

SEARCHING FOR NEUTRINO TRIDENTS IN THE NOVA NEAR DETECTOR

Reed Bowles

Submitted to the faculty of the Graduate School
in partial fulfillment of the requirements
for the degree
Doctor of Philosophy
in the Department of Physics,
Indiana University
October 2025

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

Doctoral Committee

Mark Messier, PhD

Rex Tayloe, PhD

Emilie Passemar, PhD

John Beggs, PhD

Date of Defense: 08/15/2025

Copyright © 2025

Reed Bowles

Dedicated to my wife, Taylor, who has provided unending grace, love, and patience throughout the many years required to complete this work. I could not have reached this point without you, and there are not enough words to express my gratitude for all your support.

ACKNOWLEDGEMENTS

I would like to thank the many people who have supported and encouraged me throughout my academic journey. First of all, thank you to my parents, Randy and Renee, my step-mom, Ronda, and my grandpa, Roger. You all have always encouraged me to explore my curiosity and never stop learning. Thank you for teaching me the importance of hard work, and the inherent value in pursuing the things that bring you joy in life. Your love and support have been instrumental in allowing me to achieve my dreams. Thank you to my wife, Taylor, who traveled with me both literally across the country, and metaphorically through this journey of graduate school. I cannot imagine having gone through it without you.

I owe the majority of my academic trajectory to my high school physics teacher, Stanley Bergkamp. Before taking your class I had no idea what physics was, but your passion and skill in communicating the subject ignited a fire in me that I am truly unable to describe. You have imparted a love for understanding the world onto me and thousands of other students - a feat I can only hope to achieve in my own career as an educator.

Thank you to my undergraduate mentor and advisor at Wichita State University, Mathew Muether. Your guidance provided me with a strong foundation in both research skills and advanced topics in physics. I am especially grateful for the many hours you spent answering my questions in your office and lending your expertise to make sure I understood the difficult concepts. Thanks to Holger Meyer and Nickolas Solomey, both of whom I relied on for technical guidance and academic advice while at WSU. Thank you to my REU advisor Jon Paley, who I spent the summer of 2017 working with at Fermilab. I appreciate you all for continually fueling my passion for scientific study and communication.

Thanks to my fellow students in the 2018 cohort at Indiana University who helped make the first few years of grad school fun, even during the most stressful times. Thanks to Kylie, Krystyna, and Erin for pushing me to keep going through many late nights while working on homework. Thanks also to my friends Charlie, Josh, Adam, and Jeremy for providing constant motivation while studying for the qual, introducing me to terrible movies on a regular basis, and many, *many*, trips to Chipotle.

The neutrino group at IU made me feel welcome from the very beginning - thanks to Gavin, Erica, Bruce, Justin, Micah, Teresa, and Maria for bringing me into the fold. Thanks to Mark and Jon for keeping the group running as smoothly as possible during the COVID years - it was challenging but worthwhile to maintain a sense of community. Thanks to Ashley, Erica, Maria, Teresa, and Erin for the discord work-and-vent sessions, where we attempted to hang on to any grasp of sanity we had. I want to give a specific thank you to my friend Erica Smith, the postdoc who helped develop many of the skills I gained while working on my analysis. Your willingness to chat on slack at all random hours of the day and night was one of the primary factors in my success - I feel incredibly lucky to have had you as a mentor throughout the majority of my time in grad school.

Thanks to my friends and collaborators on the NOvA experiment, without which this analysis could not have happened. A special thank you to Palash, Abdul, and Sushil for the good times and companionship. An additional thank you goes out to the members of my thesis committee: Rex Tayloe, Emilie Passemard, and John Beggs, for your questions and advice on my work over the years.

Finally, I want to express my deepest gratitude to my PhD advisor, Mark Messier, who had the biggest influence in my academic journey. I have learned so much from you throughout the countless hours spent in your office and on zoom. Your ability to communicate difficult concepts in physics, coding, experimental methods, and statistics is unparalleled. Beyond that, you have given me invaluable advice in other aspects of my career and life - from teaching methodology, to navigating working within a collaboration, to engaging with policymakers, to fatherhood. I will never forget the grace and patience you showed me throughout the years, *especially* during the height of COVID, when there were often multiple weeks in a row where I struggled to make progress. You never made me feel discouraged, and always worked to lift me up when I needed help. You carry yourself with both confidence and humility, which I respect deeply. Thank you for modeling the behavior that I aspire to have in my own interactions with students - I could not have asked for a better advisor.

Reed Bowles

SEARCHING FOR NEUTRINO TRIDENTS IN THE NOVA NEAR DETECTOR

This dissertation presents a search for neutrino trident production in the NOvA near detector through the coherent “dimuon” channel: $\nu_\mu + X \rightarrow \nu_\mu + \mu^- + \mu^+ + X$. Trident production is a rare, purely electroweak process with sensitivity to physics beyond the Standard Model. The theoretical background, motivation for studying the process, and previous experimental measurements are reviewed. The analysis uses data collected by the NOvA near detector (ND) from Fermilab’s Neutrinos at the Main Injector (NuMI) beam between November 2014 and February 2024, corresponding to an exposure of 25.5×10^{20} protons on target. The ND is a segmented tracking calorimeter located 800 m from the beam target, receiving neutrinos with a mean energy of 2 GeV. A multi-pass background reduction strategy is implemented, including the development of a novel dimuon-specific tracking technique. Trident candidates are identified using a boosted decision tree classifier trained on simulated signal and background events. Limited background Monte Carlo statistics necessitate the use of functional fits to sideband data, which are extrapolated to estimate backgrounds in the signal region. The unblinded data contain 9 trident-like events, with an estimated background of 5.66 ± 5.15 events. This yields a best fit estimate of 3.34 tridents compared to the Standard Model prediction of 4.66. A profiled Feldman-Cousins method is used to determine a 90% confidence interval of [0,9.1] on the number of signal events, corresponding to an upper limit of $1.95 \times$ the Standard Model prediction. This result represents the lowest energy search for trident events to date, and the first experimental contribution to the process in 27 years.

Mark Messier, PhD

Rex Tayloe, PhD

Emilie Passemard, PhD

John Beggs, PhD

TABLE OF CONTENTS

Acknowledgements	v
Abstract	vii
List of Tables	xiv
List of Figures	xv
Chapter 1: Neutrinos	1
1.1 Historical Basis of the Neutrino	1
1.1.1 Observing the Neutrino	3
1.2 Properties of the Neutrino	5
1.2.1 Parity Violation	5
1.2.2 Neutrino Helicity	7
1.2.3 Neutrino Flavors	8
1.3 Neutrino Oscillations	11
1.3.1 The Solar Neutrino Problem	11
1.3.2 Neutrino Oscillation Theory	13
1.3.3 Experimental Proof of Oscillations	15
1.3.4 Interpreting Neutrino Oscillations	17
1.4 Experiments Across the Energy Spectrum	20
1.5 Conclusion & Looking Ahead	26

Chapter 2: Neutrino Interactions	27
2.1 Interaction Theory	27
2.1.1 Cross Sections	27
2.1.2 Calculating Cross Sections	29
2.1.3 Scattering Kinematics	34
2.2 Neutrino Interactions in the Standard Model	36
2.2.1 Neutrino-Nucleus (Coherent) Interactions	36
2.2.2 Neutrino-Nucleon (Incoherent) Interactions	37
2.2.3 Nuclear Effects in Incoherent Interactions	41
2.3 Neutrino Scattering in Oscillation Analyses	44
Chapter 3: Neutrino Tridents	47
3.1 Neutrino Trident Theory	47
3.1.1 Trident Production Cross Sections	48
3.1.2 Trident Production as a Probe of Physics Beyond the Standard Model	52
3.2 The Experimental Landscape of Tridents	53
Chapter 4: The NOvA Experiment	59
4.1 The NuMI Beam	59
4.1.1 The Off-Axis Approach	62
4.2 The NOvA Detectors	63
4.2.1 Design of the Detectors	64
4.3 Data Acquisition	67
4.3.1 DAQ Hardware	67
4.3.2 Trigger Systems	68

4.4	Ensuring Good Data Quality	69
4.4.1	Detector Uptime	69
4.4.2	Neutrino Exposure	70
4.4.3	Online Monitoring	71
4.4.4	Nearline Monitoring	72
4.5	Simulating NOvA Data	74
4.5.1	Simulating the NuMI Beam Flux	75
4.5.2	Simulating Neutrino Interactions	76
4.5.3	Simulating Particle Propagation	77
4.5.4	Simulating Photon Transfer & Electronics Readout	79
4.5.5	Final Simulated Data	79
4.6	Calibration	81
4.7	Event Reconstruction	82
4.7.1	Event Topologies	82
4.7.2	The Reconstruction Chain	84
4.7.3	Track-Finding Algorithms	88
Chapter 5: Developing the Tridents Analysis		93
5.1	Analysis Roadmap	93
5.2	Simulation	94
5.2.1	Simulated Trident Dataset	95
5.2.2	Simulated Background Dataset	99
5.3	Initial Data Reduction	99
5.3.1	Preselection	100
5.3.2	First Round Selection	102

5.4	KLutz Custom Trident Tracker	113
5.4.1	The KLutz Tracking Algorithm	114
5.4.2	Additional KLutz Information	124
5.4.3	KLutz Performance	126
5.5	Second Round Selection	127
5.5.1	KLutz Quality Cuts	128
5.5.2	KLutz-Based BDT Selection	129
5.6	Selected Backgrounds	134
5.6.1	Systematic Background Uncertainties	135
5.6.2	Data Driven Background Estimation	138
Chapter 6: Analyzing Near Detector Data	141	
6.1	Unblinding the Background Regions	141
6.1.1	The FHC Dataset	142
6.1.2	The RHC Dataset	144
6.2	Unblinding the FHC Sideband Region	148
6.2.1	Background Estimation	149
6.3	Unblinding the FHC Signal Region	150
6.3.1	Trident Candidate Properties and Event Displays	151
6.3.2	Determining Statistical Significance	160
6.4	Results	163
Chapter 7: Conclusion	165	
Bibliography	166	

Curriculum Vitae

LIST OF TABLES

1.1	Neutrino oscillation parameters determined from 2020 global analysis [76].	25
5.1	Chemical composition of the NOvA near detector [169].	96
5.2	A table of the slice-level variables fed into a BDT.	104
5.3	A table of the one-kalman-track variables fed into a BDT.	107
5.4	A table of FHC variable effectiveness. Separation of 0 represents no separation between signal and background, and 1 represents complete separation.	131
5.5	A table of RHC variable effectiveness. Separation of 0 represents no separation between signal and background, and 1 represents complete separation.	132
5.6	Tables comparing signal and background numbers at various cuts on BDT score for both the FHC and RHC KLutz-Based BDTs. Explanation of low-B values in body paragraph.	133
5.7	Interaction types composing the background events selected by the KLutz-Based BDT.	134
6.1	Details of the fits applied to the MC and real datasets to form the background estimates.	151

LIST OF FIGURES

1.1	The distribution of β particle energies found by C.D. Ellis [6].	2
1.2	A schematic of inverse β decay taking place in the Savannah River apparatus [15]. .	5
1.3	Left: A demonstration of parity conservation and violation in a mirrored image. Right: A diagram of parity violation in the beta decay process [19].	6
1.4	A schematic of the detector hall for Fermilab's DONUT experiment [30].	9
1.5	The spectrum of measurements of Z width, proving only 3 neutrino flavors [31]. . . .	10
1.6	Angular distribution of observed neutrino events in the Kamiokande experiment. An angle of $\cos(\theta_{sun}) = 1$ represents neutrinos coming directly from the sun [36].	12
1.7	Results of the Super-K experiment proving neutrino oscillations [43].	16
1.8	Results of the SNO experiment verifying the existence of non- ν_e solar neutrinos [43].	17
1.9	The effect of the mixing angles on mass ordering and possible mass hierarchies [48]. .	18
1.10	Neutrino oscillation probability as a function of L/E, starting with an electron neutrino. Blue = ν_μ , Red = ν_τ , Black = ν_e [49].	19
1.11	The range of possible neutrino energies and the sources that produce them [55]. . . .	21
1.12	Combined $\sin^2(\theta_{13})$ results from the reactor experiments as of 2021 [63].	22
1.13	Neutrino oscillation probability as a function of L/E, starting with a muon neutrino. Blue = ν_μ , Red = ν_τ , Black = ν_e . [49]	23
1.14	Combined oscillation results of the long baseline experiments [75].	25
2.1	A simple diagram of a two-particle scattering process. One smaller particle scatters off of a single large target via a long-range force. The cross section σ of the interaction depends on the impact parameter b	28

2.2	Feynman diagram vertices for neutrino-induced weak interactions showing the charged current interaction vertex (left) and the neutral current interaction vertex (right). Here ν_ℓ corresponds to one of the neutrino flavors, with ℓ being the associated charged lepton.	30
2.3	An event display of the Gargamelle bubble chamber featuring a NC event [93].	32
2.4	Left: The Standard Model representation of inverse muon decay, mediated by a charged weak boson. Right: Fermi's representation of inverse muon decay, in which the interaction occurs directly between fermions with coupling strength G_F	33
2.5	A generic Feynman diagram of a neutrino scattering off of a nucleus.	34
2.6	A diagram of a CC coherent process which leaves the nucleus unchanged but produces a π^+ in the final state.	37
2.7	A Feynman diagram of a charged-current quasi-elastic interaction.	38
2.8	A Feynman diagram of a charged-current resonance interaction.	39
2.9	A Feynman diagram of a deep inelastic scattering interaction.	40
2.10	Cross sections of neutrino (left) and antineutrino (right) interactions with nucleons as a function of neutrino energy. The solid lines show predictions from models for QE, RES, and DIS interactions while the data points are taken from various experiments [101].	40
2.11	Nuclear potential wells for protons and neutrons bound within a nucleus of constant potential. E_F^p , E_F^n are Fermi energies of protons and neutrons, and E_B is binding energy [103].	42
2.12	The hadronic shower produced in the initial reaction must still pass through the nucleus and is subject to FSI through a variety of processes before entering the detector [107].	43
2.13	Left: Visible hadronic energy in CC ν_μ interactions in the near detector prior to introducing MEC processes. A large discrepancy between predicted and observed energy can be seen. Right: Visible hadronic energy after introducing MEC processes [112].	45
3.1	Four of the 22 Standard Model trident production processes [114].	47
3.2	A NC Coherent Trident production channel with momenta that contribute to the scattering amplitude labeled.	48
3.3	Mean values of the dimuon trident cross section on carbon determined numerically [114], compared to the mean total CC neutrino cross section [32].	51

3.4	Momentum transfer and opening angle histograms for coherent (purple) and incoherent (blue) dimuon events showing low Q^2 and small opening angles [127].	52
3.5	Generic representation of the semileptonic decay of one quark Q into another q' [134].	54
3.6	A dimuon trident event captured in the CHARM II detector [93].	55
3.7	Hadronic energy distribution of observed dimuon events in the CCFR detector [136].	56
3.8	Left: Signal selection criteria for dimuon trident events in NuTeV. Right: Distributions of hadronic energy split by invariant muon mass [138].	57
4.1	An integrated distribution of hits seen within the NOvA near detector during beam spills. The $10\mu\text{s}$ window during which the spill is happening is clearly visible.	60
4.2	Pions exiting the graphite target are focused by magnetic focusing horns. The current in these horns can be flipped to focus either π^+ or π^- . This cartoon is showing π^+ being focused and π^- being de-focused [142].	60
4.3	A diagram of the NuMI beamline [141].	61
4.4	Left: The energy of neutrinos produced at an angle θ relative to the pion direction. Right: ν_μ event rates for a distance of 810 km from Fermilab at various off-axis locations [139].	62
4.5	The ND and FD, and the beam trajectory of the NOvA experiment [143].	64
4.6	Left: A single cell featuring the wavelength shifting fiber. Right: Sixteen cells attached together to form a single extrusion [139].	64
4.7	A schematic of how layering planes in alternating vertical and horizontal orientations yields a top view and side view of energy deposited in the detectors [143, 144]. Deposited energy is resolved at the cell level, with each cell highlighted in red representing a “hit” in the detector.	65
4.8	An event display showing typical near detector activity during a NuMI beam spill. The top pane of the event display shows the XZ view, and the bottom pane shows the YZ view.	66
4.9	Left: Close up of the end of a bundle of optical fibers in a module that interfaces with the APD. Right: The 32 pixel input array of an APD utilized in NOvA [145].	68
4.10	Uptime fraction of the near detector from April 21, 2024 to May 25, 2024. Blue dots represent the daily uptime, with the green and red lines showing the weekly and monthly average, respectively. This period saw a 98% uptime fraction.	70
4.11	Daily POT exposure seen by NOvA since the beginning of the experiment [146]. . .	71

4.12	FEB Hit map of ND hits associated with NuMI beam triggers over a one-week period.	72
4.13	ND GoodRuns and failure modes recorded over a one-week period.	74
4.14	Monte Carlo predicted energy spectrum of neutrino flux at the ND. Left: Flux for beam running in FHC mode. Right: Flux for beam running in RHC mode.	76
4.15	Spectra of ν_μ and $\bar{\nu}_\mu$ CC events in the ND [152].	77
4.16	Mean energy loss rate in various materials for muons, pions, and protons. The minimum-ionizing point is shown as a vertical dash [32].	78
4.17	Event displays of simulated interactions with 2 GeV neutrinos in the ND [154]. . . .	80
4.18	Ratio of observed to true energy deposited in the detector as a function of distance from deposition to cell center [158].	81
4.19	An event display featuring a simulated dimuon trident event.	83
4.20	An event display featuring a simulated resonance event with trident-like topology. .	84
4.21	Simulated events in the NOvA ND with various steps of reconstruction applied. . . .	86
4.22	FuzzyK prongs overlayed on top of hits in the ND.	88
4.23	A diagram explaining elements of the optimum track fitting model used by BPF [166].	90
4.24	A diagram explaining elements of the optimum track fitting model used by BPF in terms of the NOvA geometry [160].	91
4.25	A simulated trident event with Kalman and BPF tracks overlayed.	92
5.1	Trident cross section vs energy for oxygen and iron generated by the trident simulation tool compared to theoretical calculations [114, 121].	96
5.2	Expected trident energy spectra for both FHC and RHC datasets.	98
5.3	Vertex locations in x and z [171].	102
5.4	Number of Kalman tracks for signal (red line), background (green line), and data (black dots). Each distribution is normalized to 1.	103
5.5	Area normalized distributions of BDT variables for signal (blue) and background (red).104	
5.6	Area normalized distribution of slice-level BDT response, cut location included. . . .	105

5.7	Event display of a dimuon event with a single reconstructed 3D kalman track. Left-over hit information for this event was distributed across additional 2D tracks.	106
5.8	Distribution of single-kalman-track BDT scores for signal (blue) and background (red).	108
5.9	Distribution of single-kalman-track BDT scores for signal (blue) and background (red).	109
5.10	Distribution of single-kalman-track BDT scores for signal (blue) and background (red).	110
5.11	Distributions of secondary track variables for signal (blue) and background (red).	111
5.12	Distribution of two kalman track BDT scores for signal (blue) and background (red).	112
5.13	A dimuon event showing a failure mode of Kalman tracks - there are two 3D tracks and two 2D tracks all capturing <i>pieces</i> of the muons, but not their entire trajectories.	113
5.14	A cartoon representation of the NOvA near detector composed of vertical cells and horizontal planes with hits representing a dimuon event.	114
5.15	Downstream and upstream windows represented by red boxes drawn around a possible endpoint candidate, highlighted red.	115
5.16	The crossed out hit does not meet the endpoint criteria, whereas the circled hit does.	115
5.17	A track seed formed from the furthest downstream endpoint candidate, highlighted in pink. Endpoint candidates not included in the seed are highlighted blue.	116
5.18	The primary track formed out of the most downstream endpoint candidate, highlighted in green. At this point the sequence starts again from the next furthest downstream endpoint.	117
5.19	Two tracks formed, shown in green and blue. Hits shared between the tracks get energy split evenly amongst the two.	117
5.20	An event display featuring the all endpoint candidates for a dimuon event. Only the furthest downstream will actually go on to form track seeds.	118
5.21	An event display showing the two track seeds grown from the endpoint candidates, highlighted in blue and pink.	118
5.22	An event display featuring the track-creation process - each hit forms one track at a time out of the furthest downstream track seed.	119
5.23	An event display showing four fully-formed 2-dimensional tracks.	119
5.24	An event display showing four 2D KLutz tracks, despite two tracks overlapping completely in the x view.	120

5.25 Differences in the slope (dQ) and y -intercept (dW) of collinear tracks. Graphs generated on a sample of 100,000 tridents, meaning muon-catcher-only tracks occur $\sim 1\%$ of the time.	121
5.26 A track formed entirely within the muon catcher (pink).	122
5.27 Remove the muon catcher track and recreate a second track without those hits.	122
5.28 Check co-linearity and stitch the muon catcher track to the newly formed track.	122
5.29 An event display showing two fully formed 3-dimensional tracks with a very narrow opening angle.	124
5.30 Results of hand-scanning events with KLutz reconstruction to measure performance.	126
5.31 The number of KLutz and Kalman tracks produced on a simulated trident dataset.	127
5.32 A clearly non-trident background event with two muon-like tracks, a small opening angle, and little hadronic energy.	128
5.33 Some variables used in the second round FHC BDT, signal scaled up 1000 times. The jump in signal events at primary track length 450 cm is an artifact of the dataset generation, in which I limited vertex locations to 450 cm from the end of the detector.	130
5.34 Output of the simulated FHC dataset BDT trained on KLutz variables.	131
5.35 Output of the simulated RHC dataset BDT trained on KLutz variables.	132
5.36 Total neutrino and anti-neutrino cross sections as a function of energy. In the relevant energy range (~ 2 to 20 GeV) the $\bar{\nu}_\mu$ cross section is roughly half that of ν_μ [32].	134
5.37 NOvA's most detailed study of flux uncertainties for $\langle E_\nu \rangle > 6$ GeV [178].	135
5.38 Sources of ν_μ flux modeled in the ND overlayed with data [179].	136
5.39 Current measurements of the coherent pion production cross section [181].	137
5.40 Coherent pion production cross section uncertainty (model uncertainty in green) [181].	137
5.41 Exponential and Power Law fits overlayed on FHC MC. Fits were performed on all available MC (up to a BDT score of 0.27) and extrapolated into the signal region.	140
6.1 KLutz-Based Trident BDT score with FHC data and MC overlayed.	142
6.2 Before and after comparisons of FHC nominal MC, weighted MC, and data.	143
6.3 KLutz-Based Trident BDT score with RHC data and MC overlayed.	144

6.4	Topology-based BDT with the cut location identified, quality cuts applied.	145
6.5	One Kalman track-based BDT with the cut location identified, quality and topology-based BDT cut applied.	146
6.6	Two Kalman track-based BDT with the cut location identified, quality and topology-based BDT cut applied.	146
6.7	Before and after comparisons of RHC nominal MC, weighted MC, and data.	147
6.8	KLutz-Based Trident BDT score sideband with FHC data and MC overlayed.	148
6.9	KLutz-Based Trident BDT score signal region with FHC data and MC overlayed. . .	150
6.10	Extrapolated function fits used for the background estimation overlayed on data. . .	151
6.11	Distributions of candidate event properties (red data points). The track lengths appear longer than predicted due to inclusion of sideband events. Signal region distributions included in the bottom two plots for a more accurate representation. .	152
6.12	Distributions of candidate event properties (red data points).	153
6.13	Distributions of candidate event properties (red data points).	154
6.14	An event display featuring a candidate trident event with BDT Score 0.280.	155
6.15	An event display featuring a candidate trident event with BDT score 0.265.	155
6.16	An event display featuring a candidate trident event with BDT score 0.335.	156
6.17	An event display featuring a candidate trident event with BDT score 0.265.	156
6.18	An event display featuring a candidate trident event with BDT score 0.265.	157
6.19	An event display featuring a candidate trident event with BDT score 0.295.	157
6.20	An event display featuring a candidate trident event with BDT score 0.290.	158
6.21	An event display featuring a candidate trident event with BDT score 0.290.	158
6.22	An event display featuring a candidate trident event with BDT score 0.265.	159
6.23	A Feldman-Cousins confidence belt representing observed events compared to mean signal event count.	162
6.24	Some BSM models [197] overlayed with experimental results from the NOvA and Charm II [133] experiments.	164

Chapter 1

Neutrinos

Neutrinos are the most abundant massive particle in the universe that we know of. They interact weakly and are lightweight, making them notoriously difficult to measure. In the nearly 100 years since they were first postulated they have played a significant role in the development of particle physics. In fact, neutrinos may hold the answers to some of the biggest questions about the universe.

1.1 Historical Basis of the Neutrino

Physics is the study of how we fundamentally understand the world. The progression of knowledge is generally a slow process - full of incremental steps that grow into an understanding of broad concepts. However, major discoveries occasionally shift the paradigm, causing physicists to look at things with a whole new perspective. The story of the neutrino begins with one such discovery.

At the end of the 19th century, the landscape of physics was changing dramatically [1]. James Maxwell had formulated a theory of electromagnetism that provided physicists with a never-before-seen comprehension of light. Additionally, the recent discoveries of cathode rays and X-rays caused excitement and vigorous debate in the field. Across the world, academics formulated experiments and theories to gain deeper understanding of these phenomena. One such experiment was performed by Henri Becquerel, who conducted a study in which he investigated the properties of “rays” emitted by uranium. Using a photographic plate, he found that the rays were emitted spontaneously and had the capability to pass through thick paper, aluminum plates, and copper foil [2]. Becquerel had discovered radiation.

His discovery captured the attention of Marie and Pierre Curie, who proceeded to study the electrical conductivity of air when exposed to uranium and various other elements [3]. They found that thorium was also radioactive and correctly observed that the elements with the highest atomic weights were the most radioactive. The Curies went on to uncover multiple new, more radioactive elements with higher atomic weights than had been known at the time. These discoveries earned the Curies and Becquerel the 1903 Nobel Prize in Physics.

At the turn of the century the flood gates had opened on the study of radiation. To better understand the newly discovered rays, Ernest Rutherford investigated the absorption spectrum of uranium in aluminum foils of various thickness [4]. From his tests he observed that the rays consisted of two separate types of radiation which he labelled as “alpha” (α) and “beta” (β) rays. With additional experimentation, Rutherford (and others) determined that the “rays” were not actually traditional rays at all, but particles emitted when unstable nuclei of atoms underwent radioactive decay. α particles were helium nuclei and β particles were singular electrons. Rutherford also observed that the processes that created these two particles were different and could be studied independently.

Experiments performed in the 1920’s by the Curies and George Briggs showed that α particles emitted in the decay of a substance had very homogeneous energies, which matched the predictions of nuclear decay at the time [5]. However, an unexpected occurrence was observed by Charles Ellis in 1927 while studying the energy spectrum of radiated β particles [6]. Ellis found that electrons coming from nuclear decay did not have a single energy as expected. Instead, they had energies distributed over a wide range, as seen in Figure 1.1. This discrepancy was a critical issue in the understanding of nuclear physics at the time.

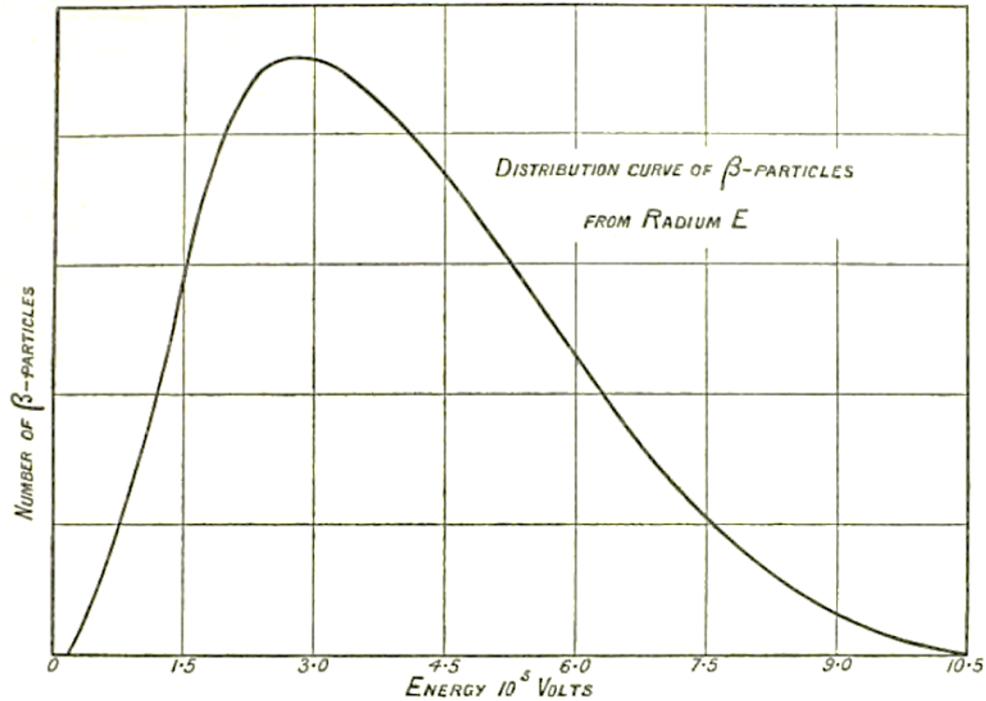


Figure 1.1: The distribution of β particle energies found by C.D. Ellis [6].

Ellis theorized several possible explanations of the phenomenon, but the correct postulation would not be made until 1930, when Wolfgang Pauli sent a letter to a physics conference being held in Tübingen, Germany [7]. In the letter, Pauli proposed the existence of a light, electrically neutral particle that had not yet been observed but could account for the missing energy in β -decay processes. He also questioned the likelihood of experimentally observing such a particle, describing that its ability to penetrate material may be “ten times larger” than that of gamma rays. In 1934 Enrico Fermi built on Pauli’s idea by proposing the “neutrino” and creating a formal theory of β -decay to describe its existence [8]. In a β -decay process, a neutron spontaneously decays into a proton, electron, and anti-neutrino. The theory was incredibly successful, providing a thorough solution to one of the biggest questions in physics at the time. However, the issue of experimentally proving the neutrino’s existence remained. According to calculations done by Hans Bethe and Rudolf Peierls in 1934, there was “no practically possible way of observing the neutrino” [9].

1.1.1 Observing the Neutrino

Bethe and Peierls were not entirely wrong in their prediction – neutrinos proved to be incredibly difficult to detect. Experiments attempting to observe the neutrino were performed for decades, and while all of the experiments were consistent with the *existence* of neutrinos, none of them provided direct observational evidence [10]. In fact, it was only after 20 years and the development of nuclear technology due to World War II that neutrinos were finally cemented as real particles.

The first nuclear chain reaction was achieved in the United States under the direction of Enrico Fermi in 1942 [11]. The technology developed in pursuit of the reaction went towards creating the atomic bomb, and went on to be widely used in the nuclear fission reactor. In a 1946 paper, Bruno Pontecorvo introduced the inverse β decay process, which described how an anti-neutrino can interact with a proton to yield a neutron and a positron [12]. Pontecorvo also estimated that the neutrino flux emitted by the newly developed nuclear fission reactors could be up to the order of 10^{14} neutrinos per cm^2/sec . He theorized that such a flux would be sufficient to make an experimental observation of neutrinos via detection of inverse β decay. Using that flux estimation and the inverse β decay theory, Luis Alvarez detailed an experimental setup that could finally be used to detect the neutrino [10].

The confirmation was eventually made in a set of experiments performed by Clyde Cowan and Frederick Reines. The procedure utilized the inverse beta decay process: $\bar{\nu}_e + p \rightarrow n + e^+$ to search for a signal of spontaneous positron generation followed by a delayed neutron capture in Cadmium [13]. Initially the duo planned to observe this signal by placing a detector 100 m away from an exploding nuclear bomb. Fortunately, they came to the “obvious conclusion” that it would be quite complicated to make a signal measurement from the explosion and turned towards a fission reactor source instead.

They performed the first experiment in 1953 at the Hanford Nuclear Reactor Site in Washington State [14]. The detector was a 10-cubic-foot cylinder, filled with 300 liters of liquid scintillator and surrounded by photomultiplier tubes (PMTs) - an unusually large and unique apparatus at the time of the experiment [15]. The detector was placed at ground level, directly near the face of the nuclear reactor. It was surrounded by a few feet of paraffin and a few inches of lead shielding in an attempt to reduce background noise in the form of neutrons. Operating the detector for a few months yielded a “marginal” difference in interaction rates between reactor-on and reactor-off running, hinting at the existence of the neutrino. Unfortunately the initial experiment was plagued by large amounts of background coming from cosmic rays and electrical noise, causing the search to end prematurely. Still, their observation was enough to motivate Cowan and Reines to give the experiment another shot.

Their second attempt at observing the neutrino utilized a new detector and enhanced shielding to minimize background rates from both cosmic rays and reactor emissions [16]. This experiment was performed in 1955 at the Savannah River Nuclear Reactor Plant in South Carolina. They placed the detector 11 meters away from the reactor and 12 meters underground, providing significantly better shielding from cosmic rays than before [15]. The detector itself was modified to be a “sandwich” of 400 liters of cadmium chloride ($CdCl_2$) doped water between large tanks filled with 4200 liters of liquid scintillator. The apparatus, pictured in Figure 1.2, was designed specifically to identify neutrino interactions much more clearly than neutrons or protons. After about 100 days of run time with the new setup, Cowan and Reines verified a reactor-power-dependent signal that agreed with the neutrino cross section. The pair had definitively confirmed the existence of the neutrino. The discovery won the Nobel Prize in physics, which was awarded to Reines in 1995 (Cowan had unfortunately passed away by then).

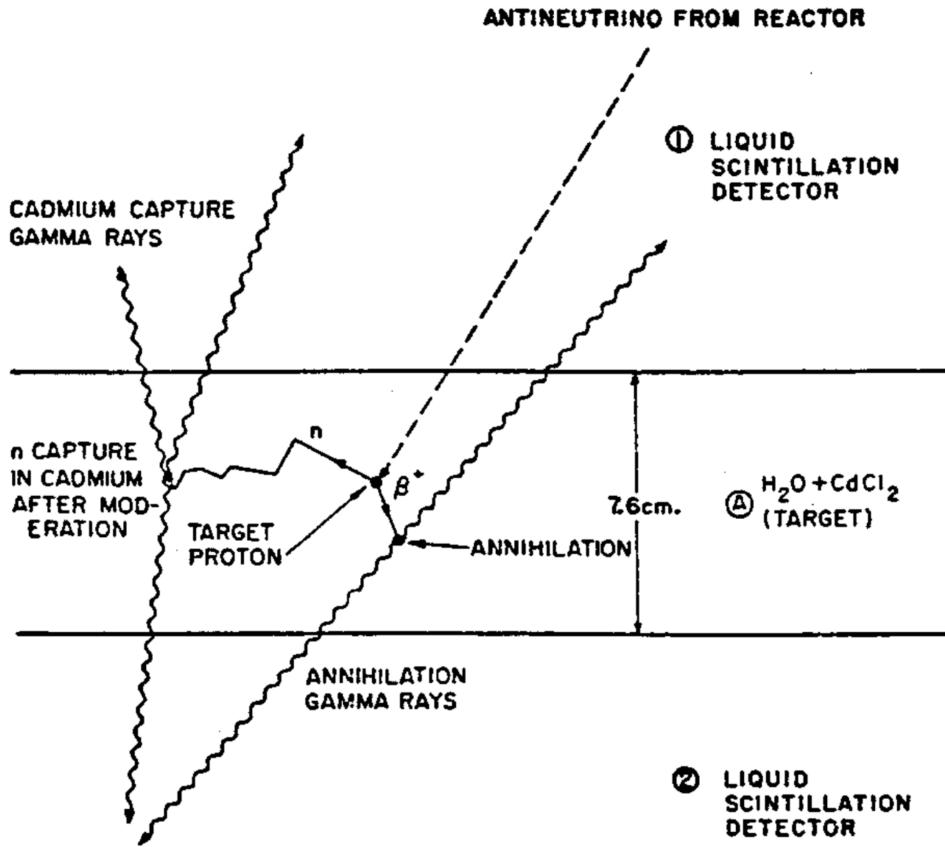


Figure 1.2: A schematic of inverse β decay taking place in the Savannah River apparatus [15].

1.2 Properties of the Neutrino

1.2.1 Parity Violation

The experimental verification of the neutrino was, while exciting, not surprising to physicists at the time. Fermi's theory of beta decay had been so successful that it was already widely accepted as predominant in the field of particle physics. However, the discovery did lead to a fervor of interest in studying neutrinos - sparking decades of theorizing about weak nuclear interactions. One such theory came in 1956 from two physicists named Tsung-Dao Lee and Chen-Ning Yang [17].

The pair postulated that β decay (and weak interactions in general) may violate the principle of parity conservation. Additionally, they postulated that if parity is violated, then a “two-component” theory of the neutrino would be valid [18]. They calculated the weak β decay process using the most general Hamiltonian containing both parity conserving (C) and parity non-conserving (C') coupling terms. They found that calculations yielded sets of terms proportional to

$|C^2|$ added to terms proportional to $|C'^2|$. They did not calculate any terms that showed interference between the couplings (CC') because experimental evidence had not yet made measurements of the properties needed to determine them. It was therefore possible that the C' coupling terms contributed to the β decay process, meaning that parity was violated. In order to better understand the possible interference between coupling terms, Lee and Yang proposed a handful of experimental tests to measure those properties. One such proposal was a measurement of the decay process $\pi \rightarrow \mu + \nu, \mu \rightarrow e + \nu + \nu$.

By letting a pion decay from rest and measuring the resulting angle θ between the muon and electron, one could determine whether parity is conserved or violated. The law of parity conservation dictates that some physical phenomenon must be indistinguishable from its mirrored counterpart [19]. For example, the scenario of the top ball in Figure 1.3 and its mirrored counterpart. As the ball spins in the direction indicated, it emits particles equally in both directions along its axis of rotation – in this case the ball is indistinguishable from its mirror image to an outside observer. However, the bottom ball *is* distinguishable from its mirror image because it only emits particles in a single direction. That makes the “handedness” of the mirrored image exactly opposite of the real ball.

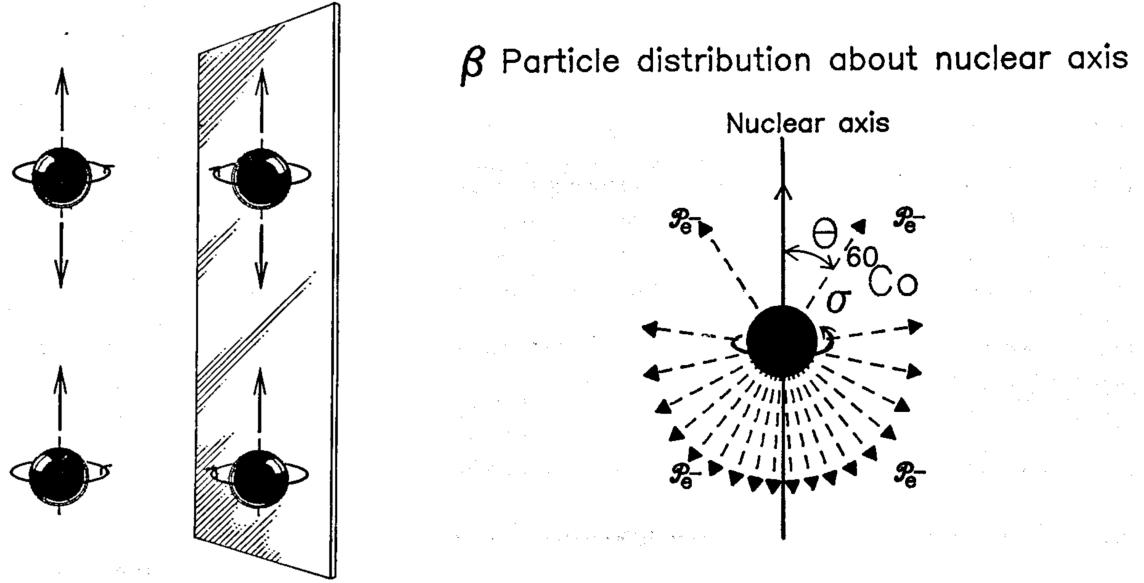


Figure 1.3: **Left:** A demonstration of parity conservation and violation in a mirrored image.

Right: A diagram of parity violation in the beta decay process [19].

In the pion decay experiment proposed by Lee and Yang, if parity is conserved one would expect to see a θ distribution that is symmetric for angles of both θ and $\pi - \theta$. However, if parity is violated the distribution would be asymmetric - meaning that the emitted beta particles would show a “preference” for a certain direction. This phenomenon is displayed in the right half of Figure 1.3. The proposal of this measurement piqued the interest of a physics professor from Columbia University named Chien-Shiung Wu.

In just a few months, Professor Wu formulated and carried out an experiment to test for violation of parity symmetry using beta decay [20]. With the assistance of Dr. Ernest Ambler and his low-temperature lab, Wu polarized a source of radioactive ^{60}Co and observed the spectrum of emitted β particles. She determined that the electrons were emitted preferentially in the opposite direction to the cobalt’s nuclear spin. In fact, the angular distribution of electrons was *so* asymmetric that the interference between parity-conserving and non-conserving terms (the CC' term mentioned earlier) was near maximal - proving that parity conservation was violated in weak interactions. Lee and Yang were awarded the 1957 Nobel Prize in physics for their theory, but Professor Wu’s experimental verification went unrecognized - a historic injustice resulting from intersecting gender and racial prejudices of the time.

1.2.2 Neutrino Helicity

Another property of neutrinos was discovered later in 1957 by a trio of physicists from Brookhaven National Lab (BNL) [21]. They performed an experiment to measure the neutrino’s “helicity” - a property that describes how a particle’s spin compares to its linear motion [22]. If the particle’s spin is parallel to its linear motion, it is said to have positive (right-handed) helicity. If the direction of the particle’s spin is anti-parallel to its linear motion, it has negative (left-handed) helicity. Helicity is a valuable property because it contains all of a traveling particle’s kinematic information.

The previous experiment performed by Wu proved that neutrinos violate parity symmetry, but it also indicated that the helicity of the beta decay electron is negative, i.e that its spin is always opposite to its emission direction [20]. Inspired by the electron helicity measurement and other experimental proposals of Lee and Yang, Maurice Goldhaber turned his sights towards confirming the result of parity violation and explicitly studying neutrino helicity [23].

Along with Lee Grodzins and Andrew Sunyar, Goldhaber formulated an experiment that could measure circularly polarized gamma rays emitted from orbital electron capture on ^{152}Eu . The measurement relied on the chain reaction $e^- + ^{152}\text{Eu} \rightarrow \nu_e + ^{152}\text{Sm}^* \rightarrow \nu_e + ^{152}\text{Sm} + \gamma$. In this process, the helicity of the neutrino is transferred to the recoiling nucleus, which is then transferred to the emitted gamma ray. Since angular momentum must be conserved, it means the resultant gamma ray has the same helicity as the original neutrino. The scientists found that the gamma rays were *always* circularly polarized with negative helicity - meaning the neutrinos were left-handed 100% of the time. This property of neutrinos is unique in particle physics; no other particles have only a single type of helicity all the time. In fact, at the time of this discovery the absence of a right-handed neutrino in the Standard Model meant there was no mechanism for the neutrino to acquire mass, and therefore must be massless [24].

1.2.3 Neutrino Flavors

Throughout the 1950s, improvements in technology made studying neutrinos more worthwhile. With every discovery, physicists were pondering unique ways to use neutrinos as probes for weak interactions at increasingly higher energies. In a 1960 paper, a professor of physics at Columbia University named Melvin Schwartz proposed the use of high energy proton accelerators as a source of high energy neutrinos. He outlined a procedure in which allowing the protons to impinge upon a target would yield high energy pions, which would then decay into neutrinos [25].

Schwartz came up with his suggestion through discussions with Lee and Yang, who had just won the Nobel Prize for their theory of parity violation. The duo was interested in studying the theoretical implications of using pions as a source of neutrinos [26]. They pointed out that in beta decay processes (such as in nuclear reactors or the sun), the resulting lepton is an electron. However, in pion decay the primary lepton from the interaction is a muon. This led them to propose the question: are the neutrinos in each of these processes truly the same, or do they differ somehow?

To answer the question of whether there are different types of neutrinos, Schwartz devised an experiment based on his pion decay premise [27]. He performed the first high-energy accelerator neutrino study in 1962 working with a team of scientists from Columbia University and Brookhaven National Lab. The team directed a 15 GeV proton beam into a beryllium target, producing a large flux of pions which decayed into muons and neutrinos.

The detector was a 10-ton aluminum spark chamber shielded from the resultant muons by a 13.5 m thick iron shield. It was exposed to 3.48×10^{17} protons, yielding 113 events that originated within the volume of the detector. Of these 113 events, 56 events were clear and energetic enough to confidently be called neutrino interactions. All 56 of the observed events were determined to come from “muon-type” neutrinos, yielding muons in the final state with no evidence of electron showers. This work, which proved that there were two distinct “flavors” of neutrino, earned Leon Lederman, Melvin Schwartz, and Jack Steinberger the 1988 Nobel Prize in Physics.

It had been proven that the two known leptons, the electron and the muon, had neutrinos associated with them: ν_e and ν_μ . So naturally when a third generation of lepton, the tau, was discovered by Martin Perl et al. in 1975 [28, 29], it was assumed that a third flavor of neutrino must exist to go with it: the ν_τ . The confirmation of the ν_τ was made 25 years later, when Fermilab’s DONUT experiment observed neutrino interactions in which a tau was the only lepton created at the interaction vertex [30]. By that time the Fermilab Tevatron facility was capable of producing an 800 GeV proton beam, which was made to collide with a large tungsten beam dump. This interaction yielded D_s mesons, which decay leptonically into τ and ν_τ . After passing through 36 m of magnets, concrete, iron, and lead shielding, the resulting ν_τ would continue to an emulsion target with scintillating fibers distributed throughout. The particles hit the main target in front of an analyzing magnet, drift chambers, and a muon catcher stretching 16m so that interactions within the emulsion target could be easily identified.

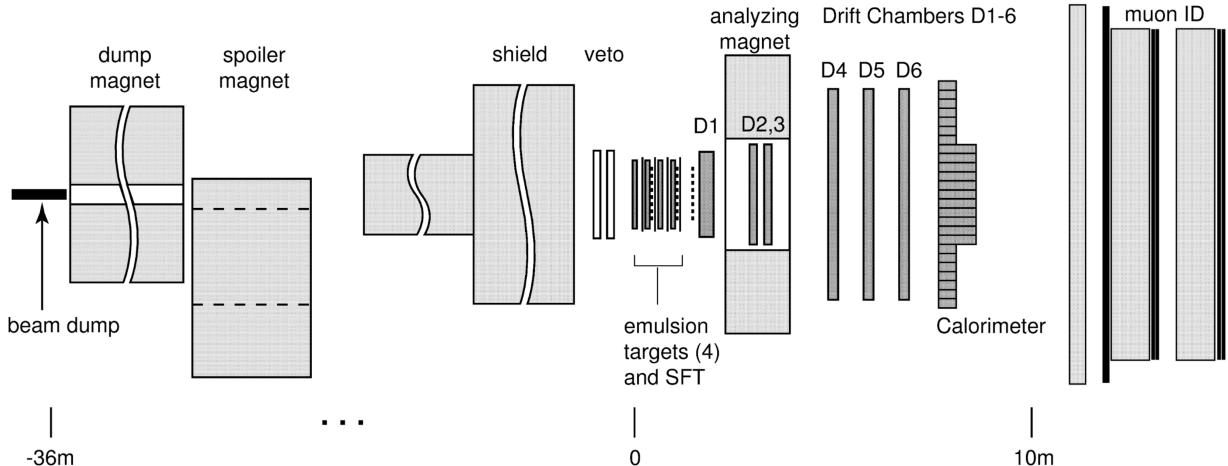


Figure 1.4: A schematic of the detector hall for Fermilab’s DONUT experiment [30].

Using this apparatus with an exposure to 3.54×10^{17} protons on the tungsten target, the DONUT collaboration identified 4 ν_τ events out of 898 neutrino interaction candidates. This observation was sufficient to make a definitive statement that the ν_τ existed, verifying the existence of three neutrino flavors.

After the discovery of the τ and ν_τ , a logical question arises: could there be additional generations of leptons and light neutrinos to go with them? Fortunately, there is experimental data to answer that question [31]. The Z boson (the neutral mediator of the weak nuclear force) can decay into leptons with smaller mass than the Z itself, and its width is dependent on the number of neutrinos. Various Large Electron-Positron Collider (LEP) experiments have measured the cross section of the $e^+ + e^- \rightarrow (hadrons)$ interaction, which yields Z bosons that will decay into neutrinos. As shown in Figure 1.5, the combined results of these experiments show that the data perfectly fits the line corresponding to only 3 flavors of neutrinos. This does leave open the possibility of neutrinos that are not decay products of the Z boson such as sterile neutrinos that do not interact with the weak force, and heavy neutrinos with $M_\nu > M_Z$ [32].

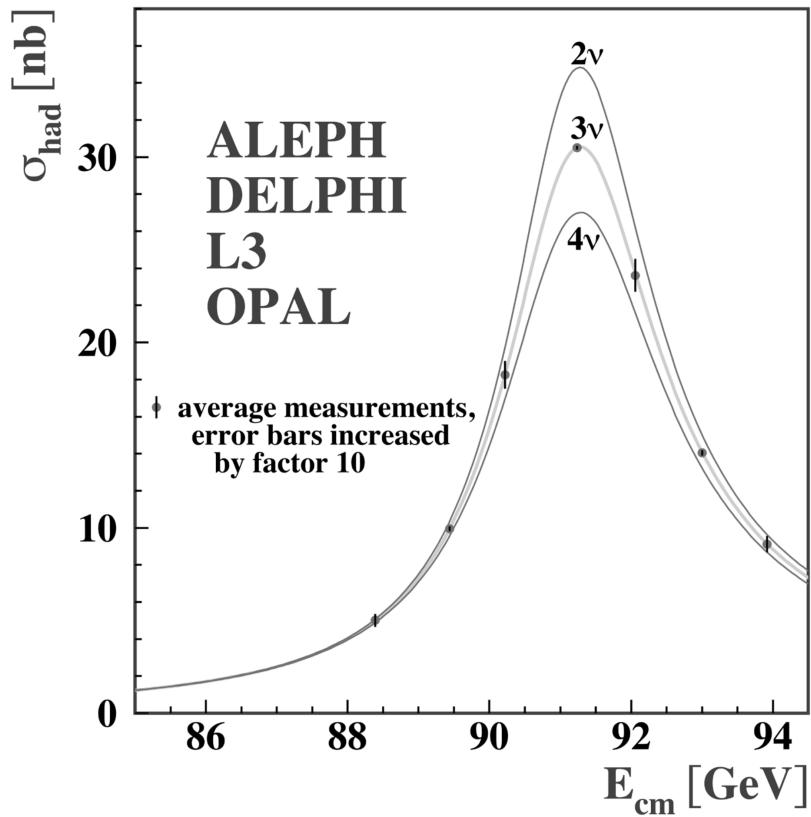


Figure 1.5: The spectrum of measurements of Z width, proving only 3 neutrino flavors [31].

1.3 Neutrino Oscillations

1.3.1 The Solar Neutrino Problem

All of the neutrino studies discussed so far have been based around neutrinos generated on Earth by either nuclear reactors or particle accelerators. However, another significant source of neutrinos is the sun. Nuclear processes within the core of the sun output an enormous amount of electron neutrinos [33]. The first experiment to utilize so-called “solar” neutrinos was designed in the 1960s by a Brookhaven physicist named Raymond Davis. Davis established the Homestake experiment to identify nuclear processes in the sun by way of measuring the solar electron neutrino flux [34].

The Homestake Solar Neutrino Detector began construction in South Dakota’s Homestake Mine in 1965. The technique used to measure solar neutrinos was the inverse beta decay reaction in chlorine: $\nu_e + {}^{37}\text{Cl} \rightarrow {}^{37}\text{Ar} + e^-$. The detector apparatus was a large cylindrical steel tank 6.1 m in diameter and 14.6 m long. The tank was placed 1478 m underground and filled with 615 metric tons of C_2Cl_4 . When neutrinos interacted with the ${}^{37}\text{Cl}$ in the tank, they would produce a very small amount of ${}^{37}\text{Ar}$. These argon atoms would be efficiently extracted from the tank and counted. The number of argon atoms counted (minus the known rate of non-solar ${}^{37}\text{Ar}$ production in the detector) corresponded directly to the number of solar neutrino interactions in the tank.

The first results of the Homestake experiment were published in 1968 with a peculiar outcome: Davis’ partner and theorist John Bahcall showed that the rates were about one third of what was expected based on solar model predictions [35]. The Homestake team double checked their argon extraction efficiency and measurements of background rates but could find no reasonable explanation for such a massive deficiency in the number of solar neutrinos. This confusing result was the beginning of what came to be known as the “Solar Neutrino Problem”.

By the early 1980s it was clear to many studying neutrino physics that the discrepancy between the Homestake experiment’s observed and predicted rate would be problematic. To verify the result, the Kamioka Nucleon Decay Experiment turned their sights towards a measurement of solar neutrinos as well [36, 37]. Located 1000 m underground in Japan’s Kamioka Mine, the Kamiokande detector was upgraded to be sensitive to solar neutrino interactions and was ready to begin taking data by 1987. The Kamiokande-II detector was a 2140-ton water Cherenkov detector capable of measuring both the energy and the direction of incoming neutrinos.

The experiment found a clear signature of electron neutrinos coming from the direction of the sun but observed less than half of the expected rate – a result consistent with that of the Homestake Experiment. This result is shown in Figure 1.6, in which the electron neutrino flux coming from the sun (where $\cos(\theta_{\text{sun}}) = 1$) is significantly lower than the prediction. The discovery and verification of the solar neutrino problem yielded the 2002 Nobel Prize in physics to Masatoshi Koshiba (director of the Kamiokande experiment) and Ray Davis.

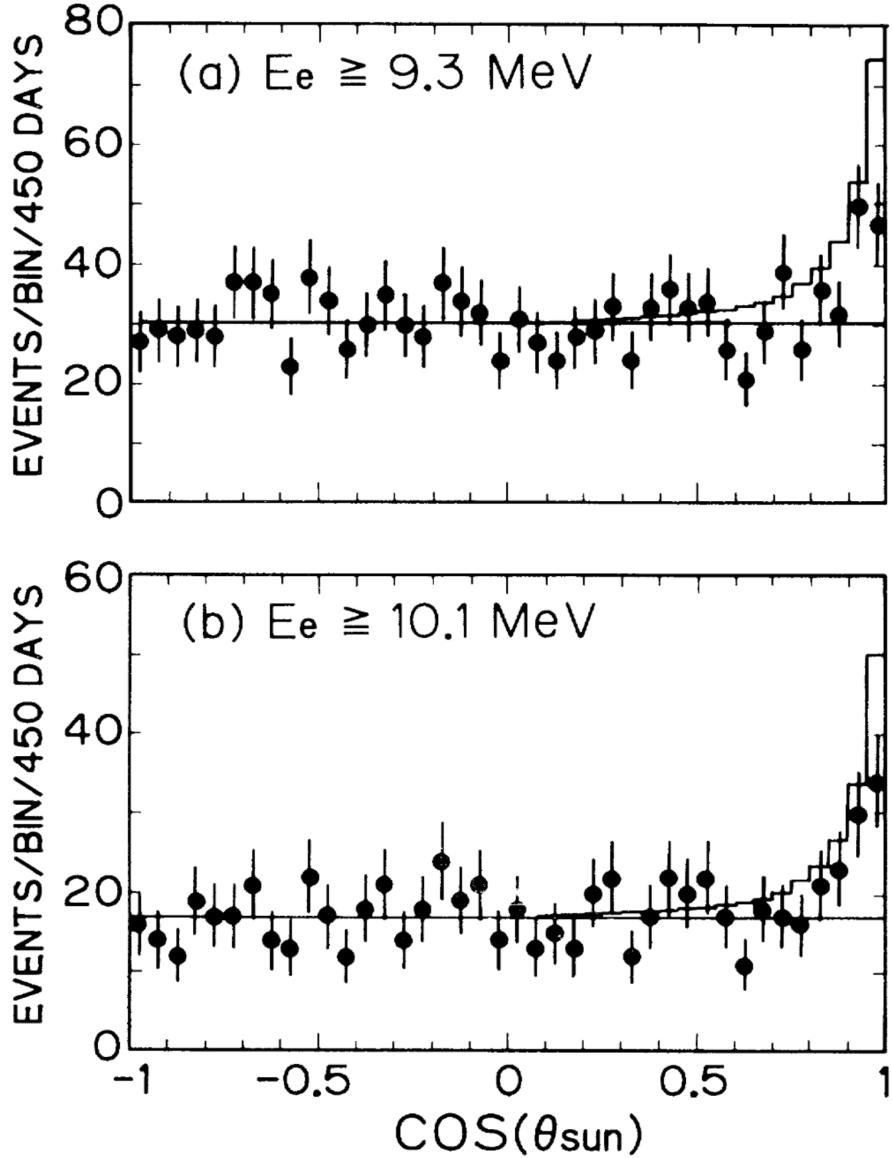


Figure 1.6: Angular distribution of observed neutrino events in the Kamiokande experiment. An angle of $\cos(\theta_{\text{sun}}) = 1$ represents neutrinos coming directly from the sun [36].

