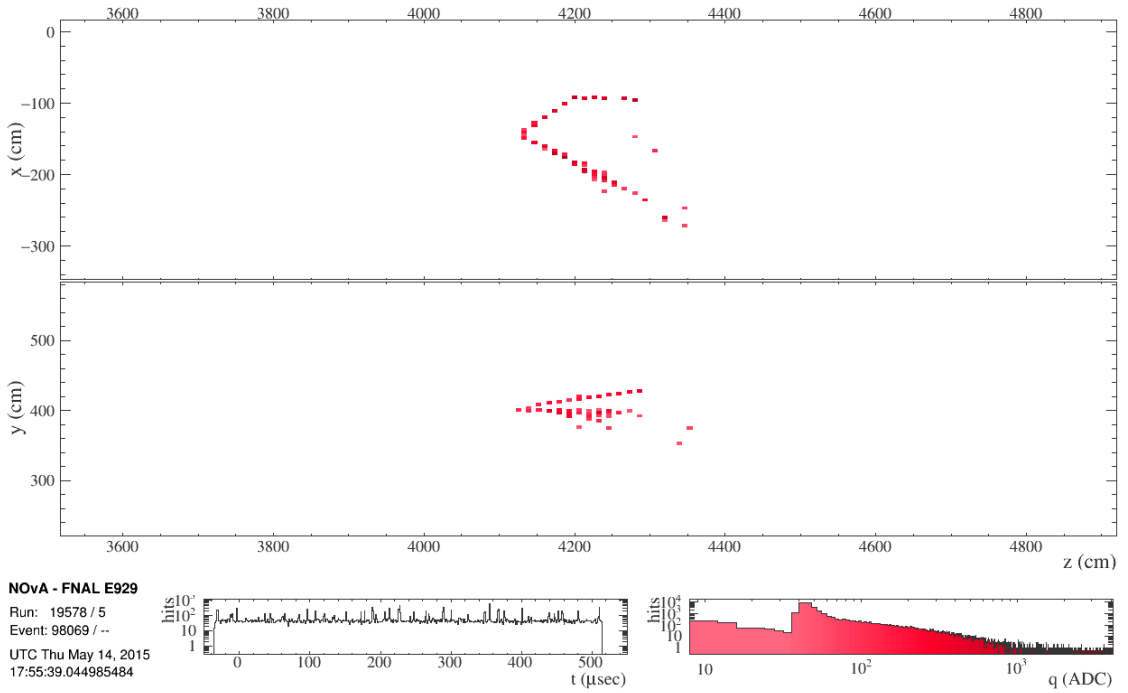

Figure C.9: Selected ν_e core event.

C.2 ν_e Peripheral Sample Selected Events

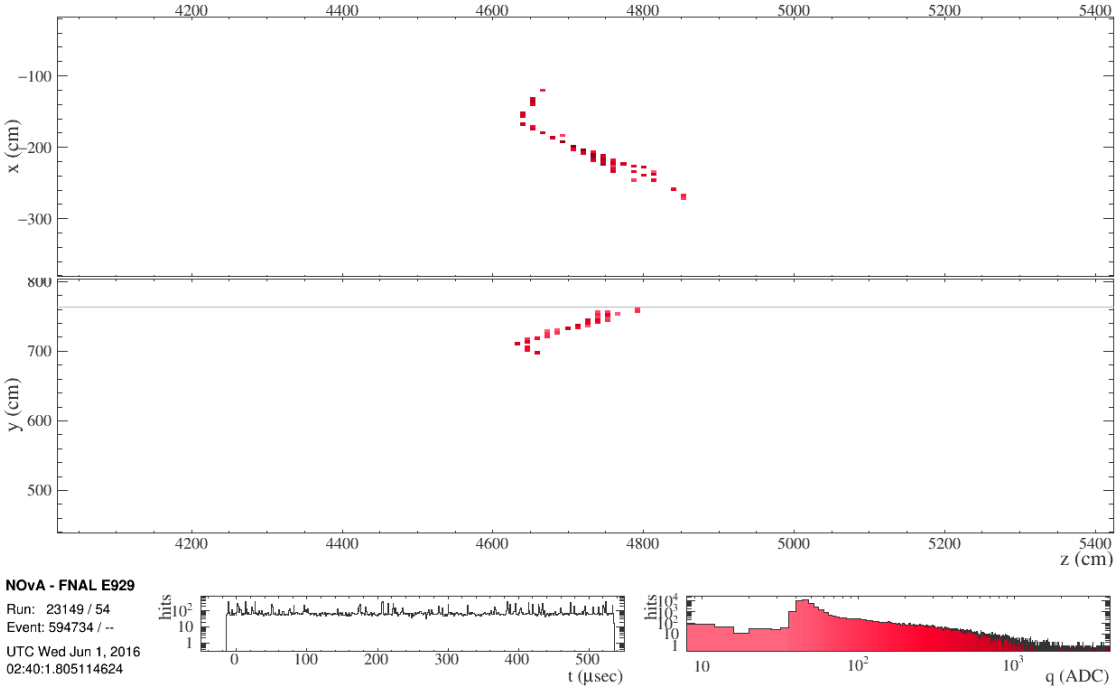

Figure C.10: Selected ν_e peripheral event.

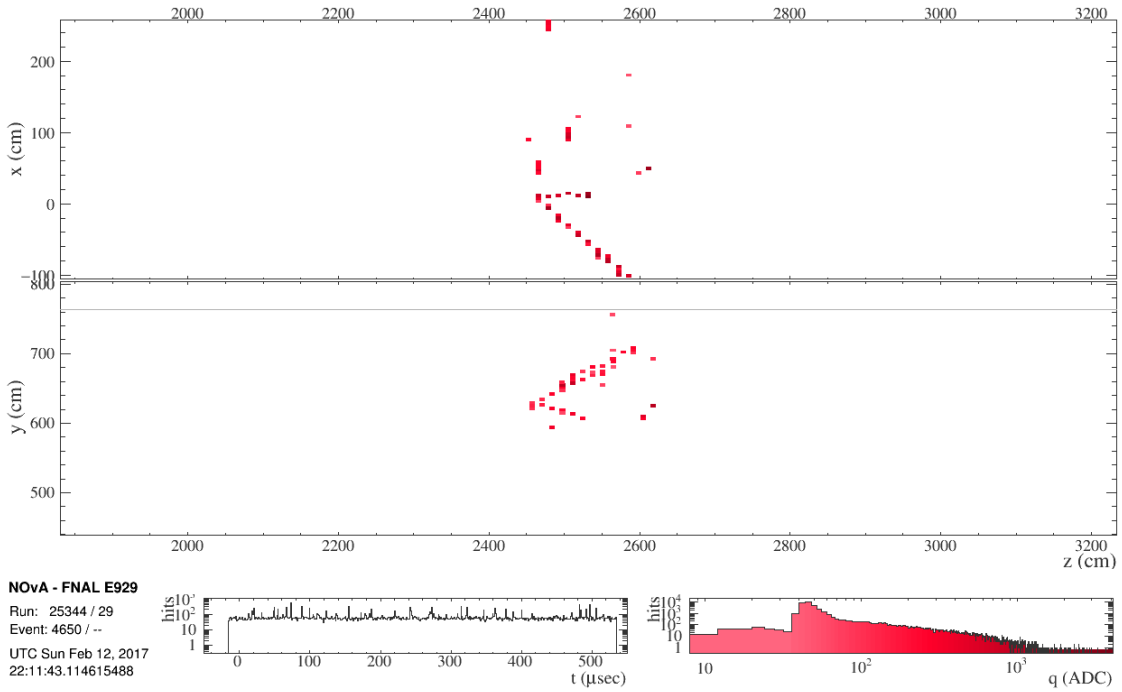

Figure C.11: Selected ν_e peripheral event.