1.3.2 Neutrino Oscillation Theory

The solar neutrino problem was perplexing because it provided evidence of some unknown physics not captured by the Standard Model at the time. However, the correct solution had been theorized before the Homestake and Kamiokande experiments even finished construction. In the late 1950s and early 1960s, two sets of theorists independently proposed neutrino “mixing” models, in which neutrinos may “oscillate” from one flavor to another. Bruno Pontecorvo with help from Vladimir Gribov [38, 39, 40], as well as the team of Ziro Maki, Masami Nakagawa, and Shoichi Sakata [41] had formalized a theory of neutrino oscillations by the end of the 1960s [42].

The basis of the theory is that there are two distinctly quantized fields with which neutrinos can be described: flavor and mass. Neutrino fields of a specific flavor (ν_ℓ) do not coincide with neutrino fields of a specific mass (ν_i) - rather, they coincide with *linear combinations* of ν_i fields. Neutrinos interact with matter in one of their three *flavor* eigenstates $\nu_l = \nu_e, \nu_\mu$, or ν_τ , but they travel through space in their *mass* eigenstates $\nu_i = \nu_1, \nu_2$, or ν_3 .

This theory addresses the solar neutrino problem, stating that the neutrinos generated in the sun are electron-type neutrinos but propagate through space as mass states, allowing the flavor states to mix as they travel to the Earth. By the time they arrive and interact at the Earth the neutrinos have oscillated into a different flavor, causing a deficit in the number of ν_e s detected. Notably, this theory depends on neutrinos having *non-zero mass*, which was directly at odds with the lack of right-handed neutrinos in the Standard Model referenced in subsection 1.2.2. Therefore, if neutrinos truly do oscillate it presented new physics to be studied and understood.

Using Fantini et al. [42] as the primary contributor to the following derivation, the basis of neutrino mixing is the linear relationship:

$$|\nu_\ell\rangle = \sum_{i=1}^N U_{\ell i}^* \nu_i, \text{ with } \begin{cases} \ell = e, \mu, \tau & [\text{flavor}] \\ i = 1, 2, 3 & [\text{mass}] \end{cases} \quad (1.1)$$

Here, $|\nu_\ell\rangle$ are the neutrino flavor eigenstates, $|\nu_i\rangle$ are the mass eigenstates, and $U_{\ell i}$ are the elements of the leptonic mixing matrix. This mixing matrix U was formalized by Pontecorvo, Maki, Nakagawa, and Sakata, leading to it being referred to as the PMNS matrix. The PMNS matrix is a unitary rotation matrix which contains three rotation angles, $\theta_{12}, \theta_{13}, \theta_{23}$ (commonly referred to as the neutrino “mixing angles”), and a potentially non-zero charge-parity (CP) violating phase, δ_{CP} .

Using $c_{ab} = \cos(\theta_{ab})$ and $s_{ab} = \sin(\theta_{ab})$, the PMNS matrix takes the form

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.2)$$

Considering a neutrino in some initial mass state $|\nu_i\rangle$ at some time t , the state can be expressed as $|\nu_i, t\rangle = |\nu_i, 0\rangle e^{iE_i t}$ as it propagates in time. This term can be plugged into Equation 1.1 to express an initial flavor state $|\nu_\ell\rangle$ in terms of a final flavor state $|\nu_{\ell'}\rangle$ as the neutrino travels.

$$|\nu_\ell, t\rangle = \sum_\ell \left(\sum_{i=1}^N U_{\ell i}^* e^{-iE_i t} U_{\ell' i} \right). \quad (1.3)$$

The transition amplitude between these two states is then

$$A_{\ell\ell'} \equiv \langle \nu_\ell, 0 | \nu_{\ell'}, t \rangle = \sum_{i=1}^N U_{\ell i}^* U_{\ell' i} e^{-iE_i t}. \quad (1.4)$$

This makes the transition probability between two flavor states

$$P_{\ell\ell'} \equiv A_{\ell\ell'} A_{\ell\ell'}^* = \sum_{i=1}^N \sum_{j=1}^N U_{\ell i}^* U_{\ell' i} U_{\ell' j}^* U_{\ell j} e^{-i(E_i - E_j)t}. \quad (1.5)$$

Now both i and j represent possible mass states of the neutrino. Finally, treating neutrinos as relativistic particles and using units where $c = \hbar = 1$, the energy E_i of mass eigenstates with mass m_i can be approximated as

$$\begin{aligned} E_i &\approx \sqrt{p_i^2 + m_i^2} \approx E + \frac{m_i^2}{2E} \quad \text{for } E \approx |p_i| \\ \implies (E_i - E_j)t &\approx \left(\frac{m_i^2}{2E} - \frac{m_j^2}{2E} \right) L = \frac{\Delta m_{ij}^2 L}{2E}. \end{aligned} \quad (1.6)$$

This result can be plugged into Equation 1.5 to yield

$$P_{\ell\ell'} \approx \sum_{i=1}^N \sum_{j=1}^N U_{\ell i}^* U_{\ell' i} U_{\ell' j}^* U_{\ell j} e^{-i \frac{\Delta m_{ij}^2 L}{2E}}. \quad (1.7)$$

This is the oscillation probability between two flavor states of a neutrino in vacuum.

1.3.3 Experimental Proof of Oscillations

After observing the deficit in ν_e s from the sun, the Kamiokande experiment set their sights on searching for definitive proof of neutrino oscillations. The upgraded Super-Kamiokande (Super-K) detector is located 1000 m underground in Japan's Kamioka Mine [43]. It is a 50-kiloton water Cherenkov detector with an imaging system made up of over 11,000 photomultiplier tubes (PMTs). It also has nearly 2000 PMTs located on the outside of the detector to veto particles entering the detector and to tag exiting tracks. This imaging system allows the detector to have very precise angular resolution, yielding detailed information about the directions neutrinos come from.

Super-K is designed to observe not only solar neutrinos [44], but atmospheric neutrinos as well [45]. By measuring the flux of atmospheric neutrinos coming into the detector at all angles, the experiment showed that the observed ratio of ν_μ/ν_e was consistent with the theory of neutrino oscillations. Super-K's first observation of neutrino oscillations was published in 1998, and according to the lead scientist of the project Takaaki Kajita, "It immediately became something really big - a major, major revolution" [43].

Around the same time as Super-K, another experiment was in the process of observing neutrino oscillations with solar neutrinos rather than atmospheric. The Sudbury Neutrino Observatory (SNO) was located 6010 m underground in the Creighton Mine in Sudbury, Canada [46]. The detector was composed of an acrylic sphere 12 m in diameter filled with 1 kiloton of heavy water (D_2O). It was surrounded by a shield of pure water contained in a 17.8 m diameter stainless steel structure, with an imaging system composed of over 9000 PMTs.

By measuring charged current, neutral current, and elastic scattering events, SNO provided direct evidence of neutrino oscillations independently of solar flux models [47]. In 2002, the experiment published oscillation results consistent with those of Super-K. They also showed that there was a non-electron contribution to the solar neutrino flux, providing a direct solution to the solar neutrino problem from nearly 30 years earlier. Results from the Super-K and SNO experiments can be seen in Figure 1.7 and Figure 1.8, respectively.

Figure 1.7 shows the arrival angle of the neutrino where the blue line represents theoretical models with *no* oscillations and red is *with* oscillations. The data matches the oscillation prediction much better. Figure 1.8 shows a non-zero contribution from ν_μ s and ν_τ s coming from the sun, proving that the solar neutrinos had oscillated into those flavors. The discovery and verification of neutrino oscillations yielded the 2015 Nobel Prize in physics to Super-K's Takaaki Kajita and Arthur McDonald, director of the SNO experiment.

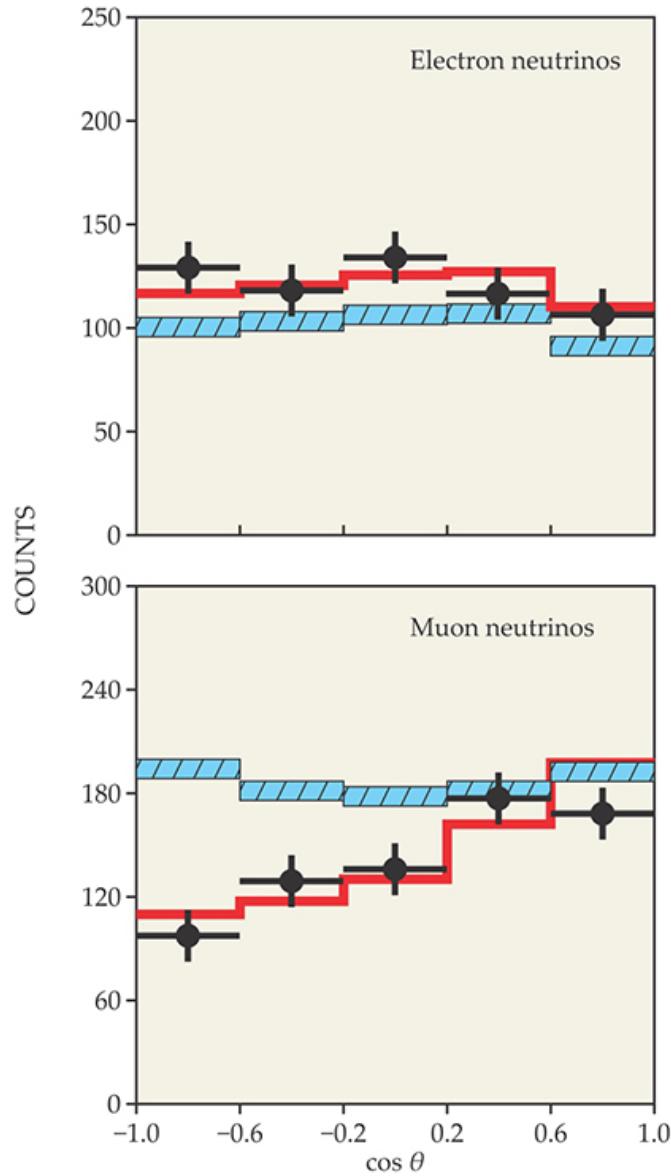


Figure 1.7: Results of the Super-K experiment proving neutrino oscillations [43].

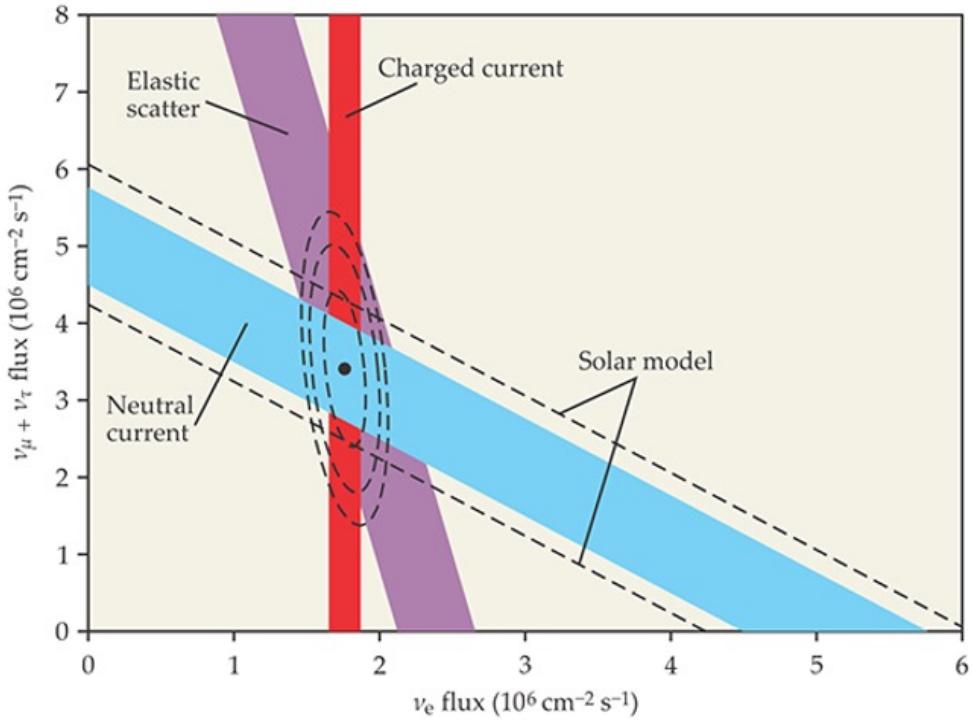


Figure 1.8: Results of the SNO experiment verifying the existence of non- ν_e solar neutrinos [43].

1.3.4 Interpreting Neutrino Oscillations

Since neutrinos have left-handed helicity, there is no Standard Model mechanism for them to acquire mass. However, because they *do* oscillate between their three flavor states as they travel, they *must* have non-zero mass states. The fact that neutrinos have mass makes them unique among Standard Model particles, and studying them can provide insight into physics beyond the Standard Model (commonly referred to as BSM physics).

To truly understand neutrinos and the oscillatory phenomenon they exhibit, one must precisely know the values present in the PMNS matrix. These are three mixing angles (θ_{12} , θ_{13} , θ_{23}) and a CP-violating phase (δ_{CP}). Additionally, the masses of the neutrinos themselves, the flavor composition of their mass states, and the “hierarchy” (that is, which neutrino mass state is heaviest, and which is lightest?) of the three mass states is still unknown, as shown in Figure 1.9. Measuring their oscillation parameters will help determine these properties.

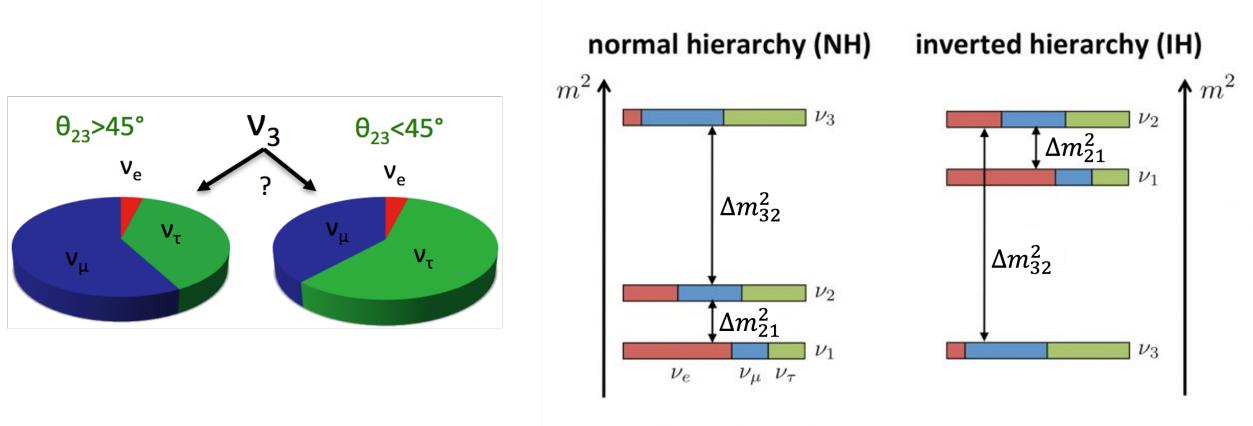


Figure 1.9: The effect of the mixing angles on mass ordering and possible mass hierarchies [48].

The oscillation probability generalized in Equation 1.7 depends on the neutrino mass splitting Δm_{ij}^2 , the neutrino's energy E , and the distance it travels L . An example commonly used to simplify this concept is the two-flavor approximation of the probability for an electron neutrino to oscillate into a muon neutrino [24], given by

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2(2\theta) \sin^2 \left(\Delta m^2 \frac{L}{4E} \right). \quad (1.8)$$

While the two-flavor approximation is not exact, it is a good representation of the oscillation phenomenon. Specific values of θ and Δm^2 depend on the L/E ratio being considered. For example, if L/E is very large, then the L/E term will yield a high oscillation frequency. The smaller of the three mass state spacings Δm_{21}^2 “slows” the frequency, providing a better opportunity of observing an oscillation under these conditions. Therefore, experiments dealing with high L/E are more sensitive to the Δm_{21}^2 mass state and its associated mixing angle, θ_{12} . Oscillation experiments are designed to target values of L/E that maximize the probability for certain oscillations to occur and provide sensitivity to different elements of the PMNS matrix.

A full three-flavor treatment with current best-fit values of the oscillation parameters yield the ν_e oscillation probabilities shown in Figure 1.10 [49].

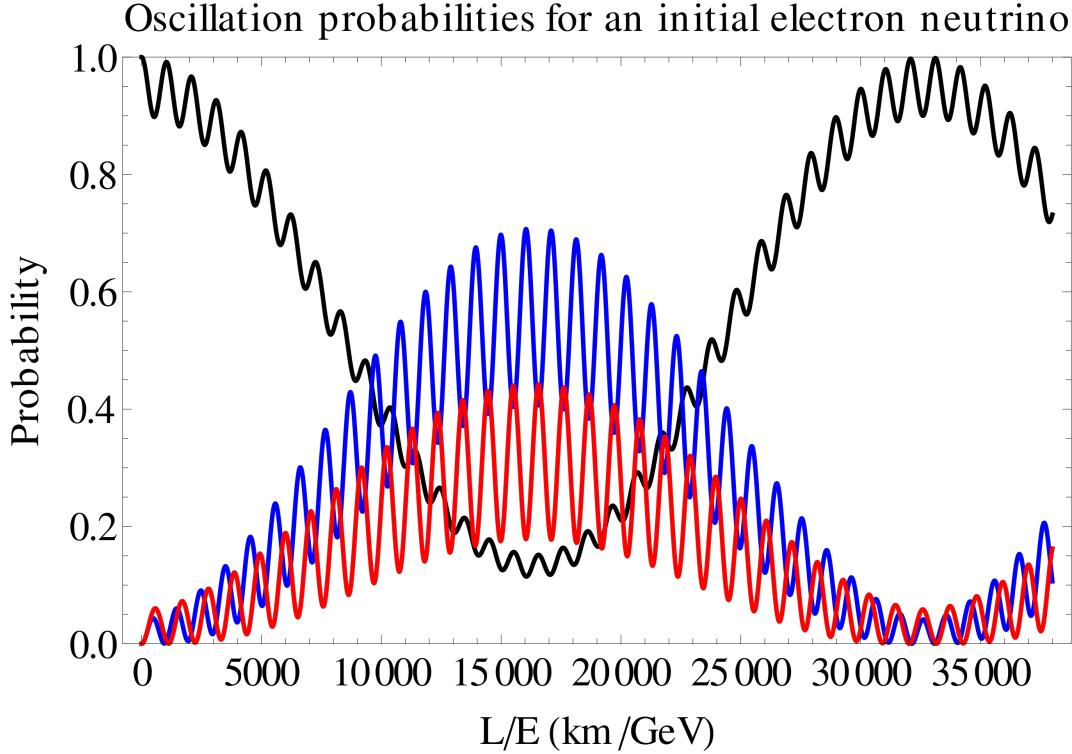


Figure 1.10: Neutrino oscillation probability as a function of L/E , starting with an electron neutrino. Blue = ν_μ , Red = ν_τ , Black = ν_e [49].

Oscillations in Matter

The mixing formalism works well for neutrinos as they travel in a vacuum. However, when neutrinos propagate through matter there are small fluctuations to their travel because of the potential to interact within the matter. These matter fluctuations require an additional potential to be added to the neutrino and antineutrino Hamiltonians. The extra term takes on the form $V = \pm \sqrt{2}G_F N_e$, where G_F is the Fermi constant (see section 2.1), and N_e is the electron number density of the matter [24].

The matter fluctuations were initially proposed by Lincoln Wolfenstein and later clarified and made rigorous by Stanislav Mikheyev and Alexei Smirnov, leading to this phenomenon being called the MSW effect [50]. The MSW effect changes the effective masses and mixing angles for neutrinos traveling in matter, which changes their oscillation probabilities. It has opposite affects on the oscillation of neutrinos and antineutrinos, since they always travel through regular matter.

Mass States, Mass Ordering, and CP Violation

By determining the oscillation parameters, some underlying symmetry may be revealed in the mixing states - such a discovery could shed new light on models of GUT-scale physics and dark matter [51]. Additionally, the ordering of the neutrino mass states could provide insight into mass generation mechanisms not currently considered in the Standard Model. There are many models of neutrino mass generation, and measuring the oscillation parameters directly related to their masses will help theorists converge on an explanation [52].

Measuring neutrino oscillations provides the possibility of discovering CP-violation in the neutrino sector, which may tell us about the origins of our existence in the universe [53]. The observable universe is primarily made up of matter, with naturally occurring anti-matter being nearly non-existent. However, at the beginning of the universe both matter and anti-matter should have been created in equal amounts. There must be some mechanism present in the universe that favors matter over antimatter; one leading theory is called “leptogenesis”, which involves CP-violating processes among neutrinos and other leptons [54].

1.4 Experiments Across the Energy Spectrum

The energy of a neutrino depends on its source, as shown in Figure 1.11, so that’s how modern neutrino experiments are often categorized. The most common sources of neutrino are the Sun, nuclear reactors, cosmic ray interactions in the atmosphere, and accelerator beams. This variety of sources allows experiments to span a wide range of neutrino energies and optimize their sensitivity to different regions of the PMNS matrix [55].

Solar Neutrino Experiments

Their long baselines and relatively small neutrino energies make solar neutrinos ideal for measuring the oscillation parameters Δm_{21}^2 mass state and its associated mixing angle, θ_{12} (as used in the example from subsection 1.3.4). The Homestake and SNO experiments were two significant observers of solar neutrinos, but there have been a number of other solar neutrino-based oscillation measurements. The SAGE [56] and GALLEX [57] experiments utilized Gallium to measure the $\nu_e + {}^{71}\text{Ga} \rightarrow e^- + {}^{71}\text{Ge}$ reaction with neutrinos in the MeV range. These experiments contributed

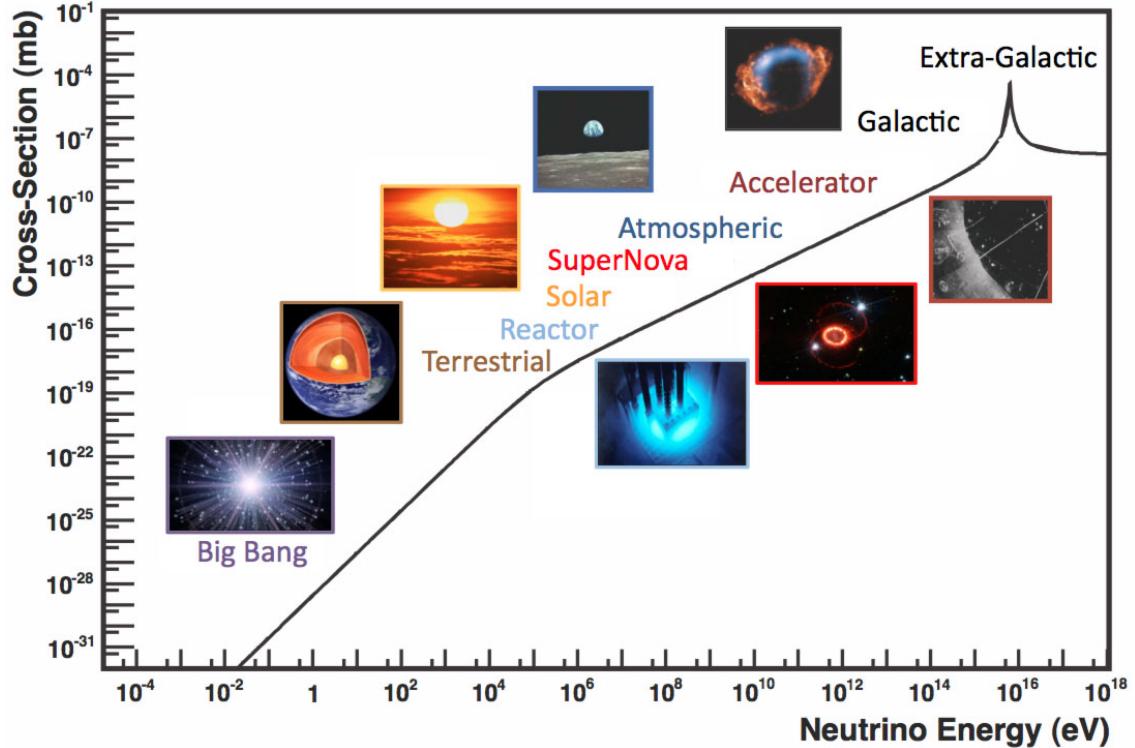


Figure 1.11: The range of possible neutrino energies and the sources that produce them [55].

additional observations of the solar neutrino problem initially found by the Homestake experiment. Another notable solar neutrino experiment is Borexino [58], which used ultra-pure liquid scintillator to make the first observation of neutrinos produced by the carbon-nitrogen-oxygen (CNO) cycle in the sun.

Reactor Neutrino Experiments

The most well-studied energy range of neutrinos comes from man-made nuclear reactors. Reactors are an incredibly intense source of electron anti-neutrinos in the single MeV energy region; it is estimated that during the fission of ^{235}U , ^{238}U , ^{239}Pu , and ^{241}Pu , reactors output $\bar{\nu}_e$ s at a rate of $2 \times 10^{20} \bar{\nu}_e$ s per second per GW of thermal power generated [55]. The $\bar{\nu}_e$ s produced in these facilities do not have the energy required to produce leptons heavier than electrons, so charged current interactions cannot happen if a reactor $\bar{\nu}_e$ changes its flavor to a $\bar{\nu}_\mu$ or $\bar{\nu}_\tau$. Therefore, reactor experiments must measure neutrino oscillation through the $\bar{\nu}_e$ disappearance channel facilitated by the inverse beta decay process $\bar{\nu}_e + p \rightarrow e^+ + n$ - the same channel originally probed by Reines and Cowan to prove the existence of the neutrino.

A few experiments that have contributed to measurements of oscillation parameters using reactor neutrinos are KamLAND in Japan [59], Double Chooz in France [60], Daya Bay in China [61], and RENO in Korea [62]. The detectors employed by each of these experiments consist of liquid scintillator surrounded by PMTs and shielded from outside radiation. Using multiple detectors with baselines (the distance between reactor and detector) on the order of 1 - 100 km each, Double Chooz, Daya Bay, and RENO precisely measured $\sin^2(\theta_{13})$ and $\Delta m_{ee}^2 \equiv \cos^2(\theta_{12})\Delta m_{31}^2 + \sin^2(\theta_{12})\Delta m_{32}^2$. The mixing angle measurements made by these experiments is shown in Figure 1.12 [63].

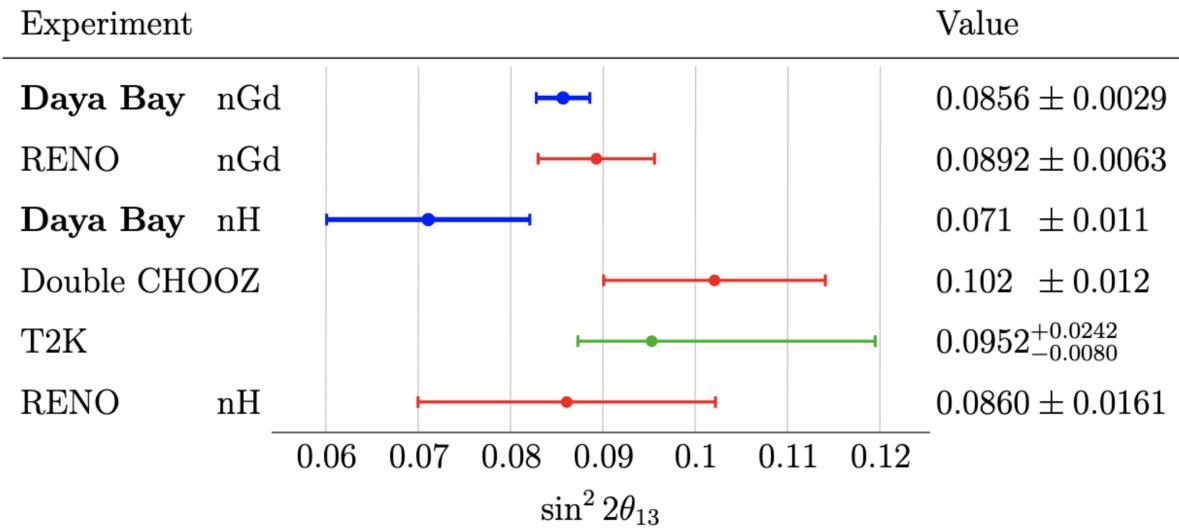


Figure 1.12: Combined $\sin^2(\theta_{13})$ results from the reactor experiments as of 2021 [63].

Atmospheric Neutrino Experiments

Atmospheric neutrinos have a slightly higher energy range than those produced by solar and reactor processes - commonly measured in the hundreds of MeVs, but potentially ranging much higher [64]. They are produced by the decay of pions and kaons generated by cosmic rays interacting with nucleons in the Earth's atmosphere. Super-K utilized atmospheric neutrinos to prove the theory of neutrino oscillations and is still running to this day. In recent years, the IceCube/DeepCore [65] neutrino experiment has made observations of muon neutrino disappearance using atmospheric neutrinos. This observation allowed them to make a measurement of the oscillation parameters Δm_{23}^2 and $\sin^2(\theta_{23})$. Both Super-K and IceCube have recently made observations of tau neutrinos [66, 67, 68], the “least studied particle” in the Standard Model [69].

Accelerator Neutrino Experiments

The last major neutrino source is particle accelerators. They yield neutrinos in some of the highest energy ranges commonly studied, with energies ranging from single MeV to hundreds of GeV. Accelerator neutrinos are generated by accelerating protons in a ring and allowing them to smash into some target. That reaction produces a large number of pions and kaons, which are focused into a beam with magnetic horns and allowed to decay into muons and muon neutrinos via the decay channels $\pi^{+(-)} \rightarrow \mu^{+(-)} + \nu_\mu(\bar{\nu}_\mu)$ and $K^{+(-)} \rightarrow \mu^{+(-)} + \nu_\mu(\bar{\nu}_\mu)$ [70]. Finally, shielding is placed at the end of the beam line to capture all of the mesons, muons, and protons left over, making the final product a relatively pure beam containing only the desired neutrinos.

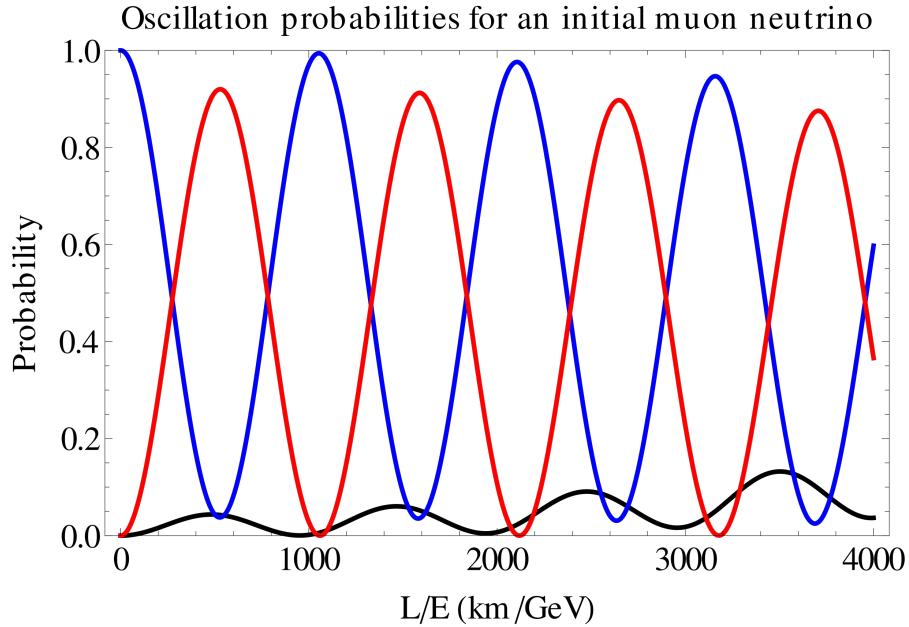


Figure 1.13: Neutrino oscillation probability as a function of L/E , starting with a muon neutrino. Blue = ν_μ , Red = ν_τ , Black = ν_e . [49]

The energy of a neutrino resulting from pion decay is

$$E_\nu \simeq \frac{[1 - (m_\mu/m_\pi)^2]E_\pi}{1 + \gamma^2\theta^2}, \quad (1.9)$$

where E_ν and E_π are the neutrino and pion energies, m_μ and m_π are the muon and pion masses, θ is the angle between the pion and neutrino direction, and $\gamma = E_\pi/m_\pi$. Knowing the energy of the neutrinos allows experiments to fine tune their baseline to maximize the oscillation probability.

Accelerator experiments vary in the types of detectors used, but the general principle is the same. Near detectors (NDs) are placed relatively close to the point of beam generation – within a few hundred meters. ND technology is typically the same technology used at the far detectors (FDs) of the experiments and are useful in reducing systematic uncertainties [55]. They allow experiments to measure the flux, energy spectrum, and interaction cross sections of their neutrinos very precisely. The far detectors employed by these experiments are typically much larger, with baselines ranging from hundreds to thousands of kilometers long. Their goal is to observe the oscillated electron neutrinos and surviving muon neutrinos.

The first long-baseline accelerator neutrino experiment started in 1999 - the Japanese KEK to Kamioka (K2K) experiment [71]. It measured a neutrino beam with an average energy of 1.3 GeV at a baseline of 250 km. It had two near detectors to monitor pion and muon kinematics in the beamline, and it utilized the water-Cherenkov Super-K detector as an FD. The Super-K detector was also used for the descendent of K2K called the Tokai to Kamioka (T2K) experiment, which began operation in 2010 [72]. It uses an upgraded proton synchrotron facility called J-PARC to produce its neutrino beam. T2K was the first off-axis neutrino experiment - this means the detector itself was placed slightly away from the main axis of the neutrino beam. Referencing the θ in Equation 1.9, being off-axis allows T2K to fine-tune the energy of the neutrinos they receive. Utilizing this feature, T2K receives neutrinos with an energy peaked at about 0.6 GeV.

The first long-baseline experiment in the United States was the Main Injector Neutrino Oscillation Search (MINOS), which started in 2005 [73]. It utilized the Neutrinos at the Main Injector (NuMI) beamline [74] to send muon neutrinos from Fermilab in Chicago, IL, to the Soudan Mine in Soudan, MN. The MINOS detectors were composed of alternating planes of steel and plastic scintillator with fiber optic cables read out by PMTs to capture interactions. The experiment had a baseline of 735 km with a neutrino beam energy peaked at about 3 GeV.

Since 2014 Fermilab's NuMI Off-Axis ν_e Appearance (NOvA) experiment has been the flagship long baseline experiment in the United States [75]. Like MINOS, it uses the NuMI beamline. However, it is placed off-axis from the beamline so that it has a neutrino energy spectrum peaked around 2 GeV. The NOvA detectors are made up of alternating planes of PVC extrusions containing liquid scintillator and read out by wavelength shifting fibers attached to avalanche photodiodes. The NOvA ND is located at Fermilab, and the FD is 810 km away near Ash River, MN.

By observing rates of muon-neutrino disappearance and electron-neutrino appearance, these accelerator experiments have measured the oscillation parameters Δm_{32}^2 and $\sin^2(\theta_{23})$, and have narrowed down the allowed region of δ_{CP} . The combined parameter results of all of the long baseline experiments discussed are visible in Figure 1.14, and the current state of the neutrino oscillation field is summarized in Table 1.1.

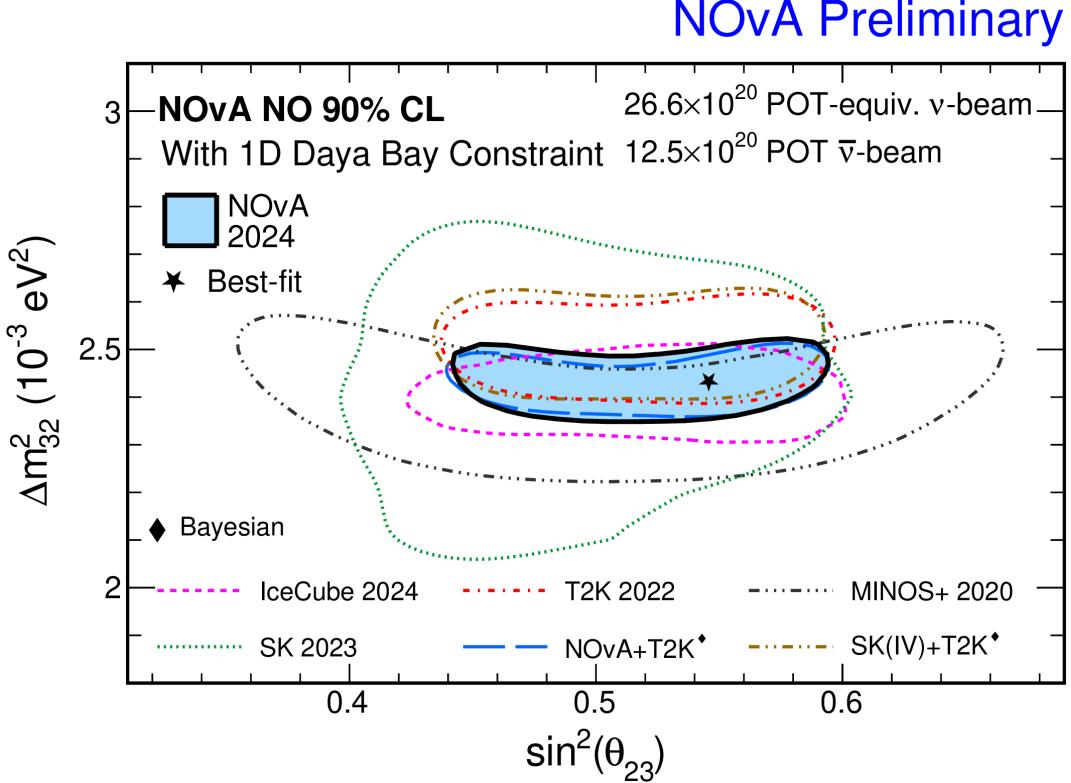


Figure 1.14: Combined oscillation results of the long baseline experiments [75].

PMNS Matrix Parameters		
Parameter	Best Fit $\pm 1\sigma$	Primary measurement techniques
θ_{12}	$34.3^\circ \pm 1.0^\circ$	Solar ν_e & Reactor $\bar{\nu}_e$
θ_{23}	$49.26^\circ \pm 0.79^\circ$	Accelerator ν_μ & Atmospheric ν_μ
θ_{13}	$8.58^\circ \pm 0.13^\circ$	Accelerator ν_e & Reactor $\bar{\nu}_e$
Δm_{21}^2	$2.45 \pm 0.21 (10^{-5} \text{eV}^2)$	Reactor $\bar{\nu}_e$
$\Delta m_{31}^2, \Delta m_{32}^2$	$2.5 \pm 0.1 (10^{-3} \text{eV}^2)$	Accelerator ν_μ & Reactor $\bar{\nu}_e$
δ_{CP}	Not Well Known	Future accelerator $\nu_e, \bar{\nu}_e$

Table 1.1: Neutrino oscillation parameters determined from 2020 global analysis [76].

1.5 Conclusion & Looking Ahead

Neutrinos are interesting to study for a number of reasons beyond learning their oscillation parameters and probing the weak force. Unfortunately, the whole gamut of neutrino experiments is too large to summarize in this literature review - there are experiments dedicated to making direct measurements of neutrino mass, searching for sterile neutrinos, and determining if neutrinos are their own anti-particles.

The remainder of this dissertation is focused on the search for neutrino trident production in NOvA. Trident production is a rare electroweak process that offers a unique probe of the Standard Model. In the chapters that follow, I present the theoretical background needed to understand the process, the motivation for studying trident production, and the Standard Model predictions. I then describe the NOvA experiment, and the ground-up development of the analysis required to perform this search. This work is the first dedicated search for trident production in NOvA and contributes to the experimental techniques required for rare process searches. It is also my hope that this search helps renew interest in investigating a process that has only been observed a handful of times, with the most recent measurement occurring more than two decades ago.

Chapter 2

Neutrino Interactions

This chapter provides the theoretical background necessary to understand neutrino interactions and cross sections. Neutrino trident production is a specific scattering process, so this chapter lays the foundation needed to understand how trident production fits into the broader framework of neutrino physics. I describe the theory of neutrino interactions in the Standard Model and how the theories were verified experimentally. I then discuss some of the interactions and modeling relevant to the NOvA experiment.

2.1 Interaction Theory

Interactions between neutrinos and other matter are a crucial part of the particle physics landscape. Studying their interactions improves the knowledge of nuclear forces and helps to lower experimental uncertainties. Neutrinos only interact via the weak nuclear force, so studying their interactions allows us to test theories within the Standard Model. This section introduces the theoretical framework and terminology used in discussing neutrino interactions.

2.1.1 Cross Sections

The primary method of understanding neutrino-nucleus interactions is comparing interaction “cross sections” predicted by theoretical models to those that are experimentally measured. In physics, a cross section is a measure of the probability that two particles will interact with each other [77]. Cross section measurements are not exclusive to neutrinos, they can be done for any pair of interacting objects. An interaction’s cross section describes the probability that some incoming particle will scatter off of a stationary target. In the case of two hard spheres (e.g., pool balls), the interaction probability is directly related to the geometric area of the spheres themselves. The effective area in which the collision can occur is the cross-sectional area (or just “cross section”) of that collision.

Cross sections are simple to measure for classically interacting objects such as pool balls, but become very complicated for small particles that interact via long-range forces. Often times, measuring an interaction cross section requires an experiment involving a large flux of incoming particles sent towards a detector filled with stationary targets. Scattered particles from the beam will change direction and be observed, allowing for a total cross section measurement of that interaction:

$$\sigma = \frac{\text{Number of Scattered Particles}}{\text{Beam Flux} \times \text{Number of Target Particles}}, \quad (2.1)$$

where σ is the total cross section of the interaction, measured in units of area.

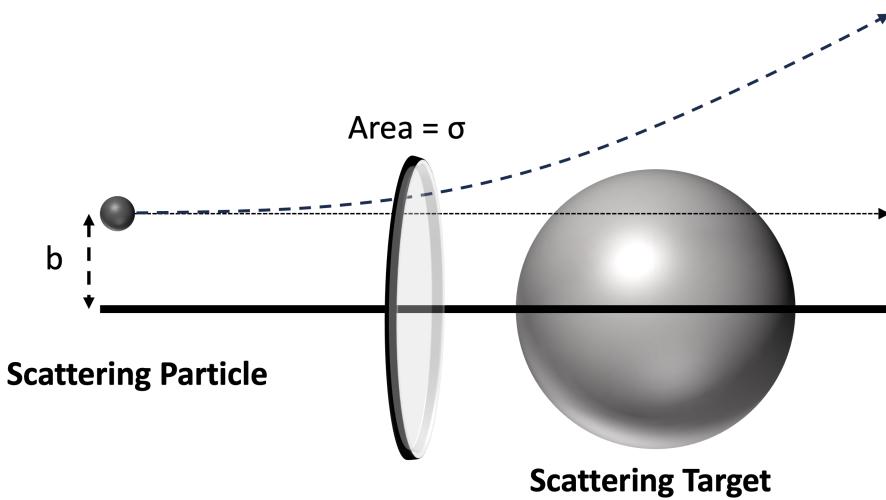


Figure 2.1: A simple diagram of a two-particle scattering process. One smaller particle scatters off of a single large target via a long-range force. The cross section σ of the interaction depends on the impact parameter b .

Despite the fact that cross sections have units of area, neutrino interactions only occur via weak channels, making the targets “transparent” to them. This means that neutrino cross sections are generally much smaller than the actual “size” of the particles involved. Of course, cross sections can depend on a number of variables other than the size of the particles, such as the energy of the incoming particle or the angle that it scatters at. Cross sections that are specified in terms of another variable are called *differential* cross sections.

2.1.2 Calculating Cross Sections

Calculating neutrino cross sections utilizes elements of the Standard Model such as the Feynman Rules to perform unique calculations for each interaction. I outline the general procedure for calculating cross sections, introduce the theory required for cross sections involving neutrinos, and demonstrate the process with the most basic CC neutrino interaction: inverse muon decay.

The Golden Rule for Scattering

The standard process by which cross sections are calculated was dubbed the “Golden Rule” by Enrico Fermi in 1950 [78]. Based on a quantum theory of radiative emission developed by Paul Dirac in the 1920’s, the process involves integrating over all possible combinations of four-momenta P_i that fulfill kinematical restraints designated by the requirements that energy and momentum be conserved in an interaction [79].

For incoming particles 1 and 2 colliding to produce outgoing particles 3, 4, . . . , n :

$$\sigma = \frac{1}{4\sqrt{(P_1 \cdot P_2)^2 - m_1^2 m_2^2}} \int |\mathcal{M}|^2 (2\pi)^4 \delta^4(P_1 + P_2 - P_3 - \dots - P_n) \prod_{i=3}^n \frac{d^3 \vec{p}_i}{2(2\pi)^3 E_i}, \quad (2.2)$$

where $E_i = \sqrt{|\vec{p}_i|^2 + m_i^2}$ is a function of the integration variable \vec{p}_i [80, 81]. The form of this integral is determined purely by the *possible* interactions that may occur for a certain energy and momentum. The dynamics of the interaction itself is wrapped up in the ‘invariant scattering amplitude’ \mathcal{M} of the interaction - this term may favor certain momentum combinations based on the interaction-mediating force. The amplitude for a certain neutrino scattering process is determined by applying the Feynman rules to the interaction with the vertex factors and propagators from the following section.

Electroweak Theory

The theory of electroweak interactions was developed by Glashow, Weinberg, and Salam (GWS) in the late 1950s [82, 83, 84]. Additionally, much of the information in this section is consolidated from a variety of sources created by theorists and fellow PhD students. [80, 81, 85, 86, 87, 88, 89, 90, 91].

Neutrinos are electrically neutral, have small mass, and only interact via weak channels; therefore, the study of neutrinos is intertwined with the development of the weak force theory. Neutrino interactions (in fact, all weak interactions) are mediated by two massive gauge bosons: the charged W-boson and the neutral Z-boson. When a neutrino interacts via W-boson exchange, it turns into a lepton of the corresponding flavor - these are called charged current (CC) interactions. When the neutrino exchanges a Z-boson, it continues on after the interaction unchanged, making the flavor of neutrino unidentifiable - these are called neutral current (NC) interactions. The GWS model contains a “weak mixing angle”, $\theta_w = 28.75^\circ$, which relates properties of the Z and W bosons. Feynman diagrams of CC and NC neutrino interaction vertices are visible in Figure 2.2

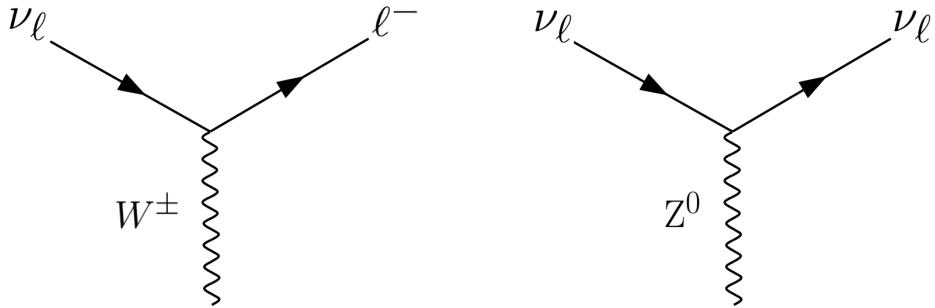


Figure 2.2: Feynman diagram vertices for neutrino-induced weak interactions showing the charged current interaction vertex (left) and the neutral current interaction vertex (right). Here ν_ℓ corresponds to one of the neutrino flavors, with ℓ being the associated charged lepton.

Following the Feynman rules for interaction calculations requires vertex factors and propagators associated with the weak bosons. The vertex factor for charged current interactions is

$$\frac{-ig_w}{2\sqrt{2}} \gamma^\mu (1 - \gamma^5), \quad (2.3)$$

where the weak coupling constant $g_w = \sqrt{4\pi\alpha_w} = 0.653$. Likewise, the vertex factor for neutral current interactions is

$$\frac{-ig_z}{2} \gamma^\mu (c_V^f - c_A^f \gamma^5), \quad (2.4)$$

where the additional coefficients c_V^f and c_A^f depend on which quarks or leptons are involved in the interaction, and $g_z = g_w / \cos(\theta_w)$. The existence of both γ^μ and $\gamma^\mu\gamma^5$ terms indicate that such interaction vertices have both vector and axial vector couplings - the theory verified by Dr. Wu in the parity-violation experiment described in subsection 1.2.1.

The propagators for each boson are

$$\frac{-i(g_{\mu\nu} - \frac{q_\mu q_\nu}{M^2})}{q^2 - M^2}, \quad (2.5)$$

where q^2 is the momentum transfer of the interaction, and M^2 is the mass of the relevant boson, $M_W = 82$ GeV and $M_Z = M_W / \cos(\theta_w) = 92$ GeV. In practice, accelerator neutrino experiments typically run at neutrino energies of $E_\nu < 100$ GeV, meaning the allowable momentum transfer in an interaction is $q^2 \ll M$. This means the propagators simplify to the more commonly used form

$$\frac{i(g_{\mu\nu})}{M^2}. \quad (2.6)$$

Experimental Verification of The GSW Theory

One NC interaction predicted in the GSW theory is neutrino-electron scattering: $\nu_\ell + e^- \rightarrow \nu_\ell + e^-$. It was the observation of this interaction in the early 1970s by CERN's Gargamelle experiment [92] that verified the GSW theory. In the experiment, neutrinos of energy 1-10 GeV produced by a proton synchrotron were directed towards the Gargamelle detector - a cylindrical “bubble chamber” measuring 4.8 m long and 2 m in diameter. The chamber weighed 1000 tons and held nearly 12 cubic meters of heavy liquid freon. The detector was inside a 2 T magnetic field, meaning that the charges of any particles within the chamber could be determined from their curvature. When a neutrino interacted in the freon the resulting particles would leave bubbles in the shape of tracks behind to be photographed.

Gargamelle observed NC events through leptonic and hadronic channels. The results were presented in July 1973 as the first direct evidence of the weak neutral current, and therefore the Z^0 boson. A NC ν_e event occurring in Gargamelle is pictured in Figure 2.3, showing a vertex originating in the detector followed by an electromagnetic shower. The formulation of electroweak theory resulted in Glashow, Salam, and Weinberg sharing the 1979 Nobel Prize in physics.



Figure 2.3: An event display of the Gargamelle bubble chamber featuring a NC event [93].

Fermi's Coupling Constant for Weak Interactions

Fermi's original beta decay theory from 1934 did not consider a W boson propagating the interaction, but rather a direct four-particle coupling between fermions [78]. He introduced a coupling constant G_F to describe the strength of the interaction. Years later, calculations of the invariant amplitude for beta decay within the GWS theory following modern Feynman rules showed that the two vertex factors and the low-energy boson propagator term combine to yield a factor with the form

$$\left(\frac{g_w}{M_W}\right)^2.$$

Fermi's theory provided a great approximation of weak interactions in the low energy regime (when $q^2 < 100$ GeV) by simply combining the propagator and vertex terms together. Therefore, it is still common to see calculations of weak interactions utilize the “shortcut” of combining the vertices and a weak boson propagator into a single point coupled by the ‘Fermi coupling constant’:

$$G_F = \frac{\sqrt{2}}{8} \left(\frac{g_w}{M_W}\right)^2 = 1.166 \times 10^{-5} \text{ GeV}^{-2}. \quad (2.7)$$

Figure 2.4 in the following section shows the shortcut applied to a Feynman diagram for one of the most theoretically straightforward weak processes - inverse muon decay.

Inverse Muon Decay

Inverse muon decay is a CC neutrino-electron scattering process which produces a muon in the final state: $\nu_\mu + e^- \rightarrow \nu_e + \mu^-$. It is one of the simplest processes to calculate the cross section of using GWS theory, making it a useful benchmark to measure neutrino properties. Inverse muon decay was used to confirm standard model predictions of weak current and the helicity of the neutrino in the 90's [94, 95], and has more recently been used in studies to constrain systematic uncertainties, such as the NuMI beam flux [96]. Figure 2.4 shows both the Standard Model and Fermi four-point interpretations of an inverse muon decay process, including the coupling terms associated with each theory.

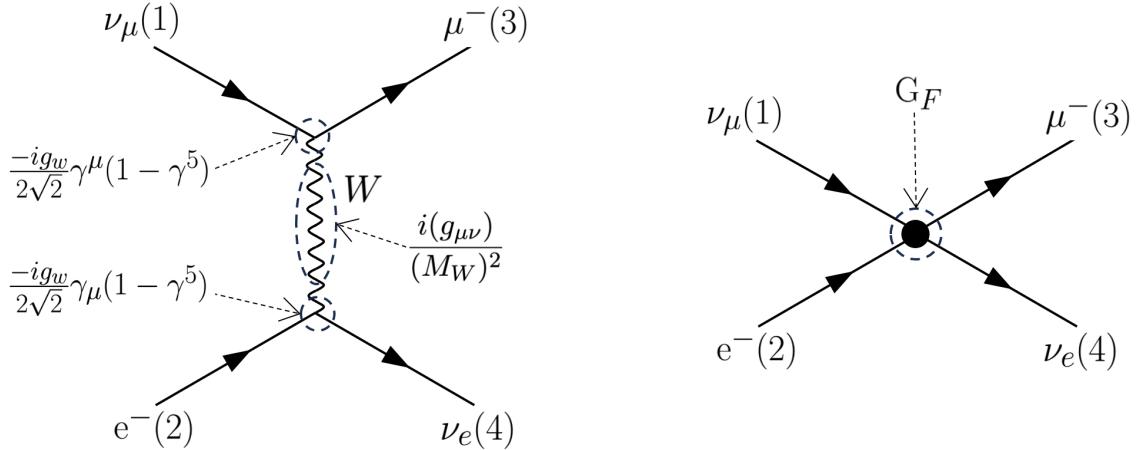


Figure 2.4: **Left:** The Standard Model representation of inverse muon decay, mediated by a charged weak boson. **Right:** Fermi's representation of inverse muon decay, in which the interaction occurs directly between fermions with coupling strength G_F .

All particles involved in this interaction are fermions, meaning they will be represented by dirac spinors u in the formation of the scattering amplitude. Stepping backwards along the interaction yields an amplitude [80]:

$$\mathcal{M} = \frac{g_w^2}{8M_W^2} [\bar{u}_3 \gamma^\mu (1 - \gamma^5) u_1] [\bar{u}_4 \gamma_\mu (1 - \gamma^5) u_2], \quad (2.8)$$

where the factor $\frac{g_w^2}{8M_W^2} = \frac{G_F}{\sqrt{2}}$. Calculating the magnitude in the center of mass frame while averaging over incoming spin states, summing over outgoing spin states, and requiring that the

neutrinos have only left-handed helicity yields

$$\langle |\mathcal{M}|^2 \rangle = 8 \left(\frac{g_w E_\nu}{M_W} \right)^4 \left[1 - \left(\frac{m_\mu}{2E_\nu} \right)^2 \right], \quad (2.9)$$

where E_ν is the energy of the incoming neutrino. This amplitude can then be plugged into Fermi's golden rule, resulting in an interaction cross section of

$$\sigma = \frac{1}{8\pi} \left(\frac{g_w^4 E_\nu^2}{M_W^4} \right) \left[1 - \left(\frac{m_\mu}{2E_\nu} \right)^2 \right]^2 = \frac{4G_F^2 E_\nu^2}{\pi} \left[1 - \frac{m_\mu^2}{4E_\nu^2} \right]^2. \quad (2.10)$$

The total cross section grows quadratically with the incoming neutrino energy, with an additional term governing the behavior of the cross section at energy scales relative to the muon mass. When E_ν approaches the threshold energy of $m_\mu/2$, the bracketed term drops to zero, making the interaction forbidden.

2.1.3 Scattering Kinematics

Events in NOvA [75] typically take the form of a neutrino interacting with a large, (mostly) stationary nucleus. Consider Figure 2.5 featuring a generic Feynman diagram in which an incoming neutrino with four-momentum P_ν scatters on a target nucleus of four-momentum P_T . The resulting particles of the interaction are some final-state lepton (which could be either a neutrino or a charged lepton) with four-momentum P_ℓ and some unspecified final-state particles (typically hadrons) with four-momentum P_H .

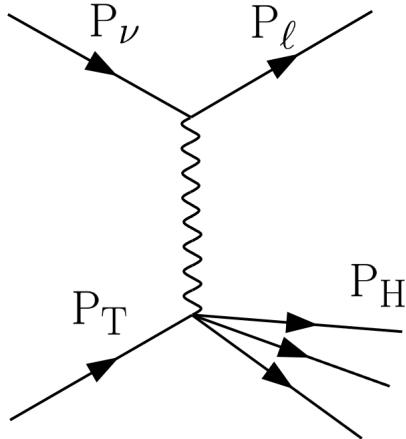


Figure 2.5: A generic Feynman diagram of a neutrino scattering off of a nucleus.