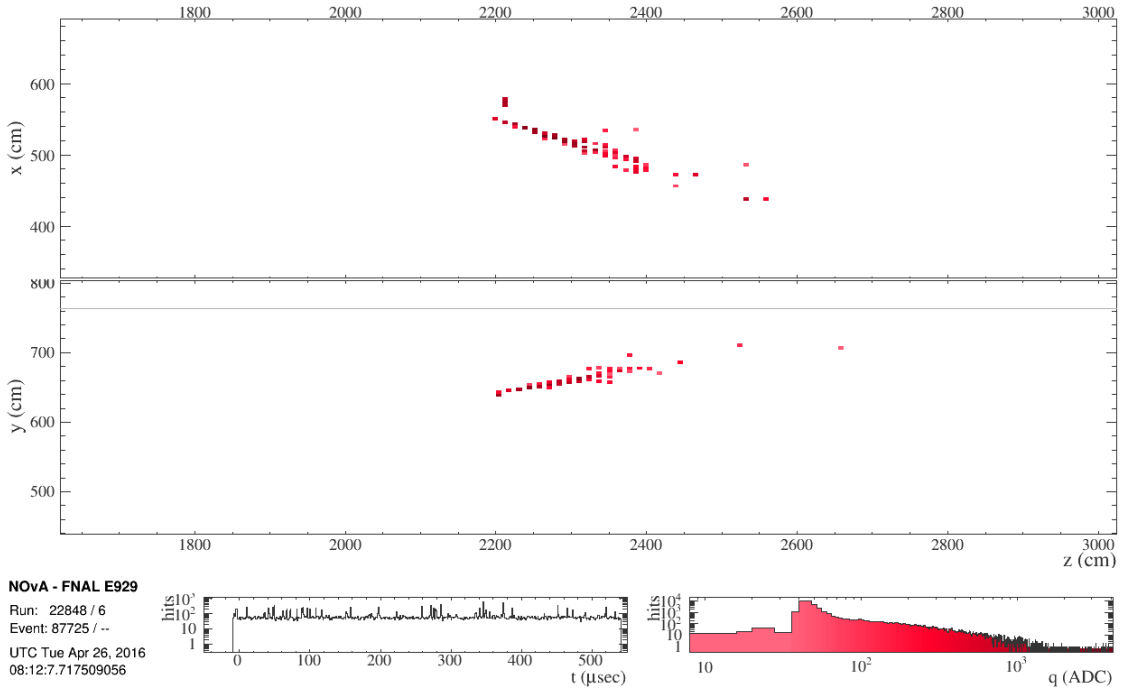

Figure C.12: Selected ν_e peripheral event.

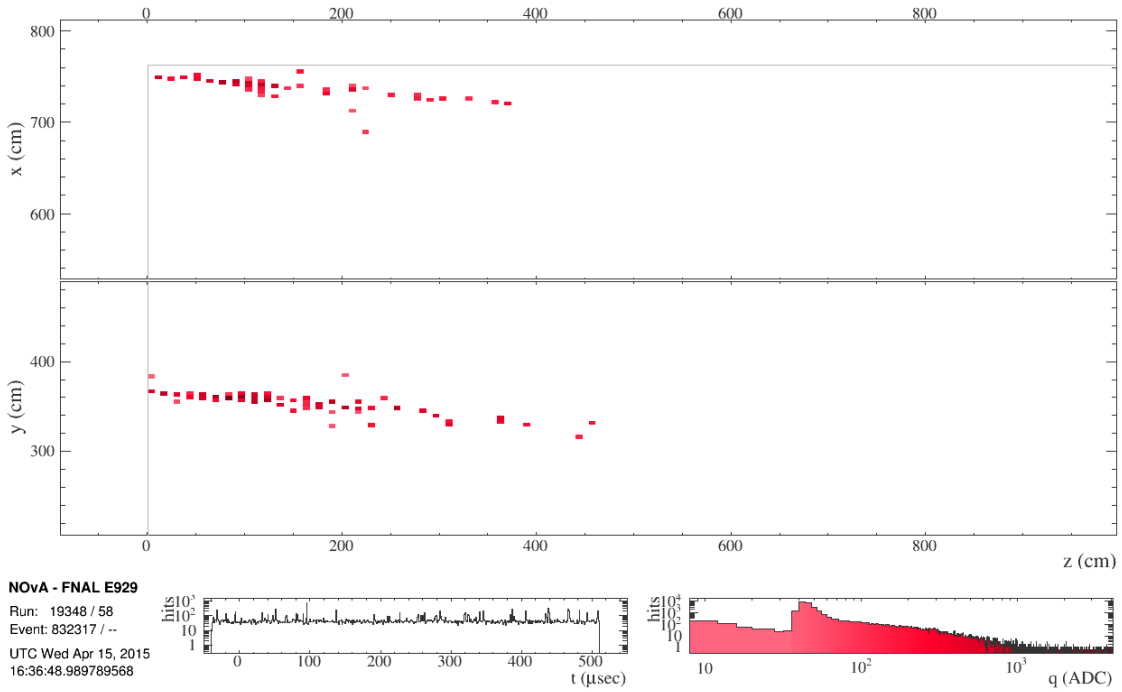

Figure C.13: Selected ν_e peripheral event.

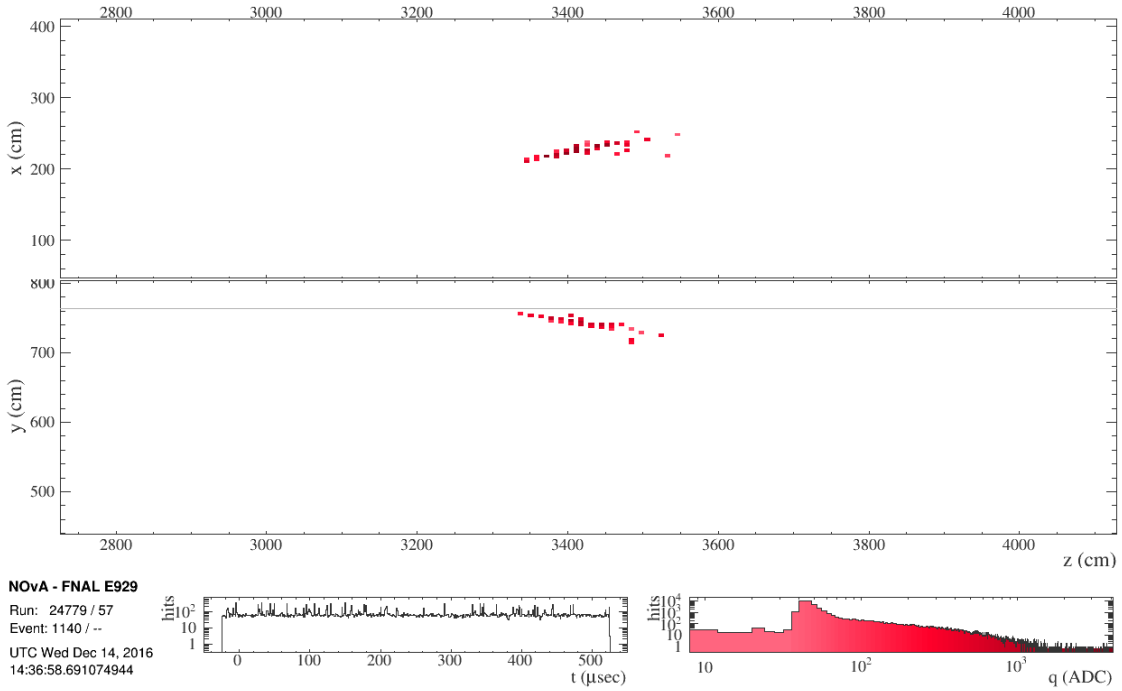

Figure C.14: Selected ν_e peripheral event.

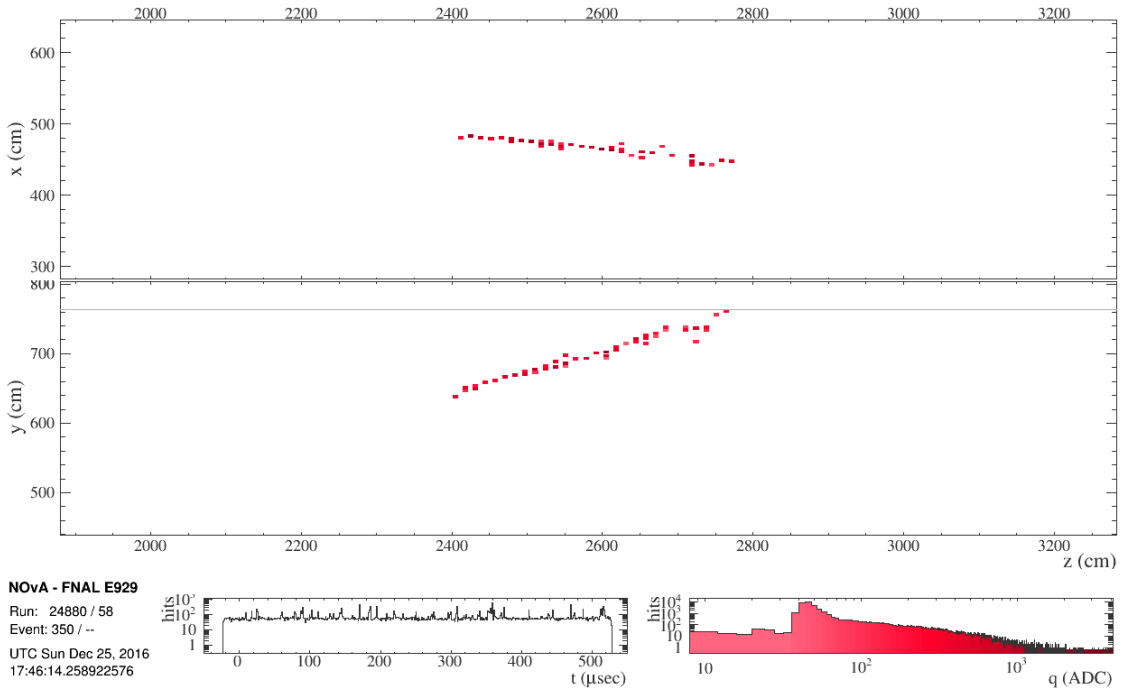

Figure C.15: Selected ν_e peripheral event.

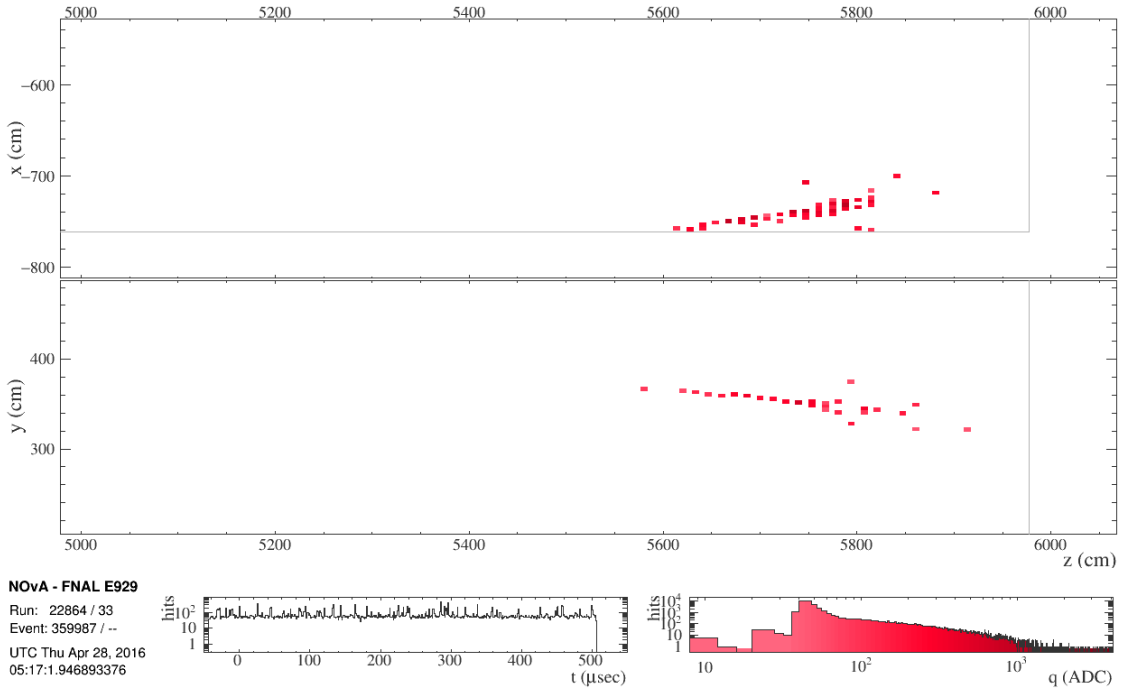

Figure C.16: Selected ν_e peripheral event.

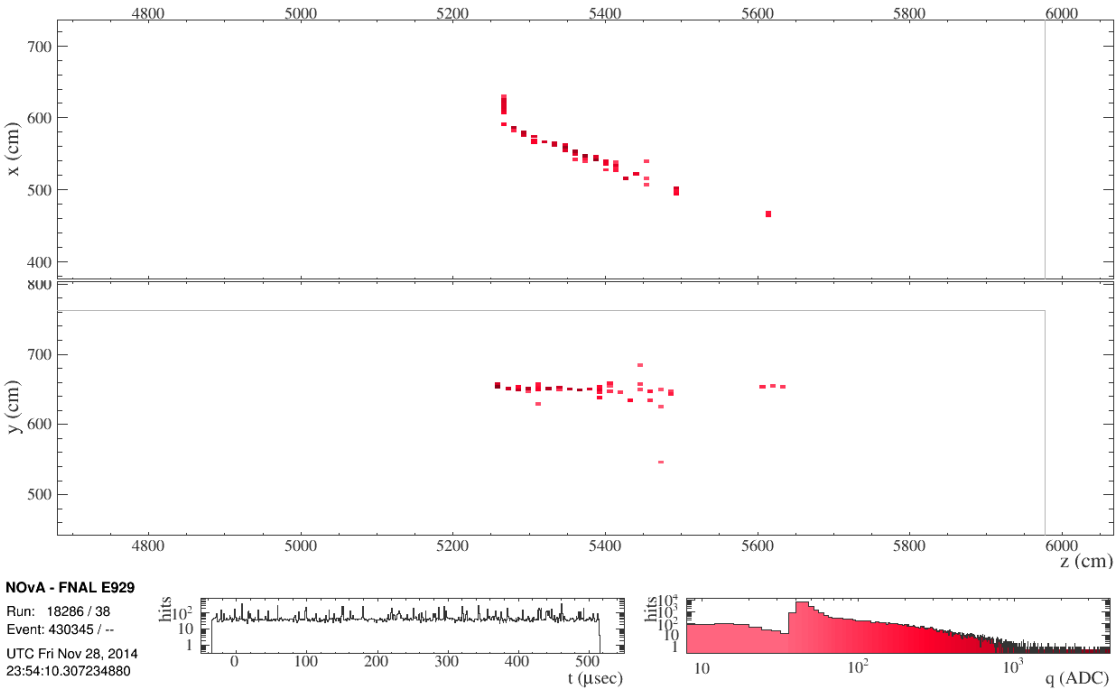

Figure C.17: Selected ν_e peripheral event.