The components of each these quantities in the lab frame can be written as four-vectors:

$$P_\nu = (E_\nu, \vec{p}_\nu)$$

$$P_T = (E_T, \vec{p}_T)$$

$$P_\ell = (E_\ell, \vec{p}_\ell)$$

$$P_H = (E_H, \vec{p}_H)$$

which stick to the convention of the 0th element corresponding to the energy of the particle. Some important kinematic terms to be familiar with for these types of interactions are: four-momentum transfer, inelasticity, invariant mass, and the Bjorken scaling variable [90].

The four-momentum transferred between the neutrino-lepton system and the target is defined as q , from which a Lorentz-invariant momentum-transfer term $Q^2 = -q^2$ can be defined:

$$Q^2 = (P_\nu - P_\ell)^2 = (\vec{p}_\nu - \vec{p}_\ell)^2 - (E_\nu - E_\ell)^2. \quad (2.11)$$

Here the energy transfer of the lepton system is $(E_\nu - E_\ell) = E_H$, which is also the energy transferred to the hadronic output of the interaction. This momentum transfer Q^2 depends heavily on the energy of the incoming neutrino, and it determines the allowable final states of an interaction.

Another important Lorentz-invariant term is the fraction of energy transferred from the initial neutrino into the interaction, defined by inelasticity y :

$$y = \frac{E_\nu - E_\ell}{E_\nu} = \frac{E_H}{E_\nu}. \quad (2.12)$$

High inelasticity ($y \approx 1$) is an indicator that the neutrino transferred a large amount of energy into the interaction, while low inelasticity ($y \approx 0$) means that the neutrino kept most of its energy.

When dealing with higher-energy neutrino interactions where neutrinos can interact directly with nucleons, quarks, and gluons, the Bjorken scaling variable, x , becomes relevant [97]. This variable describes how strongly interacting particles behave as a collection of point particles when probed at high energies.

$$x = \frac{Q^2}{2P_T \cdot q} = \frac{Q^2}{2M_T E_H} \quad (2.13)$$

where M_T is the mass of the target nucleus. The size of x ranges from 0 to 1, where large values indicate that a single quark or gluon carrying a large fraction of the nucleon's momentum was “struck”. Low values of x indicate that the neutrino interacted with the nucleon as a whole.

Lastly, the total Lorentz-invariant mass of the outgoing final state particles is represented by W , expressed as

$$W^2 = (P_T + q)^2 = M_T^2 + 2M_T E_H - Q^2. \quad (2.14)$$

W and x are convenient variables because they can be directly measured from the final-state particles of high energy interactions, which are easier to observe in detectors.

2.2 Neutrino Interactions in the Standard Model

With the theoretical foundations laid I can now discuss some specific neutrino interactions relevant to the NOvA experiment. In particular the following interactions are *backgrounds* to the neutrino trident search, as the theory for trident interactions themselves is the topic of chapter 3.

The NuMI beam generates a beam of neutrinos peaked at 2 GeV, but does contain neutrinos with energies ranging from 0.5 GeV up to 120 GeV. Therefore, a large range of interaction types can occur in the NOvA near detector (ND). The type of interaction depends on the energy of the incoming neutrino and the target it hits. Higher energy neutrinos are able to interact at smaller length scales, which range from entire nuclei down to individual quarks within a nucleon.

2.2.1 Neutrino-Nucleus (Coherent) Interactions

The first interaction type of interest is coherent neutrino scattering. Coherent scattering is a process in which the incoming neutrino scatters off of the entire nucleus as a whole. A small amount of energy is exchanged between the neutrino and the nucleus causing the nucleus to recoil, but remain otherwise unchanged from its initial state. A meson is emitted in the final state of coherent interactions, most commonly a very forward-going pion. Coherent processes can occur via both charged and neutral current channels, with the pion production processes looking like

$$\begin{aligned} \nu_\ell + A &\rightarrow \ell^- + A + \pi^+ \quad (\text{CC}), \\ \nu_\ell + A &\rightarrow \nu_\ell + A + \pi^0 \quad (\text{NC}). \end{aligned} \quad (2.15)$$

A Feynman diagram showing a CC coherent scattering process can be seen in Figure 2.6.

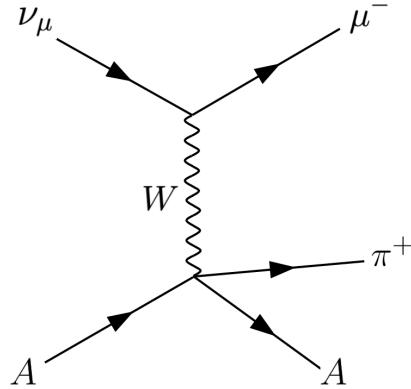


Figure 2.6: A diagram of a CC coherent process which leaves the nucleus unchanged but produces a π^+ in the final state.

Cross sections for coherent processes tend to be fairly small due to the requirement that the neutrinos have low energy. However, an important feature of these interactions is their unique dependence on atomic mass A [98]. Because the neutrino interacts directly with the nucleus, it's scattering amplitude \mathcal{M} will have A dependence, meaning coherent cross sections are proportional to A^2 : $\sigma_{\text{Coh}} \propto |\mathcal{M}|^2 \propto |A|^2$. This type of scattering is likely a significant background for the neutrino trident analysis presented in this thesis, as the primary trident production channel is via coherent scattering. Therefore, a non-trident CC coherent event with a highly energetic pion may look quite similar to a trident event in the detector.

2.2.2 Neutrino-Nucleon (Incoherent) Interactions

Coherent scattering processes describe interactions in which a neutrino scatters off of an entire nucleus because of its low energy - the remaining interaction types happen at higher values of Q^2 (and therefore, higher E_ν). At higher energies neutrinos can interact with individual *nucleons* within the target nucleus, a process known as incoherent scattering [87].

Quasi-Elastic (QE) Interactions

Quasi-Elastic (QE) interactions are the simplest of the four neutrino-nucleon interaction types. They occur when a neutrino scatters off of a nucleon, undergoing an inverse beta-decay process and ejecting the resulting nucleon intact from the target nucleus. QE interactions happen at fairly low energy ranges (on the order of 1 GeV), and they are called “quasi”-elastic because the neutrino can change into a charged lepton during this process - NC events of this nature are completely elastic.

Figure 2.7 illustrates a CC QE interaction between a muon neutrino and a neutron.

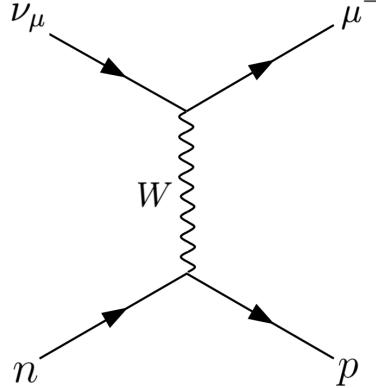


Figure 2.7: A Feynman diagram of a charged-current quasi-elastic interaction.

Despite the interaction looking simple, calculating cross sections for QE processes is quite difficult due to nuclear effects on the target and scattered hadron. Event generators currently implement the Llewellyn-Smith model of QE interactions, which considers neutrinos scattering off of free nucleons and utilizes various nuclear form factors to properly account for nuclear effects [99]. QE Interactions are relatively common to see in NOvA's neutrino energy range, with cross sections on the order of $\sigma \sim 10^{-37} \frac{\text{cm}^2}{\text{GeV}}$.

A useful outcome of measuring QE cross sections is that their two-body final state allows the interaction kinematics to be reconstructed completely, allowing for precise determination of the incoming neutrino energy. It is particularly important for the trident analysis to understand QE interactions as a potential background; in ν_μ CC QE interactions target neutrons are converted into protons, yielding a muon and a proton in the final state. Similar to coherent processes, this final state has the potential to mimic trident events, though it is less likely that a proton will travel far in the detector compared to an energetic pion.

Resonant (RES) Interactions

Resonant interactions require more energy than QE events, most commonly occurring in the energy range of 0.5 - 5 GeV. These types of interactions occur when the neutrino imparts enough energy to the target nucleon that it ends up in an excited state. This excited state produces a baryon resonance, with the primary interaction mode of a Δ resonance decaying into a pion and a nucleon. Higher energy resonances also have the capability to produce multiple pions or even kaons. RES

interactions can occur through either charged or neutral current channels, for example:

$$\begin{aligned} \nu_\ell + n &\rightarrow \ell^- + n + \pi^+ \quad (\text{CC}) \\ \nu_\ell + p &\rightarrow \nu_\ell + p + \pi^0 \quad (\text{NC}). \end{aligned} \quad (2.16)$$

Figure 2.8 shows the Feynman diagram of a CC RES interaction in which the resonance decays into a neutron and a charged pion. The current implementation of RES interactions in neutrino event generators utilizes a model developed by Dieter Rein and Lalit Sehgal [100]. The model accounts for the production of 18 resonances, and considers all combinations of neutrinos and anti-neutrinos scattering off of protons or neutrons via either CC or NC channels that obey charge conservation. Once again there are a handful of RES channels that yield two energetic charged particles in the final state, which have the ability to mimic trident events in the detector.

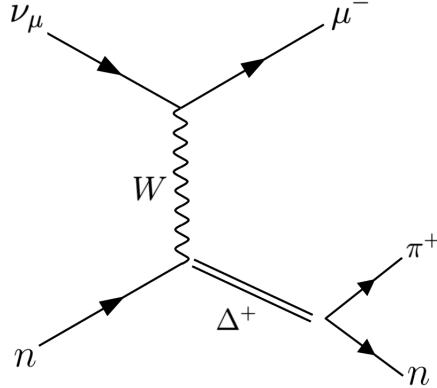


Figure 2.8: A Feynman diagram of a charged-current resonance interaction.

Deep Inelastic Scattering (DIS) Interactions

Deep inelastic scattering (DIS) interactions occur when the incoming neutrinos have significantly higher energies than the QE or RES processes. DIS interactions become the dominant channel in the NOvA experiment above about 10 GeV. In DIS processes the neutrino exchanges a W or Z boson with a singular quark within a nucleon, often causing an “explosion” of hadronic activity as the nucleus is blown apart.

These types of interactions in the high neutrino energy region have been used both as a probe of nuclear structure and as a validation of the Standard Model [101]. Though one would expect DIS events to typically result in a lot of hadronic energy, they do still have the ability to mimic trident events. In cases where the final state of a DIS interaction is a pair of highly energetic charged

particles and a neutron, the neutron could go undetected, leaving two long tracks in the detector. A Feynman diagram of a generic DIS process can be seen in Figure 2.9, with plots showing current cross section measurements of the various interaction types as a function of neutrino energy visible in Figure 2.10.

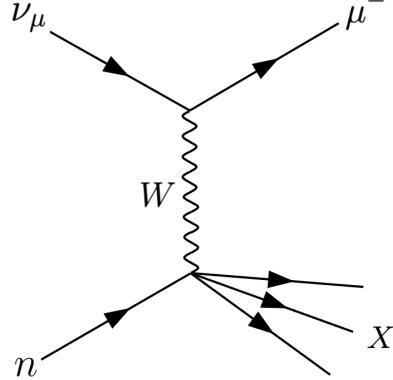


Figure 2.9: A Feynman diagram of a deep inelastic scattering interaction.

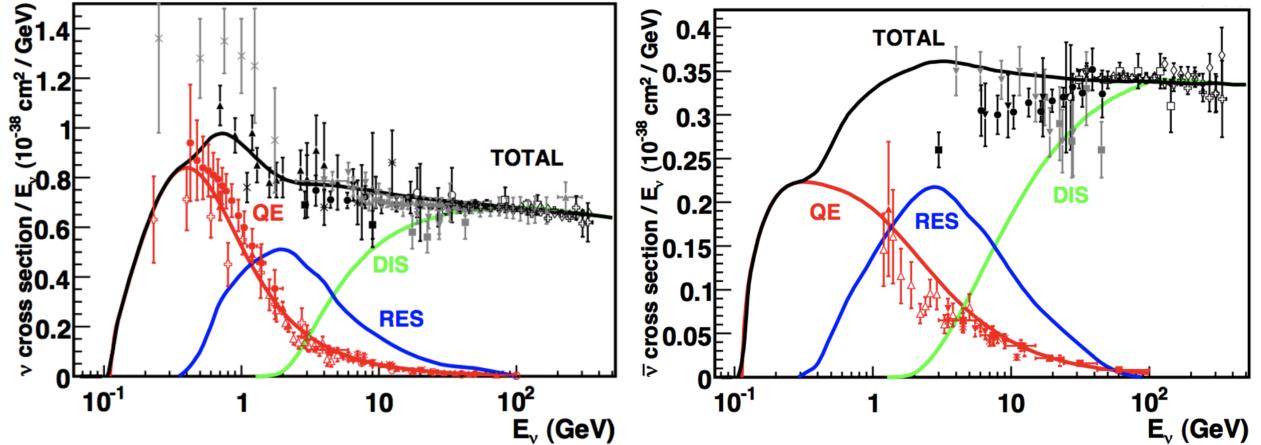


Figure 2.10: Cross sections of neutrino (left) and antineutrino (right) interactions with nucleons as a function of neutrino energy. The solid lines show predictions from models for QE, RES, and DIS interactions while the data points are taken from various experiments [101].

Meson Exchange Current (MEC)

The least well-studied type of neutrino interaction relevant to the NOvA experiment is Meson Exchange Current (MEC) - also known as multi-nucleon knockout or 2p2h interactions. These occur when the W boson is absorbed by two nucleons which are then knocked out of the nucleus,

leaving behind two holes. Frequently, these two nucleons interact with one another via some meson in the final state. The energy regime of this type of interaction lies just between QE and RES events, around 1 GeV of neutrino energy. There are a number of theoretical approaches to include this type of interaction in event generators [102], though more studies must be done to verify their validity. MEC events certainly contribute to neutrino-nucleus scattering, and the community agrees that understanding them is vital to explaining existing data.

2.2.3 Nuclear Effects in Incoherent Interactions

Modern neutrino experiments (including NOvA) use large detectors filled with a variety of target nuclei, so the neutrino-nucleon interactions described above may not always result in a clean output. Nuclear effects play a significant role in the final state of the interaction. These effects can impact nucleons either before, during, or after the interaction takes place. They can cause the outgoing kinematics of the event to change, or alter the outgoing hadronic products from the interaction as they propagate through the target nucleus. Poor understanding of nuclear effects is one of the primary causes of systematic uncertainty in cross section models - including those used by NOvA. This section is dedicated to discussing some of the intricacies of modeling inter-nucleus effects.

Fermi Gas Model

No two identical fermions can occupy the exact same quantum state according to the Pauli exclusion principle. Protons and neutrons inside of nuclei are fermions, which must follow Fermi-Dirac statistics. The statistical model used to describe nuclear states of this type is the Fermi Gas model: a framework which assumes that the nucleons within a nucleus move freely and are non-interacting, and that each energy state of the nucleus is filled sequentially. Thus, all nucleons bound within the nucleus occupy unique states in a Fermi-gas up to the maximum energy state, known as the “Fermi level” E_F [103]. The energy difference between the top of the nucleus’ potential well and the Fermi level of the gas is the “binding energy” (E_B) needed to pull a nucleon out of the nuclear potential - the principle illustrated in Figure 2.11. The Fermi Gas model can be extended to particles interacting at relativistic speeds (Relativistic Fermi Gas) and to consider the local density of nucleons within the nucleus affecting its binding potential (Local Fermi Gas).

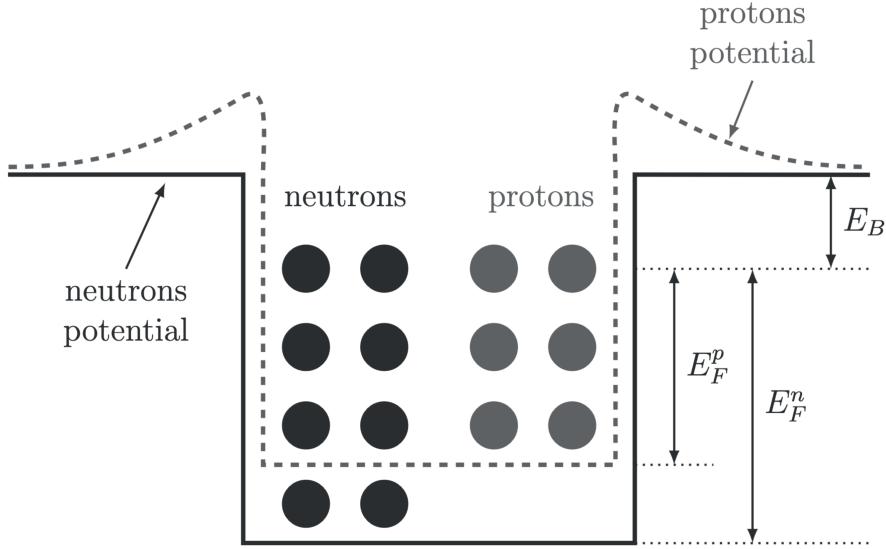


Figure 2.11: Nuclear potential wells for protons and neutrons bound within a nucleus of constant potential. E_F^p , E_F^n are Fermi energies of protons and neutrons, and E_B is binding energy [103].

Short Range Nucleon-Nucleon Correlations and Final State Interactions

As mentioned when discussing Fermi Gasses, many models do not consider interactions between nucleons within the nucleus. These must still be accounted for, as short-range correlations between nucleons can affect the momenta of those nucleons [104]. These short range correlations are primarily important to MEC interactions, which have not been particularly well studied (though models to describe multi-nucleon neutrino interactions have been proposed [105, 106]).

Neutrino-nucleus interactions can also produce a variety of non-nucleon final state particles including leptons like muons or electrons, and hadrons like pions. These ejected particles must escape the nuclear medium before they can be detected. However, it is possible for hadrons that have been knocked loose to interact again within the nucleus before they enter the detector. These final state interactions (FSI) can lead to altered kinematics, multiplicity, or charge of the final state seen in the detector - a process shown in Figure 2.12.

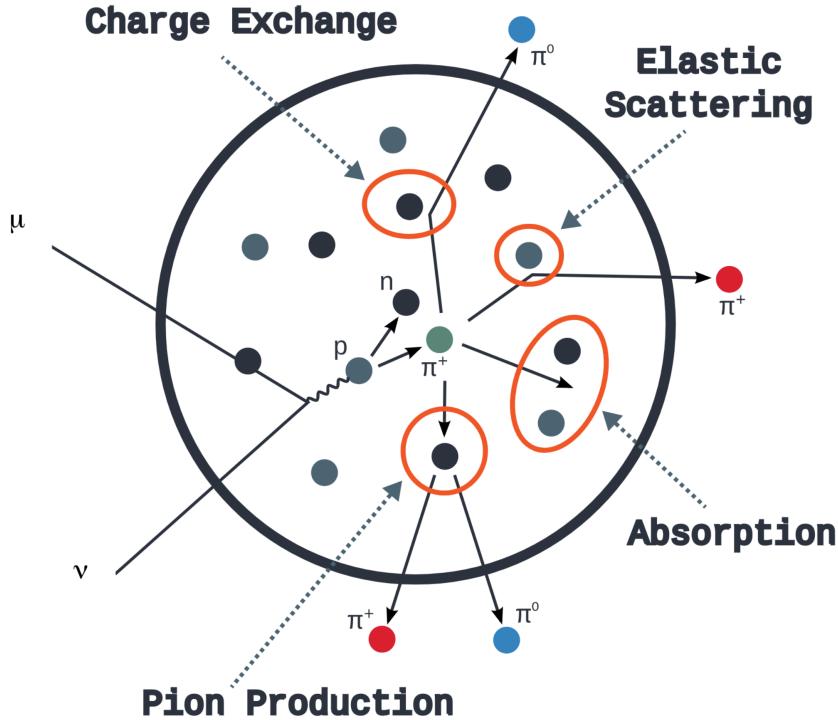


Figure 2.12: The hadronic shower produced in the initial reaction must still pass through the nucleus and is subject to FSI through a variety of processes before entering the detector [107].

Final state interactions are difficult to model properly, and they can impact the topology of events viewed in the detector. Since the energy of an incoming neutrino is based off of the reconstruction of that event, it means that the neutrino energy cannot be properly determined for events that undergo FSI. This presents a calorimetry problem for experiments to take into consideration when analyzing data. One example of such an issue may occur in the observation of what appears to be a charged current QE event. The final state of such an event would be a charged lepton and a proton being detected, as is typical for QE events. However, it is also possible that this final state could be produced by a resonance interaction in which the pion was absorbed within the nucleus, leaving only a proton behind. The possibility of FSI must be kept in mind when developing the signal criteria for an analysis.

2.3 Neutrino Scattering in Oscillation Analyses

Oscillation experiments are not able to directly measure the probability that a neutrino oscillates from one flavor into another. Rather, the event rate of neutrinos interacting after traveling some distance is measured to determine oscillation probability. Therefore, it is important to understand the components that contribute to the number of neutrinos observed in the detectors. In the case of the NOvA experiment, the NuMI beam generates a high flux of ν_μ s at Fermilab near Chicago. These neutrinos are initially sent through the near detector, then travel to northern Minnesota where they interact in the far detector. The far detector is placed at a location to maximize the probability that ν_μ s oscillate into ν_e s.

The oscillation parameters of the neutrinos can be characterized by comparing the *observed* rate of neutrino interactions in the detectors to the *expected* rate based on the neutrino flux and the probability that a ν will interact within the detectors. The expected rates (N_{ND}, N_{FD}) in the near and far detectors can be calculated using the following formulas:

$$\begin{aligned} N_{ND}(E_\nu) &\propto \phi_{ND}(E_\nu) \times \sigma_{ND}(E_\nu) \times \epsilon_{ND}(E_\nu) \\ N_{FD}(E_\nu) &\propto \phi_{FD}(E_\nu) \times \sigma_{FD}(E_\nu) \times \epsilon_{FD}(E_\nu) \times P_{\nu_\alpha \rightarrow \nu_\beta} \end{aligned} \quad (2.17)$$

The flux ϕ is the number of neutrinos produced by the accelerator per cm^2 per energy for a given number of protons on target. This differs based on the detector because the neutrino beam starts very focused and spreads out as it travels. The cross section σ of a neutrino interacting and the efficiency ϵ of each detector corrects for signal events lost during selection. The $P_{\nu_\alpha \rightarrow \nu_\beta}$ term in the far detector count represents the theoretical probability of a neutrino oscillating from one flavor into another as it travels some distance.

The cross section terms in Equation 2.17 represent the probability of a neutrino interacting within each detector. Additionally, the components depend on the neutrino energy for each interaction - these energies must be determined based on the outgoing kinematics of resultant particles. The measured neutrino energy within the detectors typically differs from the true energy of the neutrino because of detector inefficiencies. Accurately predicting the neutrino energy and detector efficiency requires additional input from neutrino-nucleus interaction models, which rely on information about interactions.

Improving Models with Neutrino Scattering

Modeling interactions properly is frequently a large source of systematic uncertainty in neutrino experiments, largely because unknown nuclear structure functions complicate the hadronic component of the interaction [108, 109, 110]. Precise neutrino cross-section measurements constrain these structure functions, feeding back into improved interaction models and, in turn, reducing experimental uncertainties.

For example, the first muon neutrino disappearance result recently published by NOvA showed a nearly 15% systematic uncertainty in the energy observed in neutrino-induced hadronic showers in the near detector [111]. This uncertainty came about due to a deficit in neutrino-nucleus interaction models: the prediction at the time did not match the hadronic energy observed. The issue with the prediction was that the default models used by the experiment did not include a variety of scattering processes that occur on *multi-nucleon* nuclei, like carbon, in the near detector. The underlying physics of this issue was not well understood at the time, so the Meson Exchange Current model was introduced better describe the ND data [112]. The difference in hadronic energy observed in the detector before and after including the MEC model is plotted in Figure 2.13.

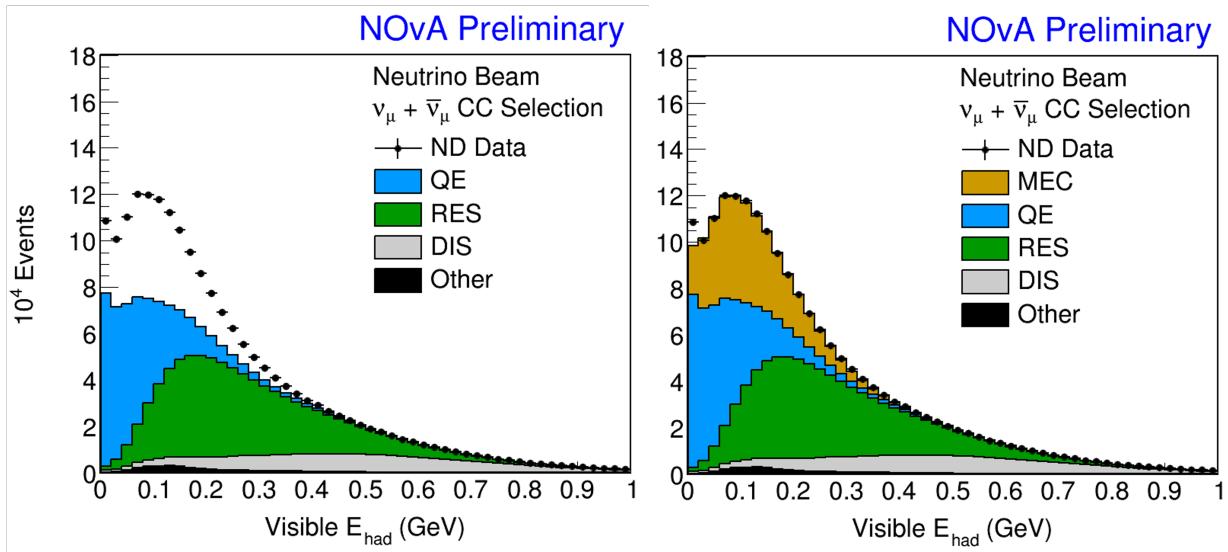


Figure 2.13: **Left:** Visible hadronic energy in CC ν_μ interactions in the near detector prior to introducing MEC processes. A large discrepancy between predicted and observed energy can be seen. **Right:** Visible hadronic energy after introducing MEC processes [112].

Measuring Cross Sections Experimentally

To make an experimental measurement of an interaction cross section, one must count the number of interactions of that type observed in the detector [87]. The goal of this process is to find the event rate for a specific interaction - how frequently it occurs in the detector. The total cross section of an interaction is calculated:

$$\sigma = \frac{N_{\text{sel}} - N_{\text{bkg}}}{\phi \times N_{\text{target}} \times \epsilon}. \quad (2.18)$$

Here, N_{sel} and N_{bkg} are the total number of selected signal events and the number of estimated background events, respectively. N_{target} is the total number of targets available to scatter off of, ϕ is the flux of incoming neutrinos, and ϵ is the signal selection efficiency.

Despite the analysis presented in this thesis not resulting in a direct cross section measurement, an understanding of how cross sections impact the observed interaction rate is necessary to properly represent the expected number of events. The general strategy for measuring a cross section in NOvA is to start with a clearly defined signal event criteria and perform an event selection on simulated data. The event selection efficiency is studied using the simulated events, which also allow estimations of the background rates. After developing the analysis procedure on simulated data, the selection is performed on actual near detector data and used to make a measurement or observation. The same general framework underlies the search for neutrino trident production. The next chapter introduces trident production properly, including the theory behind it and the current experimental status of the process.

Chapter 3

Neutrino Tridents

An interesting type of neutrino interaction that NOvA has the potential to measure is neutrino trident production, or “tridents” - the main topic of this dissertation. Trident production is one of the few *purely* electro-weak processes allowable by the Standard Model, making it a powerful probe of the weak force. In this chapter I discuss the theory of neutrino tridents and describe the current experimental status of the trident production process.

3.1 Neutrino Trident Theory

A trident event occurs when an incoming neutrino exchanges a W or Z boson with a photon from a stationary nucleus, yielding a three-lepton final state: a neutrino and two charged leptons. This interaction can occur via both coherent and incoherent channels. Current theoretical treatments identify 22 trident processes allowed within the Standard Model [113]. Figure 3.1 shows diagrams of 4 available ν_μ induced trident processes.

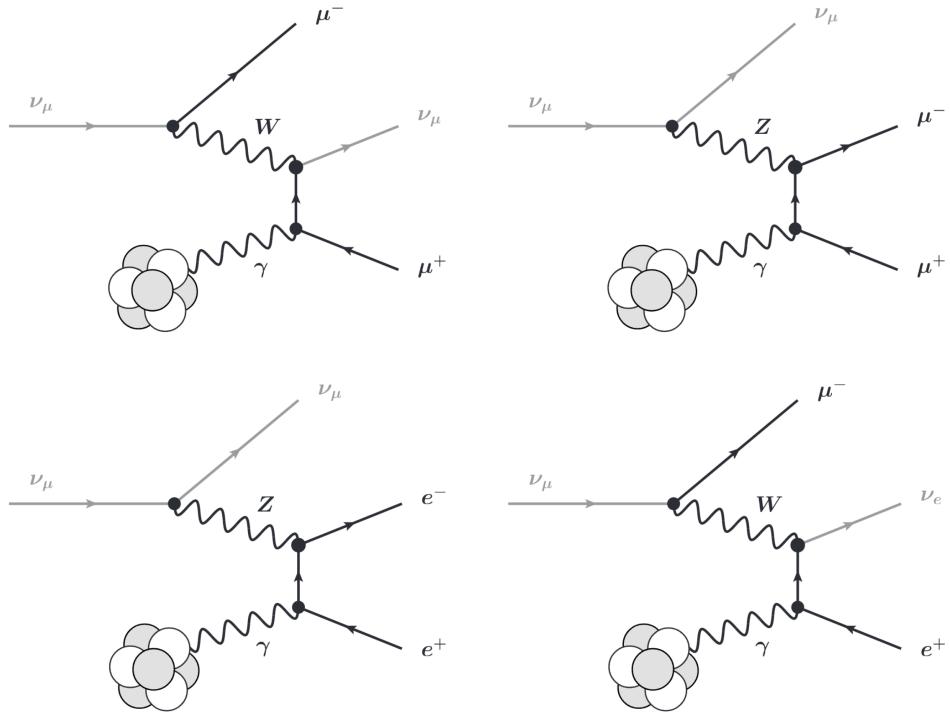


Figure 3.1: Four of the 22 Standard Model trident production processes [114].

Trident production can occur through various channels, but the main channel of experimental interest is one which yields two muons as the final-state leptons - so-called “dimuon” events. This topology can be easily distinguished from other types of neutrino interactions.

3.1.1 Trident Production Cross Sections

The first trident cross section calculation was performed in 1967 by Czyz, Sheppley, and Walecka [115]. The trio used the point four-fermion interaction “shortcut” described in subsection 2.1.2 to model the weak neutrino scattering process followed by the resulting charged lepton scattering coherently off the Coulomb field of a nucleus. Lovseth and Radomski took the theory of trident production a step further in 1971 by factoring in the available kinematic constraints of various trident processes on a number of targets [116]. A year later Brown et al. completed the first theoretical treatment of trident production within the GWS framework, accounting for the massive W and Z bosons mediating the interaction [117]. A number of trident cross section calculations have been performed since then, with a recent resurgence of interest due to an abundance of high-flux, high-energy neutrino experiments [114, 118, 119, 120, 121, 122]. Figure 3.2 shows a generic coherent trident interaction via a neutral current channel with the relevant four-momenta labeled.

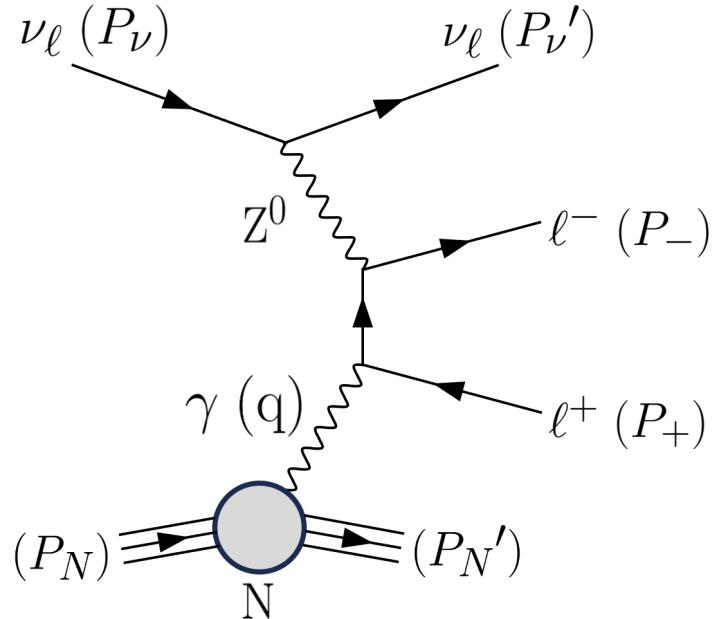


Figure 3.2: A NC Coherent Trident production channel with momenta that contribute to the scattering amplitude labeled.

Coherent Trident Production

As seen in subsection 2.1.2, the weak contribution to the cross section yields terms of order G_F^2 . The nuclear scattering resembles coherent photon pair production, shown by Bethe and Heitler to contribute terms of order $(Z\alpha)^2$ to the cross section, where Z is the atomic number of the struck nucleus and α is the fine structure constant, $\alpha = 1/137$ [123]. The fact that this process depends solely on electroweak interactions with the nucleus provides the ability to calculate it without involvement from messy nuclear structure functions.

The cross section of coherent trident scattering off of a nucleus N with mass M_N is

$$\sigma_N = \frac{(Z\alpha)^2 G_F^2}{(128\pi^6)M_N E_\nu} \int \frac{d^3 P_{N'}}{2E_{N'}} \frac{d^3 P_+}{2E_+} \frac{d^3 P_-}{2E_-} \frac{d^3 P_\nu'}{2E_\nu'} \left(\frac{H_N^{\alpha\beta} L_{\alpha\beta}}{q^4} \right) \delta^4(P_\nu - P_\nu' - P_+ - P_- + q) \quad (3.1)$$

for the particle four-momenta defined in Figure 3.2. Notably, there is a direct proportionality between the cross section and the square of the target nucleus' atomic number Z^2 . The scattering amplitude inherent in Equation 3.1 contains tensors which describe the leptonic and hadronic contributions to the interaction: $L_{\alpha\beta}$ and $H_N^{\alpha\beta}$ [91]. The leptonic tensor $L_{\alpha\beta}$ represents the lepton current summed over all available spin states of the charged leptons and neutrinos in the interaction. The hadronic tensor $H_N^{\alpha\beta}$ represents the electromagnetic current for a system scattering off of a spin-zero nucleus - it depends on the initial momentum of the nucleus, as well as the electric form factor of the struck nucleus (which are very well measured!).

Performing the phase space integration while considering nuclear electric form factors [124, 125] shows that the coherent trident cross section takes the form

$$\sigma_N \propto Z^2 \alpha^2 G_F^2 m_\mu^2 \left[\frac{E_\nu}{m_\mu} \ln \left(\frac{E_\nu}{m_\mu} \right) \right]. \quad (3.2)$$

Therefore, the cross section is directly dependent on incoming neutrino energy as $\sigma \sim E_\nu \ln(E_\nu)$. The additional electromagnetic contribution α^2 reduces the probability of trident events occurring by about five orders of magnitude compared to the total neutrino cross section for neutrinos in NOvA's energy range, making them exceedingly rare to observe.

Incoherent Trident Production

Trident interactions also occur via incoherent channels, in which the neutrino exchanges a weak boson with the coulomb field of a single nucleon: $p(n)$. The cross section for this process is similar to coherent scattering:

$$\sigma_{p(n)} = \frac{\alpha^2 G_F^2}{128\pi^6 M_{p(n)} E_\nu} \int \frac{d^3 P_{p(n)}'}{2E_{p(n)}'} \frac{d^3 P_+}{2E_+} \frac{d^3 P_-}{2E_-} \frac{d^3 P_\nu'}{2E_\nu'} \left(\frac{H_{p(n)}^{\alpha\beta} L_{\alpha\beta}}{q^4} \right) \delta^4(P_\nu - P_\nu' - P_+ - P_- + q). \quad (3.3)$$

The primary difference is the lack of a Z^2 term and the smaller target mass $M_{p(n)}$. The contribution from the hadronic tensor $H_{p(n)}^{\alpha\beta}$ now depends on the electric *and* magnetic form factors of the spin- $\frac{1}{2}$ nucleon struck. Additionally, nuclear effects contribute an uncertainty of about 30% to incoherent cross sections [114].

Incoherent trident production is suppressed relative to coherent channels due to the target's weaker electromagnetic field. Cross sections of trident production via scattering off of a proton are approximately one order of magnitude smaller than coherent scattering. Neutron channels have even smaller cross sections (about two orders of magnitude smaller than coherent channels) because they are electrically neutral [113, 114]. For the analysis described in this dissertation, we did not consider trident events from incoherent scattering due to their small cross sections.

Total Trident Cross Section

The total cross section of trident production on a nucleus composed of Z protons with atomic mass A takes the form [113]

$$\sigma_{\text{Tot}} = Z\sigma_N + Z\sigma_p + (A - Z)\sigma_n, \quad (3.4)$$

where the coherent cross section σ_N dominates due to being an order of magnitude larger than the others. Trident interactions in NOvA (where the primary target is Carbon 12) have total cross sections on the order of magnitude of $\sigma \sim 10^{-44} \frac{\text{cm}^2}{\text{GeV}}$ [117]. A treatment of all possible target nuclei within the NOvA near detector is provided in chapter 5. Figure 3.3 shows the total cross section for the trident process $\nu_\mu + {}^{12}\text{C} \rightarrow \nu_\mu + \mu^- + \mu^+ + {}^{12}\text{C}$ as determined computationally from code provided by Altmannshofer et al. [114], compared to the total CC neutrino cross section.

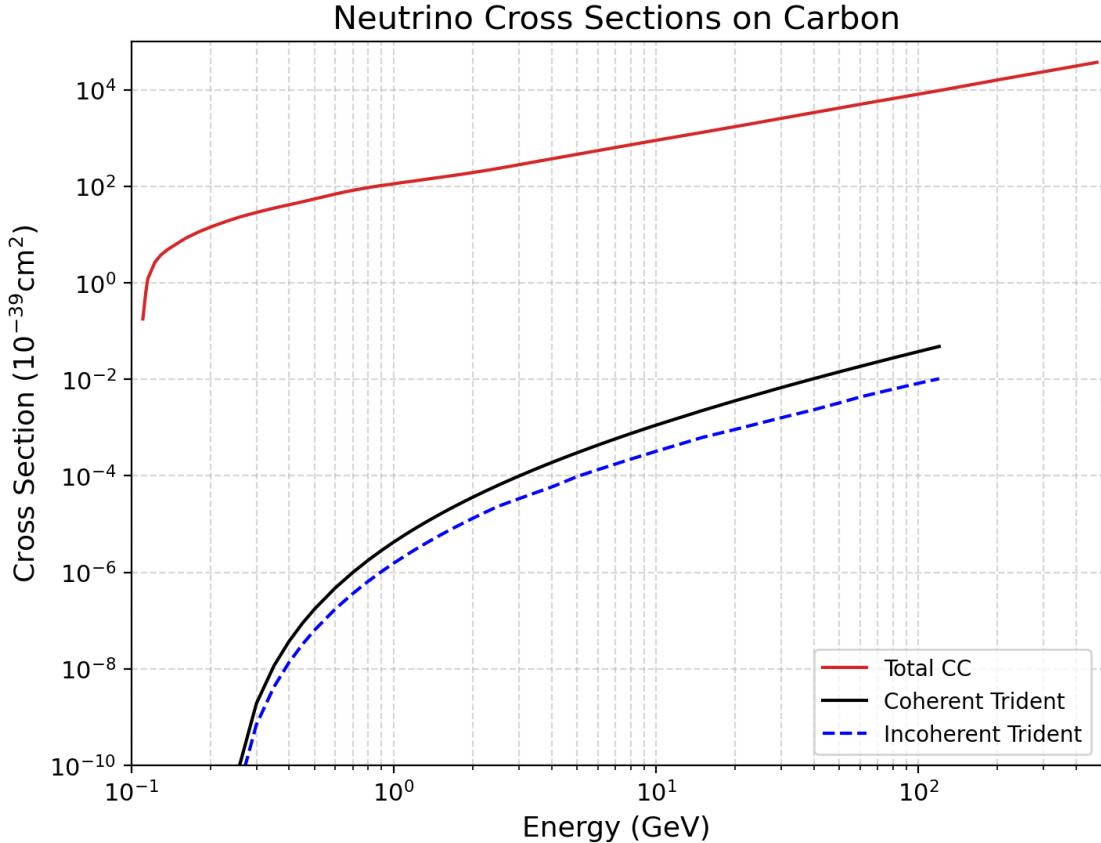


Figure 3.3: Mean values of the dimuon trident cross section on carbon determined numerically [114], compared to the mean total CC neutrino cross section [32].

Trident events containing two muons in the final state (so-called “dimuon” events) are the main focus of the search presented in chapters 5 and 6 of this thesis. This is because muons are significantly easier to identify in a detector than electrons, and events containing two muons are generally quite rare. Figure 3.4 shows some kinematic distributions of dimuon events (purple are coherent interactions, blue are incoherent) [126, 127]. The key features of these plots are very small values of Q^2 and small angles of the resulting muons with respect to the incident neutrino beam, θ_- and θ_+ . The small Q^2 mean that the neutrino does not transfer much momentum to the struck nucleus, so very little (if any) hadronic energy will be visible in the detector. The small opening angles mean that the two muons will travel very close together in the detector. Therefore, the experimental signature of both coherent and incoherent dimuon tridents is a final state containing two highly energetic, forward-going muons with little to no hadronic energy visible in the detector.

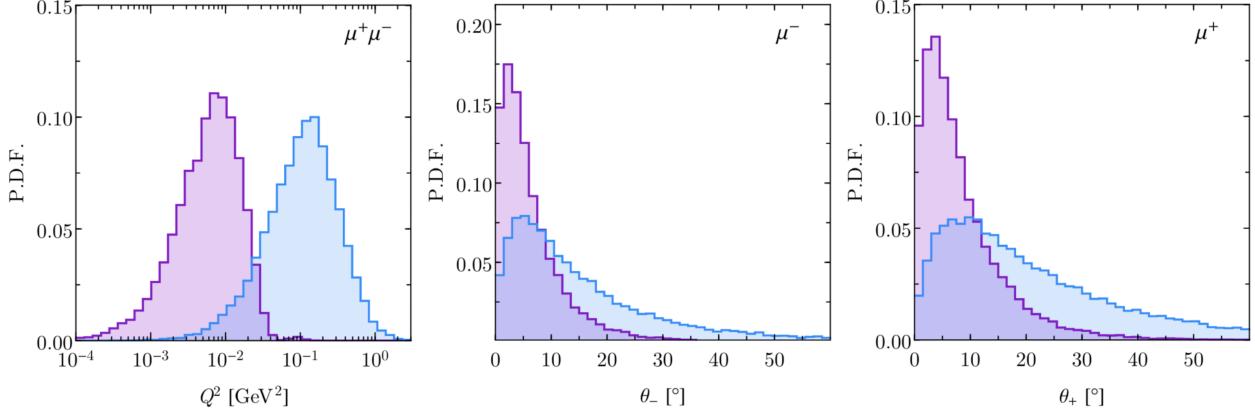


Figure 3.4: Momentum transfer and opening angle histograms for coherent (purple) and incoherent (blue) dimuon events showing low Q^2 and small opening angles [127].

3.1.2 Trident Production as a Probe of Physics Beyond the Standard Model

Trident production is interesting to study on its own merit as a rare, purely electroweak process. However, another theoretically compelling reason to study tridents is as a probe of physics not currently accounted for in the Standard Model [128, 129, 130, 131]. The most popular theory for BSM physics involving tridents takes the form of new force-mediating boson, the Z' , which would contribute a non-negligible amount to the trident production cross section.

Performing cross section calculations with the Z' mediating the neutrino-photon exchange [128] yields a trident cross section which compares to the standard model in the following way:

$$\frac{\sigma^{(SM+Z')}}{\sigma^{(SM)}} \simeq \frac{1 + (1 + 4 \sin^2 \theta_W + 2\nu_{SM}^2/\nu_{Z'}^2)^2}{1 + (1 + 4 \sin^2 \theta_W)^2}, \quad (3.5)$$

where $\nu_{SM} = 246$ GeV is the SM Higgs vacuum expectation value and $\nu_{Z'} = m_{Z'}/g'$ depends on the mass and coupling strength of the new boson. This comparison shows that the cross section of trident production should be larger than that predicted by the SM if such BSM physics exists. Existing trident cross section measurements discussed in the next section agree with SM predictions within their uncertainties. However, this does not rule out the possibility of the Z' existing - more precise trident production measurements are required to gain more understanding into the potential of BSM physics.

3.2 The Experimental Landscape of Tridents

Neutrino tridents are quite difficult to observe because of their small interaction cross sections. Nonetheless, a handful of trident observations have been made by the CHARM II, CCFR, and NuTeV experiments. So far, experimental measurements agree with the Standard Model prediction of trident cross sections.

The CHARM II Experiment

The first measurement of trident production was published by CERN’s CHARM II collaboration in 1990 [132, 133]. CHARM II was an experiment dedicated to studying NC neutrino interactions as a means of investigating the structure of the weak force - specifically the Z^0 boson. The experiment utilized a wide band beam of neutrinos and antineutrinos generated by CERN’s Super Proton Synchrotron and focused with magnetic horns. The average neutrino energies were $\langle E_\nu \rangle = 23.8$ GeV and $\langle E_{\bar{\nu}} \rangle = 19.3$ GeV.

The CHARM II detector apparatus consisted of a fine-grain, low density target calorimeter followed by a muon spectrometer. The calorimeter was composed of 420 modules made of 48 mm thick glass plates and 352 plastic streamer tubes containing digital (wire) and analogue (cathode strip) readout modes - each module had an active area of 3.7×3.7 m². The muon spectrometer consisted of six magnetized iron toroids interleaved with drift chambers, and the total mass of the target was about 800 tons. CHARM reconstructed muons with momentum of 20 GeV/c with $\pm 14\%$ resolution in the detector.

In order to measure a cross section for trident production, the experiment developed a specific dimuon trigger designed to select two-track events with track lengths of at least 10 cm covering 30 planes in the calorimeter. The analyzers required that tracks originate from the same point within the detector and penetrate the muon spectrometer, which allowed the signs and momenta of the two particles to be determined. Only muon tracks with reconstructed momenta above 4 GeV/c were kept, with events containing lower-momentum muons being discarded.

The primary background of dimuon trident events in CHARM II were dimuon events originating from the semileptonic decay of charm quarks: $\nu_\mu^{(-)} + (d^-, s^-) \rightarrow \mu^\pm + \bar{c}^- \rightarrow \mu^\pm + \mu^\mp + \bar{\nu}_\mu^{(-)}$ via the decay process shown in Figure 3.5.

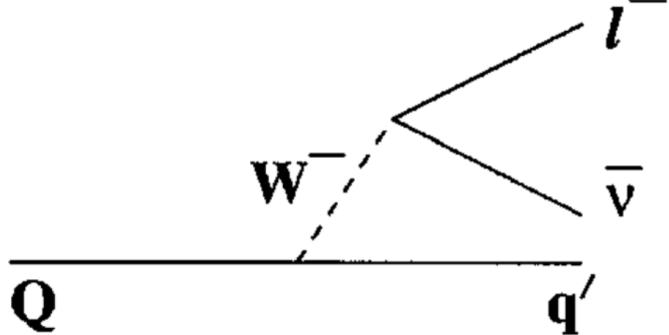


Figure 3.5: Generic representation of the semileptonic decay of one quark Q into another q' [134].

The primary indicator that this process occurred is a shower of hadronic energy near the vertex of the interaction. In order to minimize contributions from this process the analyzers measured hadronic activity by counting “additional” streamer tube hits which could not unambiguously be assigned to one of the muon tracks. If more than 8 of these additional hits occurred within the first 10 planes of the interaction, the event was discarded.

The data used for the study was obtained between 1987 and 1989. The total proton exposure of 1.5×10^{19} protons yielded $4 \times 10^7 \nu_\mu$ and $2 \times 10^7 \bar{\nu}_\mu$ events, which were analyzed via the preceding steps. After background subtraction a trident production signal of 55 ± 16 events was observed over a background of 43 ± 12 events. An event display of the CHARM II detector containing a dimuon trident event can be seen in Figure 3.6.

The trident production cross section was determined using the number of observed events, the neutrino flux, and the trigger/selection/detector efficiencies determined by studying Monte Carlo simulations. Combining those yielded a measured trident cross section (averaged between neutrinos and antineutrinos) of $\sigma_{Ex.} = (3.0 \pm 1.4) \times 10^{-41} \text{ cm}^2$ per nucleus. The analyzers calculated the theoretical value of $\sigma_{Th.} = (1.9 \pm 0.4) \times 10^{-41} \text{ cm}^2$ per nucleus for their parameters as determined by the Standard Model. The ratio of the observed and theoretical cross sections reported by CHARM II is (1.58 ± 0.64) - consistent with the Standard Model prediction at $\langle E_\nu \rangle \approx 20 \text{ GeV}$.

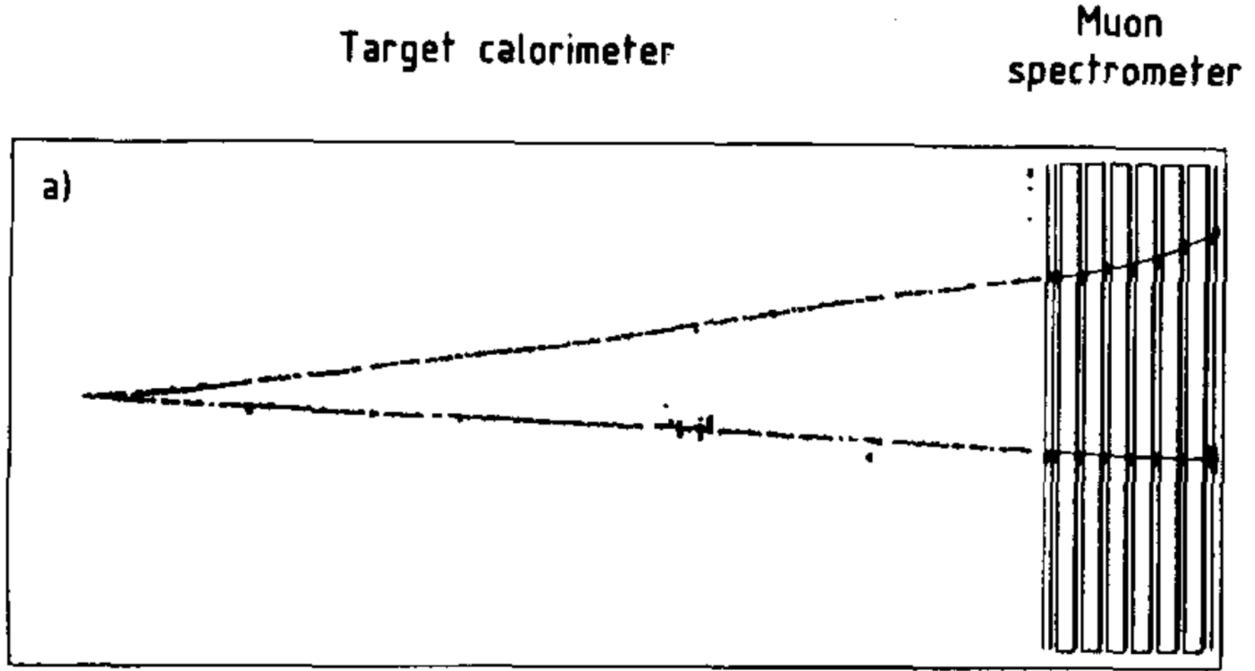


Figure 3.6: A dimuon trident event captured in the CHARM II detector [93].

The CCFR Experiment

The next measurement of trident production was made in 1991 by the Columbia-Chicago-Fermilab-Rochester (CCFR) neutrino experiment at Fermilab's Tevatron Lab E fixed target detector facility [135, 136]. The overall goal of the CCFR collaboration was to use high energy neutrino-nucleon “deep inelastic scattering” (DIS) to study proton structure and electroweak theory, as well as search for exotic particles. The neutrino beam utilized by CCFR came from Fermilab's Tevatron, which accelerated protons up to 800 GeV and collided them into a beryllium target. The resulting pions and kaons decayed through a quadrupole triplet beam line into a wide band beam of neutrinos and antineutrinos of energy 10 - 600 GeV, with $\langle E \rangle = 160$ GeV.

The CCFR detector design shared some similarities with CHARM II - it was also a calorimeter followed by a muon spectrometer. The calorimeter was formed from eighty-four $3\text{ m} \times 3\text{ m} \times 10\text{ cm}$ iron plates interspersed with scintillator segments and drift chambers. The muon spectrometer was composed of five iron toroid magnets with drift chambers throughout.

The measurement of neutrino tridents performed by CCFR analyzers was explicitly intended to verify the Standard Model channels of trident production through both charged and neutral currents

- involving both W and Z bosons. The signal selection required two muon tracks originating in the calorimeter volume with reconstructed energy above 9 GeV and 4.5 GeV for the primary and secondary muons, respectively. The analyzers distinguished trident dimuon events from charm and pion backgrounds by requiring a dimuon invariant mass of $M_{\mu\mu} \leq 2$ GeV and hadronic energy of $E_{had} \leq 1$ GeV.

The data utilized for the analysis was accumulated from 1987 - 1989 and contained a total of 3.7×10^6 muon triggers. After performing background reduction, the experiment observed a corrected trident signal of 37 ± 12.4 events. A graph of the data taken by CCFR can be seen in Figure 3.7 - the bin all the way on the left-hand side of the plot corresponds with hadronic energies below 1 GeV. A clear spike of about 15 events is present in the $M_{\mu\mu} \leq 2$ GeV data, corresponding to the observed trident signal. The analyzers reported a cross section measurement of $\sigma_{Ex.} = (4.7 \pm 1.6)E_{\nu} \times 10^{-42}$ cm 2 per iron nucleus within their energy range. The ratio of the observed and theoretical cross sections reported by CCFR is (0.82 ± 0.28) - also consistent with the Standard Model prediction.

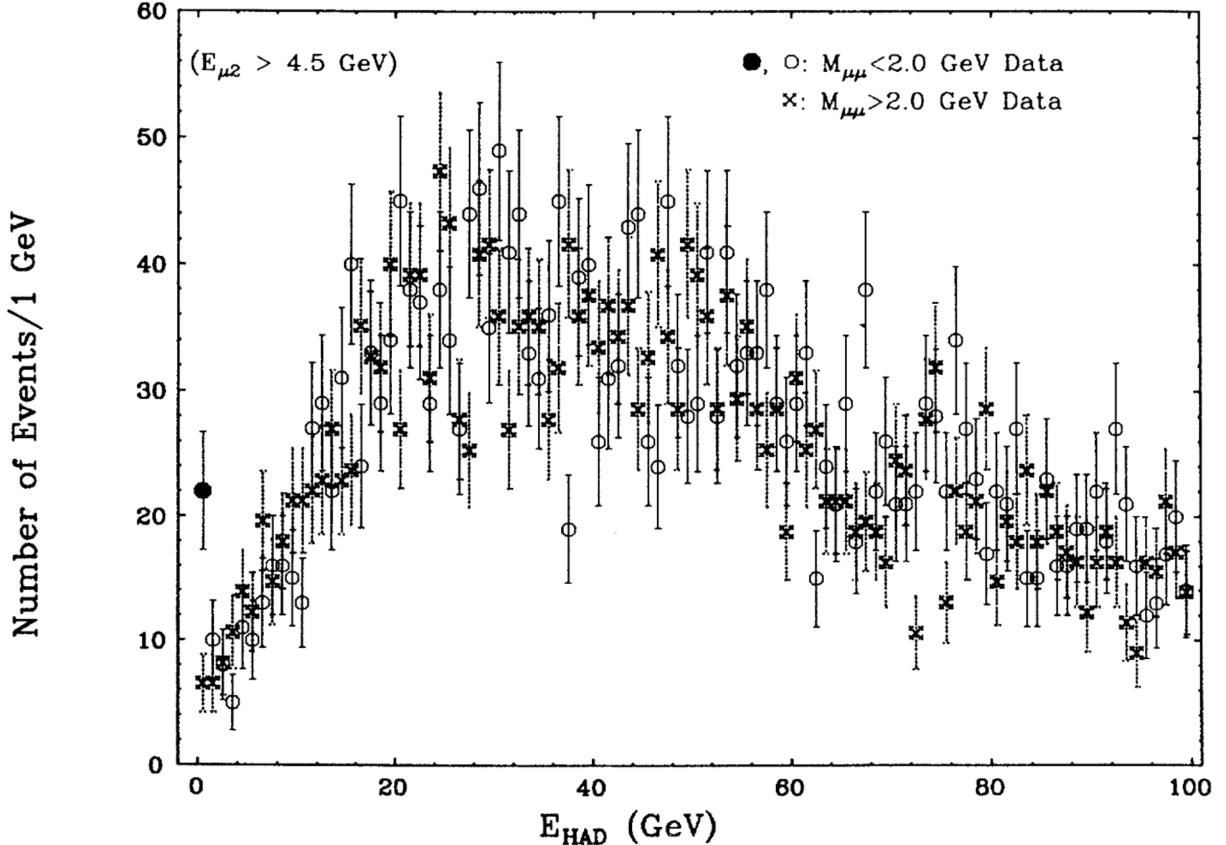


Figure 3.7: Hadronic energy distribution of observed dimuon events in the CCFR detector [136].