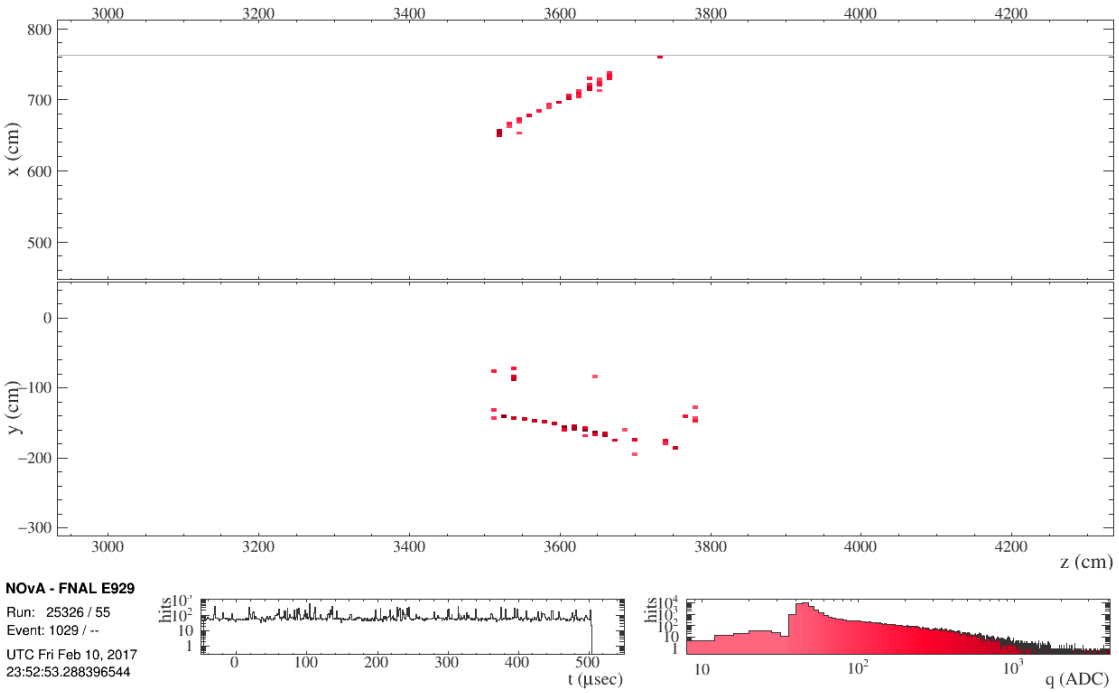

Figure C.18: Selected ν_e peripheral event.

C.3 ν_μ Selected Events

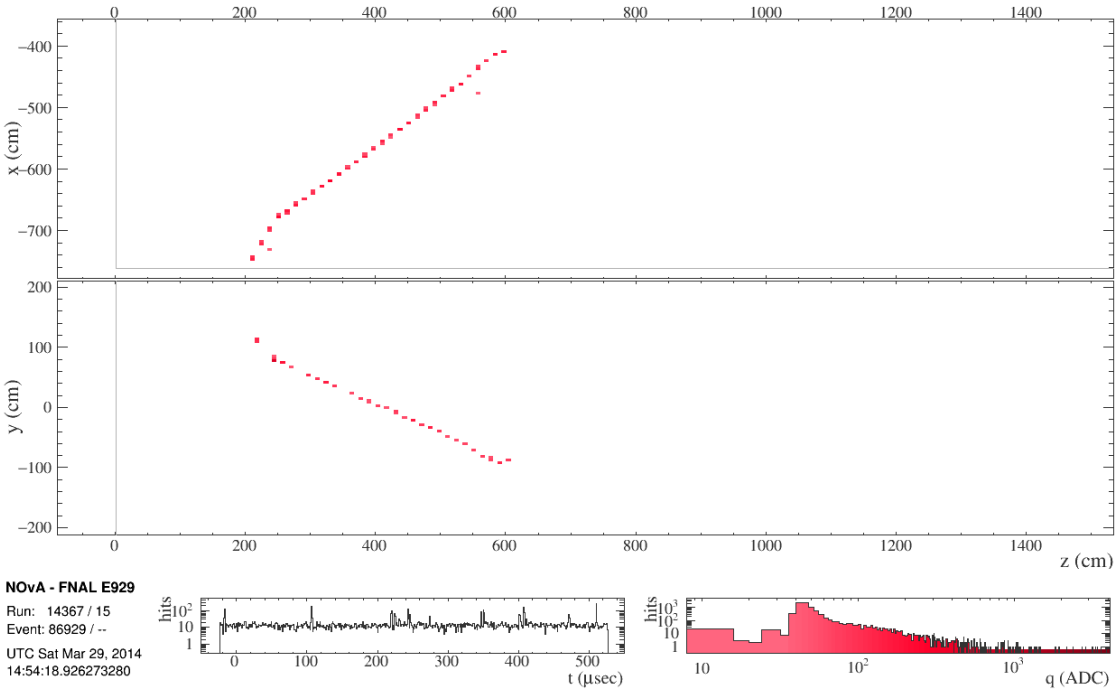

Figure C.19: Selected ν_μ event.

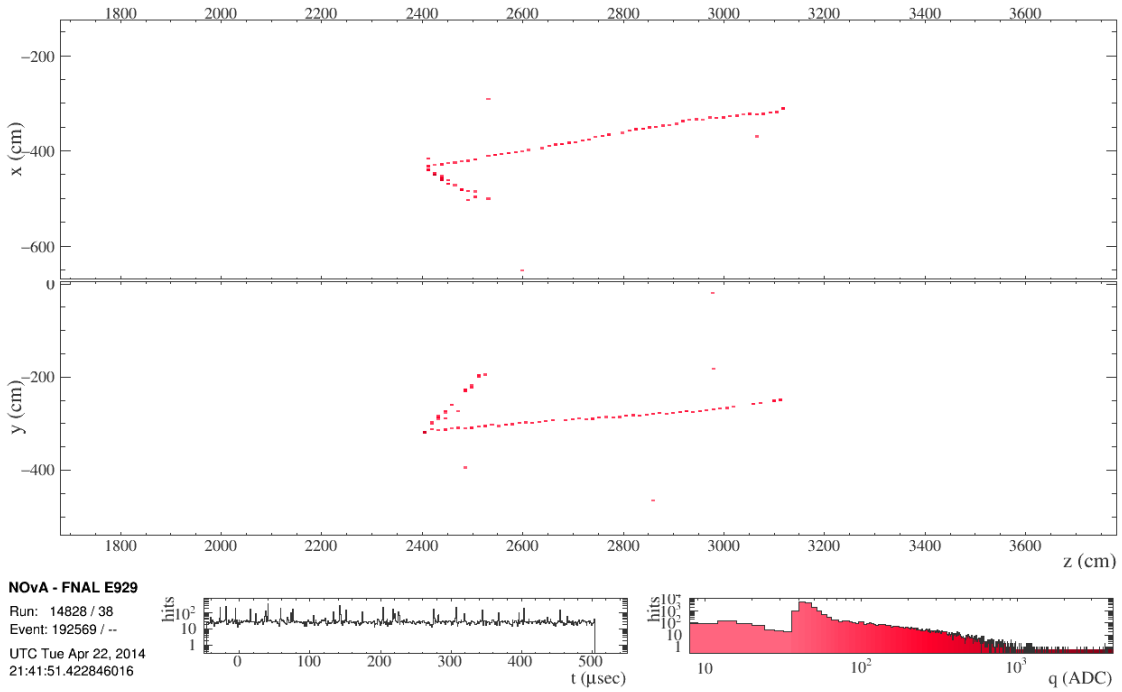

Figure C.20: Selected ν_μ event.