The NuTeV Experiment

The most recent observation of neutrino trident production was made in 1998 by Fermilab's NuTeV experiment [135, 137, 138]. NuTeV was the next-generation experiment to take place in the Lab E detector facility after CCFR. The beam composition shared many similarities to CCFR with the key addition of the ability to focus pions and kaons based on charge, yielding a beam of primarily either ν_μ or $\bar{\nu}_\mu$. The detector remained largely the same, with many of the same physics goals as CCFR.

The dimuon trident signal selection variables are presented in Figure 3.8, along with a hadronic energy plot comparing event rates between $M_{\mu\mu} \leq 2.3$ GeV (lines) and $M_{\mu\mu} \geq 2.3$ GeV (points). The elevated rate in the two lower bins of the $M_{\mu\mu} \leq 2.3$ GeV distribution indicate the presence of trident events. NuTeV observed a total of 17 events that passed trident signal selection, with the ratio of the observed and theoretical number of events reported being $(0.72^{+1.73}_{-0.72})$ - the third indication of consistency with the Standard Model.

Variable	Criterion
$E_\mu(1)$ (at event vertex)	> 9 GeV
$E_\mu(1)$ (entering toroid)	> 3 GeV
$E_\mu(2)$ (at event vertex)	> 9 GeV
$E_\mu(2)$ (entering toroid)	> 3 GeV
charge(1) + charge(2)	= 0
Both muon charges determined	
E_{HAD}	< 3 GeV
$M_{\mu\mu}$	< 2.3 GeV/c ²
Vertex Position (Transverse)	> 12 cm. from calorimeter edge
Vertex Position (Upstream)	> 20 cm of steel from calorimeter edge
Vertex Position (Downstream)	> 75 cm of steel from calorimeter edge

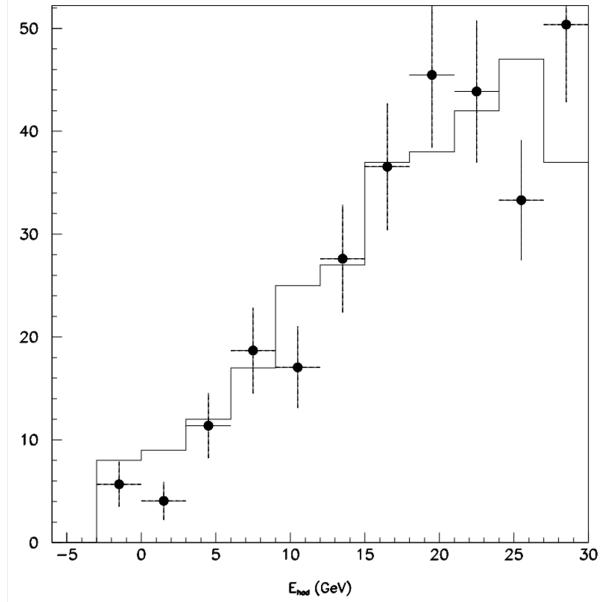


Figure 3.8: **Left:** Signal selection criteria for dimuon trident events in NuTeV. **Right:** Distributions of hadronic energy split by invariant muon mass [138].

Tridents in NOvA

Although trident production has consistently been a subject of theoretical interest, no measurements of the process have been made in nearly 30 years. Its extreme rarity makes it one of the least explored electroweak processes. This work presents the first dedicated search for neutrino trident production since NuTeV. While not a direct measurement of the cross section, this analysis demonstrates the feasibility of isolating trident-like events using a neutrino beam of $\langle E_\nu \rangle \approx 2$ GeV in the NOvA near detector. Hopefully this work helps reestablish experimental interest in trident production, and lays the foundation for a future cross section measurement with NOvA.

Chapter 4

The NOvA Experiment

The NuMI Off-axis ν_e Appearance (NOvA) experiment is a long-baseline neutrino oscillation experiment designed to measure ν_e appearance from a primarily ν_μ beam. The fundamental components of the experiment are Fermilab’s Neutrinos at the Main Injector (NuMI) beam which creates neutrinos, and two functionally equivalent detectors to observe them. The goals of the experiment are to broaden the knowledge of neutrino oscillations by constraining the value of θ_{23} and to help determine the neutrino mass hierarchy by measuring Δm_{32} . NOvA is also sensitive to CP violation in the neutrino sector δ_{CP} , which may be related to the matter-antimatter imbalance in the universe.

This chapter summarizes the details of the neutrino beam, the detectors, the hardware, and the software used to perform analyses. Most of the details and figures of the experiment in this chapter are taken from the NOvA technical design report (from which the final design of the experiment differs only slightly) [139], internal NOvA documents, or the dissertations of previous students who developed the tools discussed. Particular emphasis is given to the aspects of NOvA that I have been heavily involved in, and those relevant to the tridents analysis.

4.1 The NuMI Beam

Fermilab’s Neutrinos at the Main Injector (NuMI) [74, 140, 141] generates the most intense beam of neutrinos in the world. The beam originates in Fermilab’s accelerator complex by first stripping an electron from hydrogen, yielding a single proton. Protons are accelerated to 0.4 GeV in a Linear Accelerator, then sent through a series accelerator rings where they are stacked into “batches” of approximately 5×10^{13} protons and accelerated up to 120 GeV. These batches are released towards the NuMI beamline every 1.3 seconds in a process called a beam “spill” - each spill lasts for roughly 10 μ s, as can be seen in Figure 4.1. The intensity of the beam steadily increased since first operations from 200 kW to 900 kW.

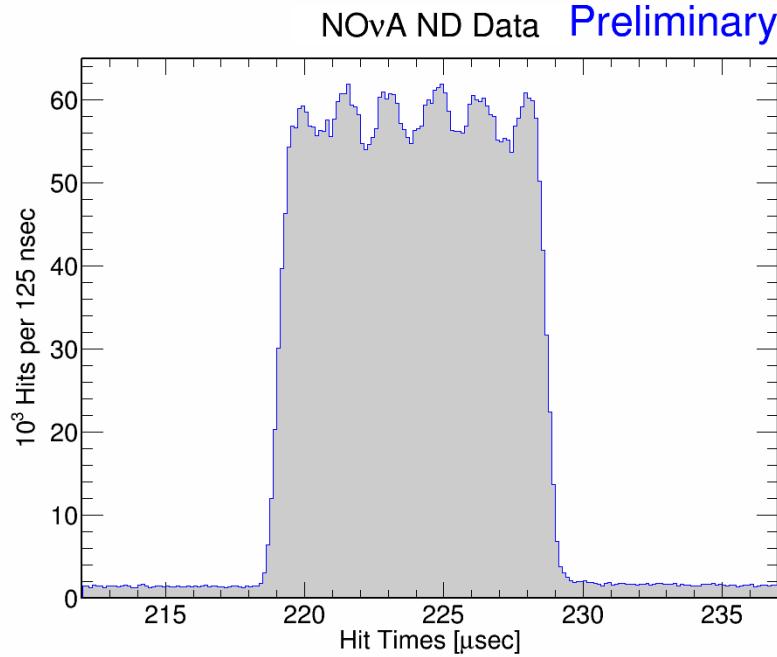


Figure 4.1: An integrated distribution of hits seen within the NOvA near detector during beam spills. The $10\mu\text{s}$ window during which the spill is happening is clearly visible.

In order to create neutrinos, the 120 GeV protons are collided into a target which is composed of a line of graphite fins. As the protons travel through the fins, they interact and produce mesons such as pions or kaons. The mesons exit the target and their trajectories are bent by two parabolic-shaped magnetic focusing horns which diminish the transverse momentum of the mesons. A cartoon of this process can be seen in Figure 4.2.

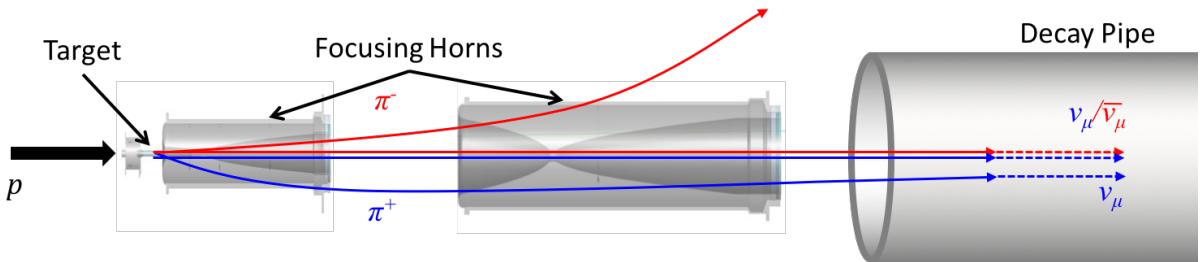


Figure 4.2: Pions exiting the graphite target are focused by magnetic focusing horns. The current in these horns can be flipped to focus either π^+ or π^- . This cartoon is showing π^+ being focused and π^- being de-focused [142].

An important feature of the magnetic focusing horns is that by changing the direction the current flows along the horns, we can choose whether to focus positively charged mesons or negatively charged mesons. Current oriented to focus positive mesons and de-focus negative mesons is called “Forward Horn Current” (FHC), whereas current oriented to focus negative mesons and de-focus positive mesons is called “Reverse Horn Current” (RHC). When the positively charged mesons decay, they yield anti-muons and muon neutrinos (ν_μ). When negatively charged mesons decay, they yield muons and muon anti-neutrinos ($\bar{\nu}_\mu$). Therefore, running in FHC or RHC mode allows the ability to choose if the primary particles being sent to the detectors are ν_μ s or $\bar{\nu}_\mu$ s, respectively. Sometimes mesons of the type to be de-focused pass through the horns, resulting in neutrinos of the opposite type in the beam, which are dubbed the “wrong-sign” component of the beam.

The mesons leave the focusing horns and enter a 675 m long helium-filled pipe, where they decay with minimal interactions in the helium to prevent defocusing by multiple Coulomb scattering. The mesons decay primarily via the channels $\pi^{+(-)} \rightarrow \mu^{+(-)} + \nu_\mu(\bar{\nu}_\mu)$ and $K^{+(-)} \rightarrow \mu^{+(-)} + \nu_\mu(\bar{\nu}_\mu)$. However, a common decay channel of kaons is $K^{+(-)} \rightarrow \pi^0 + e^{+(-)} + \nu_e(\bar{\nu}_e)$, which means that the beam also contains a small amount of impurities in the form of electron-type neutrinos. Between the wrong-sign component and the electron neutrino component, running in FHC mode yields a ν_μ beam that is 95% pure and running in RHC mode yields a $\bar{\nu}_\mu$ beam that is 93% pure.

Following the decay pipe, the beamline is composed of a hadron monitor and absorber to record and attenuate the residual hadrons that make it through the pipe. Following that are muon monitors and approximately 240 m of rock intended to absorb any muons that make it through the absorber. The only thing that makes it through the full beamline are neutrinos, which go on to be detected in various neutrino experiments. A diagram of the NuMI beamline is shown in Figure 4.3.

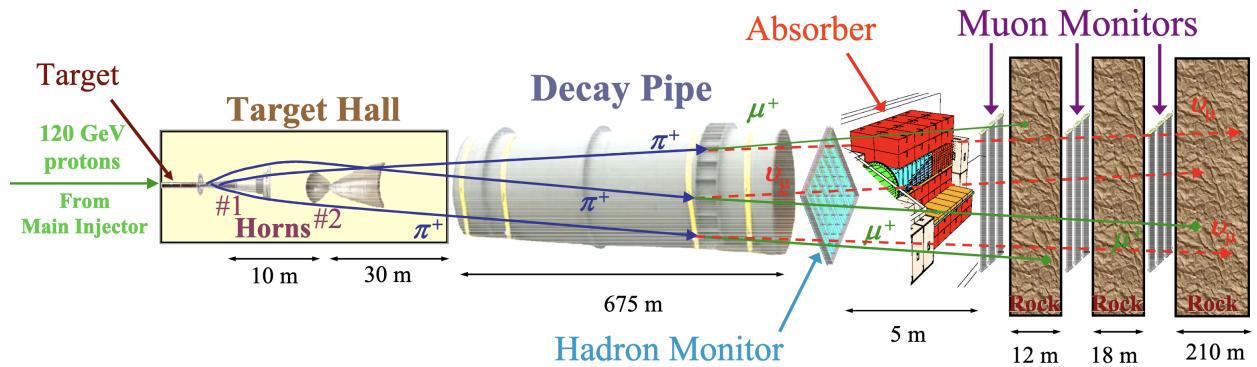


Figure 4.3: A diagram of the NuMI beamline [141].

4.1.1 The Off-Axis Approach

The NOvA detectors were placed 14.6 mrad away from the center of the NuMI beam, or 14.6 mrad “off-axis”. For the angle θ between decaying pions and resulting neutrinos, the neutrino energy is given by

$$E_\nu = \frac{0.43E_\pi}{1 + \gamma^2\theta^2}. \quad (4.1)$$

Here E_ν and E_π are the neutrino and pion energies, respectively, and $\gamma = E_\pi/m_\pi$. Because of this relationship, NOvA sees a larger amount of 2 GeV neutrinos at an off-axis angle of 14.6 mrad than it would if it were on-axis. This is useful because the first oscillation maximum (as described in section 1.4) for NOvA’s baseline of 810 km occurs around 2 GeV. The off-axis location has the added benefit of decreasing backgrounds from higher energy neutral-current events, which can mimic ν_e events in the detectors. Figure 4.4 shows the result of the 14 mrad off-axis positioning compared to placing the detectors on-axis or at various other angles. The left plot shows that the energy of the neutrino depends less on the parent pion’s energy at 14 mrad than at other angles, and the right shows that the 14 mrad location sees approximately five times more 2 GeV neutrinos than other locations would.

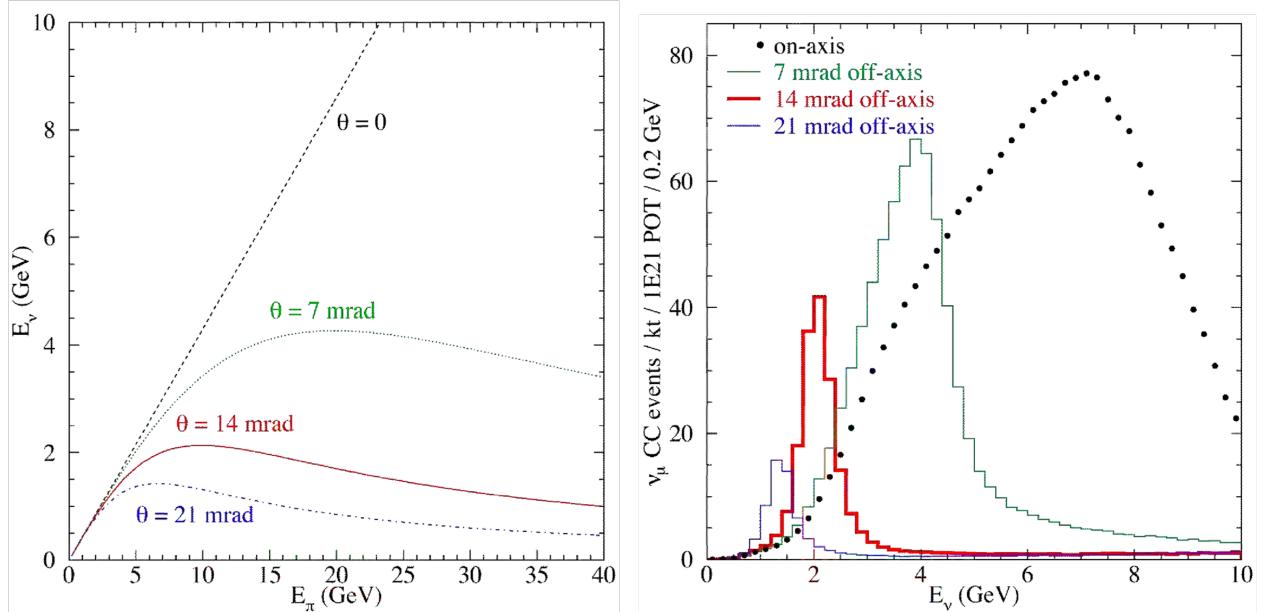


Figure 4.4: **Left:** The energy of neutrinos produced at an angle θ relative to the pion direction. **Right:** ν_μ event rates for a distance of 810 km from Fermilab at various off-axis locations [139].

4.2 The NOvA Detectors

NOvA uses two detectors - a near detector (ND) and a far detector (FD), which is typical of long-baseline neutrino experiments. The ND is relatively small, measuring approximately $4\text{m} \times 4\text{m} \times 16\text{m}$ (the beam points along the length of the detector) and weighing 300 tons. It is located on the Fermilab campus in Batavia, Illinois, just 1km away from the NuMI target. The ND lies 105 meters underground so that it is shielded from incoming cosmic rays. It is designed to measure the unoscillated composition of neutrinos coming from NuMI so the beam information can be extrapolated to the FD.

Because the ND is so close to the beam it sees a very large neutrino flux - during a typical beam spill the ND sees about 5 neutrino interactions on average. To deal with the high rate of interactions the ND electronics are set to sample very frequently, about every 125 ns. This frequent sampling and high flux makes the ND a convenient tool to perform neutrino cross section measurements and observations of rare processes, such as trident production. Despite being different sizes, the ND and FD are designed to be functionally equivalent to one another. This means they share the same physical structure, materials, and read-out electronics. The similar design was done to help minimize the impact of systematic uncertainties at the FD due to a good understanding of the ND detector response. The main difference between the detectors relevant to this analysis is a muon catcher attached at the end of the ND. The emphasis of this chapter is the ND, though I will also discuss the FD for a complete summary of the experiment.

The FD is located in northern Minnesota, 810 km from the NuMI target. It is much larger than the ND, measuring about $15\text{ m} \times 15\text{ m} \times 60\text{ m}$ and weighing 14 kilotons. The FD lies on the surface, only shielded from cosmic rays by a layer of rock and concrete on the roof of the FD complex. The primary purpose of the FD is to measure the energy spectrum of the oscillated neutrinos - namely ν_e s. The detectors are made from low Z materials to aid in discriminating between ν_e charged-current interactions and neutral-current interactions. Their design allows the ability to distinguish electron-induced showers from photon-induced showers, which is important in determining when a ν_e has interacted since the resultant lepton will be an electron. In order for the beam to reach the FD it must be shot at an angle of 3.3° down into the Earth. By sending neutrinos through the Earth, NOvA is able to leverage the MSW effect to probe the neutrino mass ordering.

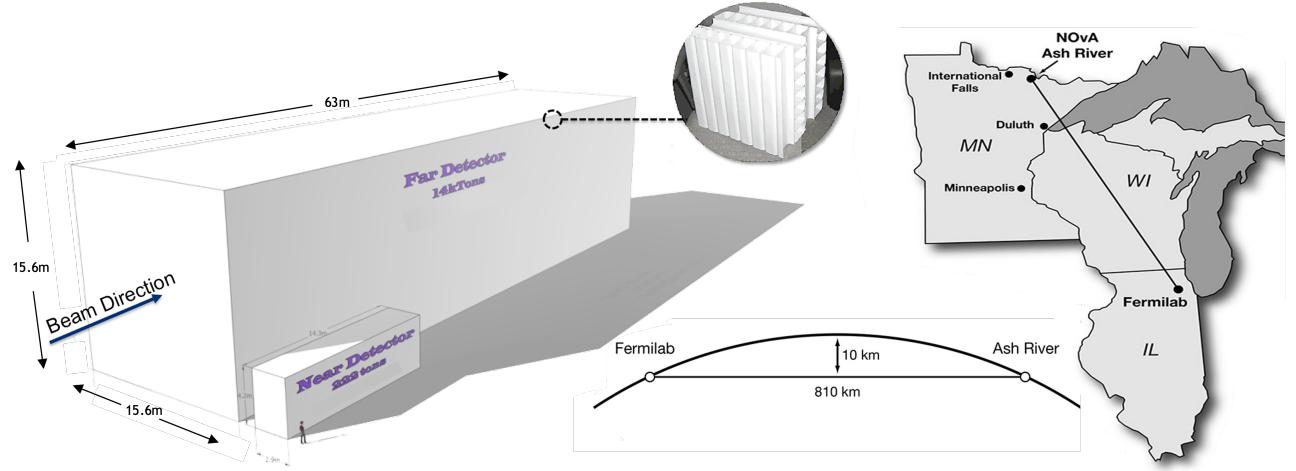


Figure 4.5: The ND and FD, and the beam trajectory of the NOvA experiment [143].

4.2.1 Design of the Detectors

The detectors are segmented tracking calorimeters formed out of reflective PVC rectangles (called “cells”) containing liquid scintillator and a wavelength-shifting fiber. The scintillator is 95% mineral oil and 5% pseudocumene mixed with other chemicals, which emit light of wavelength 400 - 450 nm when a particle interacts in it. That light is picked up and transferred down the length of the cell by optical fibers which also shift the wavelength to about 500 - 550 nm. Once the light reaches the end of the cell its wavelength is at a good length for the avalanche photo-diodes (APDs) which convert the light into an electronic signal. The APDs have an 85% quantum efficiency, meaning that 85% of the energy emitted as photons in the cell volume gets converted into electrical signal. Sixteen cells are formed into “extrusions”, 6 of which are glued end to end and called “planes” in the ND.

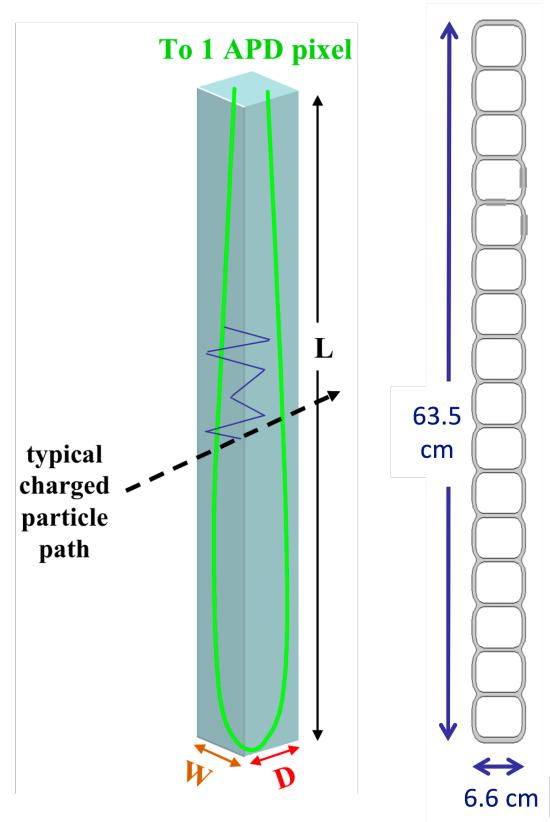


Figure 4.6: **Left:** A single cell featuring the wavelength shifting fiber. **Right:** Sixteen cells attached together to form a single extrusion [139].

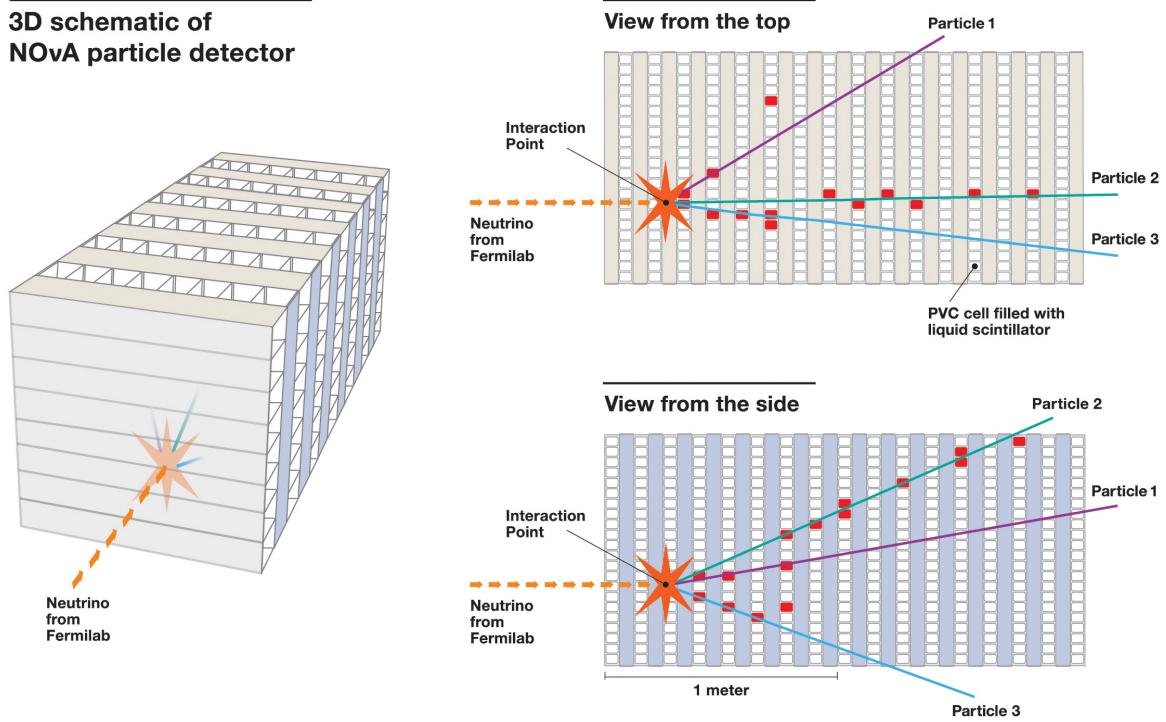


Figure 4.7: A schematic of how layering planes in alternating vertical and horizontal orientations yields a top view and side view of energy deposited in the detectors [143, 144]. Deposited energy is resolved at the cell level, with each cell highlighted in red representing a “hit” in the detector.

Figure 4.2.1 and Figure 4.7 show a single cell, cells stacked to form an extrusion, and a schematic of a NOvA detector formed out of planes stacked together. Planes are stacked in alternating horizontal and vertical orientations along the length of the detectors so that particle tracking can be done in 3 dimensions. A combination of 64 planes (32 of each orientation) is called a detector “diblock”.

When a particle deposits some energy in a cell, the APD at the end of the cell receives a signal and indicates that there was a “hit” in that cell. The distinction between cells, planes, and hits in the detector is essential to understand particle tracking and reconstruction, which is a large part of the tridents analysis. The ND contains 20,192 cells arranged into 214 planes.

A feature of the detector design is the ability to distinguish the outgoing charged lepton type after an interaction occurs. This is necessary for both the primary NOvA analysis and for the trident analysis, which depends on a signature of two resultant muons. Muons leave long, straight tracks with a characteristic energy deposition rate in the detectors, which allows them to be distinguished from other particles relatively easily.

The ND also has a segment attached to the downstream end known as the “muon catcher”. It is 3 m long and consists of ten 10 cm thick planes of steel inserted periodically into pairs of PVC planes. The purpose of the muon catcher is to stop muons from few-GeV charged current ν_μ interactions from exiting the back of the detector. In the muon catcher, the vertically aligned planes consist of three extrusion modules while the horizontal planes are made from just two extrusion modules. This is different from the main bulk of the ND, which has three vertical and three horizontal extrusion modules. Therefore the muon catcher is just as wide as the rest of the detector, but not as tall. This can be seen as the unboxed region in the top right of the bottom view in the event display shown in Figure 4.8. Also present in the event display is the distinction of three separate neutrino events taking place at slightly different times. The ND timing resolution discussed previously allows for the separation of each of these interactions, indicated by different colored hits. The time window shown at the bottom corresponds to a single NuMI beam spill.

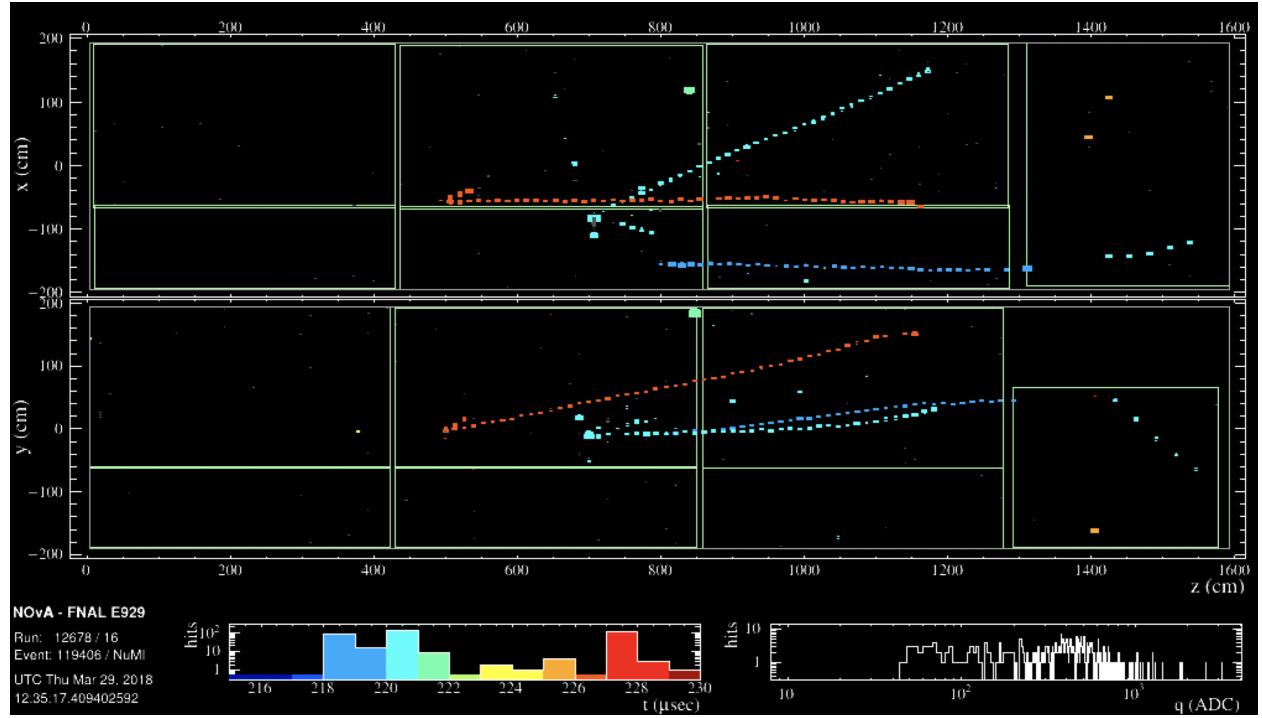


Figure 4.8: An event display showing typical near detector activity during a NuMI beam spill. The top pane of the event display shows the XZ view, and the bottom pane shows the YZ view.

4.3 Data Acquisition

Data acquisition (DAQ) is the process of turning raw energy deposited in the cells of the detector into data that can be stored and utilized for analyses. The DAQ system consists of multiple pieces of electronic hardware that operate in a continuous readout mode. All of the data collected is temporarily stored in a buffer farm and a decision is made on whether the data should be committed to long term storage or rejected.

4.3.1 DAQ Hardware

Avalanche Photo-Diodes

The DAQ system begins with the avalanche photodiodes (APDs) attached at the end of every module. The two ends of the wavelength shifting fibers running through a cell are given a “pixel” to be read out by the APD. Each APD used by NOvA is composed of an array of 32 pixels, which see the light emitted within two sets of extrusions (a segment called a “module”). Since each APD reads out 32 cells, the ND utilizes 631 APDs to convert light signals into electrical signals. Figure 4.9 shows the end of a module with all 32 optical fibers bundled together, as well as a single APD with its 32-pixel readout array which is fitted onto the end of the module.

A hit is registered in the detector when the energy output by an APD is greater than some threshold value, determined individually for each APD based on the amount of noise it experiences. To reduce thermal noise, the APDs are cooled to a temperature of -15°C by thermo-electric coolers that remove heat from the detector by a continuously flowing supply of cold water.

Front End Boards and Data Concentrator Modules

Each APD is connected to a custom-built front-end electronics board (FEB) which reads out the electronic signal from all 32 cells continuously during each beam spill. The FEB time stamps any signal pulses coming from the associated APD, digitizes it into 12 bits, then transmits the digitized data to a data concentrator module (DCM). Each DCM takes the signals from up to 64 FEBs corresponding to a localized region of the detector. The ND has 14 DCMs that collect all of the information received from the FEBs and transmit it into a “buffer farm” where it is collated and stored for up to 30 minutes.

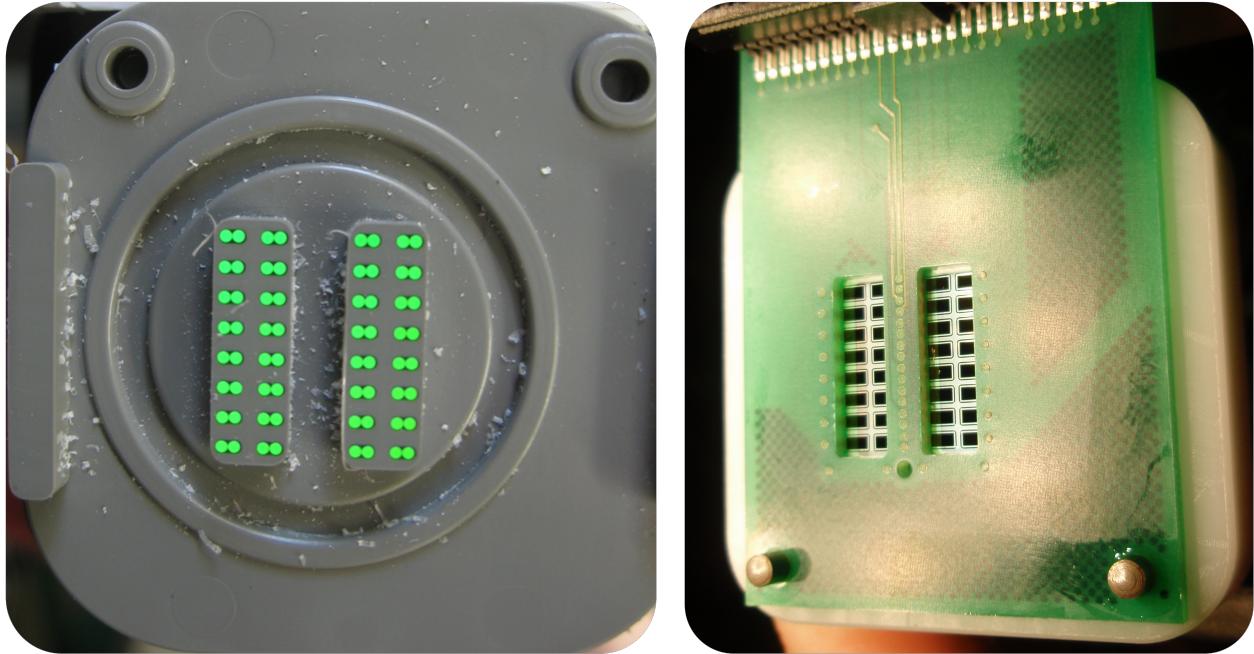


Figure 4.9: **Left:** Close up of the end of a bundle of optical fibers in a module that interfaces with the APD. **Right:** The 32 pixel input array of an APD utilized in NOvA [145].

Once the data is stored in the buffer farm, a decision must be made about whether the data will be saved to long term storage or allowed to drop out of the buffer (effectively being deleted). This decision is made by the DAQ system, which can receive one of three types of “triggers” - signals that tell the system when something significant happens. If the system is triggered, all of the data from that window of time will be recorded to disk for long term storage.

4.3.2 Trigger Systems

The three trigger types are clock triggers, signal triggers, and data-driven triggers (DDT). Clock triggers are pre-set to go off at regular time intervals to capture detector activity. For example, there is a trigger used for calibrating the FD that writes out all activity in the detector ten times per second (or just once per second at the ND).

Signal triggers happen when an external source alerts the DAQ to something “interesting” enough to save long-term. One example of this is Fermilab’s Accelerator Division sending signals to the NOvA detectors every time a beam spill happens. This triggers the DAQ to save all the data taken in a $500 \mu\text{s}$ window centered on the time the beam spill occurred. Another example of a

signal trigger is NOvA’s subscription to alerts from the Super Nova Early Warning System, which alerts the detectors to supernova occurrences picked up by other detectors around the world.

Data-driven triggers (DDT) rely on some basic reconstruction that is run on all of the data as it is stored in the buffer farm. We have the ability to cluster hits together and form tracks to look for event topologies not specific to the NuMI beam spill. If the reconstruction meets certain criteria to indicate that one of those event-types occurred in a given time window, the data corresponding to that window is saved to disk. Some interesting physics phenomena with DDTs set in the NOvA detectors include searches for supernovae, magnetic monopoles, and dark matter.

Once data is committed to long-term storage, it must be tested to ensure quality and processed through NOvA’s “online” and “nearline” monitoring systems.

4.4 Ensuring Good Data Quality

One of the most important aspects of running an experiment as large as NOvA is ensuring that the detectors are running properly and taking high-quality data as often as possible. Data Quality assurance is one of my areas of expertise and the thing I have spent the most time working on outside of the search for tridents.

4.4.1 Detector Uptime

Ensuring that all components of the detectors are turned on and recording data properly is the most basic requirement of ensuring good data quality - after all, any sophisticated monitoring tools are worthless if no data is recorded in the first place! To that end, shifts are taken by all contributing members of the NOvA experiment to monitor the detectors continuously year-round. Experts are available on-call any time a detector experiences a problem that may cause it to stop taking good data for more than a few minutes. The amount of time a detector is on and taking data is called the detector “uptime”. A plot of ND uptime can be seen in Figure 4.10 - a one month time frame is visible on the plot. The average uptime of the ND since 2014 is 98.8%.

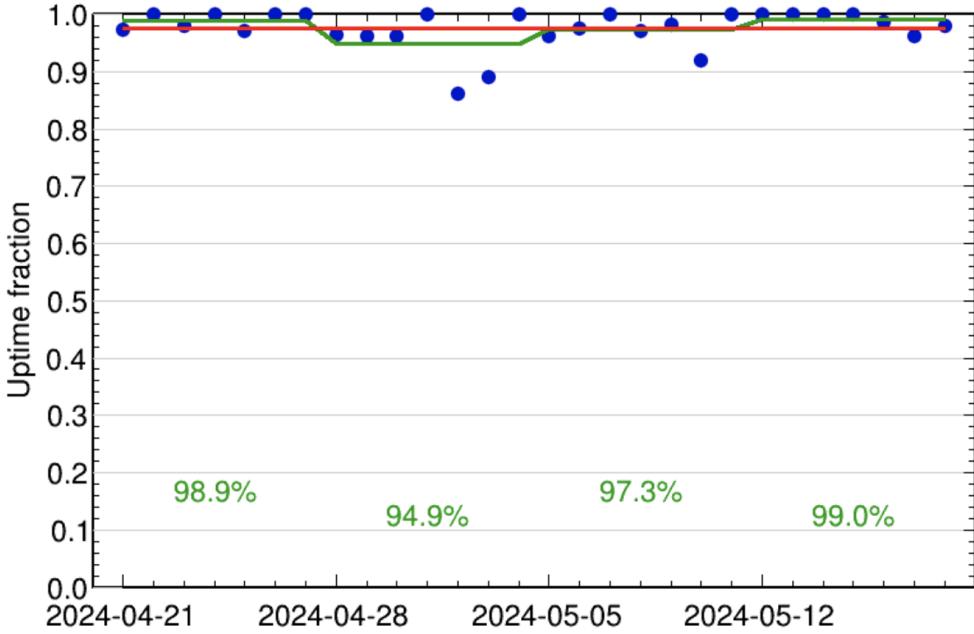


Figure 4.10: Uptime fraction of the near detector from April 21, 2024 to May 25, 2024. Blue dots represent the daily uptime, with the green and red lines showing the weekly and monthly average, respectively. This period saw a 98% uptime fraction.

4.4.2 Neutrino Exposure

As described at the beginning of this chapter, the neutrino beam impinging upon the NOvA detectors is created by colliding protons into a target and letting the resulting mesons decay into muons and neutrinos. Therefore, the neutrino flux seen at the detectors is directly proportional to the number of “Protons on target” (POT) during any given period of time - this is the primary measure of how many neutrinos the detectors are exposed to. The combination of detector uptime and time spent with the beam running is pivotal to maximizing the POT delivered to the detectors.

Figure 4.11 shows the all-time POT exposure for running in FHC (orange dots) and RHC (blue dots) mode recorded by NOvA since the beginning of the experiment through May 31, 2024. Also included are the cumulative FHC POT (orange line), cumulative RHC POT (blue line), and total accumulated POT (gray line). The dataset utilized in the trident analysis is highlighted by the light gray box enclosing data up to 2024, which provides an ND POT exposure of $25.5\text{e}20$ FHC POT and $11.38\text{e}20$ RHC POT.

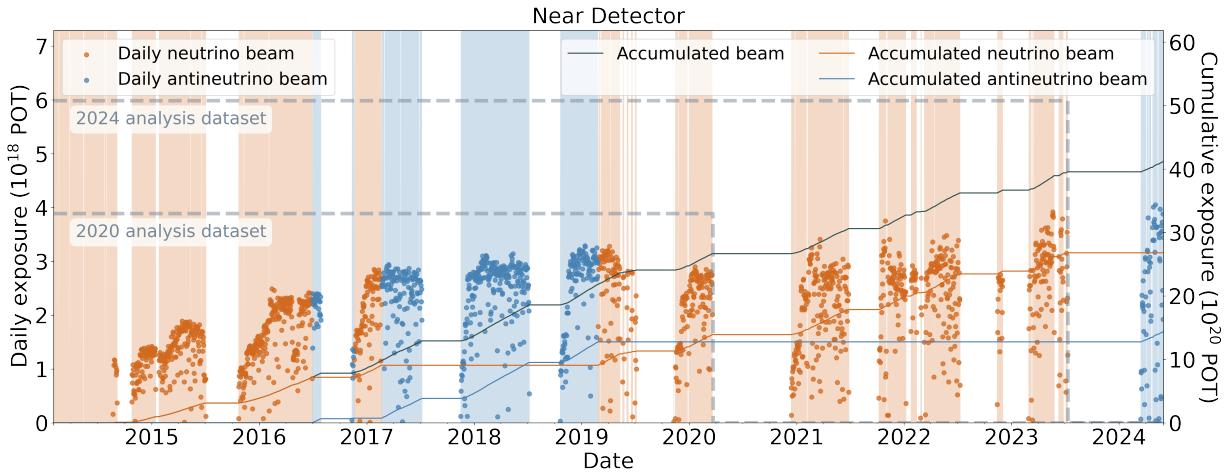


Figure 4.11: Daily POT exposure seen by NOvA since the beginning of the experiment [146].

4.4.3 Online Monitoring

The online monitoring (OnMon) software is a set of tools that allow NOvA shifters to immediately access raw data from the detectors in real time. It receives data directly from the DAQ system, organizes it into meaningful histograms, and stores the histograms for observations of detector health. It stores information such as the energy read out by each pixel, FEB, or DCM. OnMon also stores information about trigger types that occur, and organizes error messages from the DAQ hardware for use in diagnosing issues with the detector readout.

OnMon features an entire suite of plots that are useful in diagnosing issues with the detectors - one example is an “FEB Hit Map” of all the hits deposited into every FEB in a detector. Figure 4.12 shows an ND FEB hit map of only hits corresponding to NuMI triggers for the week of May 18-25, 2024. This is a two-dimensional histogram, with each box on the graph representing a single FEB and the number of hits that FEB recorded being represented by the color of the box as scaled to the axis on the right.

The majority of the boxes in Figure 4.12 are similar in color, indicating that they have seen hit rates between 1000 and 9000 over the last week. However there are a handful of boxes that don’t match - either because they see many more hits than the others (indicated by being red), or by not receiving *any* hits (white). By simply looking at the plot, somebody who is familiar with the detector readout could diagnose the 44th FEB associated with DCM 02 and diblock 03 as being

“hot” (meaning that it has been recording many more hits than the surrounding FEBs, likely due to electronic noise) and the 16th FEB in DCM04/DB02 as not recording any data. That person could then alert experts to these problem so the FEBs can be tagged as machines that may require maintenance.

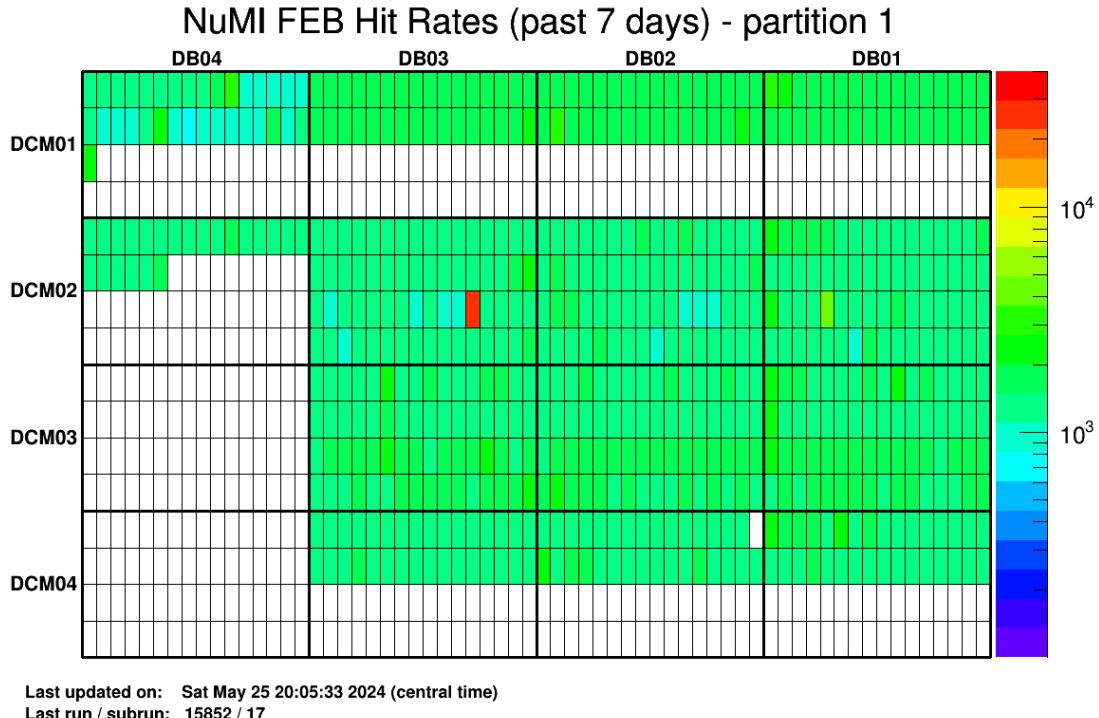


Figure 4.12: FEB Hit map of ND hits associated with NuMI beam triggers over a one-week period.

OnMon creates many useful histograms to monitor various aspects of detector health - about 400 for the ND and 4000 for the FD. These are organized in a viewing GUI for easy access by shifters and experts alike. A data quality expert will scan through many of these plots on a weekly basis looking for extreme outliers.

4.4.4 Nearline Monitoring

The nearline monitoring system shares many similarities with the online monitoring system - it is a tool that creates a suite of plots used for monitoring detector health. The difference between nearline and OnMon is the state of the data when it is processed by the system: OnMon accesses raw detector data live as it is being taken, and the nearline system only accesses data after it has been written to disk by the DAQ system. This makes the nearline a more permanent record of

detector performance and allows for more complex monitoring tools to be run, since the amount of data is reduced compared to OnMon. Most of the analyses being done on the experiment use the files that come out of a nearline process.

Nearline information is used to make general data quality decisions about the detectors, such as if data taken during a specific time period needs to be removed or partially reduced. It is also necessary to understand and profile the data that is taken. For example, processing the ND NuMI triggers through the nearline system can provide insight into things like the detector response as a function of beam power and the synchronization of the timing system with respect to the NuMI beam spills.

Much like OnMon, the nearline system creates a suite of plots used by shifters and experts to monitor the long-term health of the detector. The time between raw data being written to disk and the information showing up in the nearline plots is typically about one hour, so it is one job of a shifter or data quality expert to ensure that the plots are updating properly and look correct.

Criteria for a “Good Run”

In NOvA, a “run” is a block of detector data segmented either by time or by the amount of raw data taken. In the ND one run lasts 24 hours, and in the FD a run typically takes about one hour. These chunks of data are broken down even further into “subruns” which last about 1 hour for the ND and 5 minutes for the FD. Typically when a detector failure or interruption in the DAQ happens, the entire subrun that the problem occurred in must be removed.

The process of determining which runs are “good” (GoodRuns) relies on the output of the nearline and OnMon systems. GoodRuns thresholds are determined as a set of cuts on these files - these cuts check for criteria such as the detector being active for a certain amount of time, the pixel hit rates looking normal, and the timing of detector triggers matching the timing of the external systems that send the triggers.

An example of a plot showing the good and bad subruns for the ND is shown in Figure 4.13. This plot shows that the majority of runs recorded are good, but runs highlighted with specific colors did not pass some data quality cut threshold. For example the subruns highlighted in purple are marked as bad because the NuMI beam was not running as it should have been during those subruns.

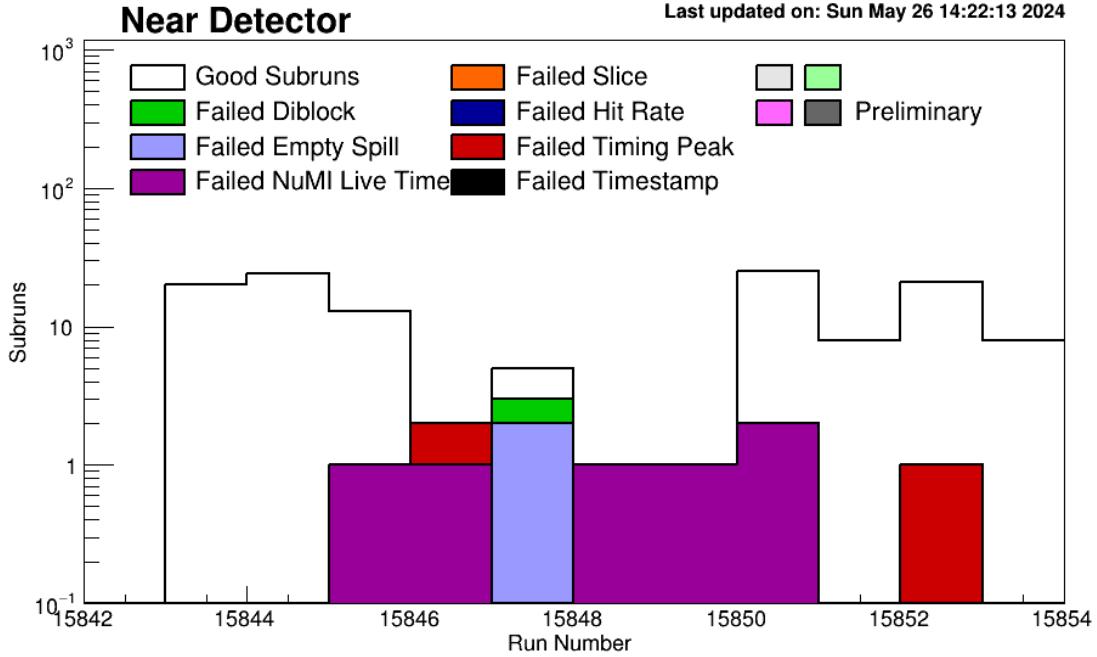


Figure 4.13: ND GoodRuns and failure modes recorded over a one-week period.

Data that makes it through GoodRuns processing is considered the “final” data that is stored long term for use by the experiment. This data is subjected to various levels of processing then used for analyses. These processes involve “calibration” and “reconstruction” of the data, which are the topics of section 3.5 and 3.6, respectively. Before discussing them; however, I must describe what goes into simulating events in the NOvA detectors. This is necessary because analyses depend on understanding what certain events look like when they occur in the detectors, and that information comes from studying simulations.

4.5 Simulating NOvA Data

Analyses in NOvA are developed using Monte-Carlo simulated data (MC) in order to reduce experimental bias. Only after an analysis has been developed, fine-tuned, and tested with MC is real detector data used to perform a final measurement. This section will summarize the steps required to generate the MC used in most NOvA analyses, including the tridents search for which it models the expected background processes. The simulation of trident signal events will be discussed more thoroughly in section 5.2.

Generating MC for NOvA is done in five steps [147], wherein the output of one step is used as the input for the next in order to reduce complexity and save computational time. The steps required to simulate data are as follows: create models of neutrinos in the NuMI beamline, allow those neutrinos to interact with matter in the detector, propagate the final-state particles from those interactions through the detector geometry and model their energy deposition, convert the energy deposited into scintillator photons in the detector, and simulate how the detector electronics respond to the photons.

4.5.1 Simulating the NuMI Beam Flux

Generating NuMI MC begins with 120GeV protons interacting with the graphite target at the accelerator complex. A combination of the GEANT4 [148] and FLUKA [149] simulation toolkits, dubbed “G4NuMI”, simulates the interactions and the resulting hadronic shower coming from the target. Modeling hadronic interactions can be difficult due to the complicated strong force processes that happen and the limited data on hadron production. An approach relying on data from many different experiments to weight the G4NuMI model is used to reduce systematic uncertainty. This is done using the “package to predict the flux” (PPFX), which was developed for Fermilab’s MINERvA experiment [150]. The PPFX weights applied to G4NuMI allow the simulation to properly match the observed data for hadron production in the NuMI energy range.

After modeling the hadron component, G4NuMI uses geometry information about the target, the cooling system, and the magnetic focusing horns to account for the bending done by the focusing horns and the decay of the hadrons into neutrinos and leptons. The information about the number of outgoing neutrinos is stored, as well as the decay point, flavor, direction, energy, and momentum of each neutrino. The hadron parent information is also saved to allow for additional tuning based on hadronic model studies. The final output of this step is a flux file which contains the simulated neutrino flux at each of the NOvA detectors. The simulated energy spectrum for PPFX-corrected NuMI beam flux (both FHC and RHC) for a single spill at the ND can be seen in Figure 4.14, broken down by neutrino type.

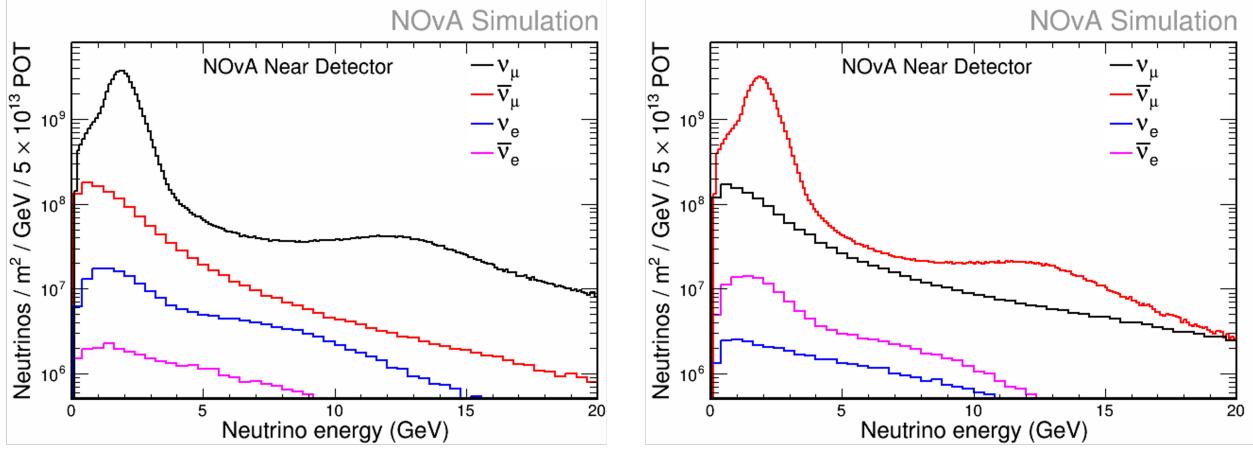


Figure 4.14: Monte Carlo predicted energy spectrum of neutrino flux at the ND. **Left:** Flux for beam running in FHC mode. **Right:** Flux for beam running in RHC mode.

4.5.2 Simulating Neutrino Interactions

The G4NuMI flux file is used as the input to the GENIE neutrino interaction generator [151]. GENIE combines information about neutrino flux and cross sections with a user-supplied detector geometry to determine if a neutrino interaction will occur. A combination of local and global relativistic Fermi gas models are used for the initial nuclear states to be interacted with.

GENIE determines the type, vertex location, and kinematics of all simulated interactions in the detectors. It uses various models for the different interaction types: Qausi-Elastic, Resonant, Deep Inelastic Scattering, Coherent and Meson Exchange Current processes (refer to section 2.2 for descriptions of the interaction types and common models used for each of them). It also handles propagating the particles created by the interaction through the nucleus and modeling final state interactions such as inter-nuclear scattering and absorption.

NOvA must tune the GENIE models in order to obtain better agreement with the observed ND data [152]. The version of genie used in this analysis had its MEC and FSI models adjusted (as summarized in section 2.3) to achieve that agreement. The results of NOvA's MEC tune as a function of four-momentum transfer $|q|$ for FHC and RHC simulations can be seen in Figure 4.15. In these plots, the dashed lines show the default GENIE configuration, the bold lines show the spectra after the models have been tuned, and the black dots represent data taken from the near detector.

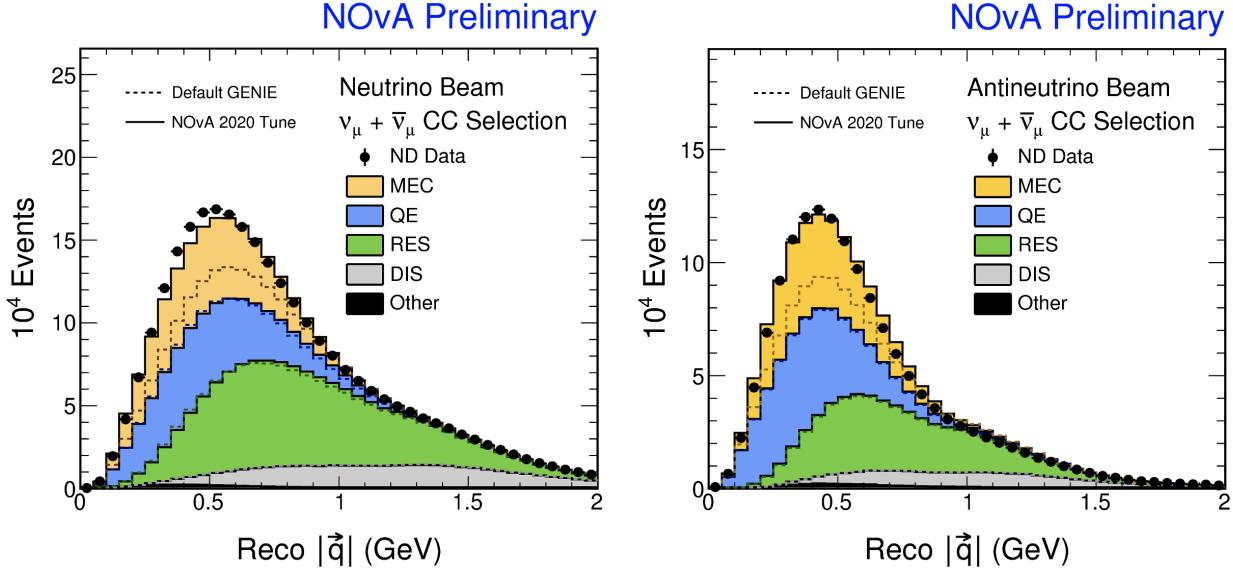


Figure 4.15: Spectra of ν_μ and $\bar{\nu}_\mu$ CC events in the ND [152].

The output from GENIE is a list of final-state particles from the neutrino interaction and their kinematic information stored as four-vectors. The beam flux at the ND is high enough that a large amount of activity from interactions with the rock upstream of the detector is visible during NuMI triggers. These “rock events” must be accounted for to accurately simulate the ND data. GENIE is used to simulate rock events, and the final-state particles of those interactions are overlaid on the other ND GENIE output. After this addition, the total number of interactions simulated to occur in the detector is consistent with what is seen during data taking.

4.5.3 Simulating Particle Propagation

The next step in the simulation chain is feeding the four-momenta for each final state particle into GEANT4, which is used to simulate the trajectory and energy deposition of each particle within the detector geometry. The average amount of energy a particle deposits as it travels a certain distance in the detector (dE/dx) is based on the Bethe-Bloch formula [32]. This formula gives the mean rate of energy loss of a relativistic “heavy” charged particle as it travels through some target medium, commonly referred to as the “stopping power” of that medium.

$$\left\langle -\frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (4.2)$$

The graph pictured in Figure 4.16 shows the Bethe-Bloch curve in the energy regions relevant to NOvA. The minima that occurs for muon momenta around 0.3 GeV is energy deposited by a “minimum ionizing particle” (MIP). That is the point where the particle deposits the least amount of energy as it travels through a material.

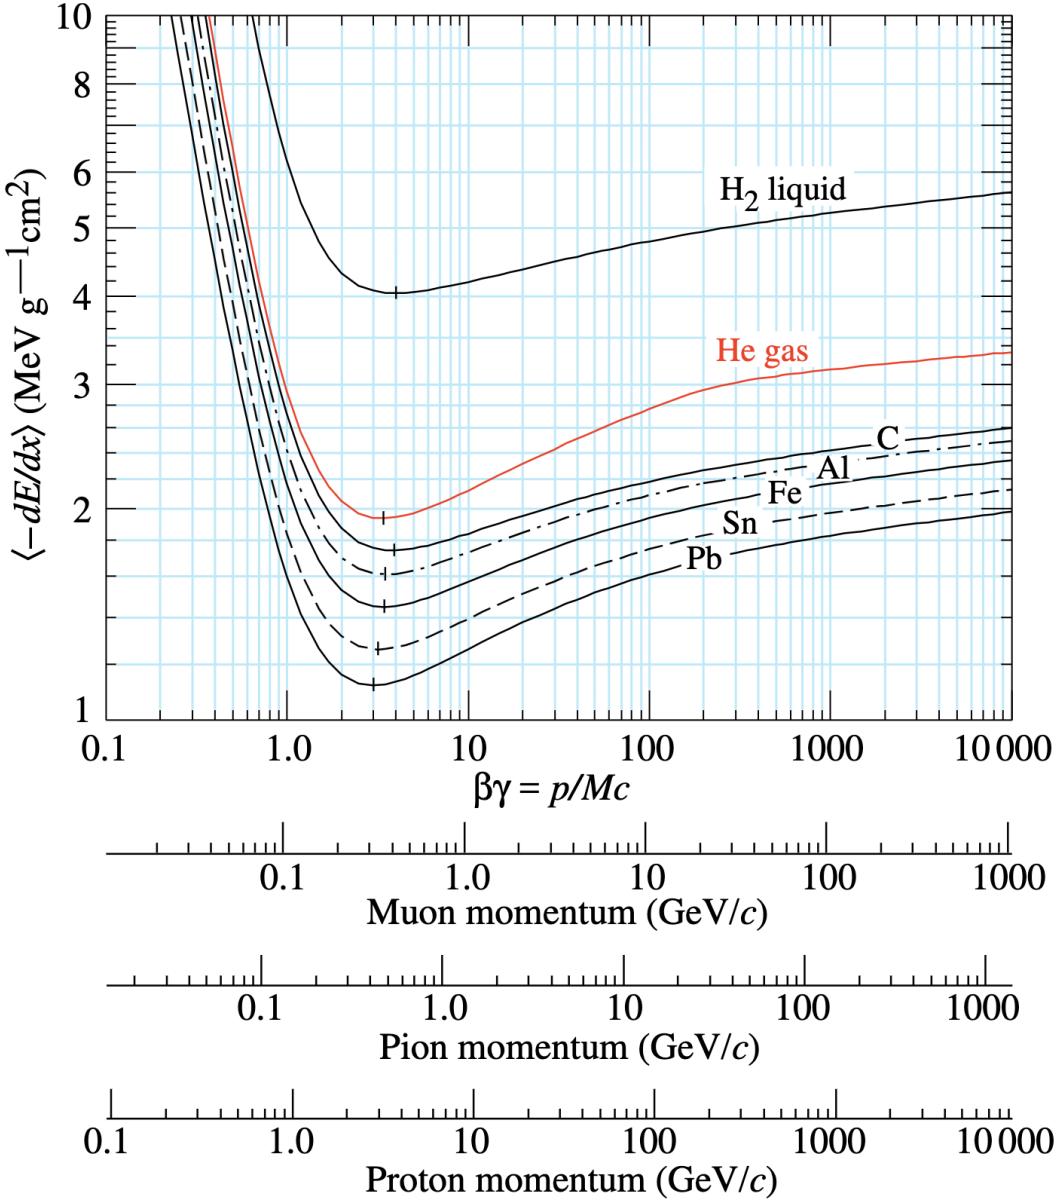


Figure 4.16: Mean energy loss rate in various materials for muons, pions, and protons. The minimum-ionizing point is shown as a vertical dash [32].

The output of running GEANT4 for particle propagation is a list of “Fiber in Liquid Scintillator Hits” (FLSHits) which represent the energy deposited in the detector by each particle involved.

4.5.4 Simulating Photon Transfer & Electronics Readout

NOvA has written a set of internal tools to simulate the response of the detectors to energy depositions starting from the list of FLSHits. The first tool is included in the *PhotonTransport* package - it simulates the process of converting energy deposited into scintillation photons, which are propagated through the wavelength shifting fibers and into the APDs. This package also models the known APD noise in the detectors. In order to match the observed data as closely as possible, the photon transport simulation considers factors such as the deposited energy’s position along the length of the cell, the arrival time based on the length of fiber traveled, and light attenuation within the fibers. It provides a number of photons that are expected to arrive at the APD as its output.

The second tool is included in the *ReadoutSim* package - it accounts for the APD efficiency and simulates the electronic pulse produced by the APDs in response to the *PhotonTransport* photons. It then performs the processes to create digitized waveforms as done by the FEBs when they receive APD signals, including accounting for electronic noise. The end product is a ROOT [153] output file in the same format as real data files with additional “truth” information about interactions describing exactly what was simulated.

4.5.5 Final Simulated Data

Figure 4.17 illustrates the final product of the simulation chain [154]. It contains three types of interactions commonly seen in the NOvA ND with neutrino of energies of 2 GeV. Each panel shows a quasi-elastic event in a single view of the detector, with the cell hits colored by the amount of charge deposited.

The top panel contains a ν_μ charged-current (CC) event with the characteristic topology of a long, straight, MIP-like track representing the muon, and a short but high-energy hadronic deposition from the proton. The middle panel shows a ν_e CC event. The electron deposits energy in a wide shower, clearly distinguishable from the long track of the previous muon. A distinct feature of electron energy deposition is that the amount of energy deposited per plane peaks towards the center of the electron’s trajectory rather than at an end. NOvA was designed to identify electron showers very well, since the primary oscillation measurement relies on observing ν_e CC events in the far detector.

The bottom panel shows a neutral current (NC) interaction with the proton and a single π^0 as the output. The π^0 decays into two photons with a 98.8% branching ratio, which it typically does very quickly. The photons travel through the detector without interacting until they pair produce into an e^-/e^+ pair, which do produce scintillation light. These overlapping particles tend to resemble ν_e CC events quite closely, with the characteristic photon-gap as one of the best ways to distinguish the two.

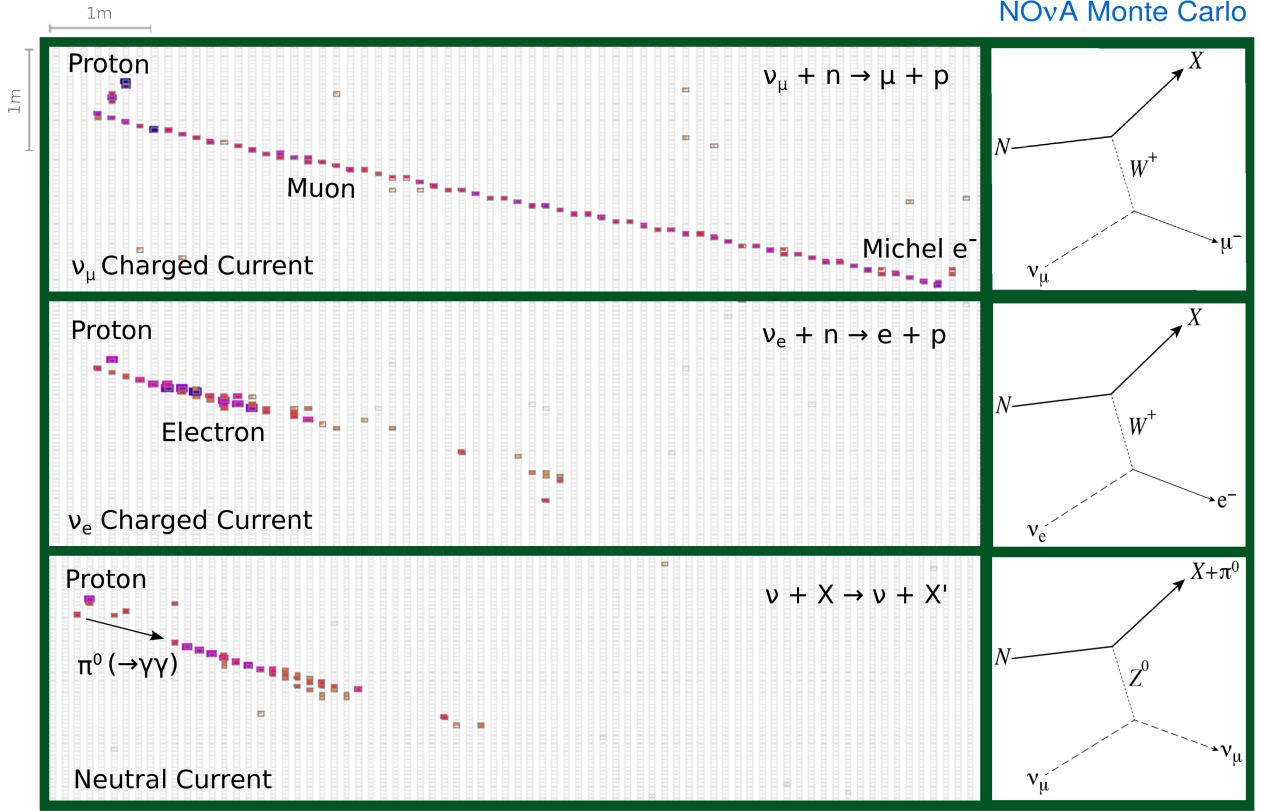


Figure 4.17: Event displays of simulated interactions with 2 GeV neutrinos in the ND [154].

These panels all feature quasi-elastic events with just two particles. Other types of interactions manifest differently in the detector: Resonance and MEC events look similar but have more hadronic energy deposited at the beginning of the event. Deep inelastic scattering events are very messy - typically featuring a muon track plus a large amount of hadronic energy. The files containing the simulated data are then sent on to have calibration and reconstruction performed on them, just the same as real detector data.

4.6 Calibration

Once data has either been taken by the detectors or simulated it must be calibrated. Cell hits in the NOvA detectors must be processed to ensure that the energy deposited into the detector is properly converted into energy read out by the DAQ, and ensuring that the timing of hits that occur are consistent across all of the read-out electronics [155].

NOvA's energy calibration consists of two phases: a relative calibration that corrects for light attenuation and aging of the detector [156], and an absolute calibration that converts energy deposition recorded by the APD into units of GeV [157]. These are both done using cosmic-ray muons because they provide uniform energy deposition across the detector. The calibration uses the energy deposition rate predicted by the Bethe-Bloch formula to match the photons seen in each cell into a measurement of energy in GeV. Detector aging effects are corrected by monitoring the mean energy response in a cell over time and continually correcting it back to the previously calibrated state. Calibration is performed on both data and MC, and the MC truth information is used to compare the energy that was truly deposited to the energy that was “observed” by the detector.

Figure 4.18 shows the profiles of reconstructed and true energy ratios of hits in the X and Y views of the ND as a function of the distance away from the center of a cell where the hit took place, W . The red curve in the plots show the ratio of observed to true energy before calibration, and the blue curve shows their ratio after [158]. A ratio of 1 means that the energy observed in the detector matches the energy that was deposited.

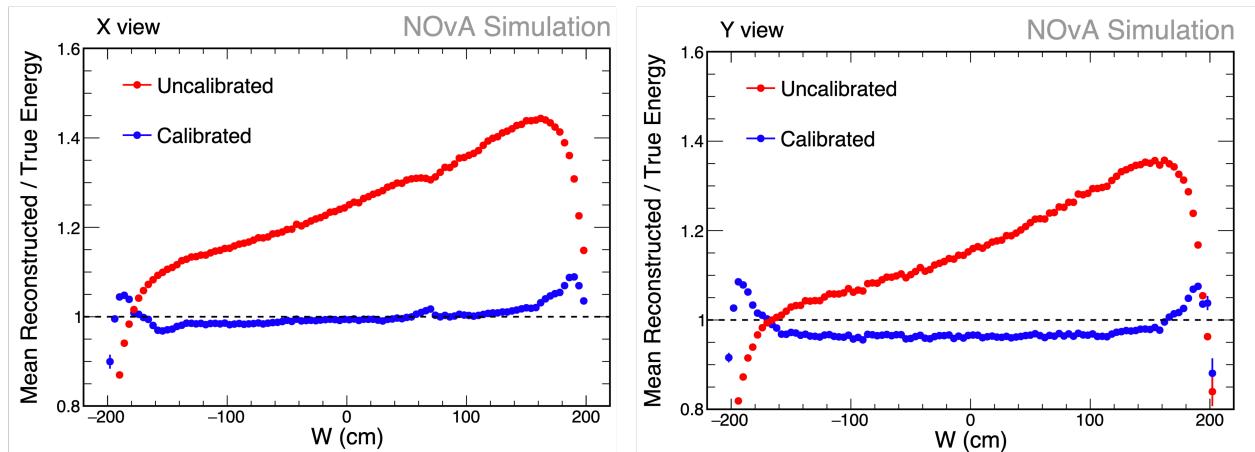


Figure 4.18: Ratio of observed to true energy deposited in the detector as a function of distance from deposition to cell center [158].

4.7 Event Reconstruction

Once data has been taken and MC has been generated, it must be sent through “reconstruction” - a process which takes the calibrated cell hit data associated with an event and turns it into a set of computational objects that have position, trajectory, energy deposition, and timing information about that event [159, 160].

The goal of most analyses in the ND is to look at the end result of an event and work backward to figure out properties about the interaction that caused it. This is done through reconstruction by determining things like the types of particles involved in the event, the direction they were going, and how much energy they had. NOvA relies on these final state properties to extract the interaction information because the incoming neutrinos themselves don’t leave tracks in the detector. This section describes event topologies in the NOvA ND and some of the standard methods used to reconstruct them.

4.7.1 Event Topologies

Trident Event Topologies

The signal of this dissertation is neutrino tridents, as described in chapter 3. Specifically, I am interested in observing “dimuon” tridents - events which result in a final state of two muons in the detector with no additional hadronic energy present. The typical dimuon event topology in the ND looks like the event presented in Figure 4.19. The following features are clearly visible:

- Two long, muon-like tracks contained within the side walls of the detector.
- A narrow opening angle between the two tracks.
- A very small (ideally zero) amount of recoil energy around the start of the interaction, denoting a lack of energy deposited from hadrons.

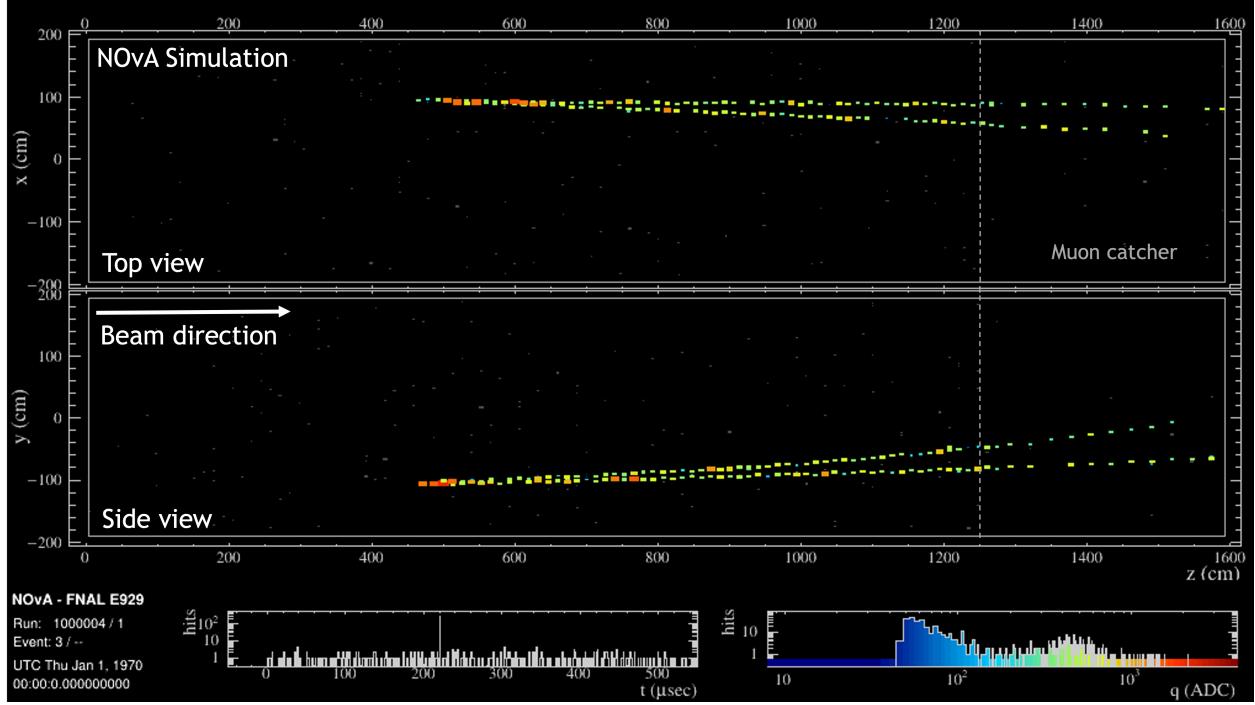


Figure 4.19: An event display featuring a simulated dimuon trident event.

Background Event Topologies

The backgrounds to observing tridents in NOvA are typically neutrino events with a final state of two highly energetic tracks and little or no hadronic energy present. As mentioned in section 2.2, any of the more standard ND interactions have the potential to mimic the trident topology.

Figure 4.20 shows an example of a background event that looks enough like a trident to be selected by some basic selection criteria. The MC truth information shows that the true interaction was a resonance event with a muon, proton, and highly energetic pion in the final state. The proton deposited its energy very quickly and its hits overlap the other particles. So the final state only has two tracks and looks very similar to a trident event, with more energy at the interaction vertex and a much wider opening angle. A more detailed discussion of specific background events will be included in chapter 5.

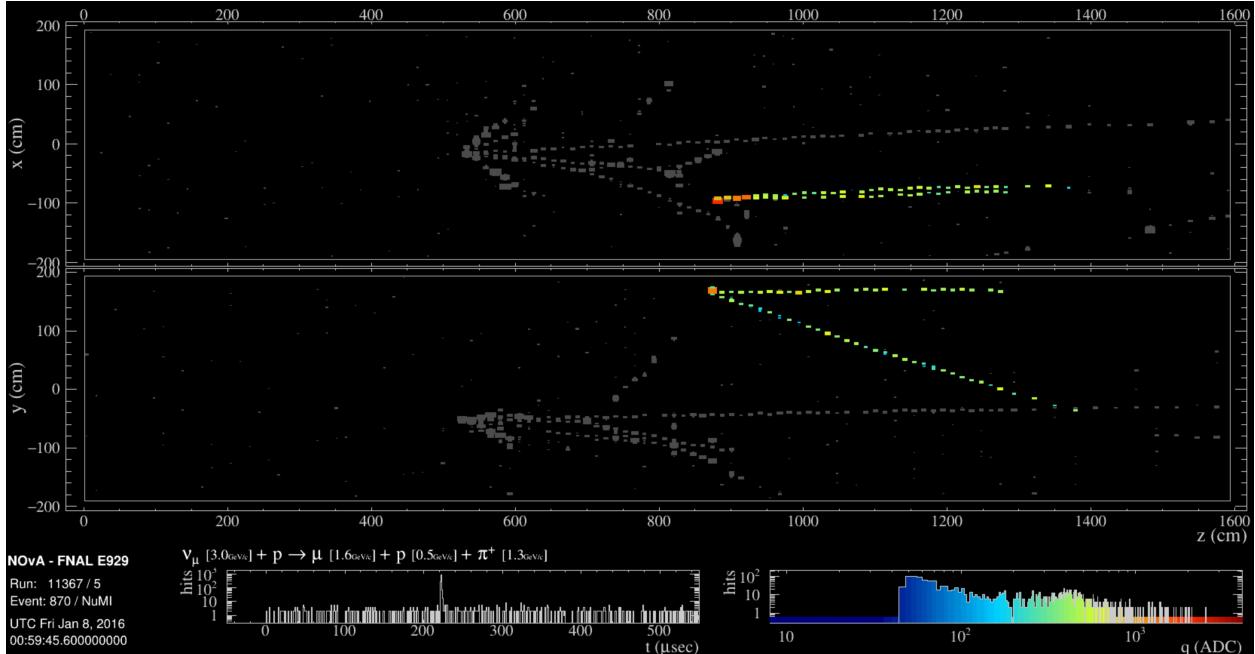


Figure 4.20: An event display featuring a simulated resonance event with trident-like topology.

4.7.2 The Reconstruction Chain

Cell hit information originating from the same event must be grouped together to reconstruct neutrino interactions. Once hits are associated with one another, the kinematic information of the event can be extracted and used to form analysis tools. The following set of reconstruction steps are run on all data that pass data quality cuts - the algorithms are developed and tested on either MC with available truth information or on very well understood data to make them as accurate as possible.

Slicing

The first step in reconstruction is to identify which cell hits in the detector originated from a single interaction - a process called “slicing”. When a NuMI spill trigger occurs at the ND, all of the data from the 550 μ s readout window is combined into a single event. However, interactions coming from NuMI neutrinos take place only during a 10 μ s window, as discussed in section 3.2.

The first step of slicing is simple - remove any events in the detector that occurred outside of the 10 μ s NuMI spill window. This will remove the majority of noise in the detector, and is sometimes sufficient to isolate singular events. However, it is possible that multiple neutrino events

happened so close together that they are indistinguishable from timing alone. Therefore, both time and space information must be used. The algorithm used to do this looks at each hit in the detector and calculates a density for surrounding hits that relies on their temporal and spatial separation. Depending on how isolated or not each hit is, it may be determined that a hit is causally linked to other hits around it, resulting in those hits being collected into a single “slice”. The slicing algorithm is run separately for the x and y views, then those slices are combined into a 3D slice using the average z-position and time values associated with possible pairs of slices between views.

Vertexing

The next step in reconstruction is to determine the actual point where the interaction assigned to a particular slice occurred, known as the “vertex” of the interaction. This process starts with the assumption that final state particles of an interaction propagate outward from a single vertex. The vertex is identified through the creation of “Hough lines”, which are lines connecting each pair of hits separated by some minimum distance. Hough lines are used to find the vertex by means of acting as a seed for an “elastic arms”-based vertex finder [161, 162]. Elastic arms searches for the optimal vertex location by finding sets of lines, or “arms”, that describe the slice hits. The set of arms that minimizes an energy cost is used to determine the estimated location of the interaction vertex.

Figure 4.21 shows a NuMI spill in the near detector with no reconstruction applied in the top event display. This event features a few neutrino interactions and a muon from a rock event streaming through the detector. In the middle event display, these events are all clearly distinct from one another in their individual slices. The bottom event display shows the Hough lines and interaction vertex (as a cross) determined by them associated with the hits in each slice. The interaction vertex in this event was on the exact same cell and plane as the true vertex.

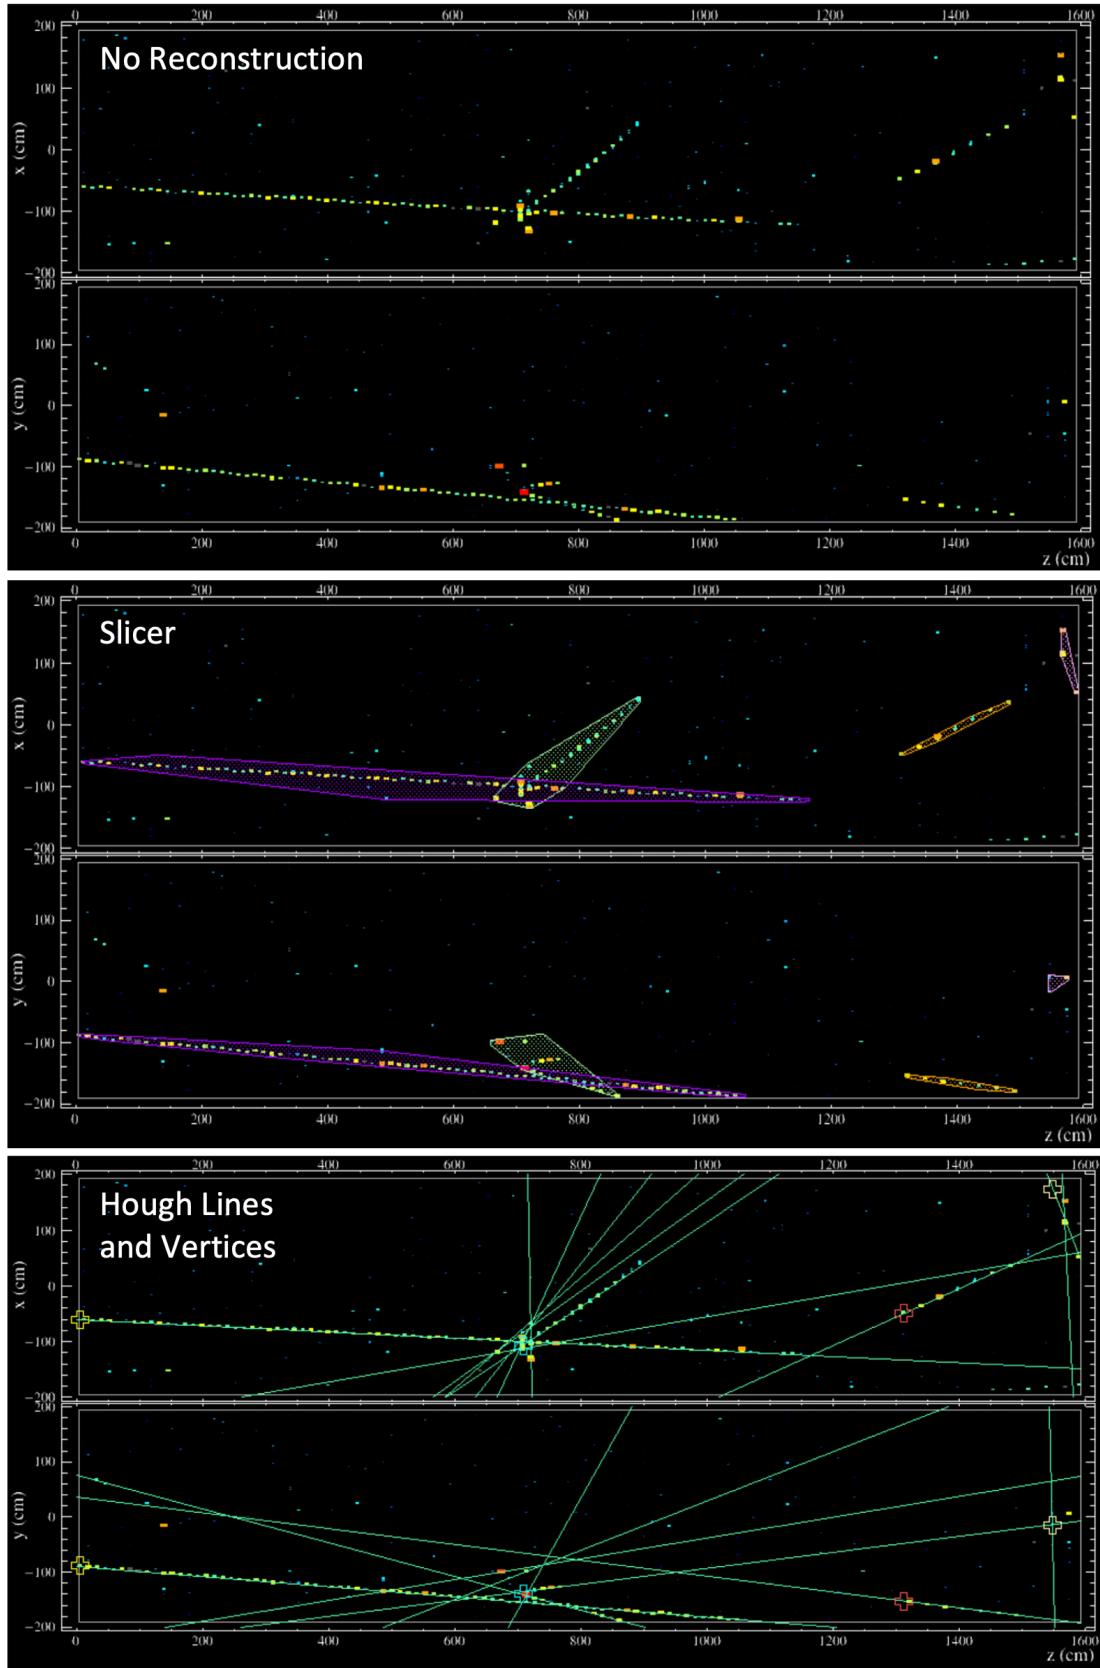


Figure 4.21: Simulated events in the NOvA ND with various steps of reconstruction applied.

Clustering

Once the slice hits and vertex of an event have been identified, the next reconstruction step is to group together (or “cluster”) hits caused by individual particles emerging from the vertex. These reconstructed cluster objects fall into two categories in NOvA: “prongs” and “tracks”.

Prongs are the more simple of the two, and sometimes act as a basis from which to form tracks. Prongs are clusters of hits with a starting point and initial direction with relation to the vertex. They allow a relatively “loose” collection of hits to be clustered together, making them powerful for clustering electromagnetic showers. The method by which prongs are formed in NOvA utilizes an adapted version of the fuzzy k-means clustering algorithm [162, 163], and is referred to as “fuzzyk”. In this method, “k” refers to the number of clusters formed and clustering being “fuzzy” means that an object can have membership in multiple clusters at one time.

Prongs are initially formed separately in the XZ and YZ detector views. To find the prongs in each view, all of the cell hits in a slice are converted to an angle with respect to the vertex. Each hit is assigned an angular uncertainty determined as a function of the distance from the vertex to the hit. Angles with high hit density are used as cluster centers, and any slice hits within a certain threshold of that region are added into the cluster. This process is repeated iteratively until no more hits pass the threshold for being included a prong.

The resulting two-dimensional prongs are matched between detector views to create three-dimensional prongs. To do so, a matrix of all possible prong-pair combinations between views is formed. The endpoints of prong-pairs must be within one plane of each other to be considered a match candidate. Any match candidates that survive are formed into temporary 3D prongs and the energy deposition rate along the prong length is calculated. A measure of similarity between the energy distributions in each view is performed via a Kuiper test. Similarity scores are assigned to each match, and permanent 3D prongs are formed out of pairs that match each other closely. Prongs for which a suitable match cannot be found remain as 2D prongs in a single detector view.

Figure 4.22 shows the results of running fuzzyk on the same NuMI spill from Figure 4.21, zoomed in to the interesting neutrino interaction towards the center of the detector. In this event display four pairs of 3D matched prongs have been generated, and can be identified as the pairs of lines with the same colors in the two detector views.

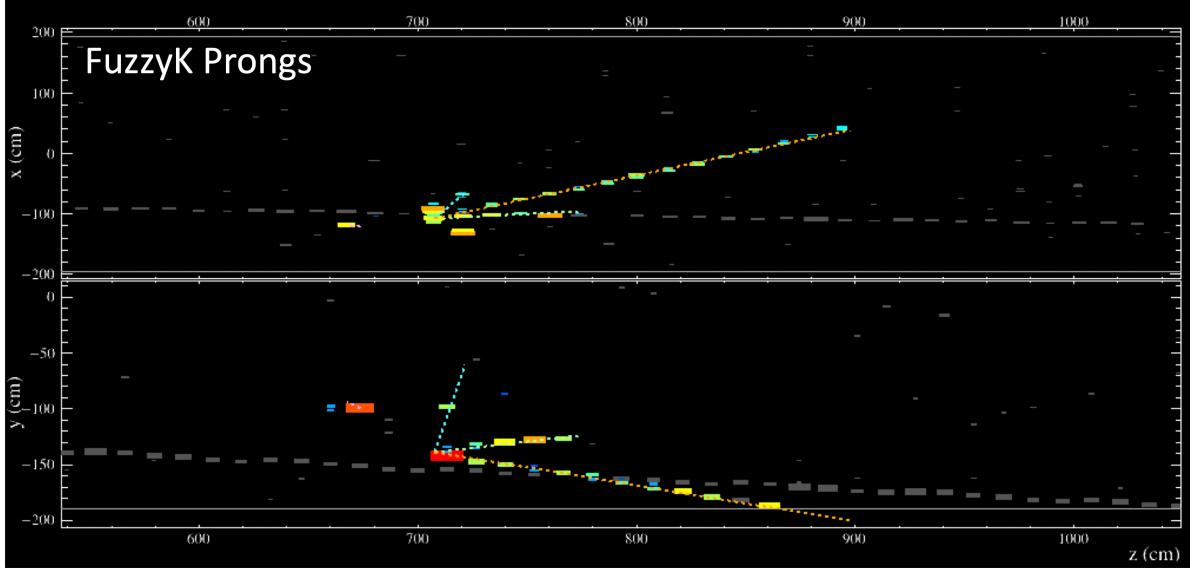


Figure 4.22: FuzzyK prongs overlayed on top of hits in the ND.

The other type of cluster objects commonly used are tracks, which are similar to prongs but contain additional information. In addition to having a starting point and initial direction, tracks have fully reconstructed trajectories describing the path a particle takes through the detector. Prongs are good for reconstructing showering particles like electrons or hadrons, and tracks are good for reconstructing long, straight hit topologies such as those caused by muons. Therefore, tracks are the more relevant tool to study dimuon trident events, since we expect to see two muons and no electrons or hadrons in the final state.

A custom track-finding algorithm was developed for the trident analysis, and is discussed thoroughly in section 5.4. However, to fully describe the workings of the custom tracker an understanding of the two existing NOvA trackers is required. The two tracking algorithms developed previously are the Kalman tracker and the Break Point Fitter.

4.7.3 Track-Finding Algorithms

Kalman

The Kalman track-fitting method implemented in NOvA [164] begins at the individual slice level, meaning it does not depend on vertices or prongs. The algorithm is based on a Kalman filter, which is a common way of estimating the true value of some measurement based on the amount of error expected in the measurement and the amount of noise in the system [165]. In this case, the desired

measurement is the precise position of some particle along its path given the recorded detector hits. The method assumes that each state of a system is linearly related to the previous state because non-showering particles are expected to follow a straight line trajectory, only deviating due to very small scattering processes.

Of particular interest to the trident analysis is the method used by the Kalman tracker to *find* tracks, as the trident track *fitting* step is done by another algorithm. Track finding is the process of determining a general list of hits in a slice that likely came from a single particle. This is done by starting at the downstream end of the slice and forming track “seeds” by connecting a straight line through combinations of two hits separated by less than 4 cells. The seeds are used to estimate the position and slope of the track, then a Kalman filter is applied to predict where the next upstream hit will be. If a hit exists close to the predicted location it is included in the track. When a hit is added to a track it updates that track’s position and direction, which are then used to repeat the process until no more hits can be added.

Break Point Fitter & Lutz Optimum Track Fitting

The Break Point Fitter (BPF) tracking algorithm was developed for NOvA as an alternative to Kalman [166]. It takes Elastic Arms vertices and FuzzyK prongs as inputs to determine particle trajectories. BPF is based on the Lutz optimal track fitter [167], which allows for Coulomb multiple scattering of a particle by breaking the particle’s path up into segments separated by scattering planes, or “break points”. It is considered an “optimal” tracker, which means it considers all available information when reconstructing a track: cell hit locations and uncertainties, scattering information, and correlations between upstream and downstream hits.

The BPF implementation of Lutz’s optimal track fitter works by projecting the hits into a “track basis” so that the z-axis points along the track direction as best it can. Measurements of the track in z are made at “scattering planes” that are orthogonal to z at fixed locations - these yield the x_i and y_i positions of the track with uncertainty σ_i . The track is allowed to scatter by an angle up to some maximum α at M scattering planes. Figure 4.23 shows a schematic of the Lutz tracking model.

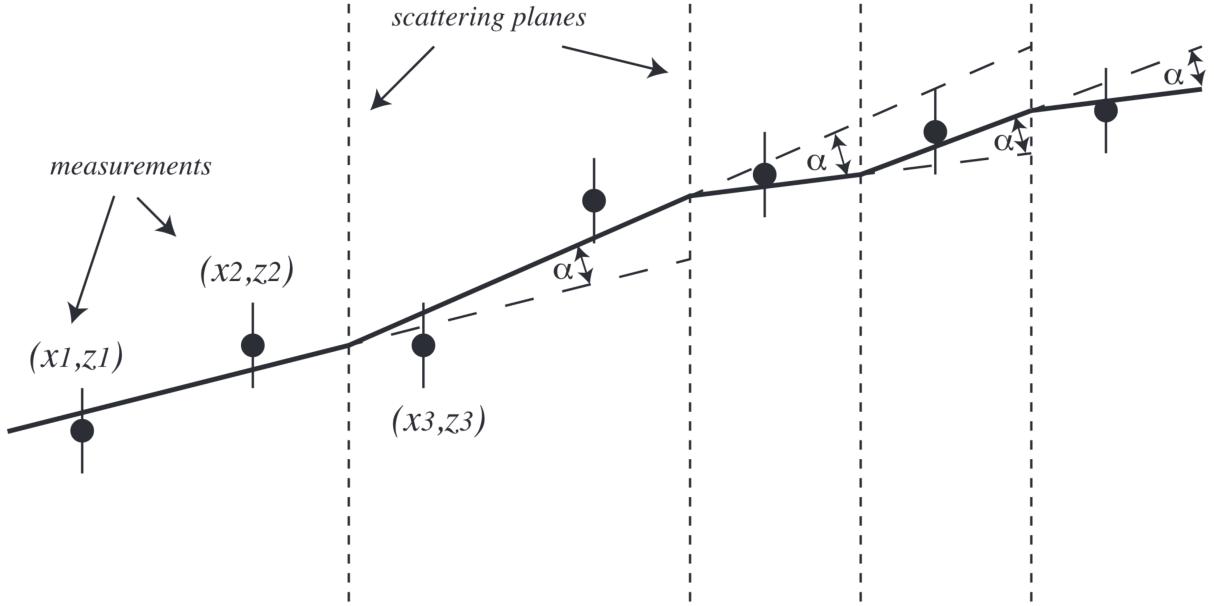


Figure 4.23: A diagram explaining elements of the optimum track fitting model used by BPF [166].

Determination of the scattering planes in the NOvA detectors utilizes a Coulomb multiple-scattering model. By starting at the downstream end where the particle's kinetic energy is assumed to be zero, the algorithm steps upstream along the particle's trajectory and determines the distance traveled by the particle in terms of detector radiation lengths, X_0 , and energy deposited, dE/dx . Scattering planes are placed along the trajectory when the possible scattering angle for the amount of energy the particle has begins to get so large as to be outside of a single cell, or if the distance between planes exceeds $2X_0$.

Once the location of the scattering planes are determined, the location of the track ξ when it coincides with the i^{th} or j^{th} scattering plane can be expressed as

$$\xi_i(z) = a + bz_i + \sum_{j=1}^M \alpha_j(z_i - z_j) \Theta(z_i - z_j), \quad (4.3)$$

where a and b are the intercept and slope of the initial track, respectively, α_j is the scattering angle at the j^{th} scattering plane, and

$$\Theta(z_i - z_j) = \begin{cases} 1 & z_i \geq z_j \\ 0 & z_i < z_j \end{cases} \quad (4.4)$$

is the Heaviside function, which ensures that only the upstream scattering angles (where $z_i \geq z_j$)

contribute to the trajectory. Figure 4.24 shows these parameters with respect to how hits show up in the NOvA geometry.

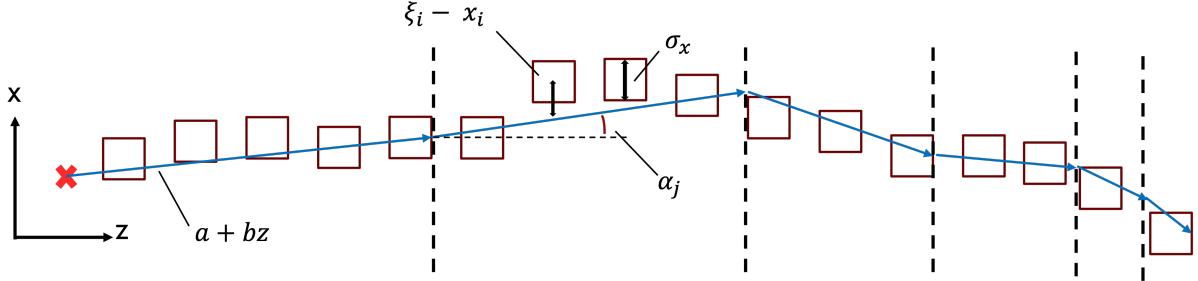


Figure 4.24: A diagram explaining elements of the optimum track fitting model used by BPF in terms of the NOvA geometry [160].

The track trajectory expressed in Equation 4.3 is optimized using a χ^2 goodness of fit test, where

$$\chi^2 = \sum_{i=0}^n \frac{(\xi_i - x_i)^2}{\sigma_{x_i}^2} + \sum_{j=0}^n \frac{(\beta_j - \alpha_j)^2}{\sigma_{\alpha_j}^2}. \quad (4.5)$$

Here σ_{x_i} is the uncertainty in x (or y) about where a hit occurred within a cell, σ_{α_j} is the *RMS* scattering angle computed for the j^{th} scattering plane, and β_j is the *expected* scattering angle at the j^{th} plane. Since NOvA does not use a magnetic field and the scatters are expected to be completely symmetrical, β is actually zero at every scattering plane in BPF. When this χ^2 is minimized by setting its second derivative equal to zero, a set of linear equations is obtained and solved to determine the optimal track parameters to match the detector hits.

Existing Tracker Performance on Simulated Trident Events

When the methods *do* reconstruct the two muon tracks properly, they perform well. However, neither of the trackers were developed to expect two long muons, so they frequently fail to reconstruct one of the tracks. The Kalman method produces two tracks on trident events approximately 60% of the time, and the BPF method produces two tracks only about 30% of the time. Even when two tracks were produced, the narrow opening angle would often cause confusion, resulting in the tracks not capturing entire muon trajectories well. Event displays showing instances of Kalman and BPF tracks applied to simulated trident events are visible in Figure 4.25.

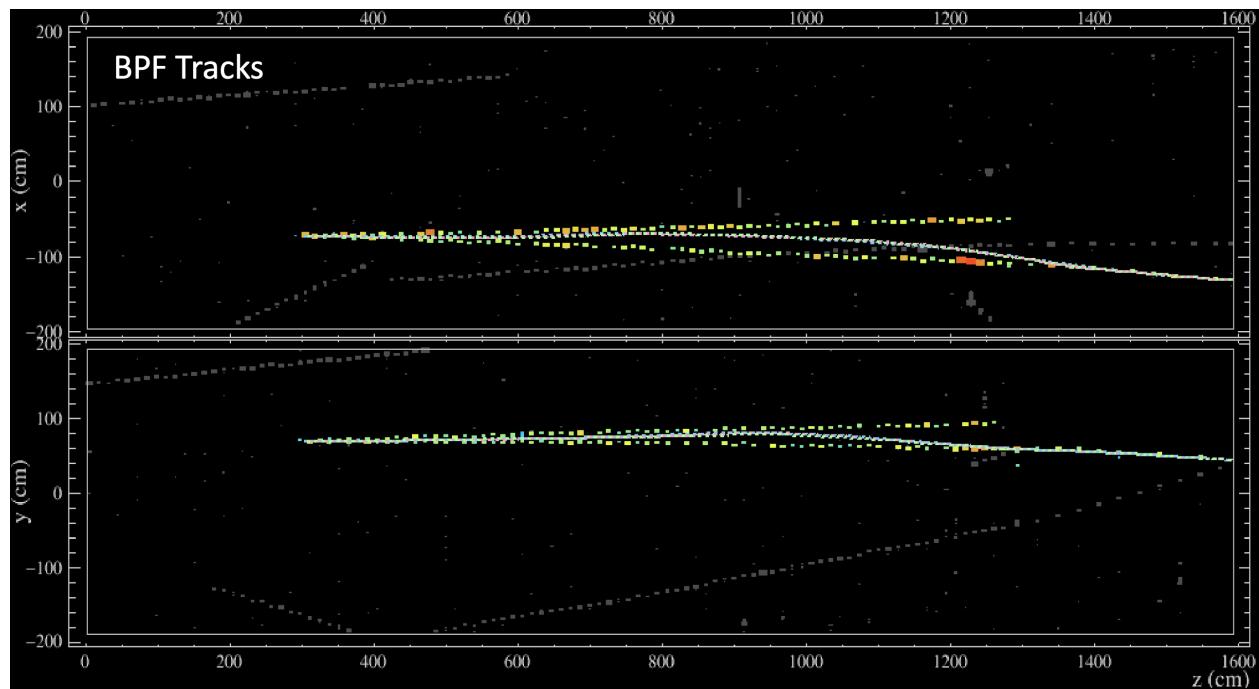
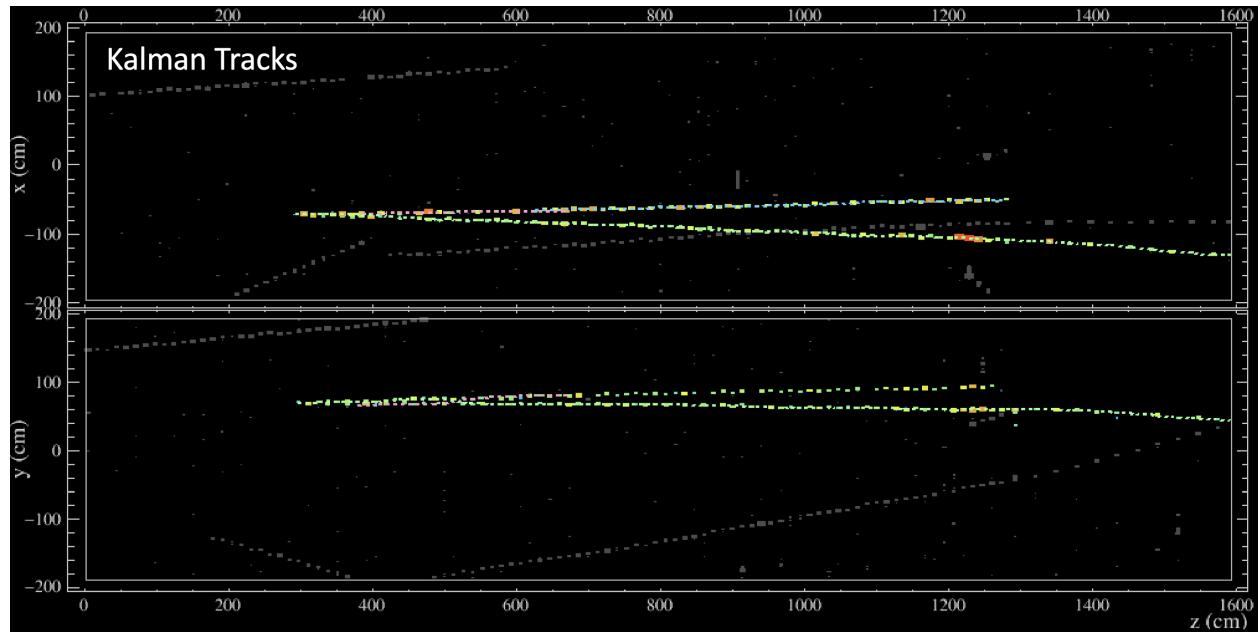


Figure 4.25: A simulated trident event with Kalman and BPF tracks overlayed.

The poor reconstruction of existing methods on simulated trident events motivated the development of a new dimuon-specific tracking algorithm, the KLutz tracker. The details of developing, testing, and utilizing the KLutz tracker for the tridents analysis are described in chapter 5.

Chapter 5

Developing the Tridents Analysis

This chapter describes the search for neutrino trident events in the NOvA near detector, laying the foundation for a future measurement of the trident cross-section. The analysis is built upon a simulation framework provided by Wolfgang Altmanshoffer and modified for use in NOvA. I describe the motivation and development of the custom dimuon reconstruction algorithm, and walk through the various steps of the signal selection used to identify trident candidates.

5.1 Analysis Roadmap

Neutrino trident production is a rare process, with a cross section about five orders of magnitude smaller than the total neutrino cross section in NOvA. Isolating singular signal events requires the background dataset to be reduced to just 0.001% of its original size. Here I outline the steps necessary to achieve such a background reduction before elaborating on each step throughout the next two chapters.

- **Trident Simulation**
 - Implement a custom trident event generator in NOvA.
 - Validate that the generator produces theoretically accurate trident events.
- **Initial Data Reduction**
 - Utilize existing tools to develop a basic selection on MC simulated data.
 - Reduce the initial background dataset to about 5% of its original size while maintaining about 98% of trident events.
- **Dimuon-Specific Reconstruction**
 - Build a custom tracker optimized for dimuon trident topologies.
 - Run the tracker on all simulated tridents and the 5% background dataset.

- **Signal Selection**
 - Develop a second selection utilizing the custom reconstruction.
 - Train and optimize boosted decision trees to isolate signal events.
 - Determine the optimal cut on the final BDT output.
- **Data Unblinding Procedure**
 - Estimate systematic uncertainties from flux prediction and cross section models.
 - Incrementally unblind real data based on BDT output sidebands.
 - Check that data/MC match within uncertainties at each step.
- **Background Estimation**
 - Use a sideband fit extrapolation to mitigate the effects of limited MC statistics.
 - Apply the extrapolation to both MC and real data to estimate final background.
 - Quantify systematic uncertainty on the fits and choice of fit function.
- **Analyze Final Result**
 - Unblind the final signal region.
 - Check the selected events for trident-like properties.
 - Calculate confidence intervals on the number of observed trident events.

5.2 Simulation

To search for tridents in the NOvA ND, a set of “selection criteria” must be developed and applied to the full dataset of neutrino interactions. These criteria are based on the characteristics of an event and act to distinguish signal events (trident interactions) from background events (everything else in the detector).

Determination of the selection criteria was performed on simulated datasets before being applied to real data so that the end result remained unbiased. The default simulation utilized by NOvA (described in section 4.5) does not account for trident interactions, so two different sets of simulated data were required in the development of selection criteria for this analysis. The simulated

background dataset was NOvA’s default event simulation based on the GENIE event simulator. This dataset includes simulation of the common neutrino interactions described in section 2.2, as well as other less common interactions that may occur in the ND.

5.2.1 Simulated Trident Dataset

The dataset of simulated trident events used for the signal of this analysis was generated with code graciously provided by Wolfgang Altmanshoffer, who originally created the generator for the DUNE experiment [114]. The generator calculates trident cross sections for a desired interaction channel and neutrino flux, then uses Monte Carlo methods to create a list of simulated interactions in the standard high energy physics (stdhep) format. These stdhep files contain information for the generated interactions such as the energy, momentum, and timing for each particle involved.

The generator was initially modified for implementation into the NOvA software framework by Kelli Michaels (a previous undergraduate student at IU) and Mark Messier [168]. Their work consisted of adding the various elements that compose the NOvA detectors, assigning a NuMI spill time to the events, and sorting out interface issues associated with the change in software environment. Mark and I then imported the NOvA beam flux from G4NuMI into the simulation.

My work, documented in the following section, was to validate that the simulated cross sections on the newly added NOvA detector elements matched theoretical predictions, and calculate weights to properly scale the number of events generated to the expected number of events in the real dataset. I then generated a full signal dataset by running the resulting stdhep files through NOvA’s particle propagation, detector response, reconstruction, and common analysis formatting tools.

Trident Cross Section Validation

Kelli and Mark modified Wolfgang’s code to generate tridents on the various target nuclei that compose the NOvA ND. I consulted NOvA technical documents and Matthew Strait, an expert on the NOvA detector operations, to determine that the overall composition of the NOvA detectors (referred to as the NOvA “soup”) is that described in Table 5.1 [169]. The stable isotopes of each element are expected to be “present in their standard ratios” within the detector, meaning that I had to keep track of trident cross sections on 23 possible target nuclei while handling the simulated data.

Composition of the NOvA Near Detector	
Element	Percent of Detector
Hydrogen	10.76%
Carbon	66.71%
Oxygen	2.98%
Nitrogen	0.03%
Sulfur	0.1%
Chlorine	16.05%
Titanium	3.22%
Tin	0.12%
Calcium	0.03%
Sodium	0.003%

Table 5.1: Chemical composition of the NOvA near detector [169].

Before generating large datasets utilizing the modified code, I had to ensure that the newly implemented cross sections matched theoretical predictions. The literature on trident cross sections is limited so I was unable to cross-check the calculations for *every* element in the NOvA soup. Fortunately, I was able to find calculations for a handful of elements to test the interpolation between carbon and argon. The results of two such cross-checks on oxygen and iron are shown in Figure 5.1. The average difference in the cross sections is 2.54% on oxygen and 5.37% on iron.