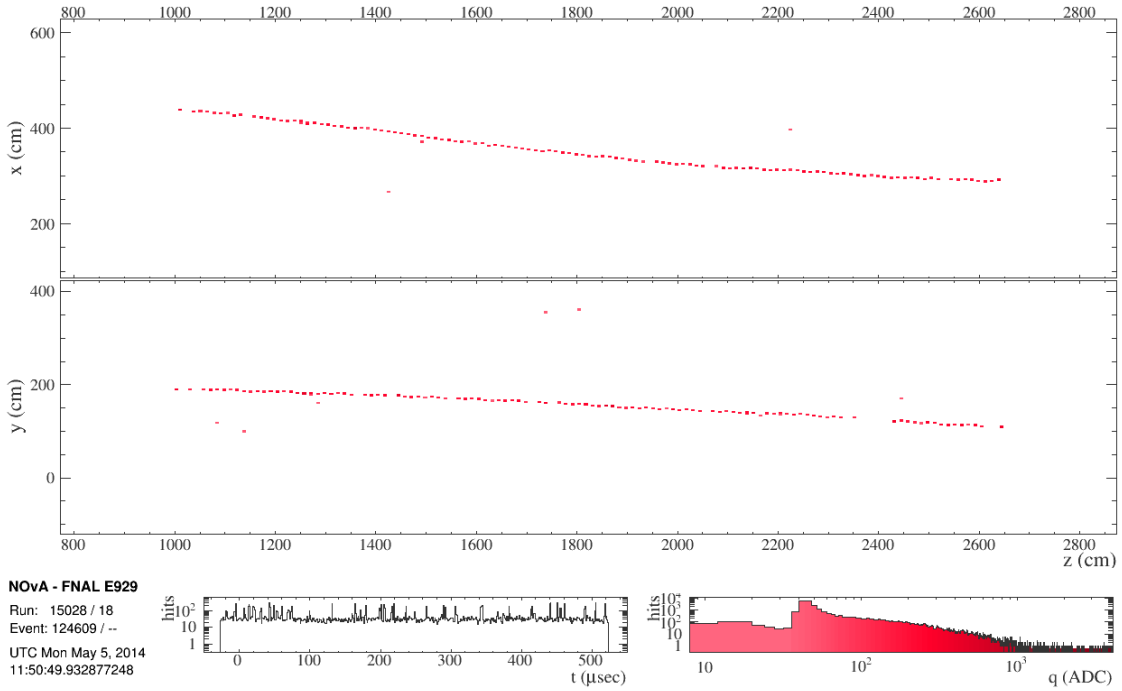

Figure C.21: Selected ν_μ event.

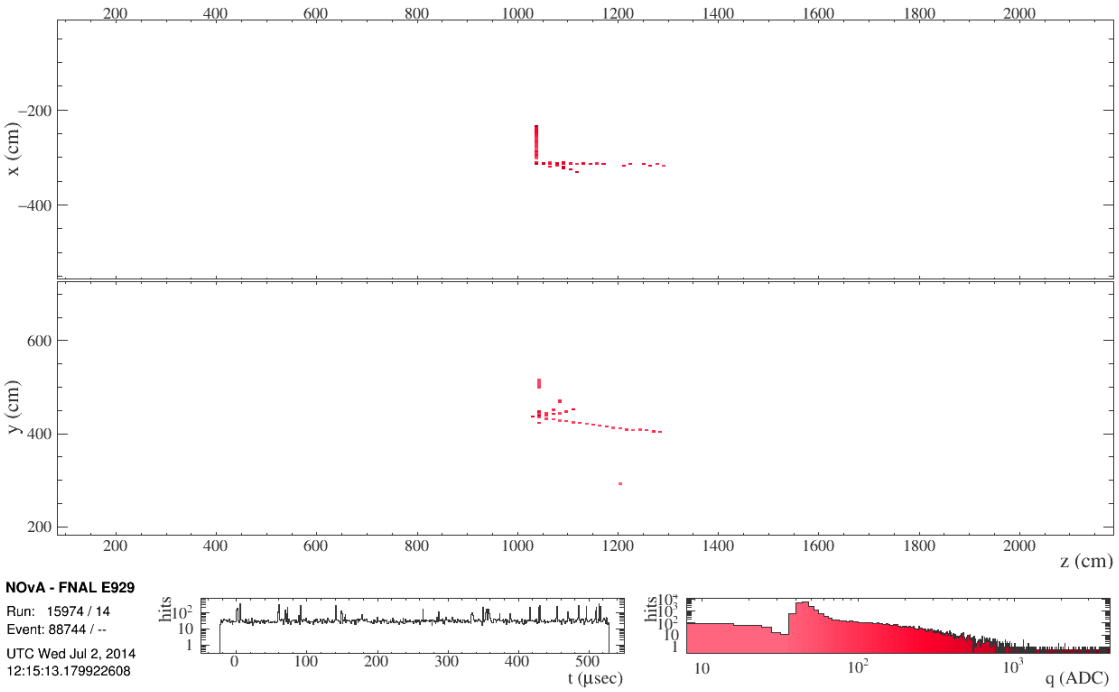

Figure C.22: Selected ν_μ event.

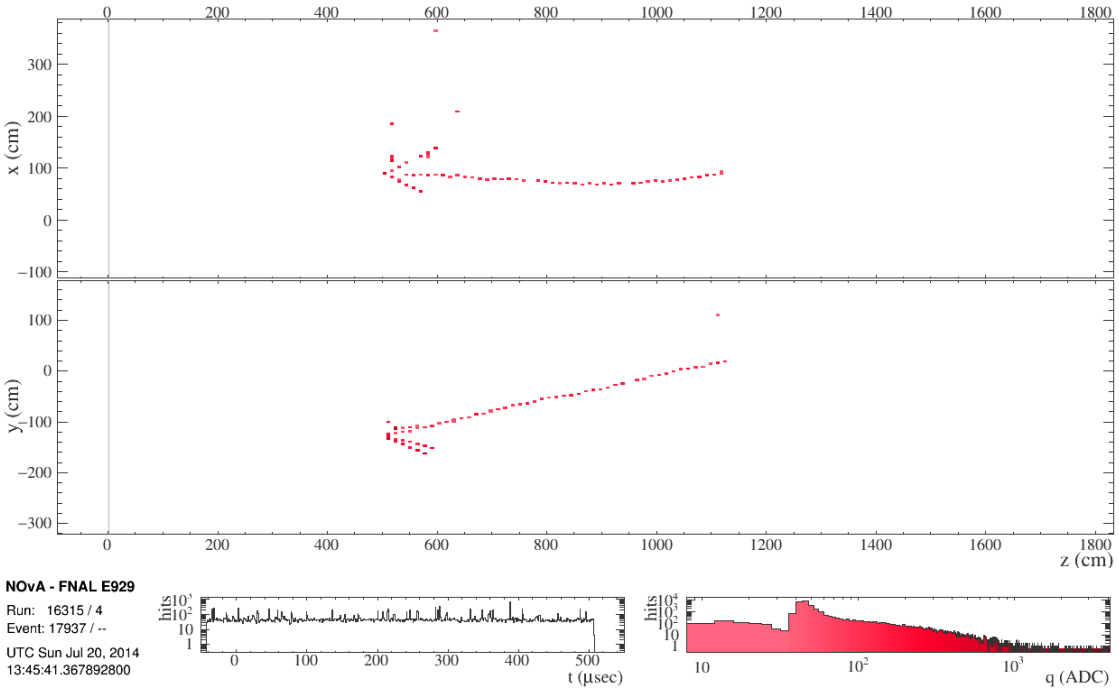

Figure C.23: Selected ν_μ event.