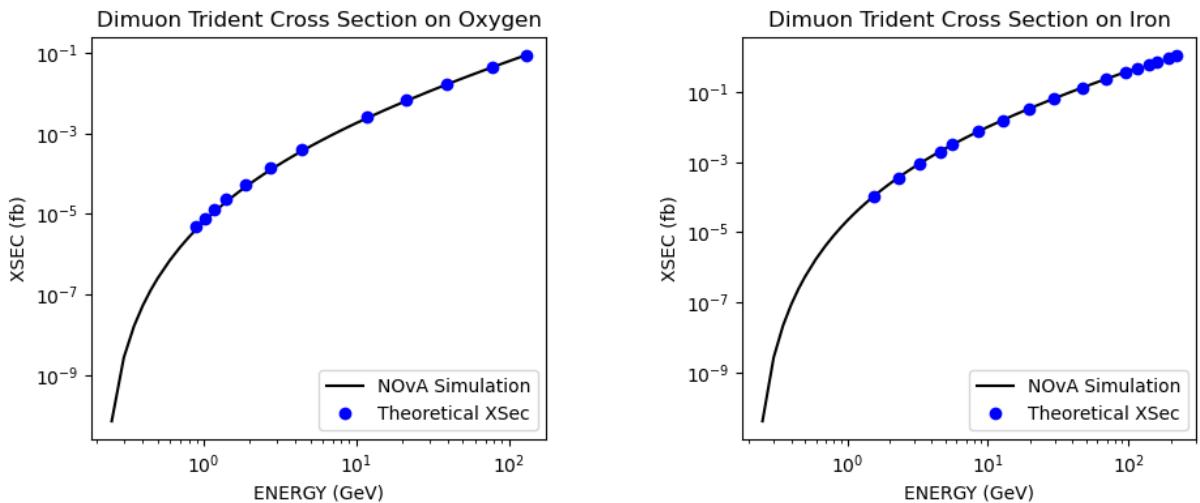


Figure 5.1: Trident cross section vs energy for oxygen and iron generated by the trident simulation tool compared to theoretical calculations [114, 121].

Event Rate Scaling

Once convinced that the cross sections forming the base of the simulation matched theoretical expectations, I moved on to the next step: reweighting the simulated dataset to properly predict the expected event rates. Doing so required accounting for the NuMI flux and its variation across the detector, and keeping track of the trident cross-section on each element of the NOvA soup.

Since trident events have such small cross sections, we did not expect to see many of them in the real detector data. However, to perform the signal selection I needed a large dataset that accurately represented the possible topologies that could manifest as a result of a trident interaction occurring. I generated a signal dataset containing one million events, despite only expecting twenty to forty in the real data. In order to accurately represent the expected number of trident events, I needed to scale the large dataset down via histogram reweighting.

As discussed in section 2.3, calculating the number of expected events N for some particular interaction with incoming neutrino energy E_ν requires the formula

$$N(E_\nu) \sim \phi(x, y, z, E_\nu) \times \epsilon(E_\nu) \times N_T \times \sigma(E_\nu), \quad (5.1)$$

where ϕ is the incoming neutrino flux, ϵ is the detector efficiency, N_T is the number of target nuclei in the detector, and σ is the cross section of the interaction. With accurate cross sections for trident events on all of the NOvA target nuclei in hand, calculating the expected number of events required inclusion of the proper NuMI beam flux across the detector.

To include the correct flux model I utilized the existing NOvA simulation infrastructure to reweight the energy spectra of other processes. Beginning with a GENIE energy spectrum of common quasi-elastic (QE) events, I weighted the spectrum by dividing out the QE cross section in every energy bin. That yielded a spectrum of events simply containing flux and detector information. I generated splines for the cross sections of all possible trident processes on each detector element with respect to incoming neutrino energy. Then I re-weighted the spectrum of QE events by multiplying the trident cross sections at all energies in the spectrum. This resulted in an energy spectrum of tridents in the ND with the number of events properly scaled to account for the trident cross section on each element of the detector and the correct beam flux, shown in Figure 5.2.

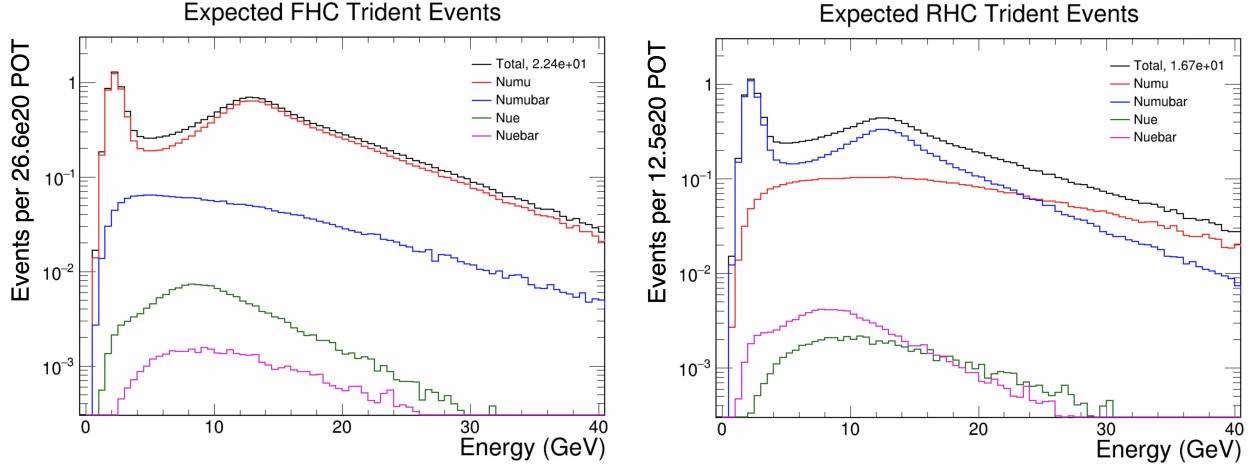


Figure 5.2: Expected trident energy spectra for both FHC and RHC datasets.

The cross section substitution yielded a weighted energy spectrum for trident events. The next step of the analysis required me to propagate that weight so that any spectrum generated for trident events would have it applied. By dividing the reweighted GENIE energy spectrum by the default trident energy spectrum I determined a new weight for each event so that projections of its variables would be properly represented based on incoming neutrino energy. That weight was calculated when I ran the stdhep files containing trident events through NOvA’s particle propagation and detector response tools and stored in the “MCTruth” information associated with an event, ready to be applied to any future histogram of interest during the analysis.

The trident datasets generated for this analysis consisted of one million trident events in the NOvA ND. The ratio of the four neutrino types in each horn current matches the values from chapter 3, with target nuclei specified to follow the ratios in Table 5.1. After simulating the detector response, I ran the events through the reconstruction chain as detailed in section 4.7 and the module that formats them into NOvA’s common analysis format (CAF). The end result was a dataset of simulated tridents with associated ElasticArms vertices, FuzzyK prongs, and Kalman tracks, formatted so that histograms of the physics variables associated with them scale to the expected number of real events.

5.2.2 Simulated Background Dataset

The simulated background datasets used for this analysis were the full NOvA production 5.1 genie FHC and RHC datasets - the default MC datasets described in section 4.5. The FHC MC dataset consisted of 5.549×10^{21} protons on target, corresponding to 2.18 times the real ND FHC dataset. The RHC MC dataset had 5.066×10^{21} protons on target - 4.45 times the real dataset [170]. Notably, the total MC only contained on average about 3.3 times more events than the real data. For an analysis that comes down to isolating singular events, like the search for tridents, the small simulated statistics was a large source of systematic uncertainty in the final measurement.

5.3 Initial Data Reduction

Because trident production is so rare, a background reduction of one million to one events was needed to confidently observe trident events. Two rounds of selection were used to achieve the necessary background reduction - the first round was developed on pre-existing NOvA reconstruction (and thus required no additional processing time), and the second round was developed with additional dimuon-specific reconstruction. This section documents the steps taken in the first round of selection, intended to achieve the largest possible background reduction while maintaining a high signal efficiency. This approach was taken to reduce the size of the dataset requiring special reconstruction tools. The initial round of selection was developed in large part by IU postdoc Erica Smith [171], from whom I took over by finishing the selection and filtering the total datasets based on the final results.

The focus of this search was the observation of “dimuon” tridents – events which result in a final state of two muons in the detector with no additional hadronic energy present. The identifying features of the signal topology are:

- Two long, muon-like tracks contained within the side walls of the near detector.
- A narrow opening angle between the two tracks.
- A very small (ideally zero) amount of energy around the interaction vertex, denoting a lack of energy deposited from hadrons.

An event display featuring an “ideal” dimuon event with these features clearly visible was shown in Figure 4.19.

5.3.1 Preselection

Preselection is a set of selection criteria (frequently called cuts) that are defined before looking at the actual physical characteristics of an event. These cuts remove from the sample events with poor detector running conditions, events with the bulk of the energy deposited outside of the detector, or events that started outside of the detector. Preselection cuts are made solely based on event slicing and a simple hit-based reconstruction of the event vertex location with no advanced reconstruction. The variables that go into the preselection are described in the following sections:

Data Quality

All interaction analyses performed in the NOvA ND have a basic set of data quality cuts applied at the beginning of preselection. These cuts ensure the quality of the data taking conditions for both the detector and the beam meet a certain standard at the time of data-taking.

Beam Quality Cuts: Ensure that the beam meets quality standards during runtime [172].

- Fraction of detector hits occurring outside the beam time peak < 0.45 : Ensures that the majority of energy deposited occurred within the time window of the beam spill. Hits outside of this window are primarily caused by electronic noise in the ND.
- Time Delay between NuMI time stamp and spill < 0.5 ns: Requires the assigned time stamp of a spill properly matches the true time of the beam spill.
- $-202 \text{ kA} < \text{Horn current} < -198 \text{ kA}$: Requires the magnetic horn current to be within the appropriate range to focus pions such that the expected neutrino flux is impinging upon the ND.
- $2.0 \text{ mm} < \text{Beam x and y position on target} < 2.0 \text{ mm}$: Ensures that the beam was positioned correctly on the graphite target to produce the expected neutrino flux.
- $0.57 \text{ mm} < \text{Beam width in x and y} < 1.58 \text{ mm}$: Ensures that the beam was focused correctly to produce the expected neutrino flux.

Detector Condition Quality Cuts: Ensure that the detectors meet the condition for a good run during each spill [173] (see Sections 4.2 and 4.3 for terminology details).

- Number of Active Diblocks = 4: Require all 4 of the ND diblocks to be active during data-taking.
- Fraction of Good FEBs per DCM > 0.80: Require at least 52 good FEBs per DCM to minimize gaps in the data.
- $12 \text{ Hz} < \text{MIP Hit Rate} < 20 \text{ Hz}$: Require the hit frequency to be in a window that ensures that components are working properly but not overly noisy.
- $1 < \text{Number of Reconstructed Slices} < 6$: Require that a reasonable number of slices is created from a single beam spill.
- Fraction of 3D Tracks > 0.95: Ensures that any tracks generated can be formed into 3D tracks (i.e. energy deposited in the two detector views align properly in space and time).

Event Reconstruction Quality

These cuts ensure that events in the detector are properly reconstructable:

- Number of Hits in Slice > 20: Fewer than 20 hits often results in bad reconstruction.
- Number of Planes in Slice > 10: Ensures that hits deposited energy over a range of the detector rather than all in a single location.
- Number of Reconstructed Vertices > 1: Require at least a single reconstructed vertex per event.

Fiducial Volume

Fiducial volume cuts ensure that the reconstructed vertex of the interaction is within the boundary of the detector. This removes a large amount of energy entering the detector from events that take place in the surrounding rock.

- $-170 \text{ cm} \geq x\text{-location} \leq 170 \text{ cm}$
- $-170 \text{ cm} \geq y\text{-location} \leq 170 \text{ cm}$
- $50 \text{ cm} \geq z\text{-location} \leq 1200 \text{ cm}$

Plots showing reconstructed vertex locations and fiducial cuts in x and z are shown in Figure 5.3. Three interesting features of these graphs are the peak at the start of the detector, the asymmetry in the x direction, and the drop off of vertices after $z = 1250$ cm. These features are due to muons generated in the rocks before the detector, the beam entering the detector at an angle rather than straight on, and the muon catcher stopping most particles before they enter into the back region of the detector, respectively.

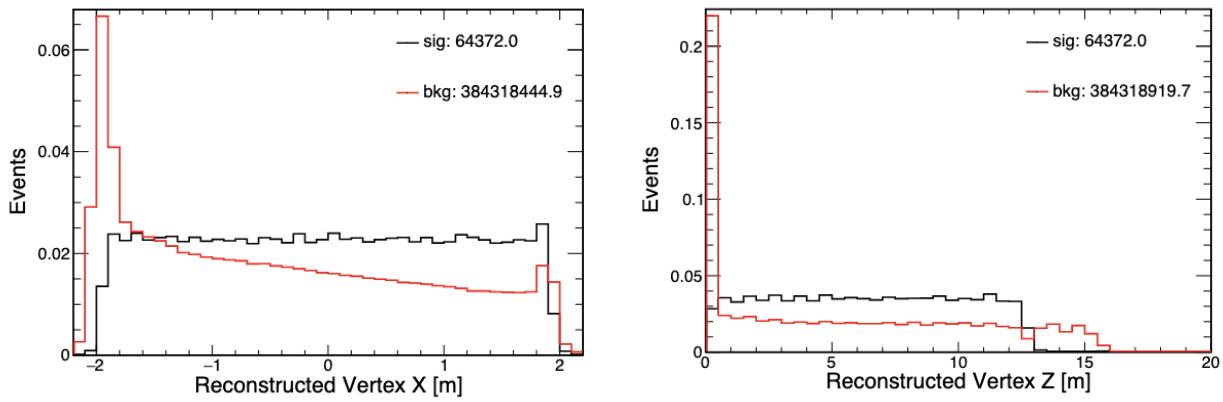


Figure 5.3: Vertex locations in x and z [171].

5.3.2 First Round Selection

We investigated the number of Kalman tracks created by the standard NOvA reconstruction after applying the preselection cuts. Figure 5.4 shows that over 95% of the simulated trident events had either one or two reconstructed three-dimensional tracks. Generally, when two tracks are formed it means the Kalman algorithm reconstructed one track for each muon in the detector. When only one 3D Kalman track is reconstructed, it is either because the opening angle between the muons was too small or because hits from one muon overlap with hits from the other muon in at least one view of the detector.

We cannot expect similar distributions of track-specific quantities between the two distinct track-level samples. However, slice-level quantities that reflect the overall topology of the entire event should be similar for both populations. Therefore, the approach in this first round of selection was to first perform a slice-level selection that takes topological variables as inputs, then split the events that remain into separate one-track and two-track populations to select on further.

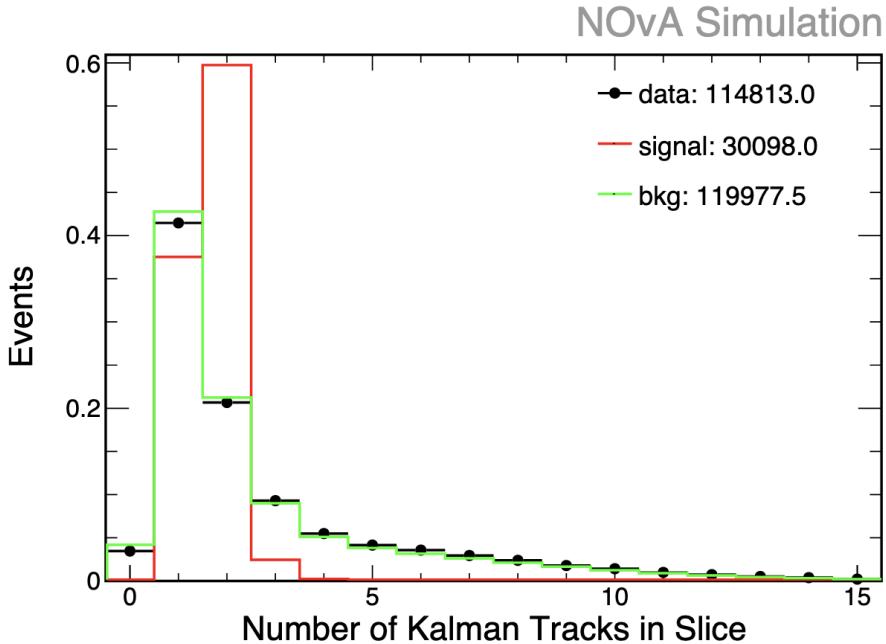


Figure 5.4: Number of Kalman tracks for signal (red line), background (green line), and data (black dots). Each distribution is normalized to 1.

Slice-Level Selection

Slice-level cuts were utilized to remove many of the neutrino events with topologies consistent with backgrounds to the dimuon signal. The variables for this level of the selection are listed in Table 5.2, and area-normalized plots showing the signal and background distributions are shown in Figure 5.5. We fed these variables into ROOT’s Tools for Multi-Variate Analysis (TMVA) framework [174]. TMVA comes equipped with a boosted decision tree (BDT) algorithm which takes input variables and determines their effectiveness at separating signal and background. The algorithm then ranks the variables based on their separation power. Correlation matrices are generated to compare variables with one another - uncorrelated variables are desirable because making cuts on them does not change the effectiveness of the other variables.

After running the BDT on the chosen variables, a new variable called the “gradient boosted decision tree output” (BDTG output) was created and cut upon to maximize the amount of signal retained while removing background. Figure 5.6 shows the BDTG output of the signal and background samples. We decided on a BDTG cut value of -0.258, which removes 81.6% of background while retaining 99.8% of dimuon signal events.

Slice-Level BDT Variables	
Variable	Description
kXWidthToLength	Ratio of x-width to length of hit distributions in the slice
kYWidthToLength	Ratio of y-width to length of hit distributions in the slice
kHPP	Number of hits in a single plane
kNContPlanes	Number of planes with continuous hits (i.e. no large gaps)
kVtxE40	Calorimetric energy within 40 cm of the interaction vertex
kNKal2D	Number of reconstructed 2D Kalman tracks
kNKal	Number of reconstructed 3D Kalman tracks

Table 5.2: A table of the slice-level variables fed into a BDT.

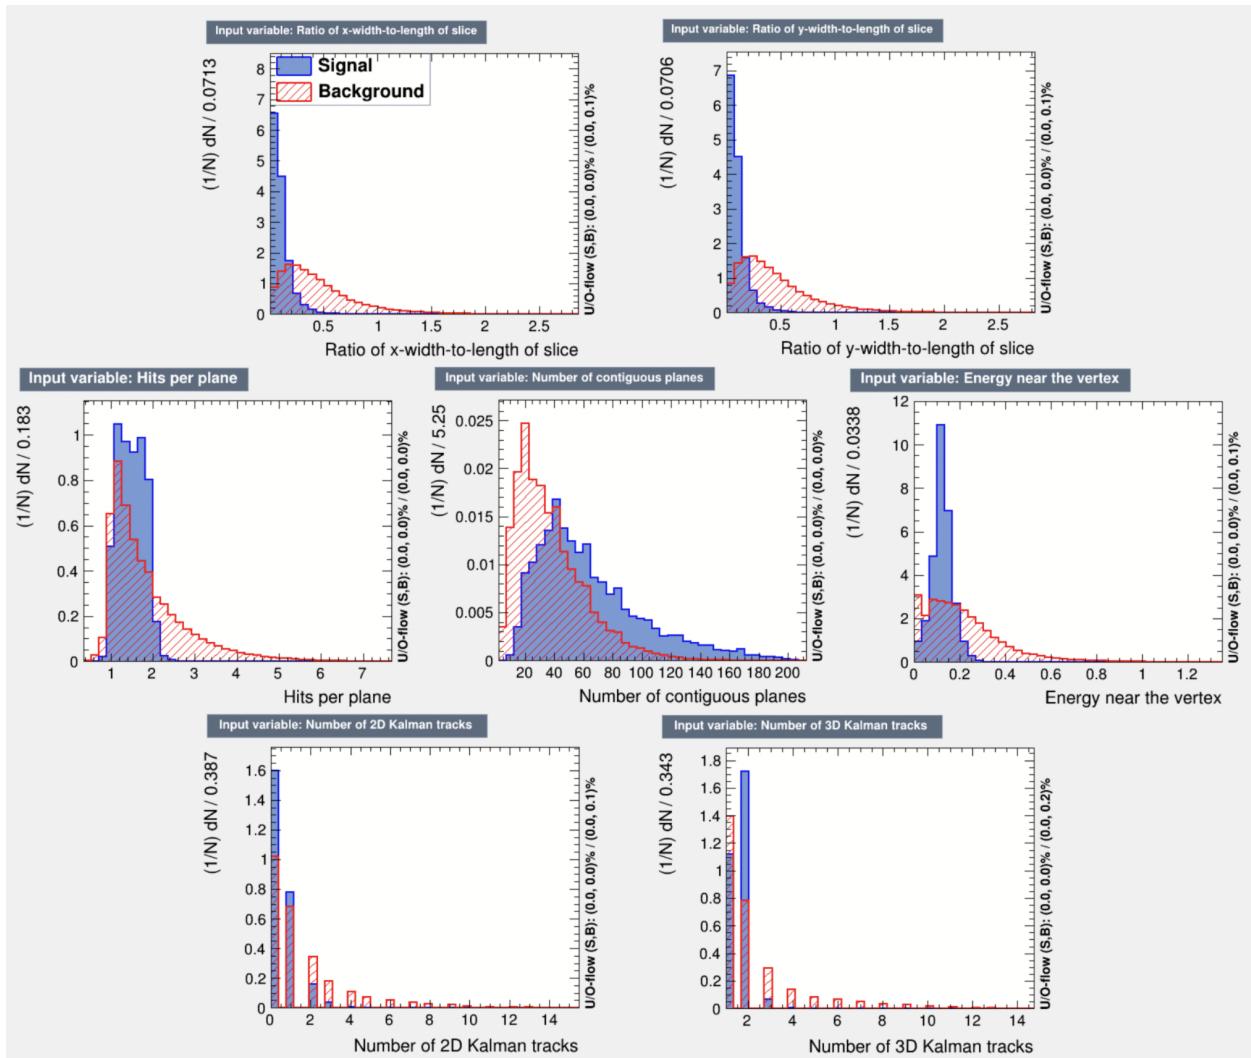


Figure 5.5: Area normalized distributions of BDT variables for signal (blue) and background (red).

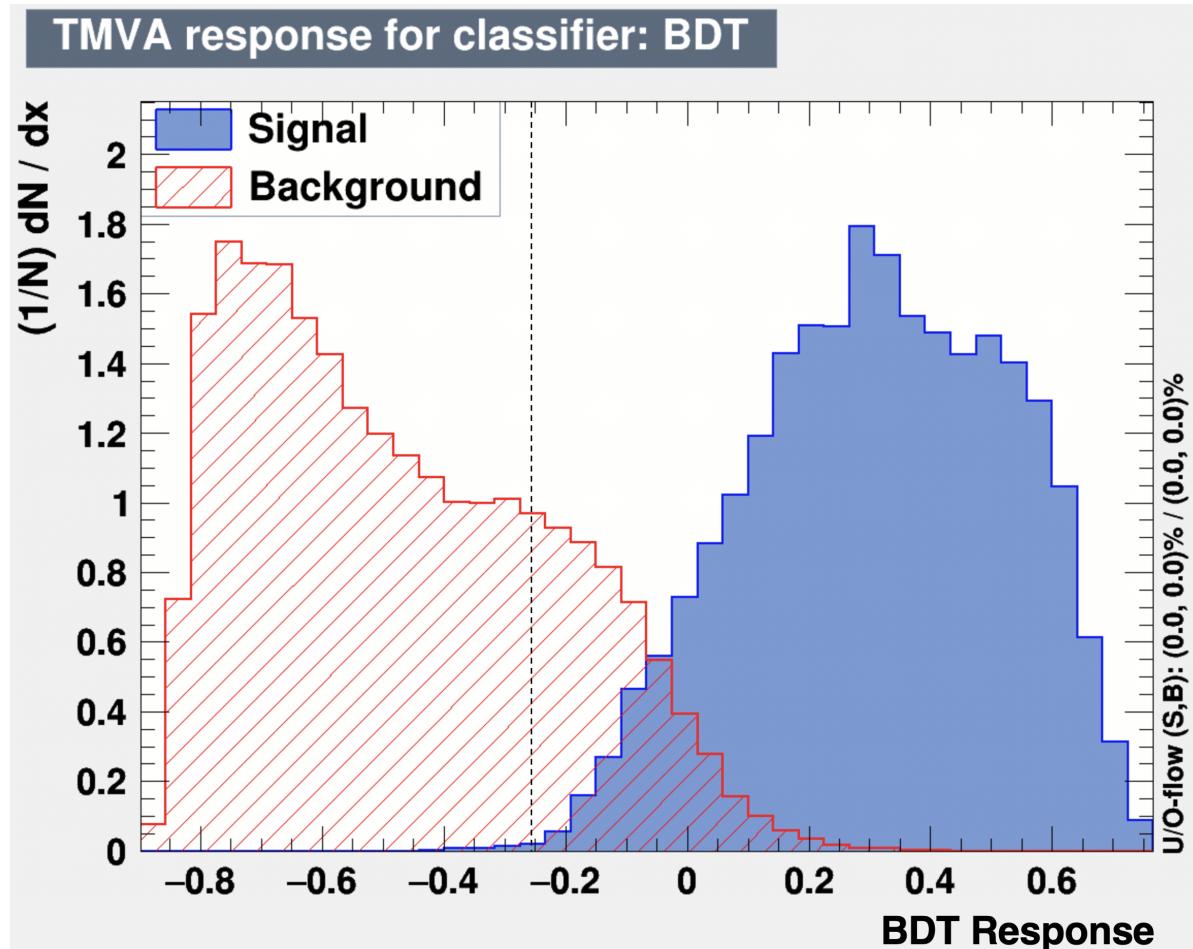


Figure 5.6: Area normalized distribution of slice-level BDT response, cut location included.

Track-Level Selection

After applying the cut on the slice-level BDT, we analyzed more specific track-level variables on the remaining data. These consist of precise information calculable when information such as the starting point, stopping point, and trajectory of the energy deposition is well known. We developed two selections in tandem to handle events with single and double kalman track reconstructions. Both selections followed the same process: identify track-level variables that provide good separation between signal and background, and feed them into a BDT similar to that of the slice-level selection.

One 3D Kalman Track Selection

As shown in Figure 5.4, only a single 3D kalman track was reconstructed on trident events about 40% of the time. An event display featuring an example of a trident event with a single 3D reconstructed kalman track is shown in Figure 5.7. The single track selection relies on 24 input variables that identify physical properties of the event, described in Table 5.3 and shown in Figures 5.8 and 5.9.

Unfortunately, when only one 3D track is formed it is very difficult to identify features that indicate an event has two separate muons. That issue makes the single-track selection the least efficient of the three first-round BDTs. Here we decided on a BDT cut value of -0.30, which removes 62% of background while retaining 98% of dimuon signal events. The BDT output is shown in Figure 5.10.

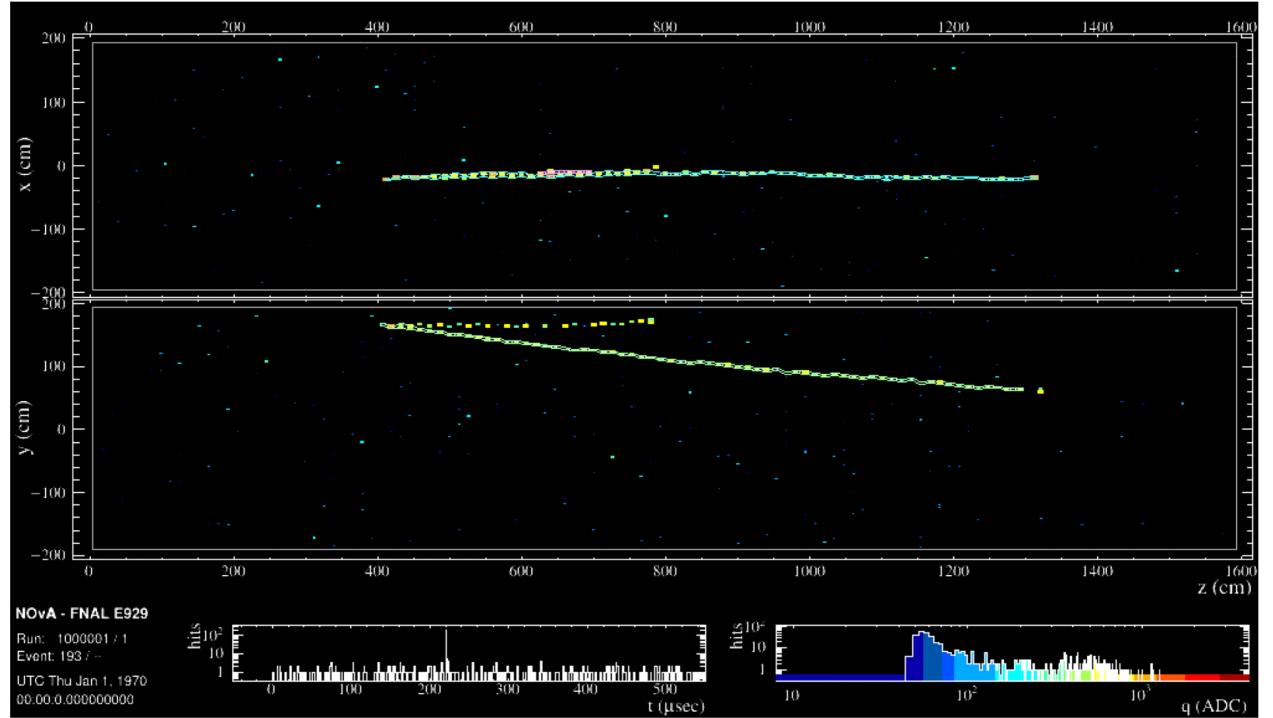


Figure 5.7: Event display of a dimuon event with a single reconstructed 3D kalman track. Leftover hit information for this event was distributed across additional 2D tracks.

One Kalman Track BDT Variables	
Variable	Description
kXWidthToLength	Ratio of x-width to length of hit distributions in the slice
kYWidthToLength	Ratio of y-width to length of hit distributions in the slice
kHPP	Number of hits in a single plane
kNKal2D	Number of reconstructed 2D Kalman tracks
kVtxE40	Calorimetric energy within 40 cm of the elastic arms vertex
kKalAvedEdx0	Average deposited energy/length ratio for the longest 2D track
kKalScatLL0	Log-Likelihood that the particle scattered somewhere along its trajectory
kKaldEdxLL0	Log-Likelihood of the dE/dx , useful in identifying muons
kKalDirX0	Initial x-direction of the track
kKalDirY0	Initial y-direction of the track
kKalDirZ0	Initial z-direction of the track
kKalTrkCalE0	Calorimetric energy stored in the track
kKalTrkCalEPerNHit0	Overall energy per hit in the track
kKalOverlapE0	Overlapping energy of any 2D tracks created
kKalTrkNhits0	Number of hits in the track
kKalTrkLength0	Length of the track
kKalNhitLength0	Overall number of hits per unit length
kKalReMId0	Reconstructed Muon ID score of the track
kKalLast10cm0	Average dE/dX deposited over the last 10cm of the track
kKalLast20cm0	Average dE/dX deposited over the last 20cm of the track
kKalLast30cm0	Average dE/dX deposited over the last 30cm of the track
kKalLast40cm0	Average dE/dX deposited over the last 40cm of the track
kCVNmlooseptp	CVN score identifying transverse momentum to be muon-like
kNTotalPng2D	Total number of 2D fuzzyk prongs generated

Table 5.3: A table of the one-kalman-track variables fed into a BDT.

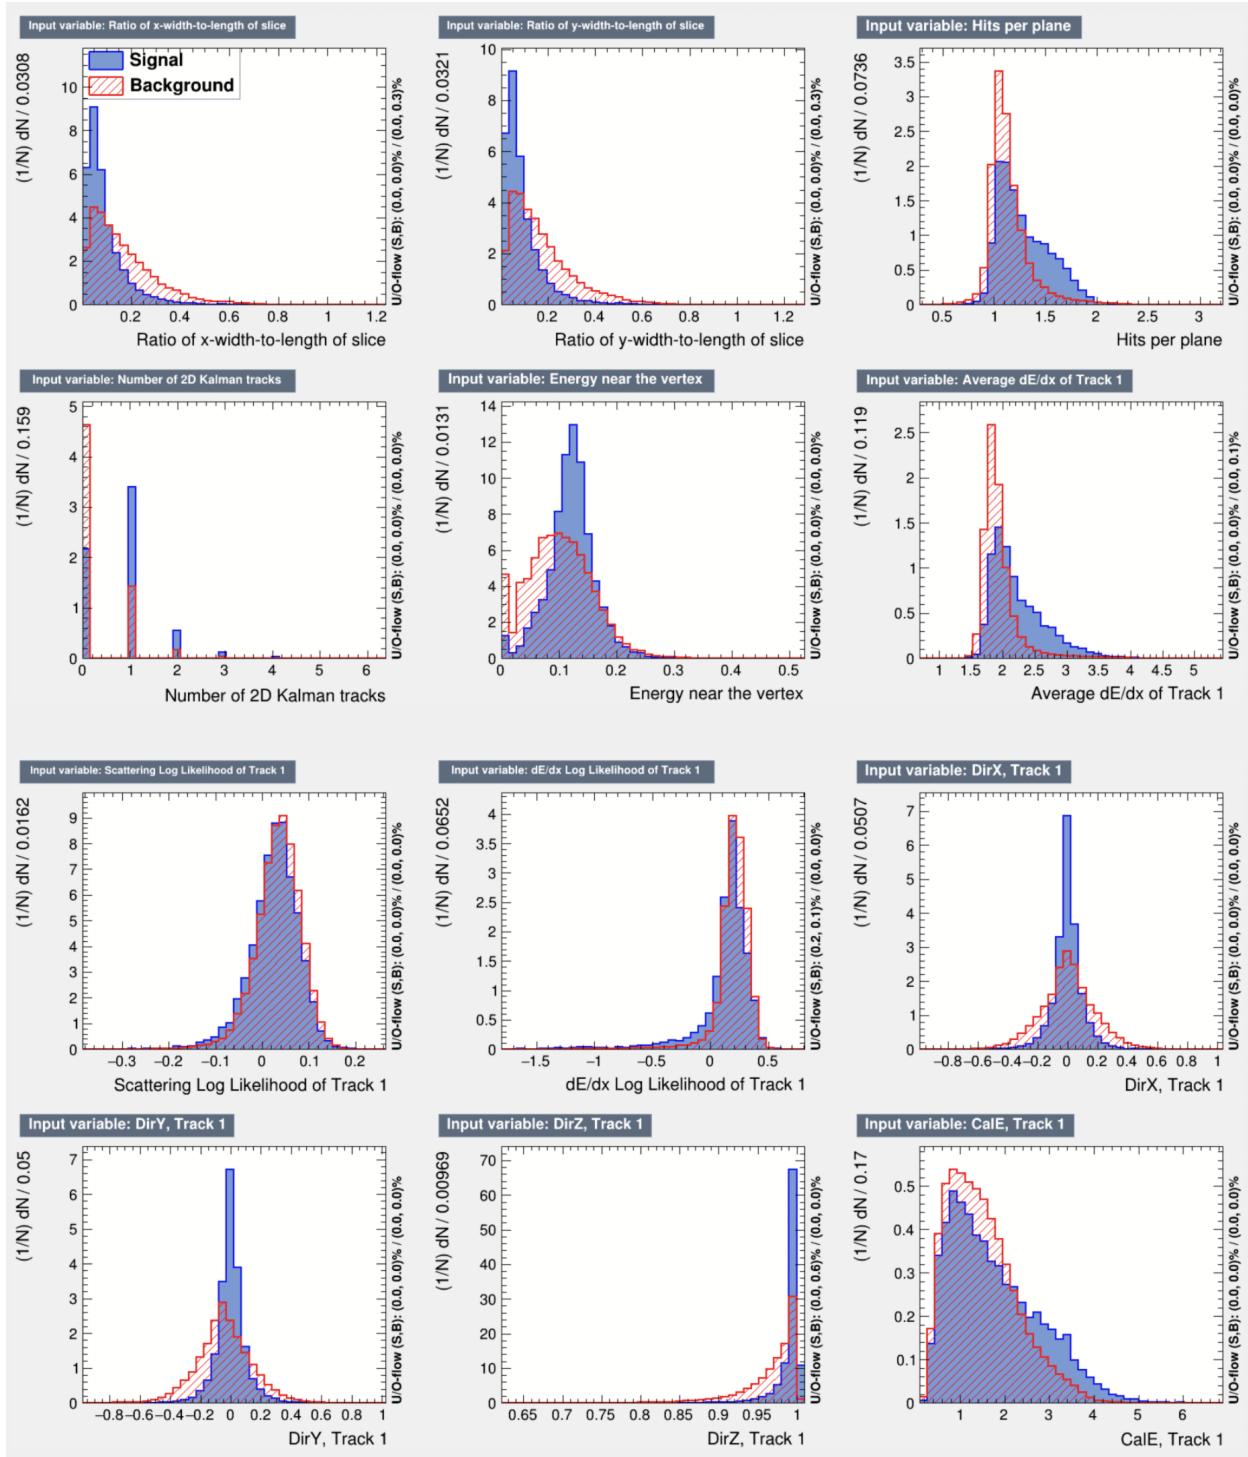


Figure 5.8: Distribution of single-kalman-track BDT scores for signal (blue) and background (red).

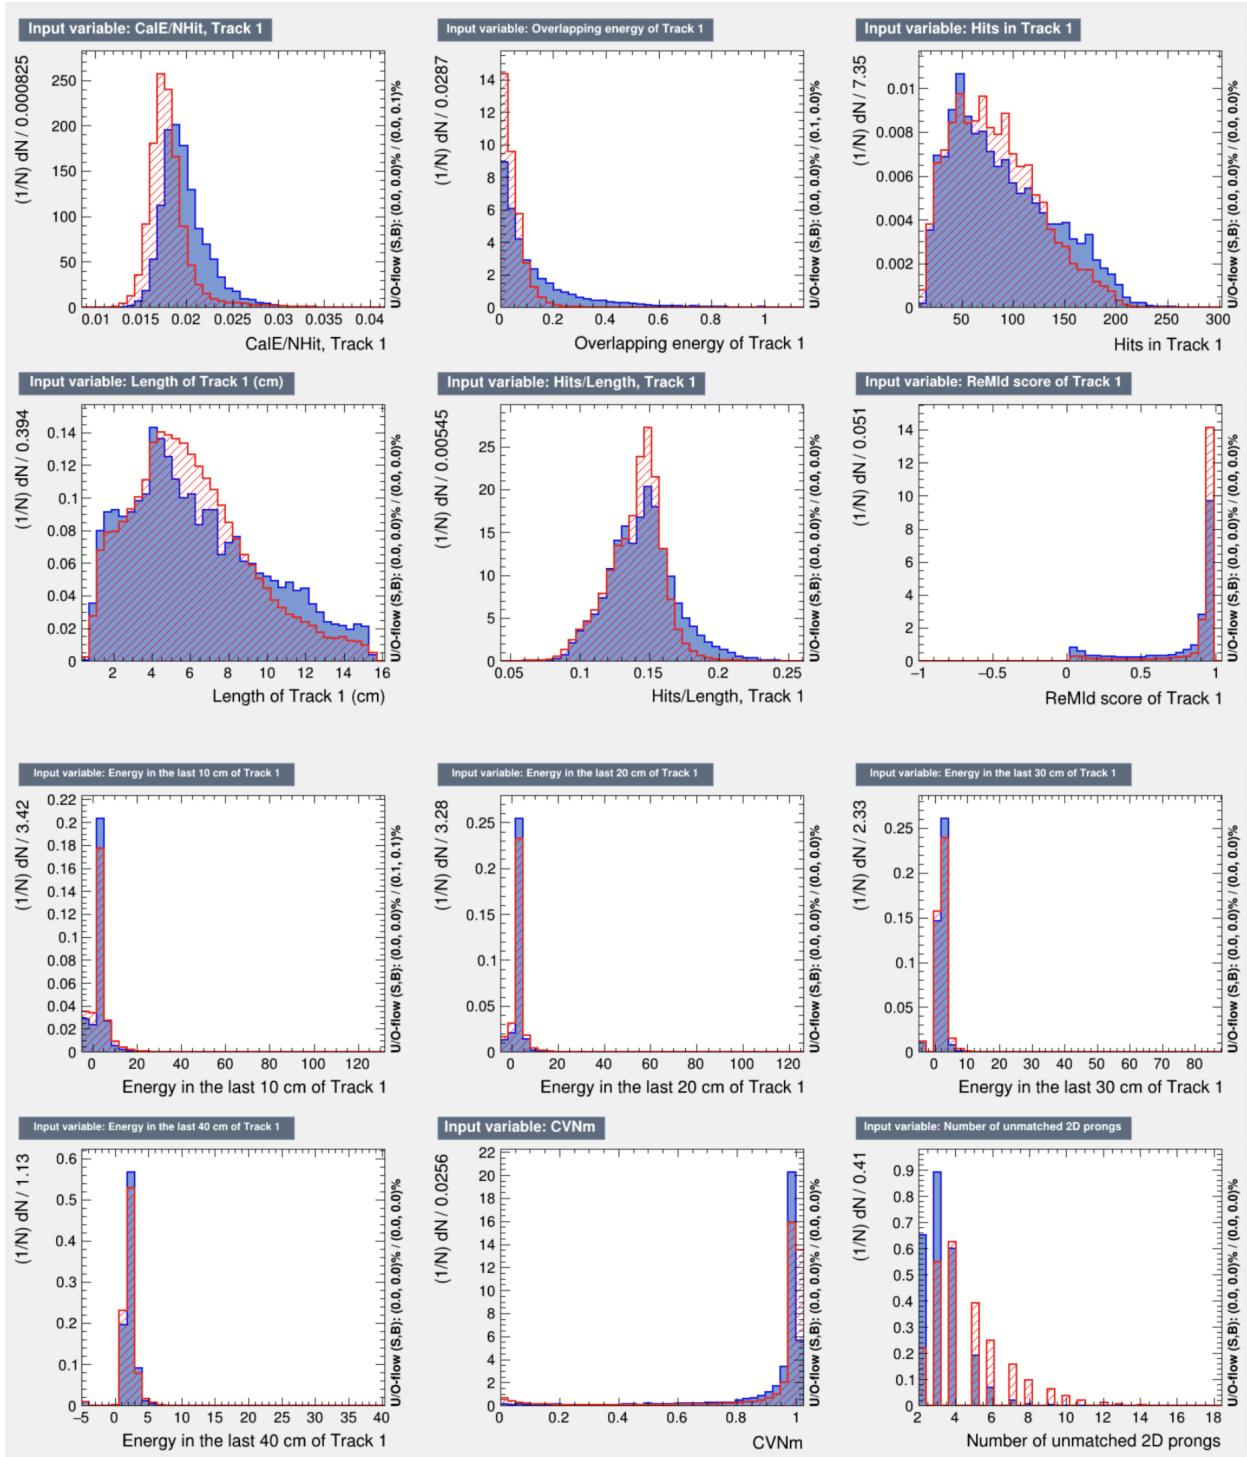


Figure 5.9: Distribution of single-kalman-track BDT scores for signal (blue) and background (red).

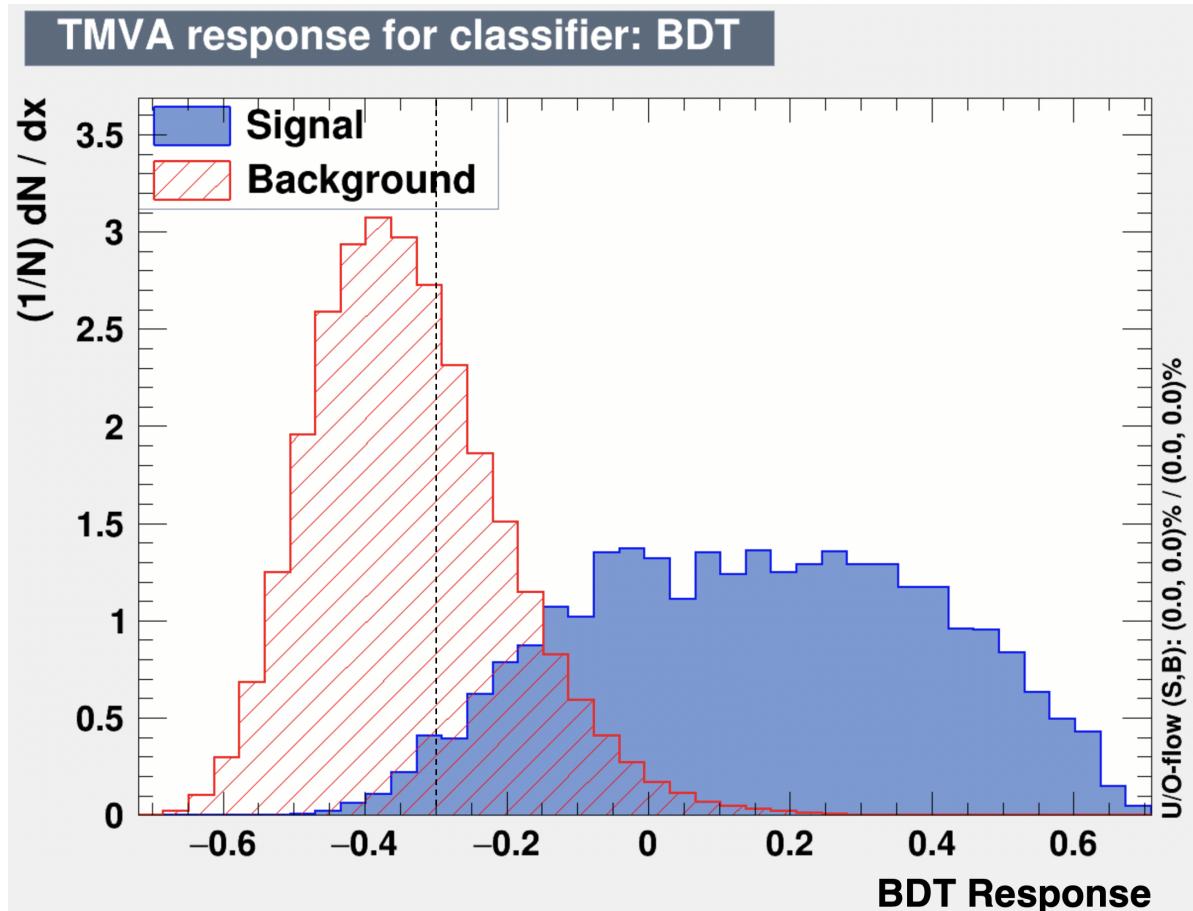


Figure 5.10: Distribution of single-kalman-track BDT scores for signal (blue) and background (red).

Two 3D Kalman Track Selection

The Kalman method reconstructed two 3D tracks on simulated trident events about 60% of the time, which allows for more powerful signal identification. For example the ReMId score for the shorter track shown in Figure 5.11 shows a clear distinction between signal and background.

We identified 38 variables with strong separation ability to feed into TMVA for the two-kalman track dataset. They are the same variables listed in Table 5.3 for the one-track BDT, but any track-specific variables are included for both 3D tracks. Figure 5.11 only shows distributions for secondary (shorter) track variables, since many of the primary (longer) track distributions are similar to their one-track counterparts. Here, the chosen BDTG cut value is -0.175, which removes 98% of background while retaining 95% of signal. The output of the two kalman track BDT is shown in Figure 5.12.

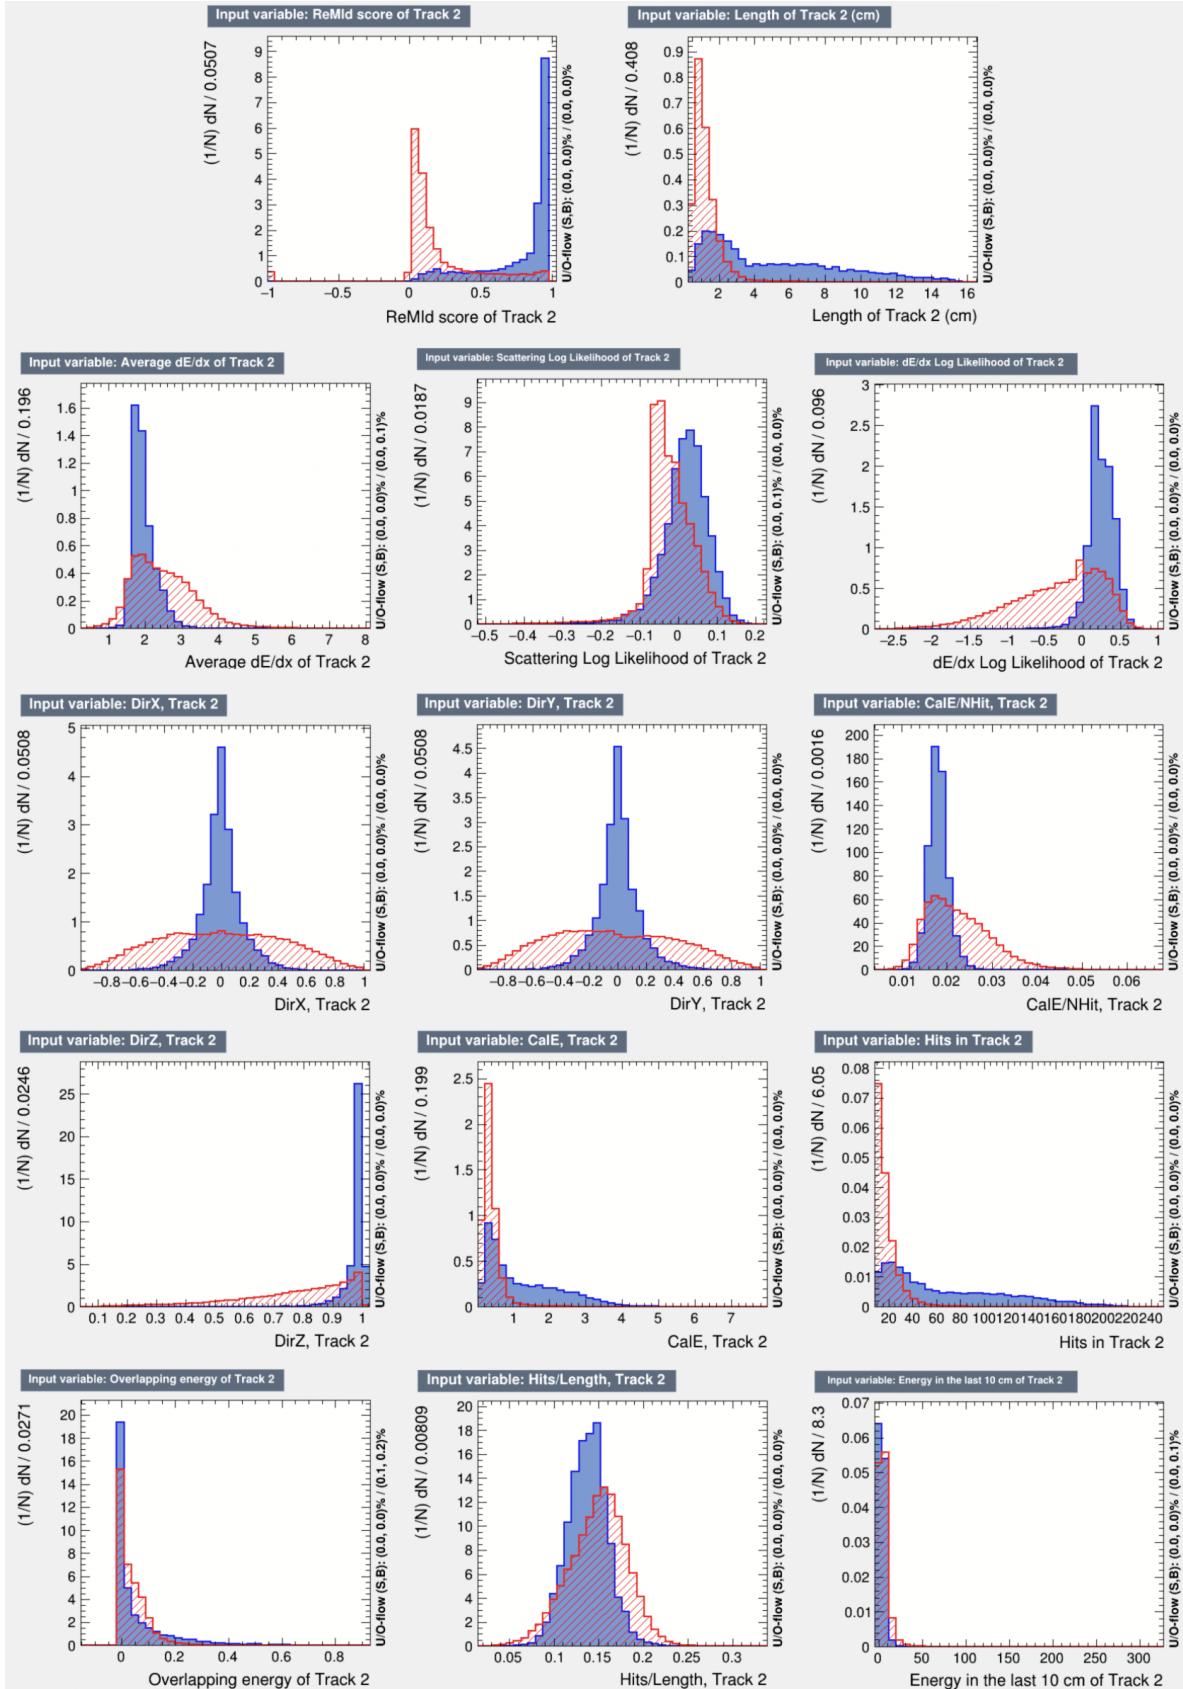


Figure 5.11: Distributions of secondary track variables for signal (blue) and background (red).

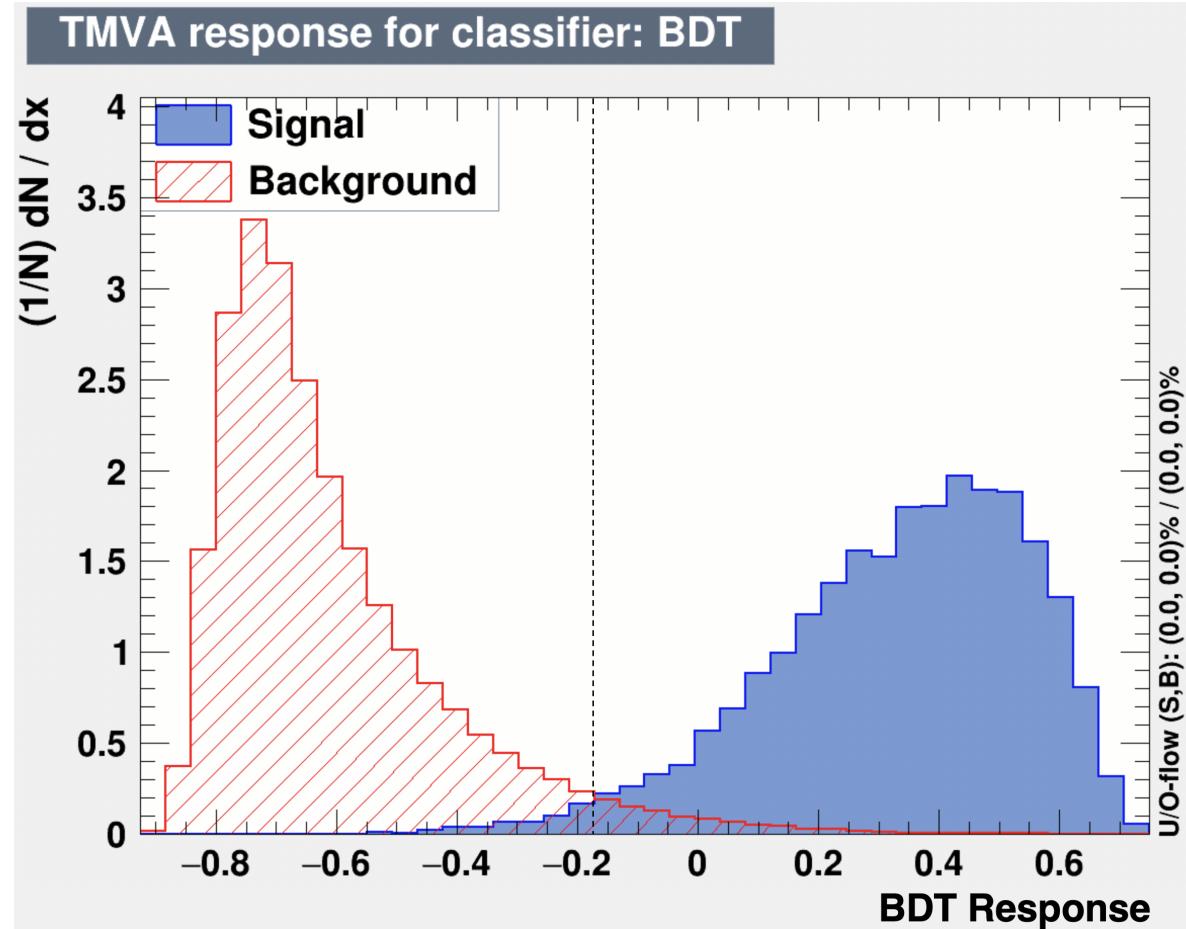


Figure 5.12: Distribution of two kalman track BDT scores for signal (blue) and background (red).

Overall Dataset Reduction

Applying the cuts from these BDTs to the FHC and RHC datasets brought the number of events predicted by the simulation from about 39.1 tridents on 10 million background down to about 38.2 tridents on 550k background. Overall signal events were preserved with 97.7% efficiency, while background was reduced to just 5.5% of the original size. However, the remaining background dataset was still about four orders of magnitude too large to reach the signal to background ratio needed for a trident observation. We developed a second round of selection to reduce the background more - a process requiring a new tool. Fortunately, after applying these selections we reduced the full 300TB dataset down to 16TB – an amount that would make running custom reconstruction feasible. I developed a new reconstruction tool that improves our ability to distinguish the two muons and reduce the background further: the “KLutz” dimuon tracker.

5.4 KLutz Custom Trident Tracker

The Kalman tracking algorithm worked fairly well, but it was not developed with the identification of two long tracks in mind, leading to a handful of issues. The very small opening angles between muons often caused Kalman to produce just a single track, as shown in Figure 5.7. Events that *did* get reconstructed with two tracks commonly had issues with “jumping the tracks”, where the track hops back and forth between energy depositions from the two muons. There were also cases of two 3D Kalman tracks being formed without capturing all of the muon energy and additional 2D tracks being formed to pick up the remaining energy. An example featuring both of these issues is visible in the event display of Figure 5.13. The trident event presented looks ideal, but faulty reconstruction would cause this event to be thrown out of the analysis.

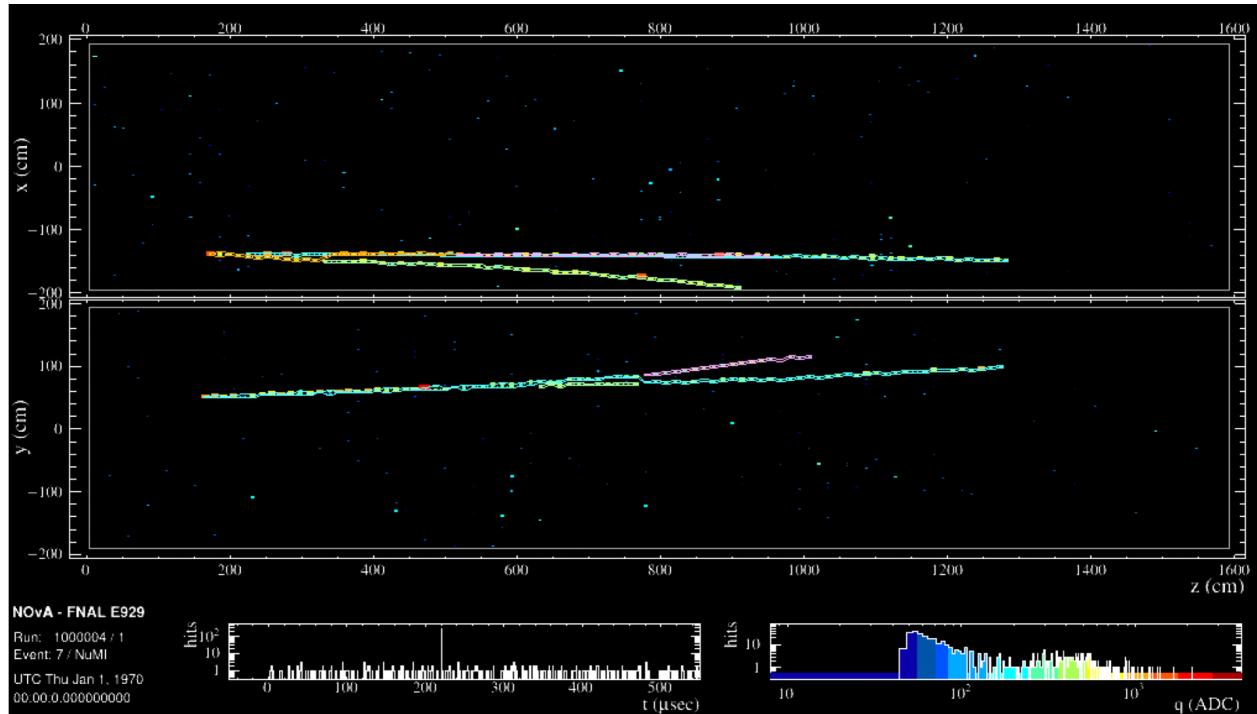


Figure 5.13: A dimuon event showing a failure mode of Kalman tracks - there are two 3D tracks and two 2D tracks all capturing *pieces* of the muons, but not their entire trajectories.

These issues occurred frequently enough that continuing the analysis required the development of a new dimuon-specific tracking tool that could capture both muon tracks cleanly. I developed the algorithm to do this with inspiration from both the Kalman and Lutz methods described in

subsection 4.7.2, resulting in the name “KLutz”. The method uses the endpoint and track seed generation of Kalman to identify tracks, and the Lutz optimum fitting method to form the tracks. Combining these two methods allows KLutz to be very sensitive to small deviations in tracks, which was necessary when the tracks were often so close together.

5.4.1 The KLutz Tracking Algorithm

This section details the KLutz track reconstruction algorithm. The starting point is an event composed of raw energy deposited in specific planes and cells of the near detector, called a hit (see section 4.2 for details). Each hit occupies one plane and one cell in the grid that makes up a single view of the detector. Figure 5.14 is a cartoon version of the ND as a grid, with hits in a given cell identified by a red box.

Initial Track Formation

The first step is to find the end of the muon’s path, where it either stopped or exited out the back of the detector. We assign this as the point where the muon has zero energy, which acts as the starting point to generate a track from. To find endpoint candidates, we generate windows in the direction the beam approaches from (“upstream” - left in these diagrams) and in the direction the resulting particles move in (“downstream” - right in these diagrams) for every hit in a slice. The downstream window around each hit is 3 planes long and 2 cells wide, and the upstream window around each hit is 5 planes long and 5 cells wide, as seen in Figure 5.15.

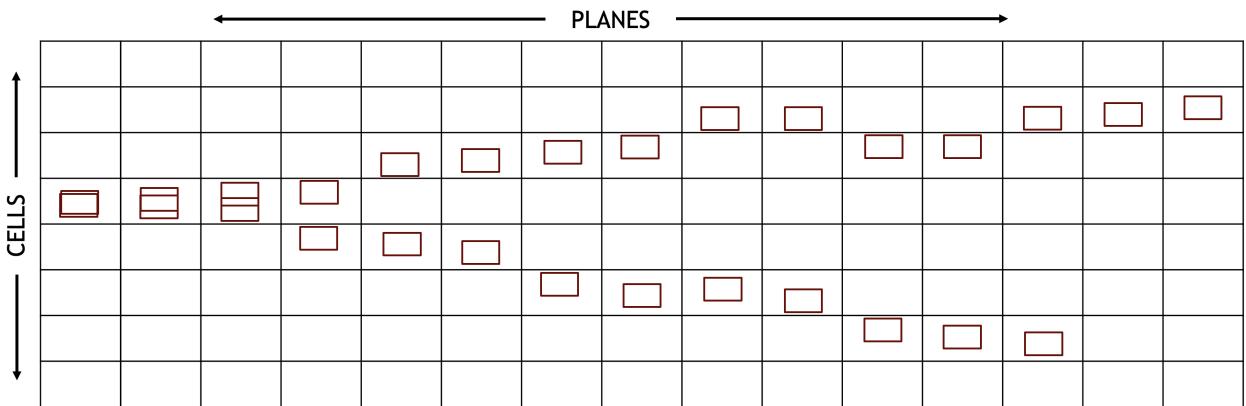


Figure 5.14: A cartoon representation of the NOvA near detector composed of vertical cells and horizontal planes with hits representing a dimuon event.

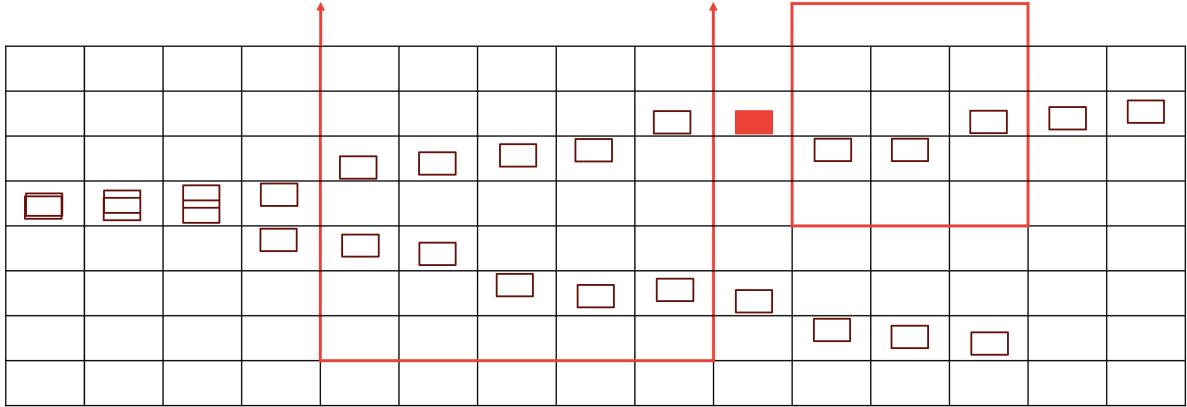


Figure 5.15: Downstream and upstream windows represented by red boxes drawn around a possible endpoint candidate, highlighted red.

We sum the number of hits in each of these windows and set criteria that identify hits as possible endpoint candidates. After meticulously testing a variety of window sizes and endpoint criteria to account for possible gaps at the end of the track, the final criteria for an endpoint are:

- Only a single hit is *allowed* in the downstream window of the candidate endpoint.
- At least three hits are *required* in the upstream window of the candidate endpoint.

These criteria require that a hit is at the farthest downstream end of a track, with some leeway for deviations. If a hit passes those thresholds, it is tagged as an endpoint candidate. For example, in the cartoon shown in Figure 5.16, the hit in red has a sufficient number of hits upstream of it, but too many hits downstream, since only one is allowed in the downstream window. Therefore, the hit in red is not tagged as an endpoint. On the other hand, the hits circled in blue do meet the criteria to be tagged as endpoints, because they have three hits upstream and either one or zero hits downstream of them.

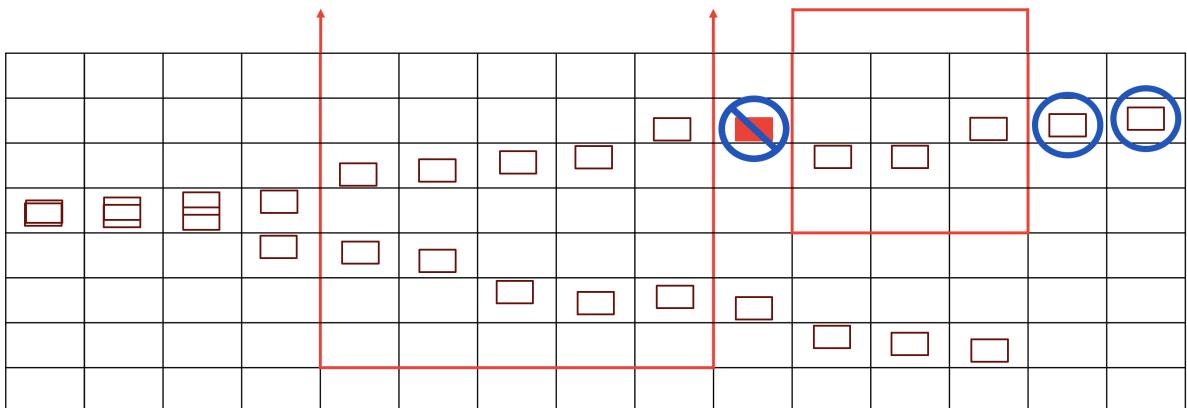


Figure 5.16: The crossed out hit does not meet the endpoint criteria, whereas the circled hit does.

Once all endpoint candidates have been identified, we “grow” tracks from them by creating an object called a “seed” track. Starting at the furthest downstream candidate endpoint in the detector, we step upstream and grab any hits within 1 plane and 2 cells of the starting point, combining them into a single cluster of hits. That process is repeated for every hit in the slice until there are 5 hits in the seed track. The cartoon in Figure 5.17 shows hits in one seed track formed from the furthest downstream hit in pink, and the two remaining endpoints in blue.

The resulting track seed is the initial parameter fed into the Lutz optimum fitting algorithm described in subsection 4.7.2. Using the seed as a starting point provides a very strong foundation to form full tracks since the position and direction of the seed tracks are known, making KLutz tracks less prone to errors due to small deviations than tracks formed entirely by the Kalman or Lutz methods alone.

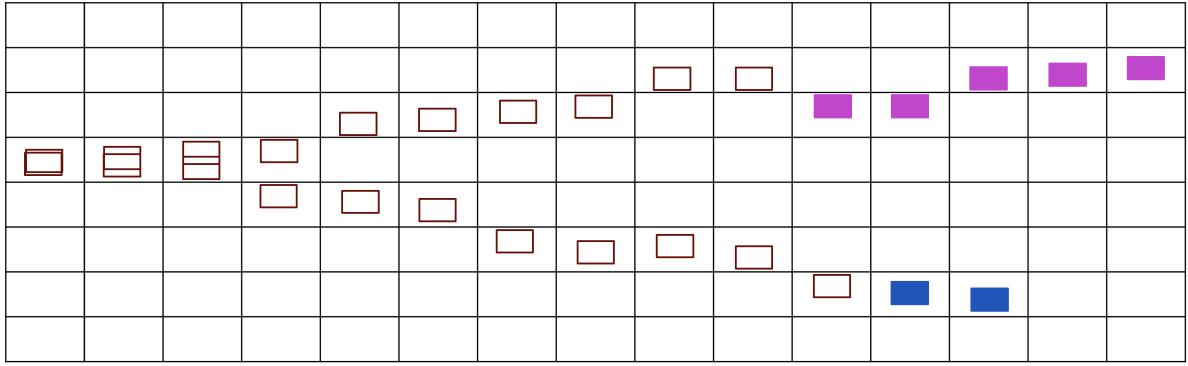


Figure 5.17: A track seed formed from the furthest downstream endpoint candidate, highlighted in pink. Endpoint candidates not included in the seed are highlighted blue.

At this point I chose to place scattering planes every 30 cm from the downstream end of the event - a value chosen based on the radiation length and average distance traveled by muons before scattering in the detector. The Lutz fitting process is applied to hits considering the scattering planes, and a set of linear equations is solved to minimize the χ^2 and yield parameters for a fully formed track. Any hits with a chi squared value above 43 with respect to this initial track are removed from contention in the fit, and the whole process is repeated again starting from the next furthest downstream endpoint. Figure 5.18 shows the state of hits in the slice at this point in the process - one set of hits is clustered into a fully formed track, and the other set of hits still only has an endpoint associated.

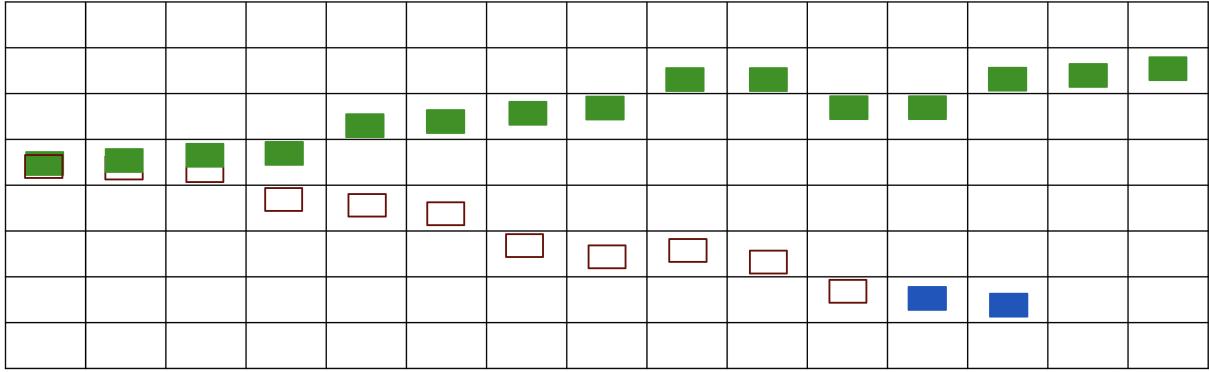


Figure 5.18: The primary track formed out of the most downstream endpoint candidate, highlighted in green. At this point the sequence starts again from the next furthest downstream endpoint.

When the sequence runs again each view of the detector has two independent tracks associated with it, so there are four total tracks. Because the opening angles between the tracks are so small, the two tracks typically have some upstream hits deposited in the same cell, causing the tracks to overlap at the start. KLutz handles these events by searching for any shared cells between the two tracks and splitting the energy deposited in those cells evenly between the two tracks. Figure 5.19 shows the cartoon representation of two tracks with hits associated with a single track highlighted in one color, and shared hits highlighted in both colors.

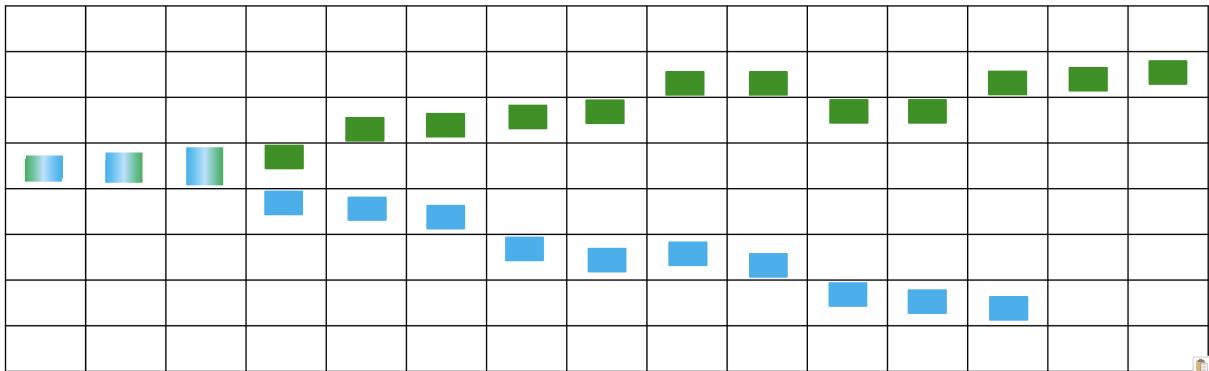


Figure 5.19: Two tracks formed, shown in green and blue. Hits shared between the tracks get energy split evenly amongst the two.

The cartoons in Figures 5.14 - 5.19 are accurate representations of an ideal scenario, and a set of real event displays featuring these steps are shown in Figures 5.20 - 5.23.

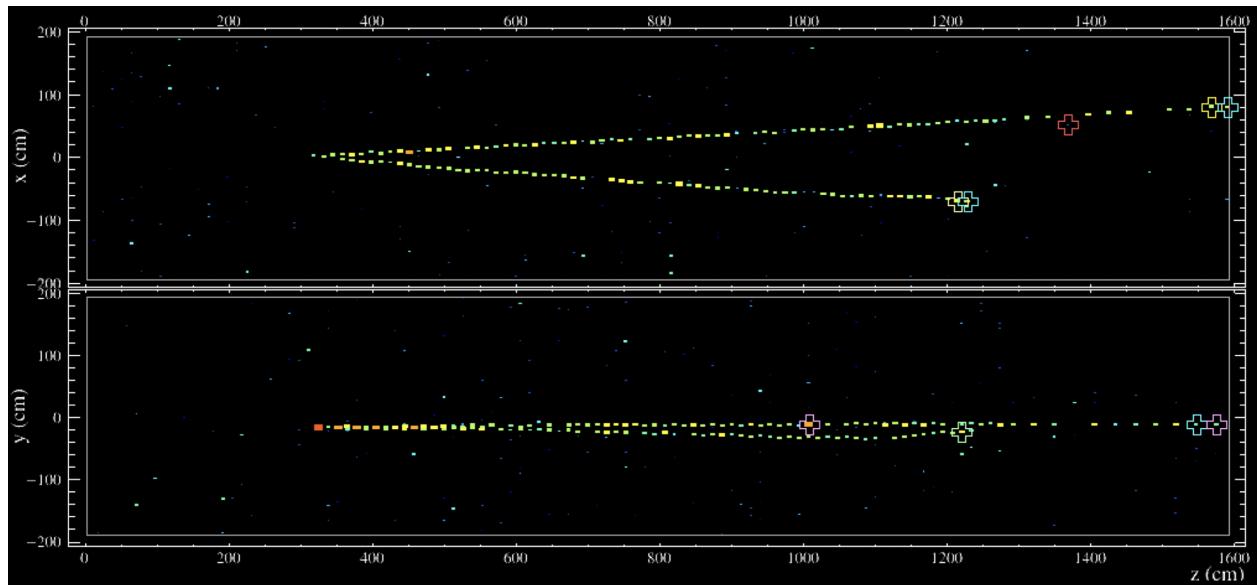


Figure 5.20: An event display featuring the all endpoint candidates for a dimuon event. Only the furthest downstream will actually go on to form track seeds.

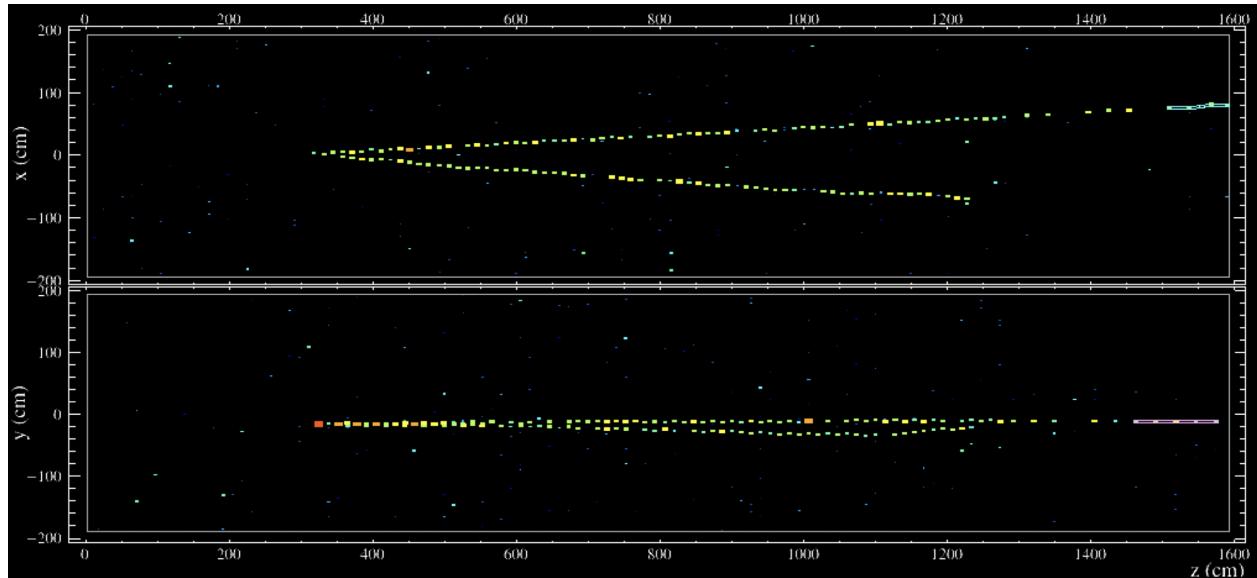


Figure 5.21: An event display showing the two track seeds grown from the endpoint candidates, highlighted in blue and pink.

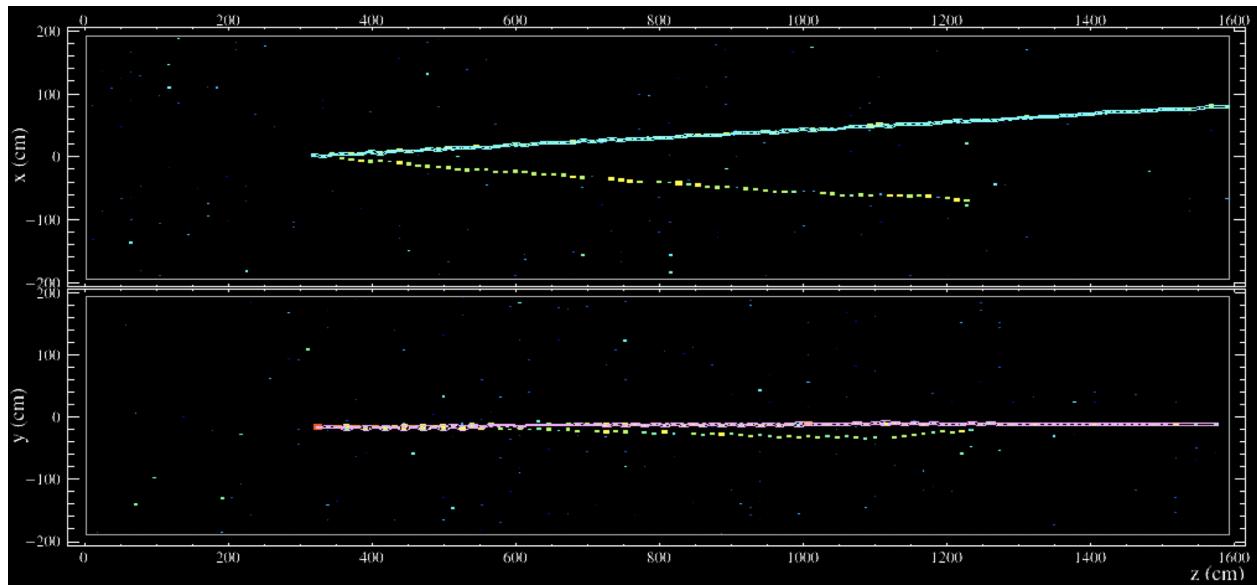


Figure 5.22: An event display featuring the track-creation process - each hit forms one track at a time out of the furthest downstream track seed.

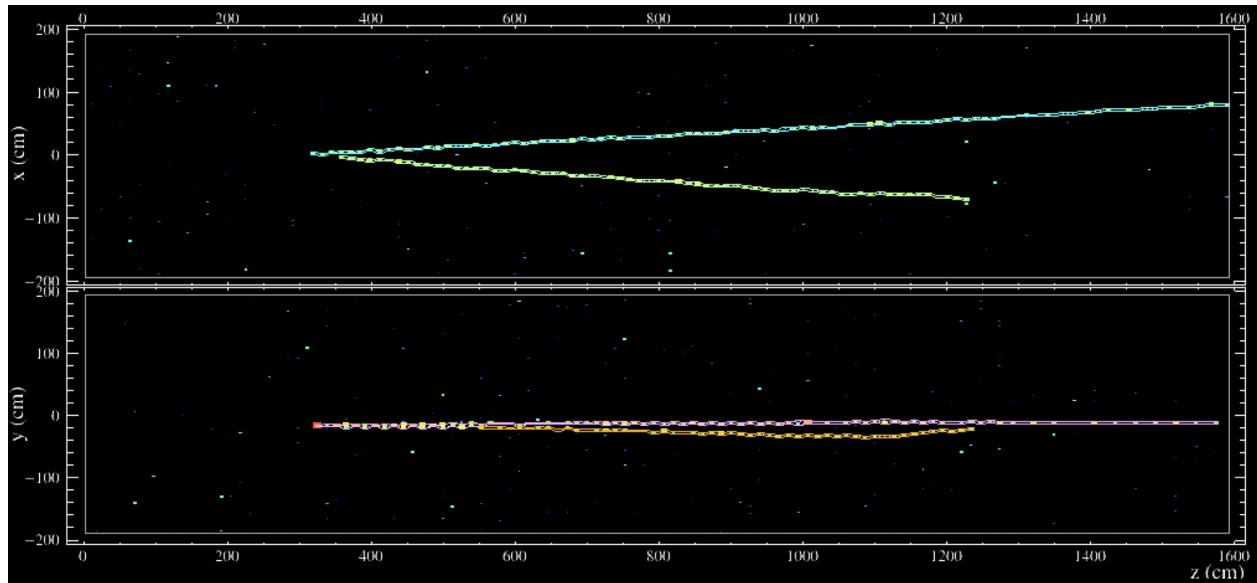


Figure 5.23: An event display showing four fully-formed 2-dimensional tracks.

Overlapping Tracks

Occasionally the energy deposited by the two muons overlap one another entirely. If they overlap completely in *both* views, nothing can be done to distinguish the two muons and their reconstruction must be skipped. However if the two tracks only overlap in a single view, I utilize the two tracks in the other view to split any hits in the overlapping region into two tracks. The energy is split evenly between the two tracks in the region where they overlap.

This creation of two tracks for overlapping hits is one of the most powerful distinguishing abilities that KLutz has over existing methods in NOvA. Figure 5.24 shows an event display featuring the same event that caused Kalman to only create a single track because of overlapping hits in a single detector view. KLutz tracks are drawn on the display, with the two sets of two tracks clearly visible, even in the section with overlapping energy depositions.

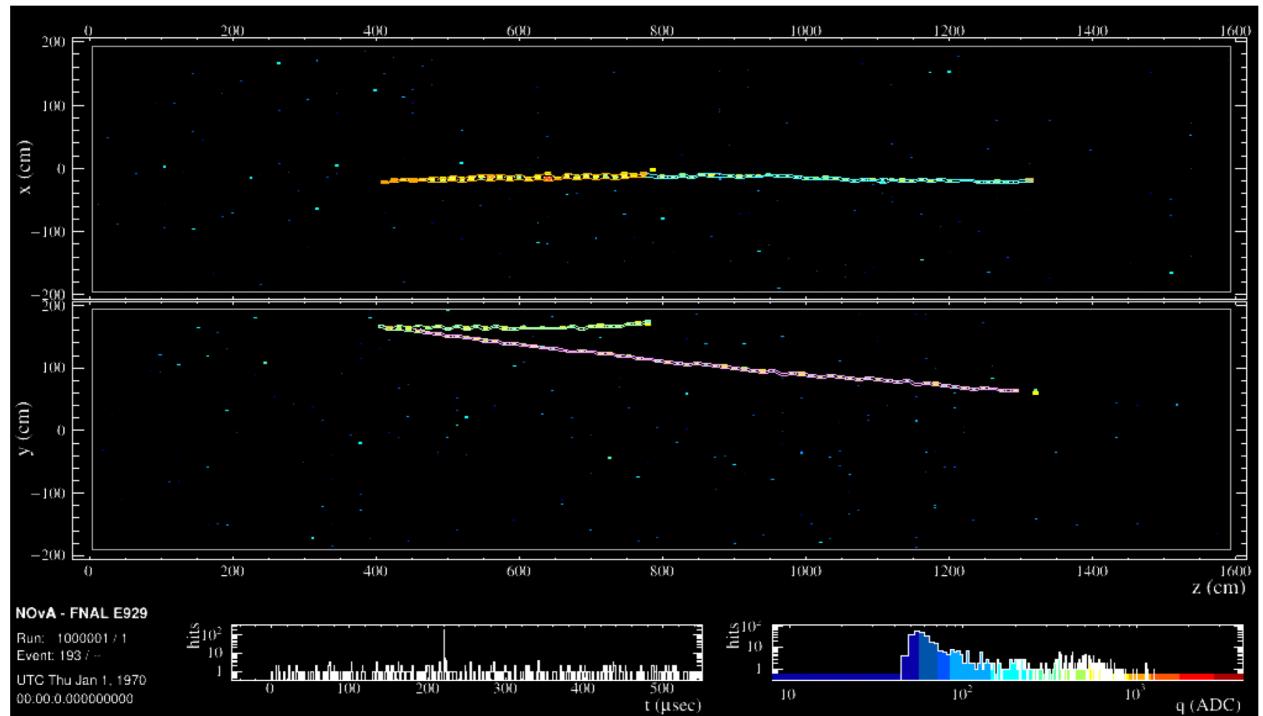


Figure 5.24: An event display showing four 2D KLutz tracks, despite two tracks overlapping completely in the x view.

Muon Catcher Tracks

The last thing needed to complete the set of four independent tracks is to ensure that hits from any muons traveling into the muon catcher are included in the tracks. The plane spacing in the muon catcher is larger than it is in the rest of the detector, so the initial algorithm sometimes fails to connect tracks in the muon catcher to hits in the primary detector volume. To address this issue, any track that are made of hits entirely within the muon catcher are stored and the hits that compose them are removed from the slice. I then proceed with the algorithm to create tracks as normal from the remaining hits deposited in the main volume of the detector.

After all of the hits in the primary detector volume have been utilized, I refer back to the stored track composed of only muon catcher hits. I check the co-linearity of the muon catcher track with the other primary volume tracks by comparing the slope and y-intercept of two linear functions fitted to each of them (denoted by dQ and dW , respectively). I check that the slope and y-intercept are within the thresholds of $dW < 7$ cm and $dQ < 7^\circ$, determined from the graphs shown in Figure 5.25. Any muon catcher tracks that meet the criteria for being co-linear with primary volume tracks are stitched together to form a single long track that travels through the detector and into the muon catcher. The process of identifying muon-catcher-only tracks, removing them from contention in the slice, and re-adding them back in to match co-linear tracks is illustrated in the event displays shown in Figures 5.26, 5.27, and 5.28 on the next page.

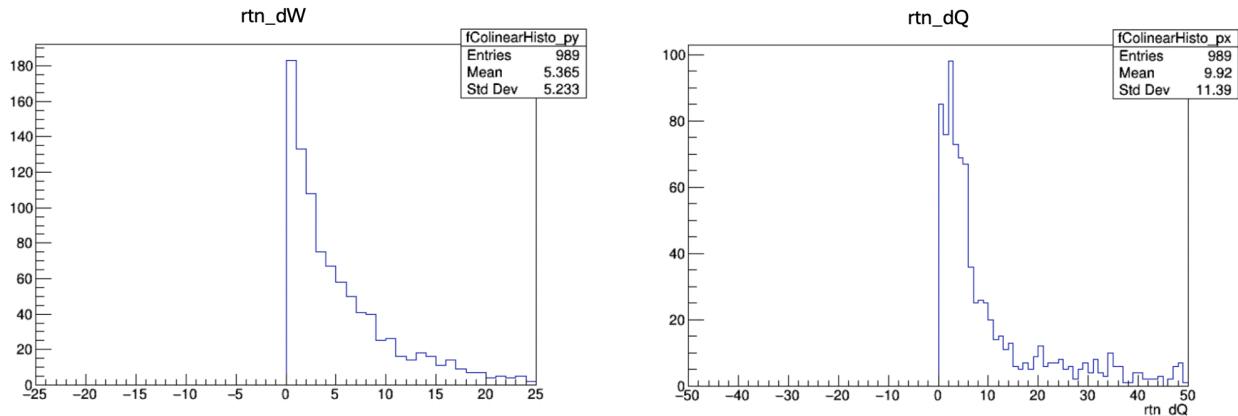


Figure 5.25: Differences in the slope (dQ) and y-intercept (dW) of collinear tracks. Graphs generated on a sample of 100,000 tridents, meaning muon-catcher-only tracks occur $\sim 1\%$ of the time.

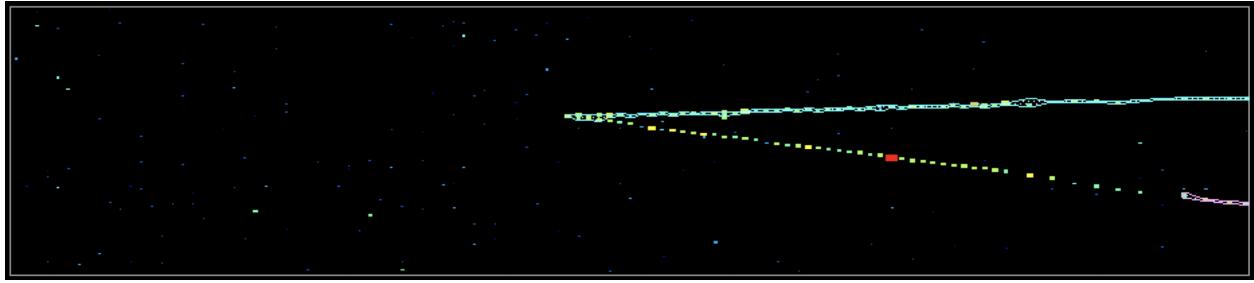


Figure 5.26: A track formed entirely within the muon catcher (pink).

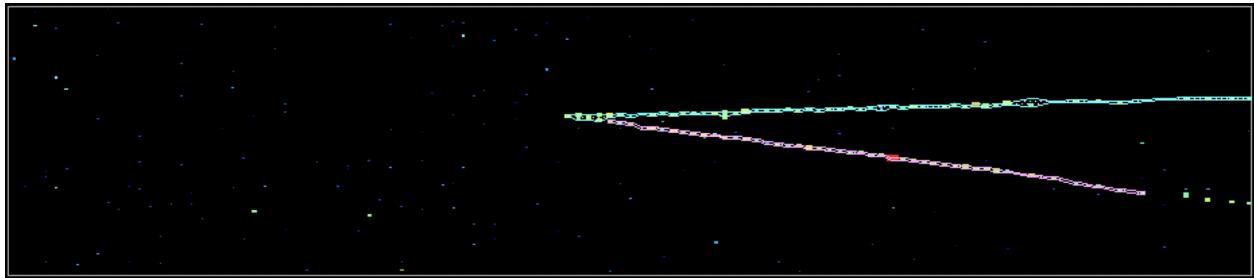


Figure 5.27: Remove the muon catcher track and recreate a second track without those hits.

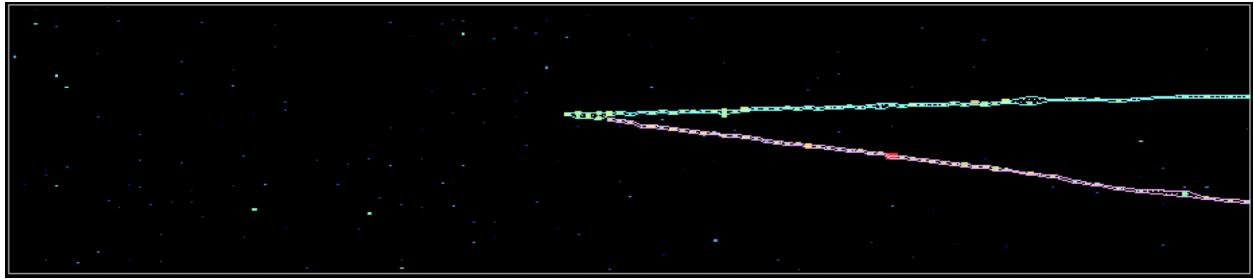


Figure 5.28: Check co-linearity and stitch the muon catcher track to the newly formed track.

Global Track Fit

Once all four independent tracks are formed, I perform a “global fit”, where all of the hits in the entire original slice and the four 2D tracks are provided to Lutz for one last set of track formation in each view. I grow four new global tracks by a process of “simulated annealing” - a method used in the Elastic Arms track-fitting algorithm [161]. In this annealing process, tracks are initially free to explore all the hits in a slice but are slowly solidified (or “frozen”) into place as hits are included on or excluded from that track.

I provide the global tracks with weights V_{ia} identifying which hit i should be associated with which track a , called the “Potts factor”. The Potts factor is equal to one if the a^{th} track goes through the i^{th} hit, and is zero otherwise. For our globally fit tracks, hits exclusively on one of the initially formed tracks are assigned a default value of $V_{ia} = 1$ for the associated global track, and hits that are not on the initially formed track are assigned a value of $V_{ia} = 0$. Any hits shared between the initial tracks are assigned $V_{ia} = 0.5$.

The annealing is accomplished by varying the parameter β in Equation 5.2 from low initial values (representing high temperatures) to high final values (representing low temperatures). By iterating track fits and χ^2 values calculated by Lutz, this process lets the global tracks “explore” hits with various weights before definitely including them or excluding them by setting $V_{ia} = 1$ or 0. The Potts factor is:

$$V_{ia} = \frac{e^{-\beta M_{ia}}}{e^{-\beta \lambda} + \sum_b e^{-\beta M_{ib}}}. \quad (5.2)$$

Here, M_{ia} is the contribution to the χ^2 of the a^{th} track by the i^{th} hit, λ is a chosen value that determines the χ^2 value for a noise hit, and β is the inverse “temperature” of a track, which influences how far a track can search for hits. Looking at the denominator of this function, one can see that if the χ^2 contribution of a hit is much larger than the chosen threshold λ , the noise term dominates and causes V_{ia} to approach zero. However if the hit’s χ^2 contribution is small then the Potts factor approaches one - indicating that the hit belongs on a track. As β grows larger, these V_{ia} values become more difficult to change, “locking in” the hit’s status. The result of this process is a set of four tracks that have had the chance to explore all hits in a slice before locking in a fit.

View Matching

The last step of tracking the two individual muons is to identify which tracks in the two views came from the same muon. This is done with a “View Match Algorithm”, initially developed by previous IU graduate student Evan Niner [162] and modified by me to take KLutz tracks as inputs. The algorithm takes multiple tracks in separate views, and matches tracks with like starting points, ending points, and energy deposition rates. The result is two reconstructed 3-dimensional tracks, each representing one of the muons that deposited energy in the detector.

The formation of two 3D muon tracks every single time KLutz runs successfully is my algorithm's largest advantage over existing reconstruction methods - this means that many of the events that previously had only one 3D track associated with them now have two tracks accurately capturing the energy from both muons. Figure 5.29 shows an event display of the simulated trident with poor Kalman reconstruction from the previous Figure 5.13, but with my two 3D KLutz tracks accurately capturing the path of each muon.

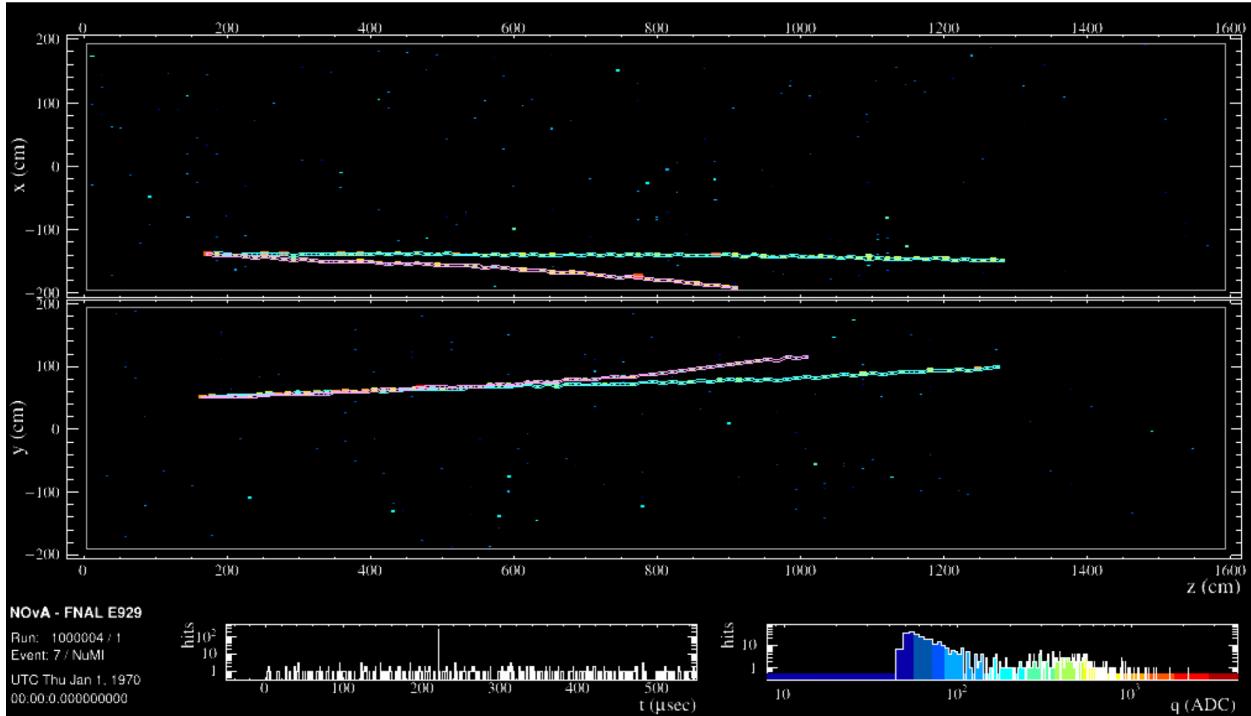


Figure 5.29: An event display showing two fully formed 3-dimensional tracks with a very narrow opening angle.

5.4.2 Additional KLutz Information

Two 3D muon tracks are the final product of the KLutz module. I label the longer of the two the “primary” track, and the shorter the “secondary” track. After running the resulting tracks through NOvA’s common analysis file making module and investigating some distributions, it became clear that additional information was needed at the individual cell-hit level that reconstructed is performed at. I added that additional information into the KLutz tracking module as extra clusters tacked onto the end of the two muon tracks.

Vertex Energy

We observed a number of events with high energy vertices due to hadronic activity in the same direction as the muons. To keep track of that information I determined a method to measure the vertex energy of an event by forming a cluster of all energy within the first 3 planes from the most downstream hit within the muon tracks.

Overlapping Hits

Background events that made it past the initial data reduction process frequently looked like single muons that emitted some bremsstrahlung radiation in the middle of their travel. The downstream end of that radiation would get picked up by KLutz and reconstructed into a second muon track that primarily overlapped with the first track in both views except for a small deviation at the end. To identify events like these I created a cluster formed out of all hits shared between the muon tracks. This allowed me to form a distribution of “the fraction of track 2 that overlaps with track 1”, which we expected to be a relatively small value for two true muons.

Remaining Slice Energy

The global track fit captures hits that belong to the two muon tracks, but any energy remaining in the slice initially went unused - an issue we resolved by classifying the energy into three types: hadronic energy close to the event vertex, hadronic energy far away from the event vertex, and excess slice energy. These classifications are determined based on the nuclear collision length λ of the detector - the average distance a hadron travels before interacting. Based on the nuclear composition of the near detector, its collision length is about 60 cm, or 10 detector planes. The probability of a hadron interacting within some detector distance z is

$$P(z) = 1 - e^{-z/\lambda} \quad (5.3)$$

based on the random Poisson nature of an interaction occurring [175]. From this idea, the classifications for hadronic energy are as follows:

- Collect 90% of hadronic energy into the “close” hadron energy cluster. Setting Equation 4.3 equal to 0.90 and solving for z yields a distance of approximately 2.3 interaction lengths, or 23 detector planes. So the close hadron energy cluster is built out of any energy deposited within 23 planes of the interaction vertex.
- Similarly, we wanted to collect 99% of hadronic energy within the “far” hadron energy cluster. This cluster includes hits that occur between 23 and 46 detector planes away from the vertex.
- Any hits further than 46 planes away from the vertex and not in a muon track are added into the “excess slice energy” cluster, as it has a 1% probability of being hadronic energy and is more likely to be simple noise hits.

5.4.3 KLutz Performance

I hand-scanned through 1000 event displays of simulated trident events with KLutz tracks overlayed to get a sense of how the final product performs. The category criteria and results of the scan are shown in Figure 5.30:

1000 Total Events				
Event Quality				
Easily identifiable as a trident event				
Trident-like, but not totally clear				
Indistinguishable from other events/just one muon				
Reco Quality				
Event Quality				
		712	60	0
		31	22	2
		9	28	136
		752	110	138
Reco Quality				
Tracked both muons nearly perfectly				
Decently matched muons, with some imperfections				
Didn't create any tracks/totally failed to track				

Figure 5.30: Results of hand-scanning events with KLutz reconstruction to measure performance.

The primary takeaway from the hand scan is that of 862 “trident-like” events, KLutz identifies two individual muons well in 827 events - a reconstruction efficiency of nearly 96%. Another lesson learned from the hand-scan is that approximately 14% of simulated trident events are “low quality”, meaning that they are indistinguishable from non-trident events by eye. The hopes of identifying such events in the final selection are virtually non-existent.

Figure 5.31 compares the number of reconstructed 3D Kalman tracks to the number of 3D KLutz tracks for 100,000 simulated trident events. KLutz runs successfully by creating 7 total tracks (2 tracks for muons and 5 for overlapping muons, vertex energy, and hadronic energy described previously) about 70% of the time. In contrast Kalman produces a track for each muon only about 45% of the time.

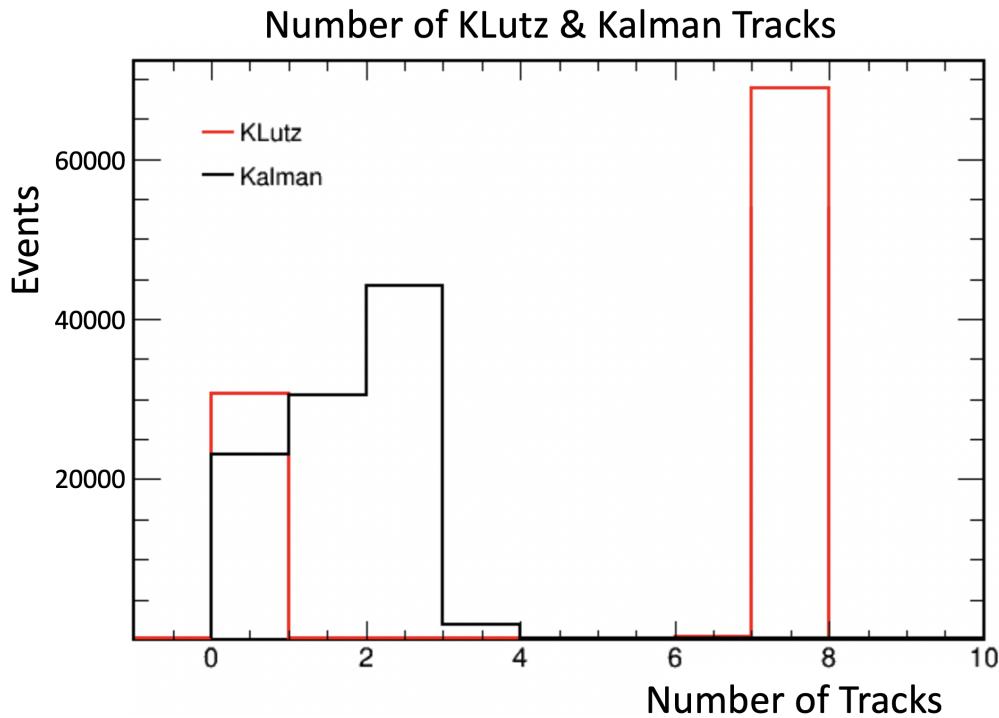


Figure 5.31: The number of KLutz and Kalman tracks produced on a simulated trident dataset.

5.5 Second Round Selection

I applied the KLutz tracker to the filtered background dataset described at the end of section 5.3 consisting of approximately 550k background events. I then started on the second round of selection by investigating how the KLutz variables looked in some basic distributions and identifying a few obvious non-trident event topologies that frequently passed through the first round selection. I made some basic quality cuts to remove them, then identified useful KLutz variables to feed into a BDT intended to isolate tridents.

5.5.1 KLutz Quality Cuts

Two obvious background topologies that frequently pass through the first round selection are single muons that tricked the algorithm into creating two tracks (typically caused by bremsstrahlung radiation), and two muons that match the criteria for being a trident event but are nowhere near each other in the detector. An example of the latter is shown in Figure 5.32, where one can see KLutz tracks overlayed on a background event with trident-like topology.

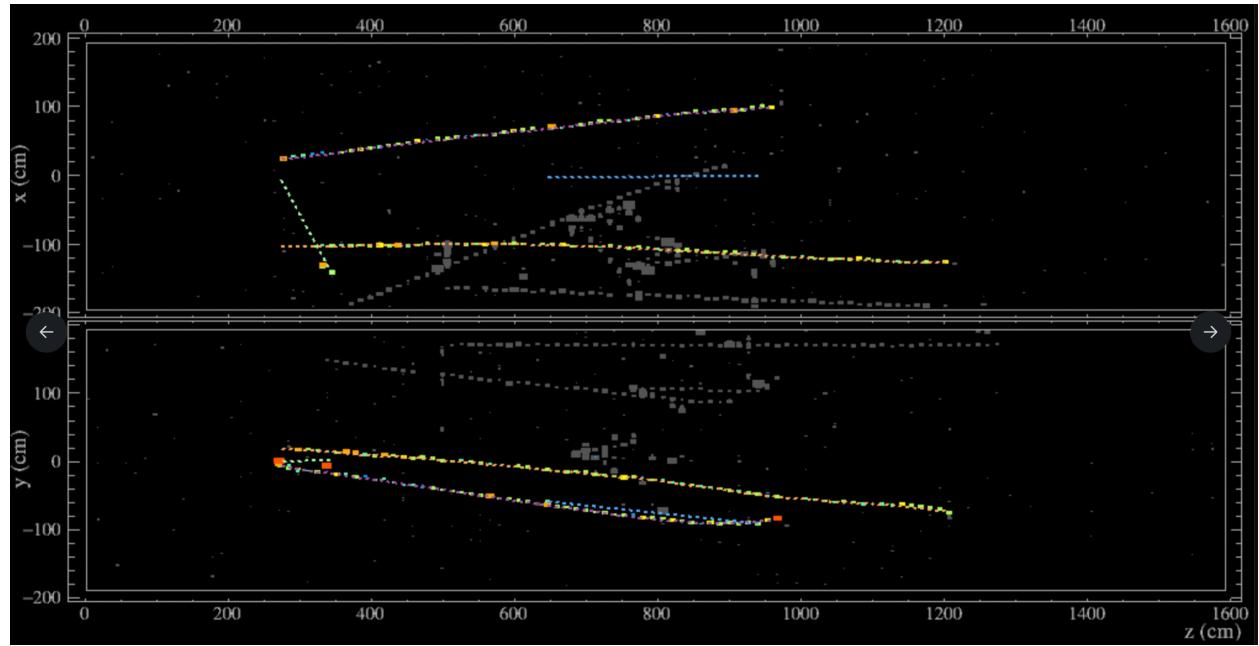


Figure 5.32: A clearly non-trident background event with two muon-like tracks, a small opening angle, and little hadronic energy.

The KLutz quality cuts designed to remove similar events are:

- Muon track overlap fraction < 1.0 : Require two distinct muon-like tracks.
- Muon track start point distance < 12 cm: Require any two muon-like track's furthest upstream hits to be near one another.
- Primary track ReMId score > 0.80 : Require that the longest track clearly looks like a muon.

Applying these cuts reduced the number of events from 21.97 signal on 461k background to 21.57 signal on 160k background in FHC, and 16.62 signal on 100k background to 16.40 signal on 30k background in RHC.

5.5.2 KLutz-Based BDT Selection

Developing the BDTs

I investigated KLutz track variables to find good BDT training candidates. Some of the most powerful distinguishing variables are the close-range hadronic energy, secondary track ReMId score, track vertex energy, opening angle between tracks, primary track length, and secondary track length. Figure 5.33 shows these key variables applied to FHC datasets after KLutz Quality cuts with the number of signal events scaled up by a factor of 1000 so that they are visible compared to background events.

The FHC and RHC datasets each had their own BDT trained to optimize the separation between signal and background. I investigated combinations of variables and BDT hyper-parameters to maximize the number of trident events maintained by various BDT cuts while keeping the following “figure of merit” (F.O.M) as high as possible.

$$\text{F.O.M.} = \frac{\text{Signal}}{\sqrt{\text{Signal} + \text{Background}}}. \quad (5.4)$$

A F.O.M. is a single value used to estimate the statistical significance of a selection; it does not guarantee a specific level of confidence in a measurement, but is indicative of the selection’s discriminating power [176]. It compares the observed rate of signal events to the uncertainty in observing a certain number of background events. A higher F.O.M. suggests that the analysis is *more likely* to yield a statistically significant excess of signal events, but the true confidence level of the measurement is not determined until the selection is applied to real data and a full statistical inference is performed - a discussion for chapter 6.

After multiple iterations of training and testing the BDTs, I settled on a set of optimal parameters that yielded the distributions shown in Figures 5.34 and 5.35. Tables 5.4 and 5.5 show all of the variables fed into each BDT, ranked by importance.