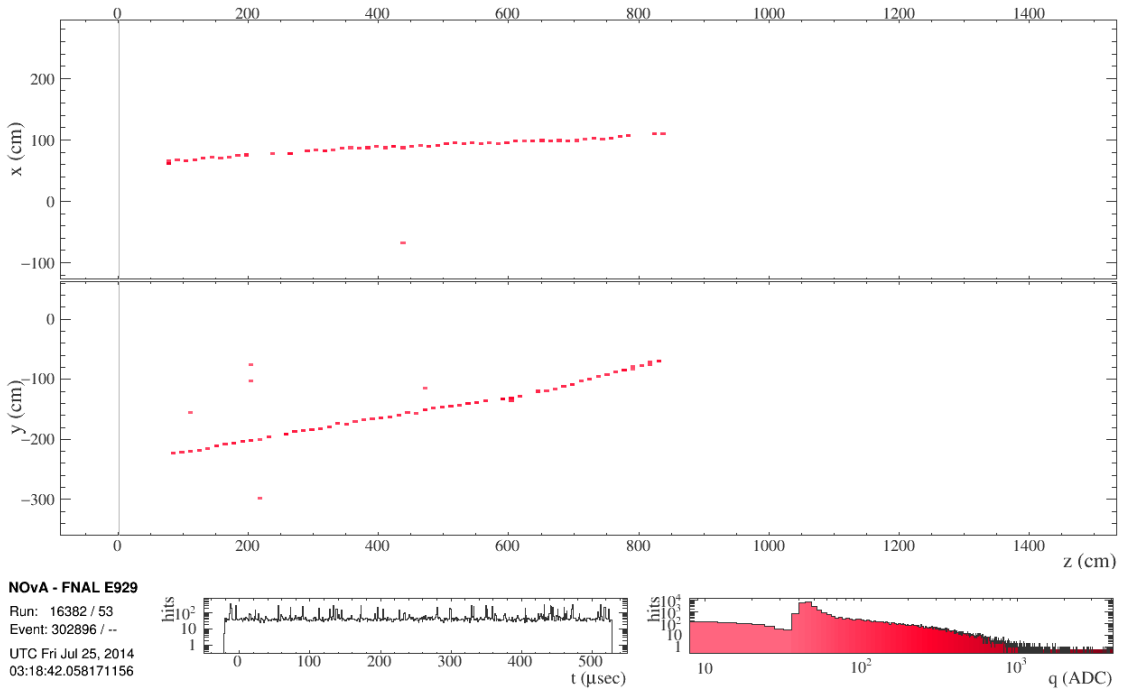

Figure C.24: Selected ν_μ event.

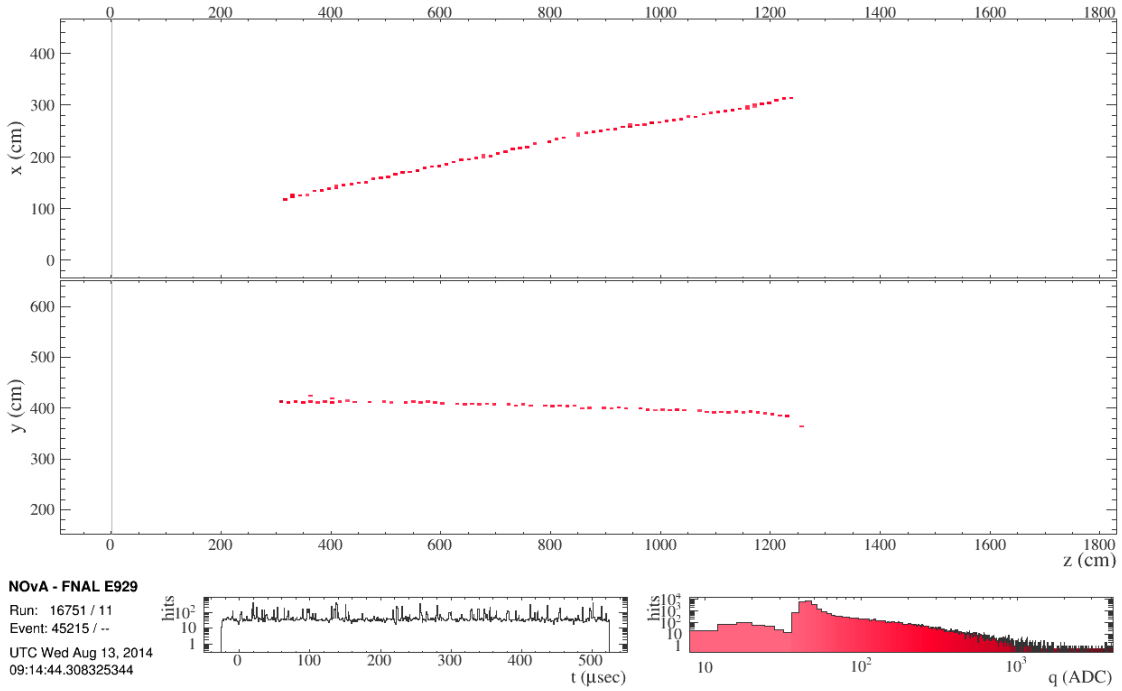

Figure C.25: Selected ν_μ event.

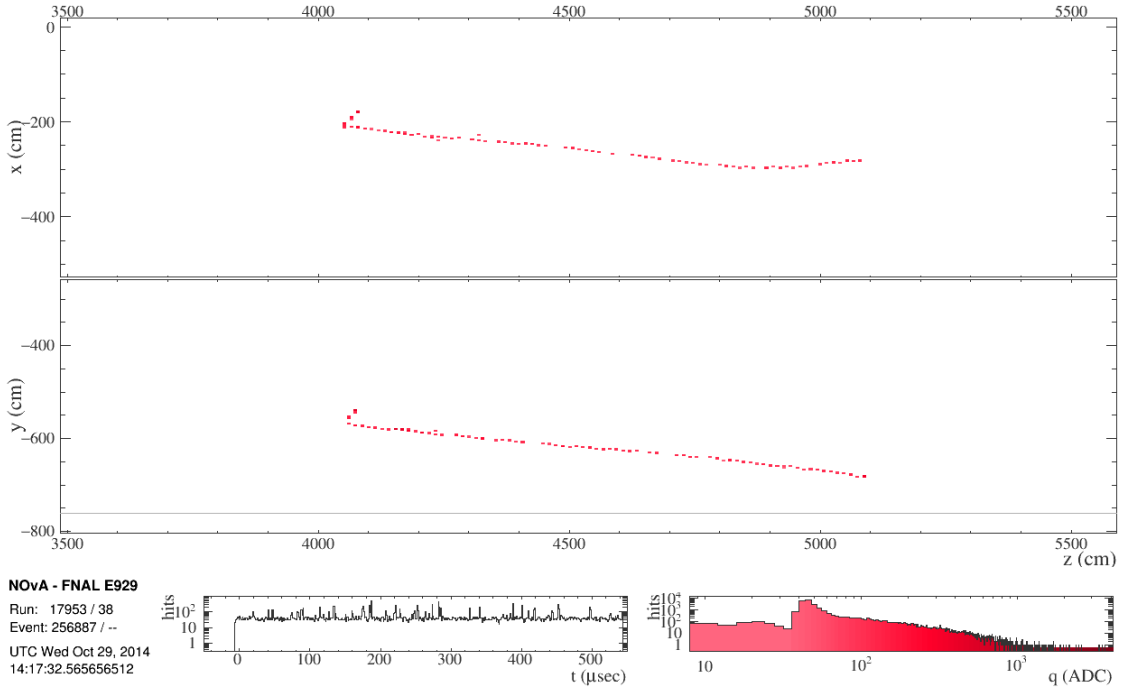

Figure C.26: Selected ν_μ event.

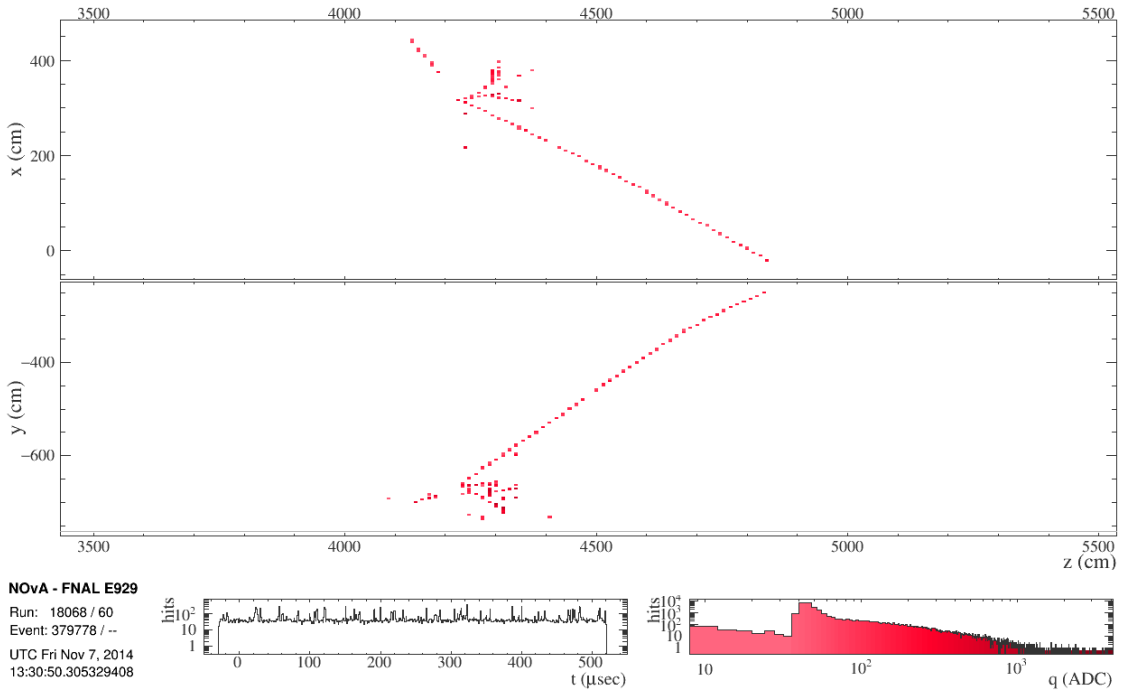

Figure C.27: Selected ν_μ event.

BIBLIOGRAPHY

[1] D.S. Harmer R. Davis and K.C. Hoffman. Search for neutrinos from the sun. *Phys.Rev.Lett*, 20:1205–1209, 1968.

[2] W. Heisenberg. Physics and Philosophy. The Revolution in Modern Science. 1958.

[3] K.S. Hirata et al. Observation of ^8B solar neutrinos in the Kamiokande-II detector. *Phys.Rev.Lett.*, 63:16, 1989.

[4] A.I. Abazov et al. Search for neutrinos from sun using the reaction Ga-71 (electron-neutrino e^-) Ge-71. *Phys.Rev.Lett.*, 67:3332–3335, 1991.

[5] P. Anselmann et al. Solar neutrinos observed by GALLEX at Gran Sasso. *Phys.Lett.*, B285:376–389, 1992.

[6] B. Aharmim et al. Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory. *Phys.Rev.*, C72:055502, 2005.

[7] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[8] A. Bellerive, J. R. Klein, A. B. McDonald, A. J. Noble, and A. W. P. Poon. The Sudbury Neutrino Observatory. *Nucl. Phys.*, B908:30–51, 2016.

[9] Giorgio Giacomelli. Neutrino physics and astrophysics with the MACRO experiment at the Gran Sasso lab. *Braz. J. Phys.*, 33:211–217, 2003.

[10] M. G. Aartsen et al. The IceCube Neutrino Observatory: Instrumentation and Online Systems. *JINST*, 12(03):P03012, 2017.

[11] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.

[12] J. K. Ahn et al. RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang. 2010.

[13] Y. Abe et al. Reactor electron antineutrino disappearance in the Double Chooz experiment. *Phys. Rev.*, D86:052008, 2012.

[14] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam. *Phys. Rev.*, D64:112007, 2001.

[15] Heather L. Ray. The MiniBooNE Experiment. 2007.

[16] A. Suzuki et al. Design, construction, and operation of SciFi tracking detector for K2K experiment. *Nucl. Instrum. Meth.*, A453:165–176, 2000.

[17] P. Adamson et al. Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS. *Phys. Rev. Lett.*, 110(25):251801, 2013.

[18] N. Agafonova et al. Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment. *Phys. Rev. Lett.*, 115(12):121802, 2015.

[19] K. Abe et al. The T2K Experiment. *Nucl. Instrum. Meth.*, A659:106–135, 2011.

[20] D. S. Ayres et al. The NOvA Technical Design Report. 2007.

[21] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2016.

[22] K. Abe et al. Letter of Intent: The Hyper-Kamiokande Experiment — Detector Design and Physics Potential —. 2011.

[23] Donald A. Glaser and David C. Rahm. Characteristics of bubble chambers. *Phys. Rev.*, 97:474–479, Jan 1955.

[24] W. Pauli. Dear radioactive ladies and gentlemen. *Phys. Today*, 31N9:27, 1978.

[25] J. Chadwick. Possible Existence of a Neutron. *Nature*, 129:312, 1932.

[26] C.L Cowan and F. Reines. Detection of the Free neutrino. *Phys. Rev.*, 92:830–831, 1953.

[27] H.W. Kruse A.D. McGuire F. Reines, F.B. Harrison and C.L. Cowan. Detection of the Free neutrino: a Confirmation. *Science*, 124:103–104, 1956.

[28] Georges Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett.*, B716:1–29, 2012.

[29] Eva Halkiadakis, George Redlinger, and David Shih. Status and Implications of Beyond-the-Standard-Model Searches at the LHC. *Ann. Rev. Nucl. Part. Sci.*, 64:319–342, 2014.

[30] Hitoshi Murayama. The origin of neutrino mass. *Physics World*, 15(5):35, 2002.

[31] Andre de Gouvea. Neutrino mass models. *Annual Review of Nuclear and Particle Science*, 66(1):197–217, 2016.

[32] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[33] Boris Kayser. Are neutrinos their own antiparticles? *J. Phys. Conf. Ser.*, 173:012013, 2009.

[34] Stefano Dell'Oro, Simone Marcocci, Matteo Viel, and Francesco Vissani. Neutrinoless double beta decay: 2015 review, 2016.

[35] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172, 1958.

[36] M. Nakagawa Z. Maki and S. Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870, 1962.

[37] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Sov. Phys. JETP*, 26:984–988, 1968. [Zh. Eksp. Teor. Fiz.53,1717(1967)].

[38] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys.Rev.*, D17:2369–2374, 1978.

[39] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics.* Oxford University Press, Oxford, UK, 2007.

[40] L. Kolupaeva, O. Samoylov, and I. Shandrov. Matter effect in neutrino oscillations for long-baseline experiments. *Phys. Part. Nucl. Lett.*, 14(7):975–980, 2017.

[41] Francesco Capozzi, Eleonora Di Valentino, Eligio Lisi, Antonio Marrone, Alessandro Melchiorri, and Antonio Palazzo. Global constraints on absolute neutrino masses and their ordering. *Phys. Rev.*, D95(9):096014, 2017.

[42] K. N. Abazajian et al. Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure. *Astropart. Phys.*, 63:66–80, 2015.

[43] K. A. Olive et al. Review of Particle Physics. *Chin. Phys.*, C38:090001, 2014.

[44] Junpei Shirai. Results and future plans for the KamLAND-Zen experiment. *J. Phys. Conf. Ser.*, 888(1):012031, 2017.

[45] Jonathan Engel and Javier Men√©ndez. Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review. *Rept. Prog. Phys.*, 80(4):046301, 2017.

[46] Stefano Dell'Oro, Simone Marcocci, Matteo Viel, and Francesco Vissani. Neutrinoless double beta decay: 2015 review. *Adv. High Energy Phys.*, 2016:2162659, 2016.

[47] Eleonora Di Valentino, Alessandro Melchiorri, and Joseph Silk. Reconciling Planck with the local value of H_0 in extended parameter space. *Phys. Lett.*, B761:242–246, 2016.

[48] P. MACHADO. Recent highlights from neutrino theory. *Presented at the APS Division of Particles and Fields DPF 2017 conference.*, 2017.

[49] A. D. Sakharov. Violation of CP Invariance, c Asymmetry, and Baryon Asymmetry of the Universe. *Pisma Zh. Eksp. Teor. Fiz.*, 5:32–35, 1967. [Usp. Fiz. Nauk161,61(1991)].

[50] J. A. Formaggio and G. P. Zeller. From eV to EeV: Neutrino Cross Sections Across Energy Scales. *Rev. Mod. Phys.*, 84:1307–1341, 2012.

[51] Super-Kamiokande. Online gallery. http://www-sk.icrr.u-tokyo.ac.jp/sk/index-e.html. [Accessed November-2017].

[52] MINOS. Online gallery. http://www-numi.fnal.gov/PublicInfo/PhotoGalleries/index.html. [Accessed November-2017].

[53] P. Adamson et al. The NuMI Neutrino Beam. *Nucl. Instrum. Meth.*, A806:279–306, 2016.

[54] Kuldeep K. Maan, H. Duyang, Sanjib R. Mishra, and Vipin Bhatnagar. *Neutrino Flux at NOvA: Empirical Constraint and Sytematics Errors*, pages 311–315. Springer International Publishing, Cham, 2016.

[55] Alex Himmel. The NuMI beam simulation with Flugg. MINOS internal docdb-6316.

[56] G. Battistoni et al. The fluka code: description and benchmarking. *AIP Conf.Proc.*, 896:31–49, 2007.

[57] Fermilab. Online gallery. http://news.fnal.gov/wp-content/uploads/2015/10/NOvA$_e vent_i llustration - hr.jpg. [Accessed January - 2018]$.

[58] Pietro Antonioli et al. SNEWS: the SuperNova Early Warning System. *New Journal of Physics*, 6(1):114, 2004.

[59] Justin A. Vasel, Andrey Sheshukov, and Alec Habig. Observing the Next Galactic Supernova with the NOvA Detectors. In *Meeting of the APS Division of Particles and Fields (DPF 2017) Batavia, Illinois, USA, July 31-August 4, 2017*, 2017.

[60] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth.*, A389:81–86, 1997.

[61] Biswaranjan Behera, Gavin Davies, and Fernanda Psihas. Event Reconstruction in the NOvA Experiment. In *Meeting of the APS Division of Particles and Fields (DPF 2017) Batavia, Illinois, USA, July 31-August 4, 2017*, 2017.

[62] Michael David Baird. *An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter.* PhD thesis, Indiana U., 2015.

[63] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters a density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, pages 226–231. AAAI Press, 1996.

[64] Mark Messier. Vertex reconstruction based on elastic arms. NOvA internal docdb-7530.

[65] M. Baird, J. Bian, M. Messier, E. Niner, D. Rocco, and K. Sachdev. Event Reconstruction Techniques in NOvA. *J. Phys. Conf. Ser.*, 664(7):072035, 2015.

[66] Leandro A.F. Fernandes and Manuel M. Oliveira. Real-time line detection through an improved hough transform voting scheme. *Pattern Recognition*, 41(1):299 – 314, 2008.

[67] Evan David Niner. *Observation of Electron Neutrino Appearance in the NuMI Beam with the NOvA Experiment.* PhD thesis, Indiana U., 2015.

[68] J.C. Dunn. A fuzzy relative of the isodata process and its use in detecting compact, well-separated clusters. *J. Cybern.*, 3:32–57, 1974.

[69] J.C. Bezdek. *Pattern Recognition with Fuzzy Objective Function Algorithms.* Plenum Press, New York, 1981.

[70] R. Krishnapuram and J.M. Keller. A possibilistic approach to clustering. *IEEE Trans. Fuzzy Syst.*, 1:98–110, 1993.

[71] N.H. Kuiper. Tests concerning random points on a circle. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen, Series A*, 36:38–47, 1960.

[72] Michael Baird. Tech Note for the Multi-Hough Transform. NOvA internal docdb-8241.

[73] Silvia Fernanda Psihas Olmedo. Muon Energy Reconstruction Through the Multiple Scattering Method in the NOνA Detectors. Master's thesis, Minnesota U., Duluth, 2015.

[74] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010.

[75] Evan Racah, Seyoon Ko, Peter Sadowski, Wahid Bhimji, Craig Tull, Sang-Yun Oh, Pierre Baldi, and Prabhat. Revealing Fundamental Physics from the Daya Bay Neutrino Experiment using Deep Neural Networks. 2016.

[76] A. Schwartzman, M. Kagan, L. Mackey, B. Nachman, and L. De Oliveira. Image Processing, Computer Vision, and Deep Learning: new approaches to the analysis and physics interpretation of LHC events. *J. Phys. Conf. Ser.*, 762(1):012035, 2016.

[77] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going Deeper with Convolutions. 2014.

[78] R. Acciarri et al. Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber. *JINST*, 12(03):P03011, 2017.

[79] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv:1408.5093*, 2014.

[80] P. Adamson et al. Constraints on Oscillation Parameters from ν_e Appearance and ν_μ Disappearance in NOvA. *Phys. Rev. Lett.*, 118(23):231801, 2017.

[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. *JINST*, 11(09):P09001, 2016.

[82] Dominick Rosario Rocco. *Muon Neutrino Disappearance in NOvA with a Deep Convolutional Neural Network Classifier*. PhD thesis, Minnesota U., 2016.

[83] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

[84] NOvA. G4numi and the numi beam. https://cdcvs.fnal.gov/redmine/projects/numi-beam-sim/wiki/. [Wiki; Accessed January-2018].

[85] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl.Instrum.Meth.*, A506:250–303, 2003.

[86] Leonidas Aliaga Soplin. *Neutrino Flux Prediction for the NuMI Beamline*. PhD thesis, William-Mary Coll., 2016.

[87] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, et al. The GENIE Neutrino Monte Carlo Generator. *Nucl.Instrum.Meth.*, A614:87–104, 2010.

[88] A. Aurisano, C. Backhouse, R. Hatcher, N. Mayer, J. Musser, R. Patterson, R. Schroeter, and A. Sousa. The NOvA simulation chain, 2015.

[89] A. Aurisano. NOvA internal document, docDB 13577.

[90] R. E. Kalman. A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 1960.

[91] Susan M. Lein. *Muon Neutrino Contained Disappearance in NOvA*. PhD thesis, 2015.

[92] Luke Vinton. *Measurement of Muon Neutrino Disappearance with the NOvA Experiment*. PhD thesis, Sussex U., 2018.

[93] Erika Catano Mur. BEN Decomposition Tech Note. NOvA internal docdb-23174.

[94] Ryan Nichol and Bruno Zamorano. Executive summary of calibration for the thrid analysis. NOvA internal docdb-23558.

[95] C.N. Chou. The Nature of the Saturation Effect of Fluorescent Scintillators. *Phys.Rev.*, 87:904, 1952.

[96] Adam Aurisano. Tech Note: 2017 Light Model. NOvA internal docdb-23228.

[97] A. Kaboth. Systematic Uncertainties at the T2K Experiment for Oscillation Measurements. 2013.

[98] Tyler Alion. NOvA internal document, Third Analysis: Absolute Calibration Technote, docDB 23372.

[99] Tyler Alion. Calorimetric Energy Scale in the NOvA Detectors. https://indico.fnal.gov/event/13745/session/6/contribution/67.

[100] Alex Radovic. Latest Oscillation Results from NOvA. http://theory.fnal.gov/jetp/.

[101] K. Abe et al. Combined Analysis of Neutrino and Antineutrino Oscillations at T2K. *Phys. Rev. Lett.*, 118(15):151801, 2017.

[102] M. G. Aartsen et al. Measurement of Atmospheric Neutrino Oscillations at 6–56 GeV with IceCube DeepCore. *Phys. Rev. Lett.*, 120(7):071801, 2018.

[103] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev.*, D57:3873–3889, 1998.

[104] P. Adamson et al. Search for active-sterile neutrino mixing using neutral-current interactions in NOvA. *Phys. Rev.*, D96(7):072006, 2017.

[105] J. Renner et al. Background rejection in NEXT using deep neural networks. *JINST*, 12(01):T01004, 2017.

[106] Chang-Wei Loh, Zhi-Qiang Qian, You-Hang Liu, De-Wen Cao, Rui Zhang, Wei Wang, Hai-Bo Yang, and Ming Qi. Deep Learning the Effects of Photon Sensors on the Event Reconstruction Performance in an Antineutrino Detector. 2017.

[107] Hao Qiao, Chunyu Lu, Xun Chen, Ke Han, Xiangdong Ji, and Siguang Wang. Signal-background discrimination with convolutional neural networks in the PandaX-III experiment. 2018.

LIST OF FIGURES