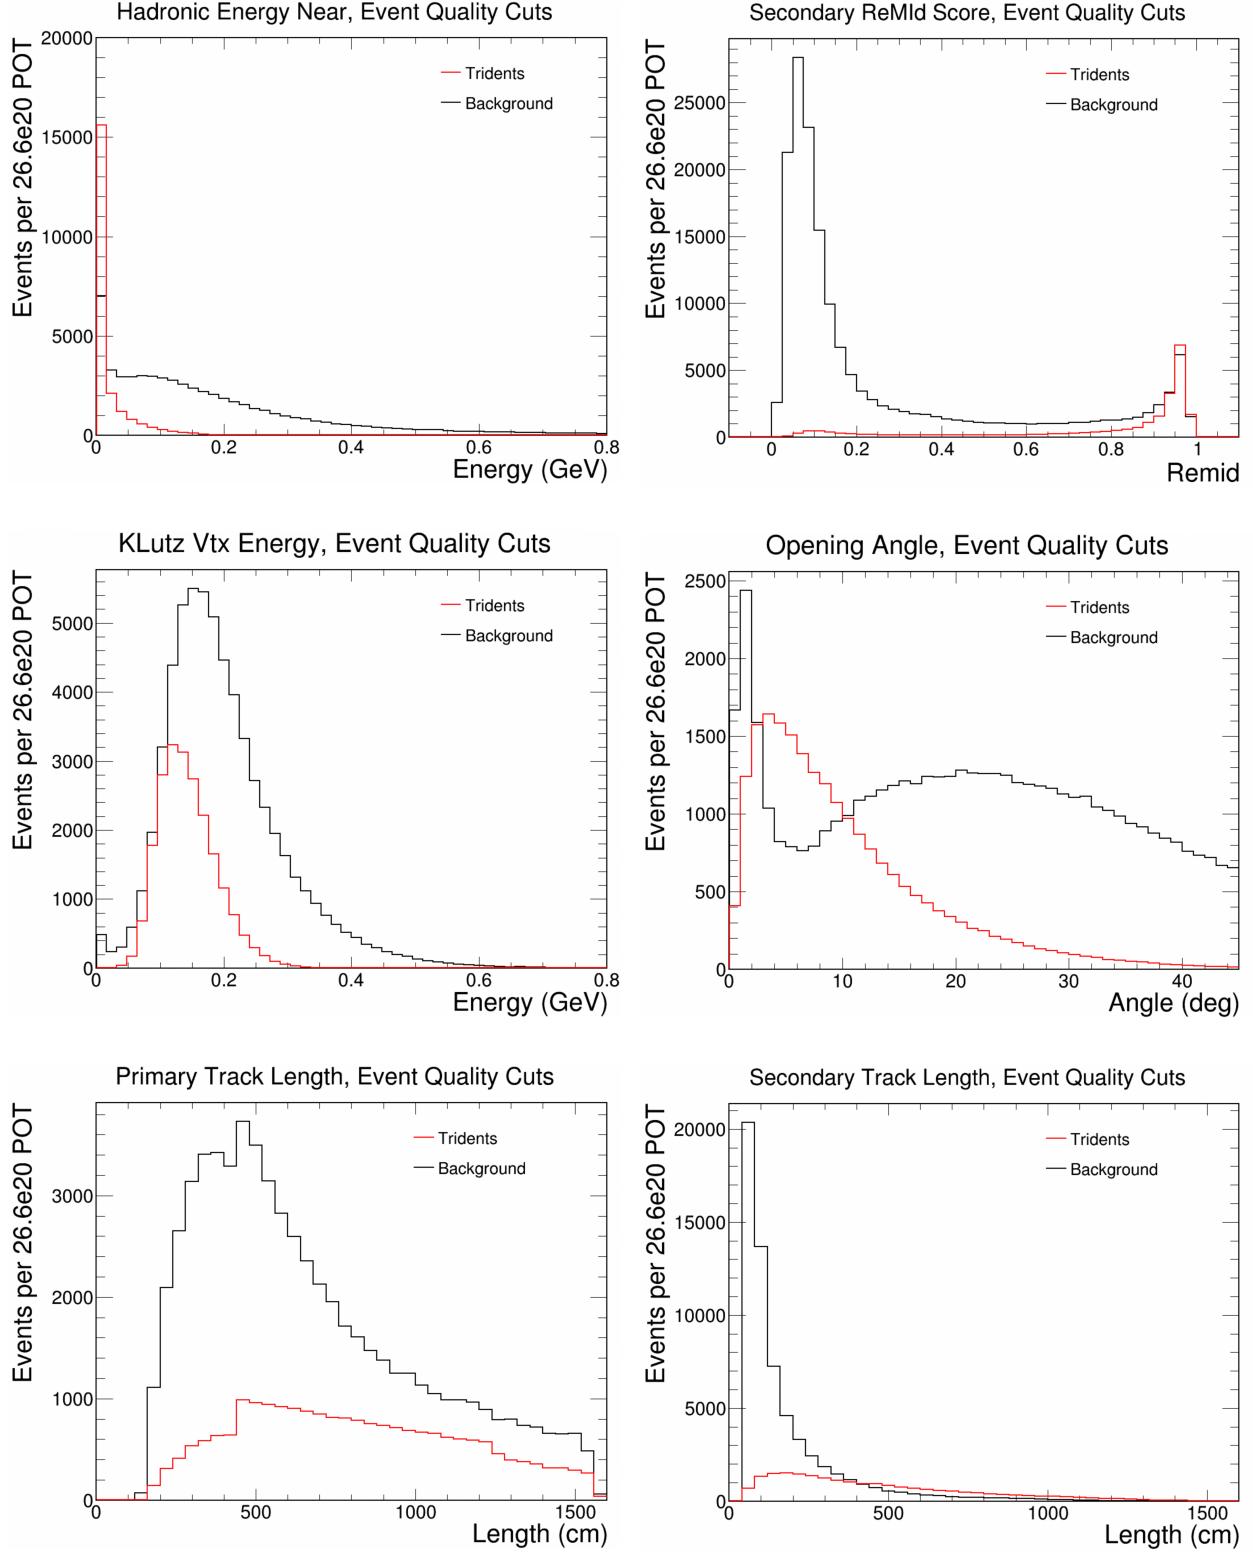


Figure 5.33: Some variables used in the second round FHC BDT, signal scaled up 1000 times. The jump in signal events at primary track length 450 cm is an artifact of the dataset generation, in which I limited vertex locations to 450 cm from the end of the detector.

FHC BDT Variable Importance		
Rank	Variable	Separation
1	Close-Range Hadronic Calorimetric Energy	0.08552
2	Secondary Track Length	0.08394
3	Track Opening Angle	0.08001
4	Secondary Track ReMID Score	0.07907
5	Primary Track Length	0.07480
6	Track Overlap Percent	0.07452
7	Vertex Calorimetric Energy	0.07428
8	Track Start Point Distance	0.07269
9	Secondary Track vs NuMI Opening Angle	0.06998
10	Primary Track vs NuMI Opening Angle	0.06747
11	Primary Track ReMID Score	0.06162
12	Reconstructed Invariant Mass	0.05890
13	Medium-Range Hadronic Calorimetric Energy	0.04624
14	Far-Range Hadronic Calorimetric Energy	0.03733
15	Number of Michel Electrons in the Slice	0.03362

Table 5.4: A table of FHC variable effectiveness. Separation of 0 represents no separation between signal and background, and 1 represents complete separation.

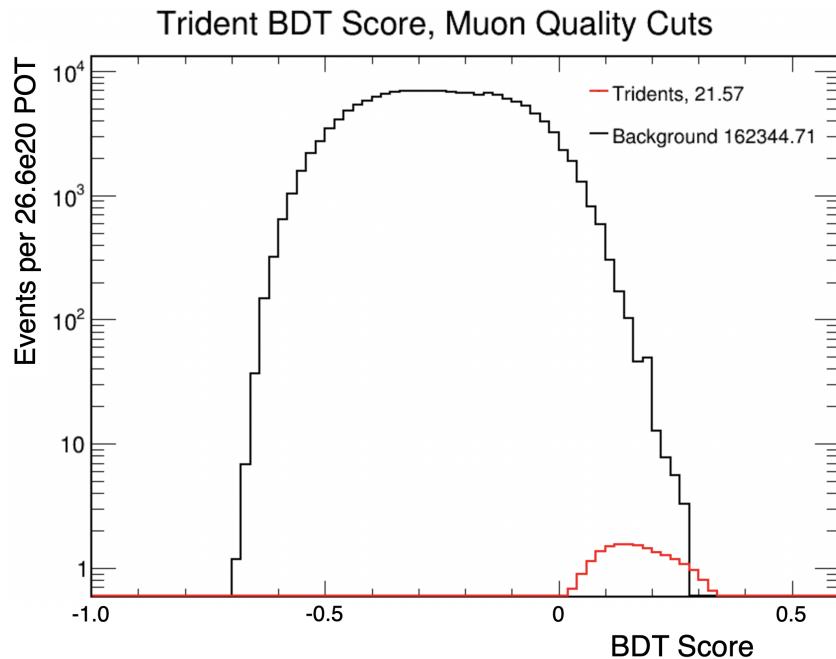


Figure 5.34: Output of the simulated FHC dataset BDT trained on KLutz variables.

RHC BDT Variable Importance		
Rank	Variable	Separation
1	Secondary Track ReMId Score	0.1050
2	Primary Track Length	0.1004
3	Secondary Track Length	0.09287
4	Close-Range Hadronic Calorimetric Energy	0.09272
5	Primary Track ReMId Score	0.09229
6	Track Opening Angle	0.08779
7	Track Overlap Percent	0.08535
8	Track Start Point Distance	0.08462
9	Vertex Calorimetric Energy	0.08243
10	Primary Track vs NuMI Opening Angle	0.07996
11	Reconstructed Invariant Mass	0.06347
12	Number of Michel Electrons in the Slice	0.03338

Table 5.5: A table of RHC variable effectiveness. Separation of 0 represents no separation between signal and background, and 1 represents complete separation.

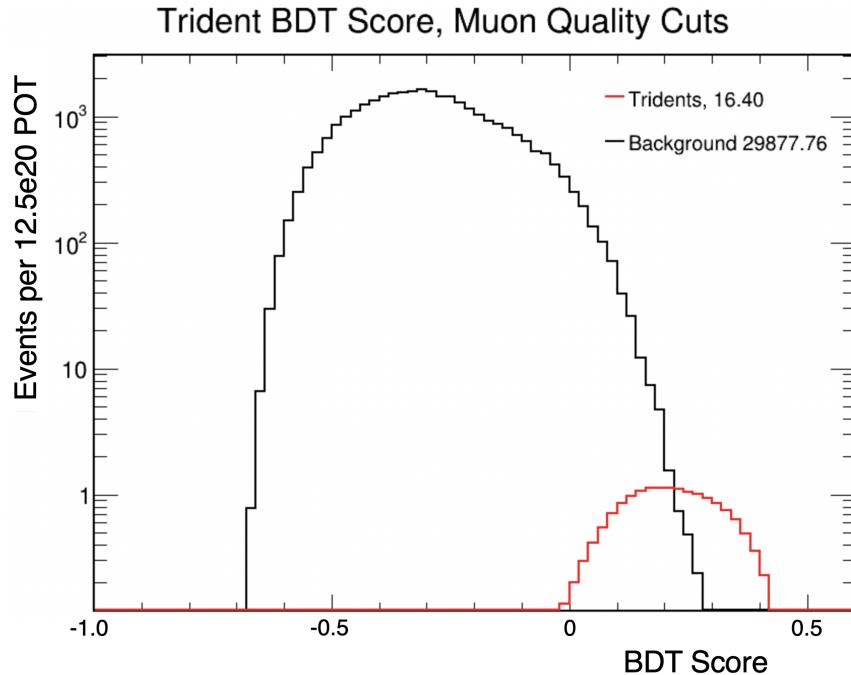


Figure 5.35: Output of the simulated RHC dataset BDT trained on KLutz variables.

FHC BDT Cut Values				RHC BDT Cut Values			
BDT >	S	B	F.O.M.	BDT >	S	B	F.O.M.
0.160	11.40	89.69	1.13	0.120	13.00	56.23	1.56
0.180	9.89	56.99	1.21	0.140	12.03	28.70	1.88
0.200	8.46	21.24	1.55	0.160	10.97	15.07	2.15
0.220	7.11	12.03	1.63	0.180	9.84	7.27	2.38
0.240	5.84	6.40	1.67	0.200	8.71	3.12	2.53
0.260	4.66	2.37	1.76	0.220	7.57	1.56	2.31
0.280	3.58	0.00	1.89	0.240	6.47	0.78	2.40
				0.260	5.42	0.26	2.27
				0.280	4.41	0.00	2.10

Table 5.6: Tables comparing signal and background numbers at various cuts on BDT score for both the FHC and RHC KLutz-Based BDTs. Explanation of low-B values in body paragraph.

BDT Performance

The numbers in Table 5.6 show the predicted signal and background numbers for various cuts on these BDT output distributions. I determined reasonable cut values by calculating the optimal F.O.M. based on these numbers. Cuts of FHC BDT ≥ 0.270 and RHC BDT ≥ 0.210 optimized the individual F.O.M.s as seen in the table, yielding a combined F.O.M. of 3.05. It is worth noting again that the relatively small background MC sample presents a large amount of uncertainty in the actual values estimated in these tables. The zeros in the last rows of each table do not truly reflect an expectation of zero background, but rather that no MC events survive into that region of the BDT. As a result it is not immediately obvious if cuts at 0.290 on both BDTs would improve the F.O.M. relative to the chosen cuts. Further investigation of these distributions using real data is discussed in chapter 6.

Another point to note here is that the RHC BDT performs better than the FHC BDT. This could be attributed to multiple factors, but one good explanation is that the trident production cross section is the same for both neutrinos and anti-neutrinos, whereas the total overall muon anti-neutrino cross section is only about 50% the size of the overall muon neutrino cross section (as seen in Figure 5.36 [32]). Another important consideration is that the cross sections of common background events may not be modeled as well for RHC as they are for FHC.

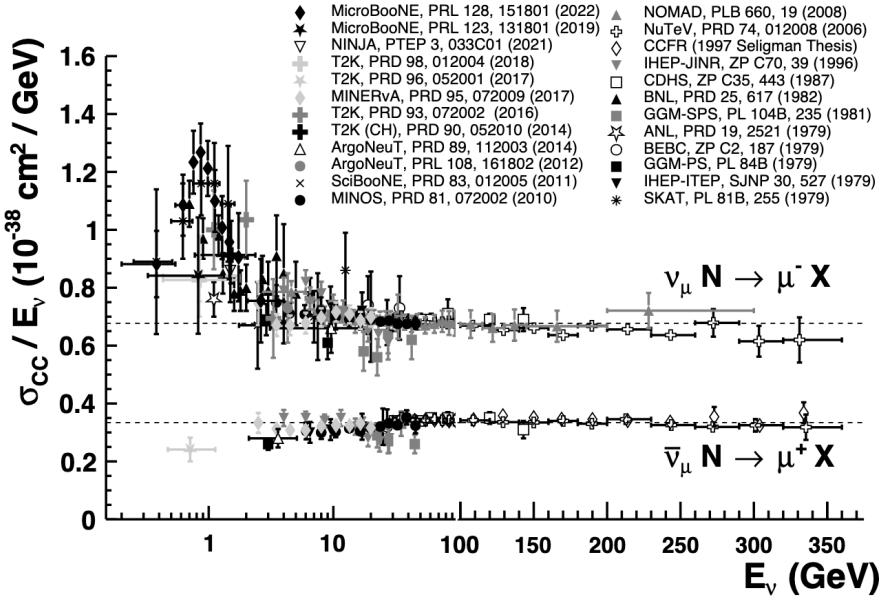


Figure 5.36: Total neutrino and anti-neutrino cross sections as a function of energy. In the relevant energy range (~ 2 to 20 GeV) the $\bar{\nu}_\mu$ cross section is roughly half that of ν_μ [32].

5.6 Selected Backgrounds

Choosing the cuts with the optimal F.O.M. yields a combined expected event rate of 13.37 tridents on 5.87 backgrounds with MC alone. The average true neutrino energy of events that passed the selection was approximately 18 GeV. A background rate of 5.87 corresponds to 35 simulated events that actually passed the selection with data POT scaling applied.

I scanned the 35 simulated background events to better understand their topology and interaction modes. The majority of events consist of coherent interactions that produce a pion. The pion frequently decays quickly into a muon, resulting in a final state containing a pair of muons and very little hadronic energy - a topology identical to that of trident events, but resulting from a different interaction mechanism. The breakdown of expected background types is in Table 5.7.

Rank	Interaction Type	% (Number)
1	Coherent Pion Production, Decay in Flight	54% (19)
2	Coherent Pion Production, No Decay	20% (7)
3	DIS with a Non-Interacting Neutron	17% (6)
4	Other Unlucky Topologies	9% (3)

Table 5.7: Interaction types composing the background events selected by the KLutz-Based BDT.

5.6.1 Systematic Background Uncertainties

The systematic uncertainties associated with this background estimation come primarily from the NuMI beam flux at the energies selected by the BDT and the modeling of coherent cross sections at those energies.

NuMI Beam Flux

As mentioned previously, the average true neutrino energy selected by the KLutz-based BDTs is 18 GeV. The NuMI beam peaks at approximately 6 GeV, with the NOvA near detector exposure peaked at 2 GeV because of its location 14 mrad off the primary beam axis. Therefore, the NuMI flux at energies above 5 GeV is rarely studied in NOvA [177]. Figure 5.37 shows the most detailed study of flux uncertainties on neutrino energies greater than 6 GeV performed so far [178].

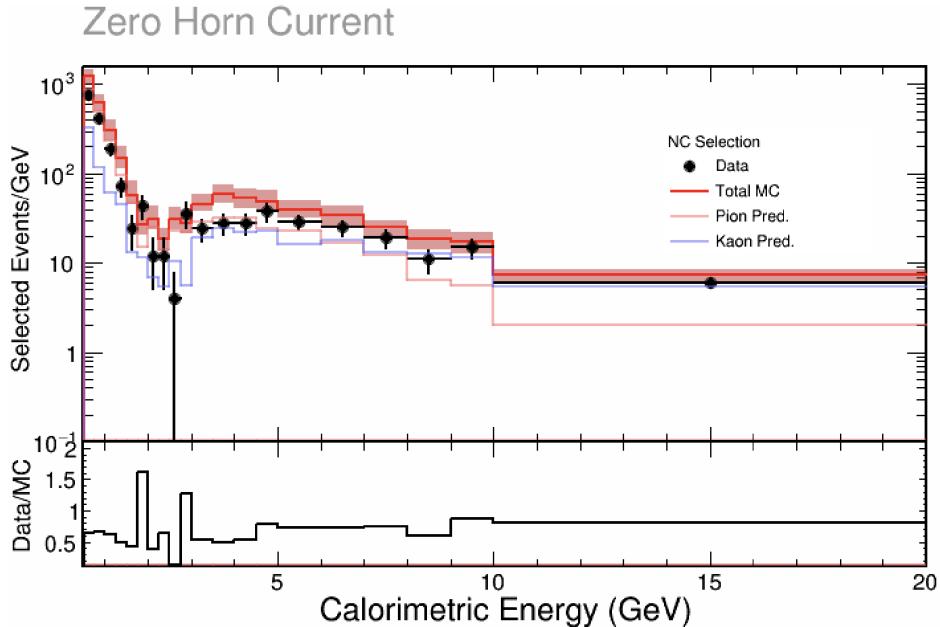


Figure 5.37: NOvA’s most detailed study of flux uncertainties for $\langle E_\nu \rangle > 6 \text{ GeV}$ [178].

Typically, the uncertainty on NuMI flux used in NOvA ND analyses is calculated using PPFX (see subsection 4.5.1), and considers the fiducial volume of the ND the beam is impinging upon. The uncertainty up to 5 GeV ranges from about 12-20% and the total flux uncertainty is determined based on the relative frequency of events at certain energies. For this analysis, we assume a conservative 20% systematic uncertainty on the NuMI flux for $\langle E_\nu \rangle \approx 18 \text{ GeV}$.

Kaon Flux

An additional reason for the mismatch in this region may be the flux of muon neutrinos originating from kaon decays. Events that pass the first round of data reduction tend to have higher energy than is typical of NOvA's neutrino spectrum because of the preference for events with multiple long tracks.

As shown in Figure 5.2, the high-energy region of the neutrino beam consists primarily of Kaon-induced neutrinos [179, 180]. Since NuMI's kaon flux is less understood than the pion flux, the NOvA standard for handling kaon modeling uncertainty is to apply an additional 10% systematic uncertainty to the analysis.

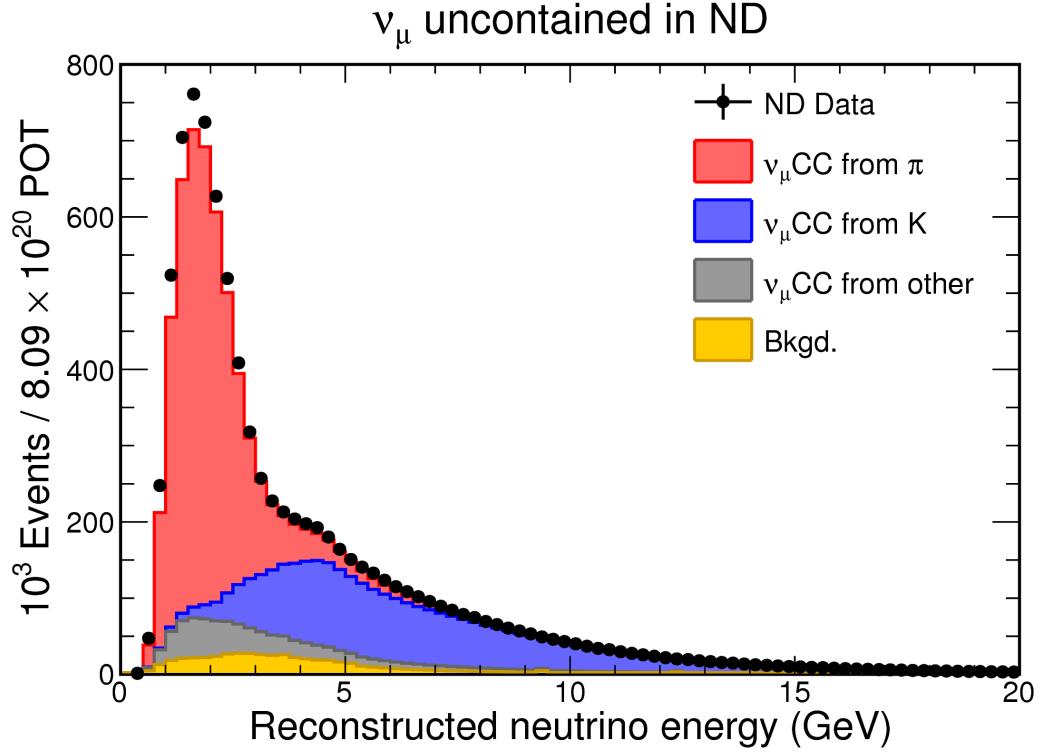


Figure 5.38: Sources of ν_μ flux modeled in the ND overlayed with data [179].

Coherent Cross Section Modeling

The primary background type for this analysis is coherent pion production - a very well studied process [181]. However, there is disagreement between data and models in the regime of neutrino energy greater than 10 GeV. Figures 5.39 and 5.40 show that for neutrinos in this range, the measurement uncertainty on Carbon is about 25%.

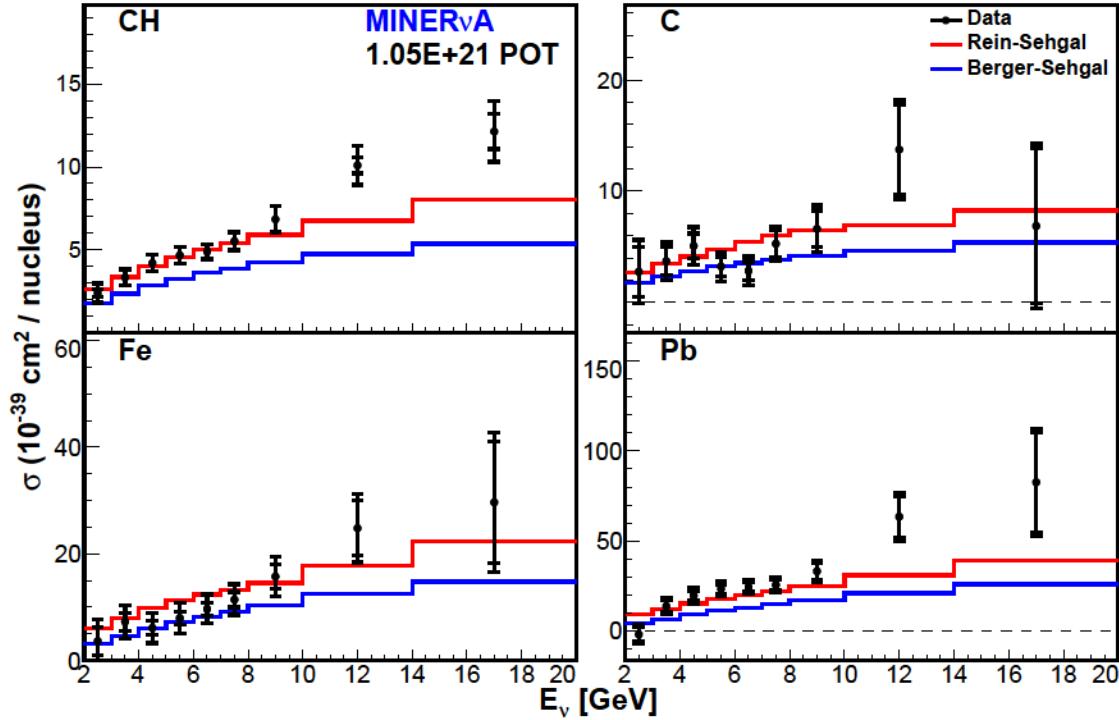


Figure 5.39: Current measurements of the coherent pion production cross section [181].

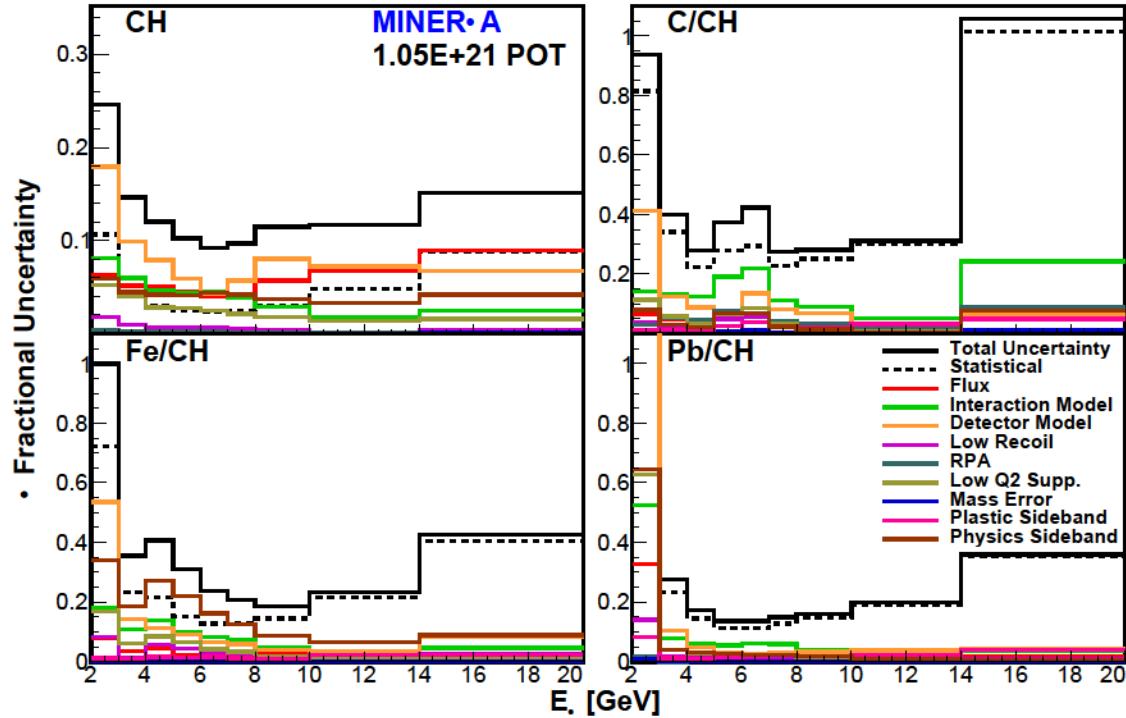


Figure 5.40: Coherent pion production cross section uncertainty (model uncertainty in green) [181].

A full cross section measurement of trident production would require each of these uncertainties to be investigated in detail. However, since the goal of this analysis is a count of signal over background events, adding the approximate flux, kaon, and model uncertainties in quadrature is sufficient to estimate our systematic background uncertainties. This process yields a total systematic uncertainty of roughly 33.5% on the background simulation.

5.6.2 Data Driven Background Estimation

The estimate on systematic uncertainty described in subsection 5.6.1 works well in the low-BDT score “background” region where the MC statistics are plentiful. However, in the high-BDT score “signal” region, the limited MC statistics lead to large fluctuations in predicted background rates, resulting in statistical uncertainties that exceed the systematics. For example, an FHC bin containing three MC background events represents approximately one expected real data event. Since particle interactions follow Poisson statistics, the uncertainty on the MC prediction is $\sqrt{3} \approx 1.73$ events. This corresponds to a 58% relative statistical uncertainty on that bin, meaning the MC prediction in the signal region is unreliable. This also does not account for bins containing no MC events despite having non-zero probability of background events occurring there, which may skew the estimate even further.

To mitigate the effects of the high statistical uncertainty, I adopt a data-driven approach to estimate backgrounds in the signal region. The first step in the process is to determine the “shape” of the background - the functional form the distribution adheres to most closely. I fit that function to data in a mid-BDT score “sideband” region, then extrapolate it into the signal region. The integral of the extrapolated function in the signal region is used to estimate the background count, essentially simulating the “infinite MC” limit where the background distribution is known perfectly (up to model uncertainties).

I developed and validated the procedure on FHC MC before applying it to real data. An initial extrapolation using an exponential best fit predicts a background of 5.39 events in the signal region, while the MC itself predicts 2.37 events. The extrapolated estimate is higher due to the fact that the extrapolation is smooth and extends up to a BDT score of 0.5, whereas the available MC only ranges up to a score of 0.29.

Uncertainty on the Extrapolated Fit Background Estimate

Although it avoids the issue of low MC statistics, there are still systematic uncertainties associated with this method of background estimation. They come from the uncertainty in the fit itself, and the uncertainty that the functional form accurately represents the true shape of the distribution.

Statistical Fit Uncertainty

The statistical uncertainty of a fit reflects how well a model corresponds to the observed data and how the data fluctuate around the chosen functional form [182]. ROOT provides tools to evaluate the goodness-of-fit for each parameter of a function, as well as the correlations between parameters [183]. These results are stored in a “covariance matrix” - a matrix containing the variance of each parameter along the diagonal, and covariances between the parameters in the off-diagonals. It captures the extent to which parameters are constrained by the data and how they vary together [184]. By propagating the covariance matrix through the extrapolation, I determine the statistical uncertainty on the integral of the fit function in the signal region. Applying this procedure to the exponential function from the previous section yields a prediction of 5.39 ± 1.16 background events in the signal region.

Fit Shape Uncertainty

Uncertainty in the underlying functional form of the background distribution is handled using a “discrete profiling” estimation method [185]. The method considers the choice of functional form as a “nuisance parameter” - a variable which is allowed to change in order to explore the phase space of plausible models. Rather than relying on a single function, this method forms an “envelope” spanning a range of reasonable estimates for the expected number of background events.

Commonly suggested functional forms for this type of profiling include exponential functions, Bernstein and Laurent polynomials, and power-law functions [186]. I initially tested all of these, but after finding that polynomials perform poorly for extrapolation I settled on just exponential and power law estimates in the final analysis. The background predictions in the signal region from these functions are:

- Exponential: $B_{\text{exp}} \pm \sigma_{\text{exp}} \approx 5.39 \pm 1.16$ events
- Power Law: $B_{\text{pow}} \pm \sigma_{\text{pow}} \approx 10.22 \pm 1.51$ events

The fits matched to FHC MC data estimates are shown in Figure 5.41.

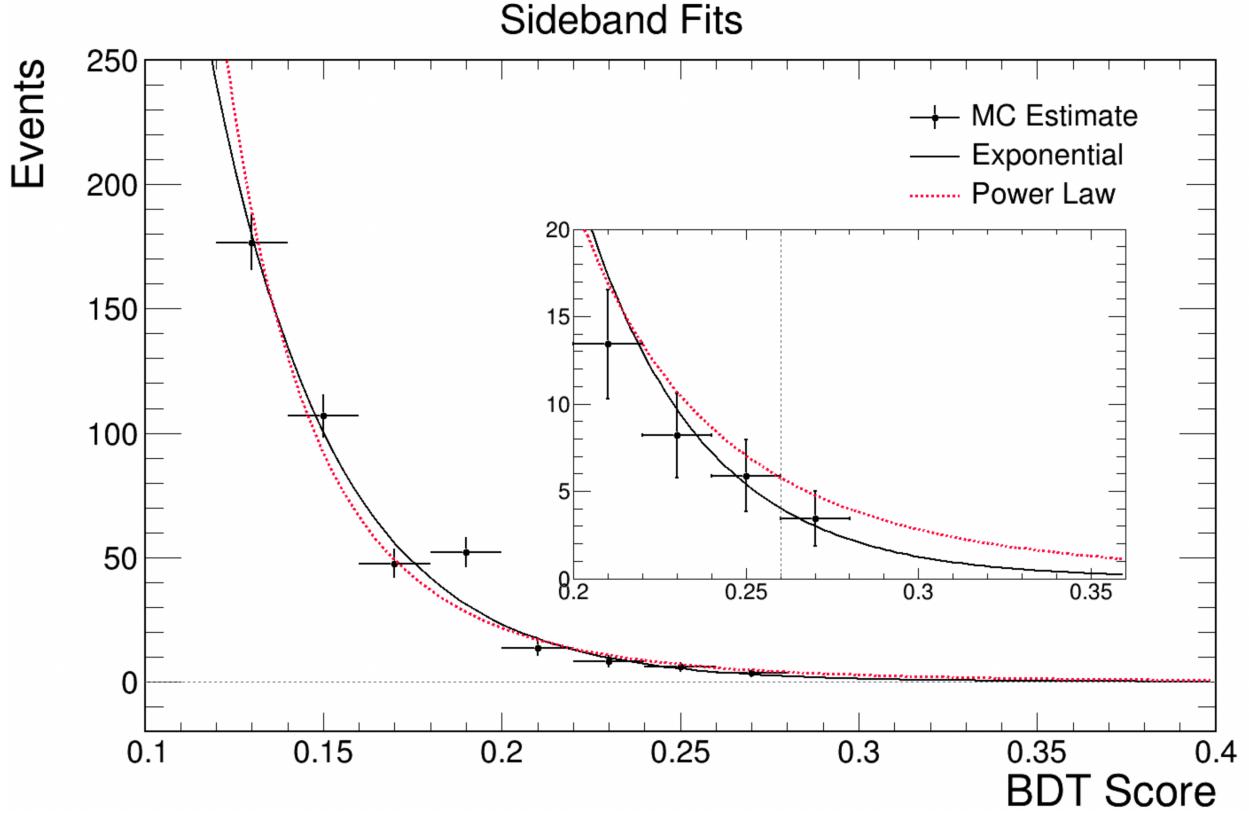


Figure 5.41: Exponential and Power Law fits overlayed on FHC MC. Fits were performed on all available MC (up to a BDT score of 0.27) and extrapolated into the signal region.

The exponential function is the fit with the smaller χ^2 , so I use it as the central value estimate. I incorporate the power-law prediction by treating it as a source of systematic uncertainty on the exponential fit. The power law fit uncertainty is not considered, as it likely has 100% correlation with the exponential fit uncertainty. I estimate the distribution shape uncertainty by finding the relative difference between the two fits and adding that in quadrature with the statistical uncertainty of the exponential fit. For example with the numbers reported here:

$$\begin{aligned} \sigma_{\text{fractional pow}} &= \frac{|B_{\text{pow}} - B_{\text{exp}}|}{B_{\text{exp}}} = 0.896, \quad \sigma_{\text{absolute pow}} = \sigma_{\text{fractional pow}} \times B_{\text{exp}} = 4.83 \\ \implies B_{\text{exp}} &= 5.39 \pm \sqrt{4.83^2 + 1.16^2} = 5.39 \pm 4.97 \text{ events.} \end{aligned}$$

In chapter 6 I will discuss how this result is used to create a weighted average estimate on the expected number of real data background events.

Chapter 6

Analyzing Near Detector Data

The analysis methods of chapter 5 were developed using simulated data to avoid bias when looking at real detector data. In this chapter, I describe the application of this procedure to the real data, and how the data was analyzed in each step of the process. I also present the final result of the search. I used NOvA’s nominal production 5.1 datasets for this search: the FHC ND dataset corresponds to 25.5×10^{20} POT, and the RHC ND dataset corresponds to 11.38×10^{20} POT.

6.1 Unblinding the Background Regions

“Unblinding” refers to the process of looking at real data after developing an analysis based on simulation [187]. In general, it is important to avoid making changes to the analysis procedure during unblinding, within reason. To approach this process carefully I adhered to a three-phase unblinding procedure. In each phase I performed a set of checks to ensure consistency between data and simulation within reasonable uncertainty.

- **Phase 1:** Unblind the BDT score ≤ 0 “background” region where the MC predicts very little signal to occur.
 - Check that the data and simulation distributions agree to a reasonable extent, considering systematic uncertainties.
 - Determine normalizations for the MC interaction components that best match the data and apply them as an MC weight.
- **Phase 2:** Unblind the mid-BDT score “sideband” region, where MC predicts approximately 6000 background events and 20 signal split between the datasets.
 - Confirm that the distributions still agree within systematic uncertainties, with the understanding that low MC statistics will start to impact the comparison.
 - Perform a data-driven estimation of expected background events in the signal region based off of extrapolation from this sideband.

- **Phase 3:** Unblind the high-BDT score “signal” region, where MC predicts approximately 6 background events and 13 tridents split between the datasets.
 - Use the background estimate and the number of observed events to perform a statistical inference on the potential number of trident events.

6.1.1 The FHC Dataset

The results of unblinding the background region of the FHC dataset are shown in Figure 6.1

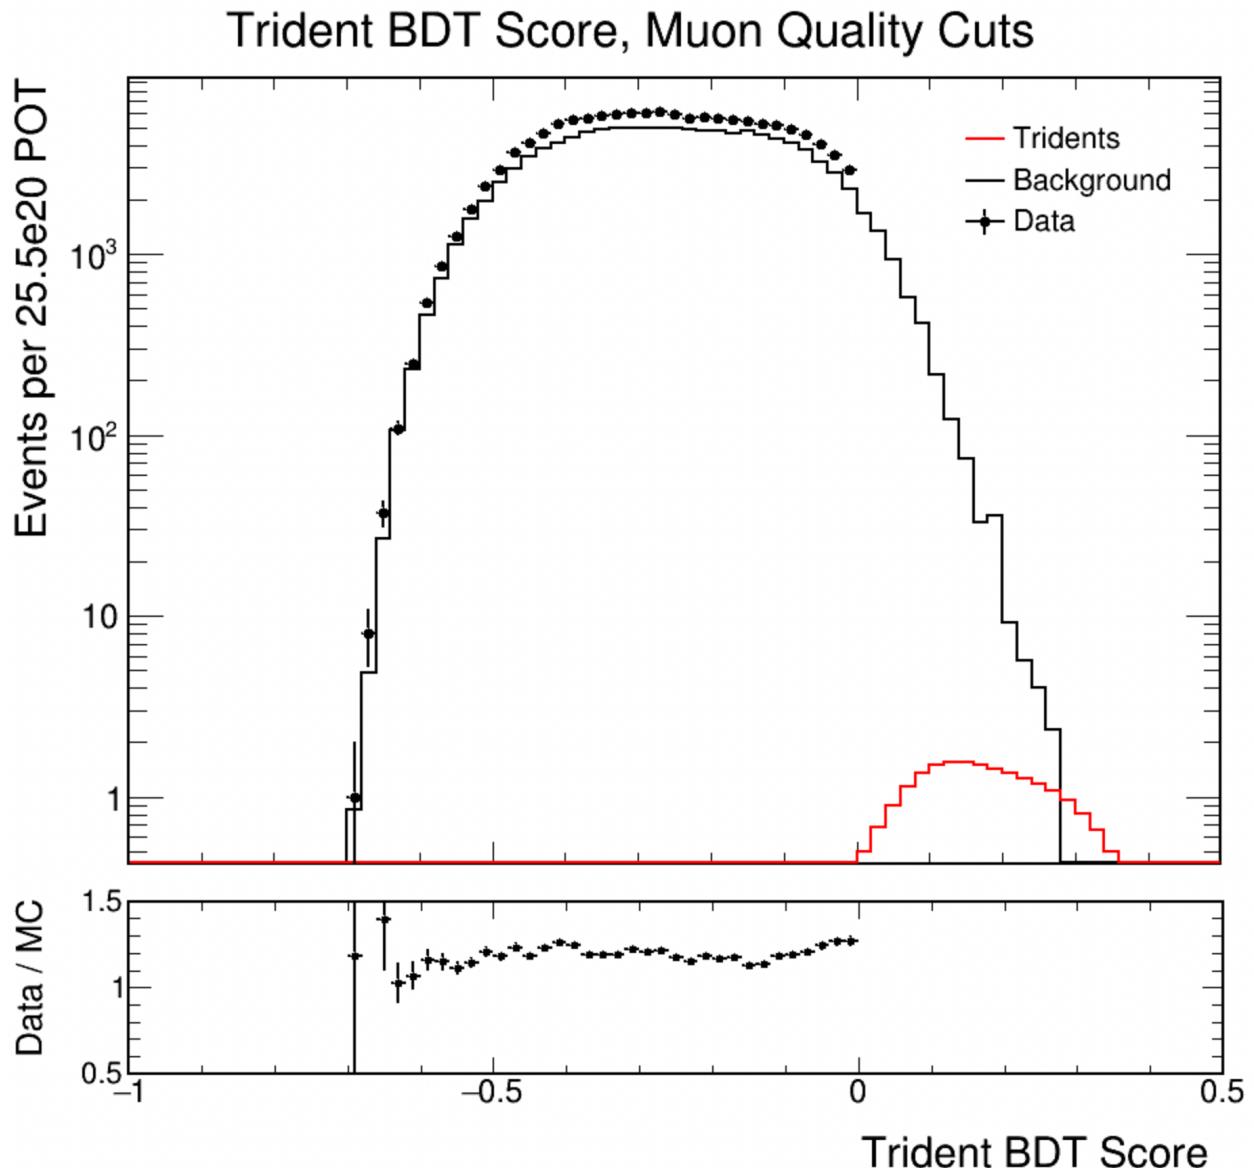


Figure 6.1: KLutz-Based Trident BDT score with FHC data and MC overlayed.

The ratio of data to MC ranges from about 1.0 to 1.3 where statistics are plentiful. As discussed in subsection 5.6.1, the MC estimate carries a 33.5% systematic uncertainty due to flux and cross section modeling, so the data and MC do agree within expectations.

Exploring MC Weights

Although the MC and data agree within uncertainties, we were interested in studying potential causes of the discrepancy. We investigated the contribution of individual interaction channels to the overall distribution and reweighted the interaction components of the nominal MC by minimizing the best-fit χ^2 as determined with the MINUIT fit optimization tool [188]. Each component shares a common flux normalization factor, and can vary proportionally to the QE component. Variations are constrained via additions to the fit χ^2 . The results of this study are shown in Figure 6.2.

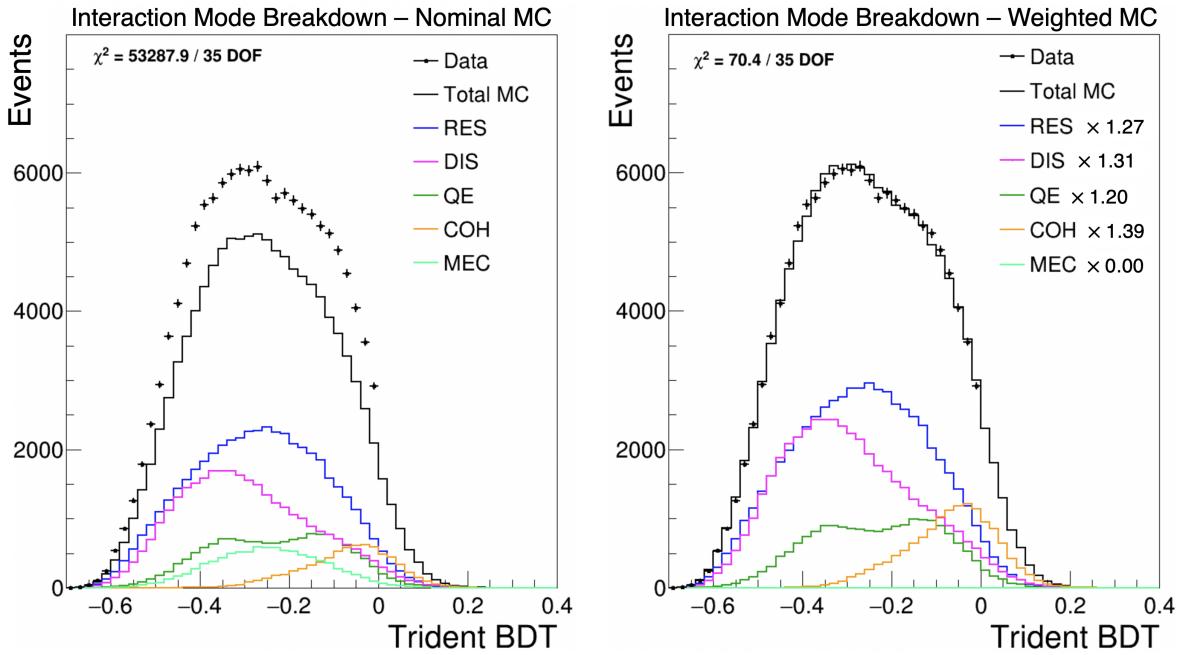


Figure 6.2: Before and after comparisons of FHC nominal MC, weighted MC, and data.

The RES, DIS, and QE contributions do not change more than the expected systematic uncertainty, but the coherent contribution is increased by 40%. Additionally, the MEC events contribute 28 points to the χ^2 , and the best-fit reduces their contribution to zero, indicating that there may be an issue with NOvA's MEC process modeling. No further investigation was done beyond this point, but the results of the study could act as motivation for further consideration of coherent and MEC interactions, particularly at high energies.

6.1.2 The RHC Dataset

The results of unblinding the background region of the RHC dataset are shown in Figure 6.3. A clear mismatch develops as the BDT score rises, indicating a potential issue with the MC modeling. I investigated possible causes of the discrepancy, starting with going back to the very beginning of the selection process and seeing how the distribution changes with each cut applied.

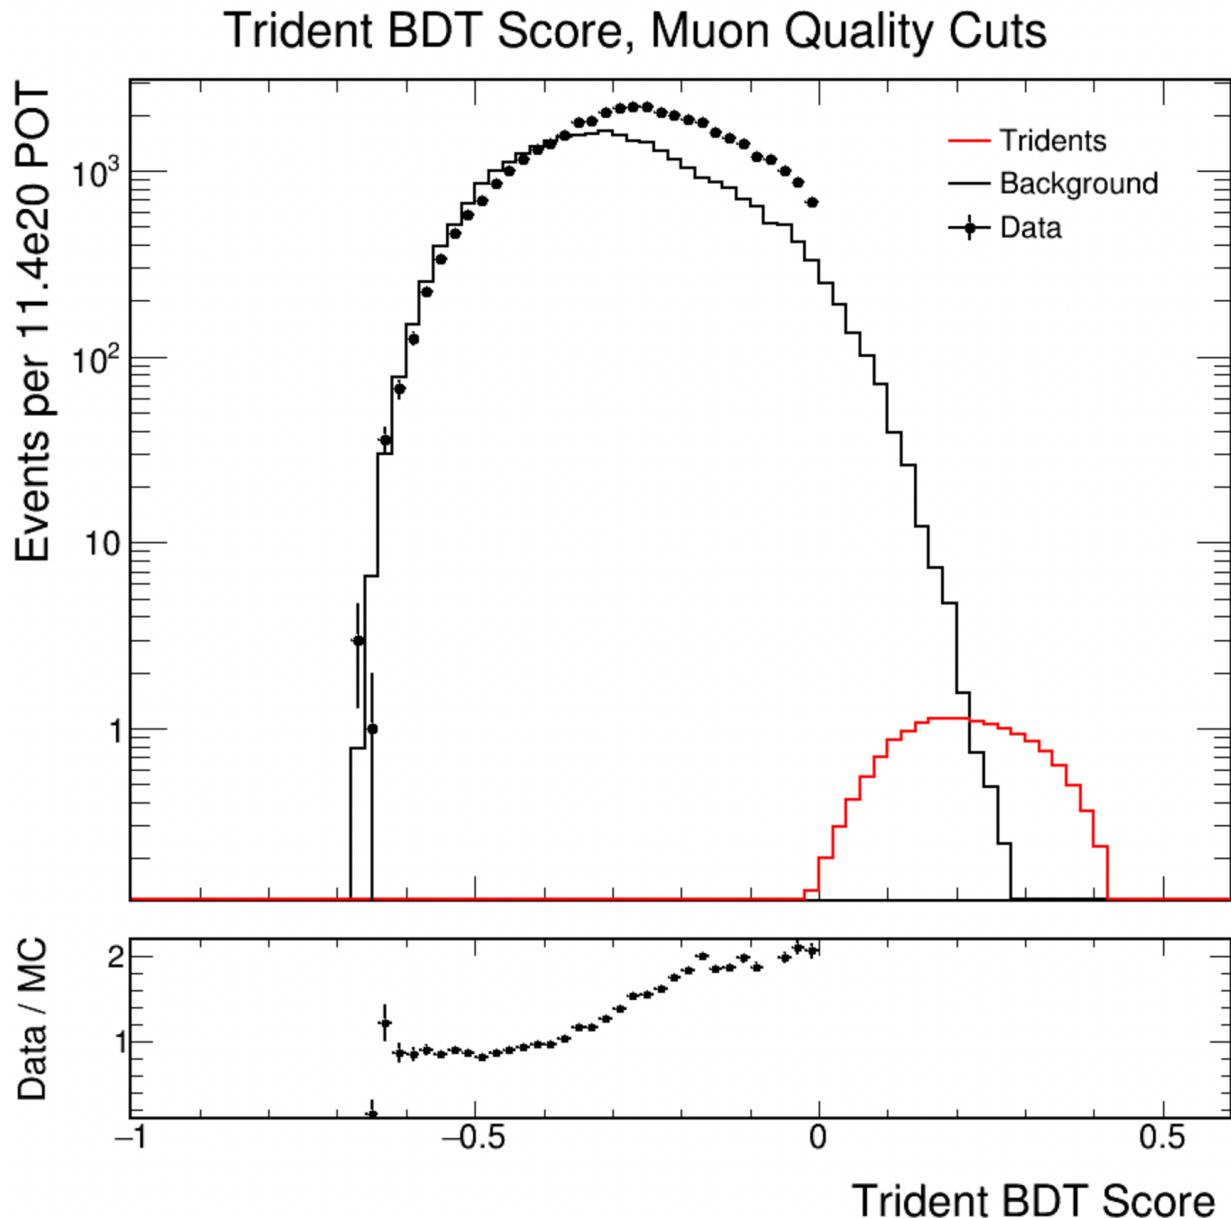


Figure 6.3: KLutz-Based Trident BDT score with RHC data and MC overlayed.

Cut-By-Cut Investigation

Starting with the standard set of NOvA data quality cuts, I generated distributions of the three original BDTs with MC and data overlayed. I applied the appropriate topology BDT cut when investigating the one-track and two-track BDTs. The results of the investigation are shown in Figures 6.4, 6.5, and 6.6.

The mismatch is not too dramatic at the start, but trends towards what is seen in the KLutz-based BDTS as cuts are applied. It seems that as the events approach a more “trident-like” topology, the modeling discrepancies get larger - most likely because these events occur at such high neutrino energies. I scanned through all of the individual KLutz variables to ensure that no single variable was skewing the distribution, and did not find any obvious culprits. Instead, what I found was a consistent mismatch across the full range of all the input variables.

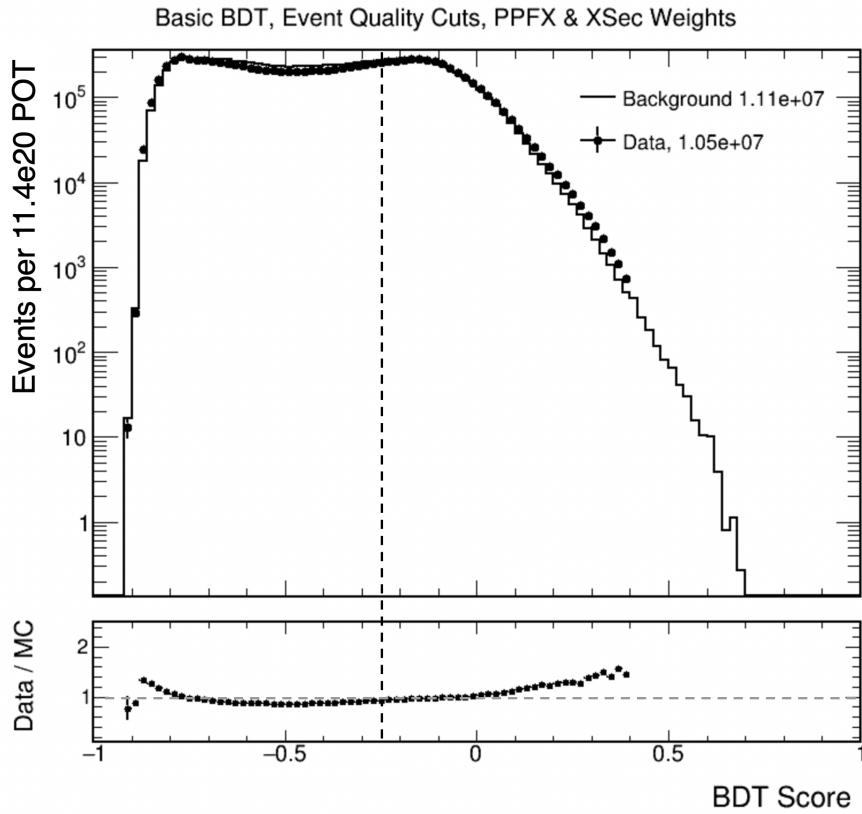


Figure 6.4: Topology-based BDT with the cut location identified, quality cuts applied.

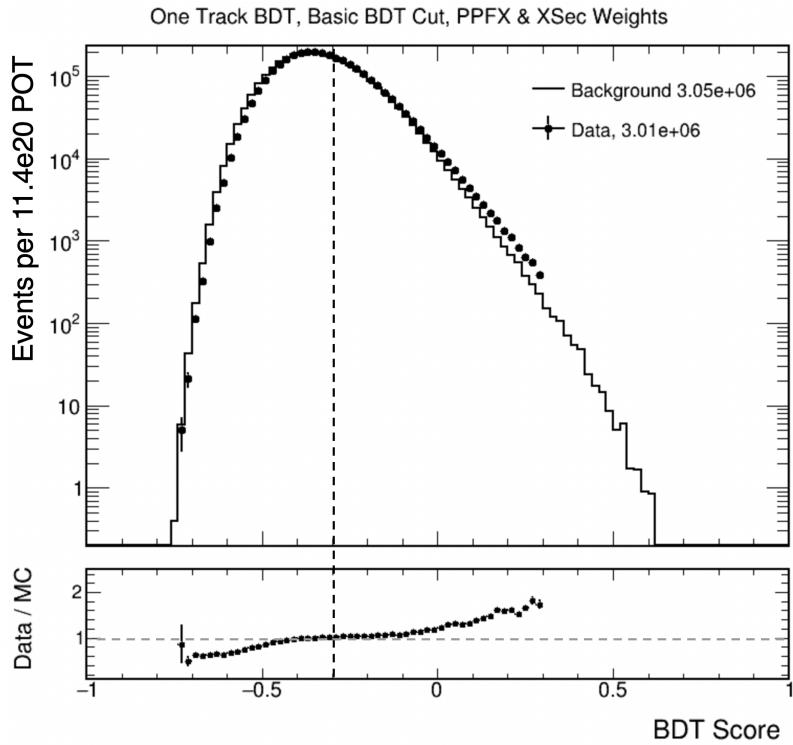


Figure 6.5: One Kalman track-based BDT with the cut location identified, quality and topology-based BDT cut applied.

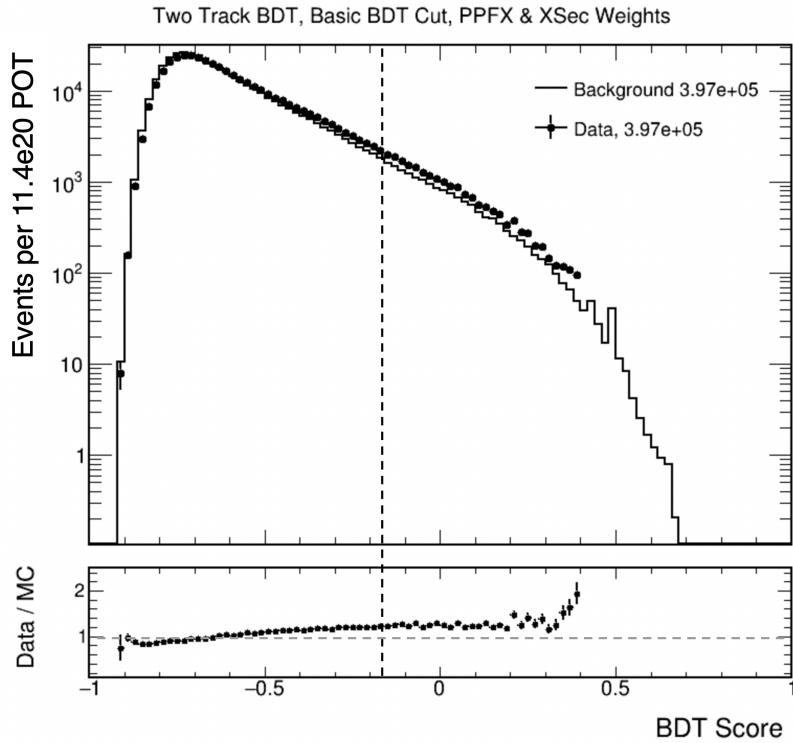


Figure 6.6: Two Kalman track-based BDT with the cut location identified, quality and topology-based BDT cut applied.

Potential Model Corrections

Following the same χ^2 -minimizing best fit procedure as the FHC distributions, a reasonable fit *can* be found, as shown in Figure 6.7. However, the resulting weights are well outside our expectations based on current knowledge of neutrino cross sections. After consulting the collaboration, we decided to set the RHC dataset aside for exploration at another time. At this point we proceeded with unblinding the sideband region of the FHC dataset.

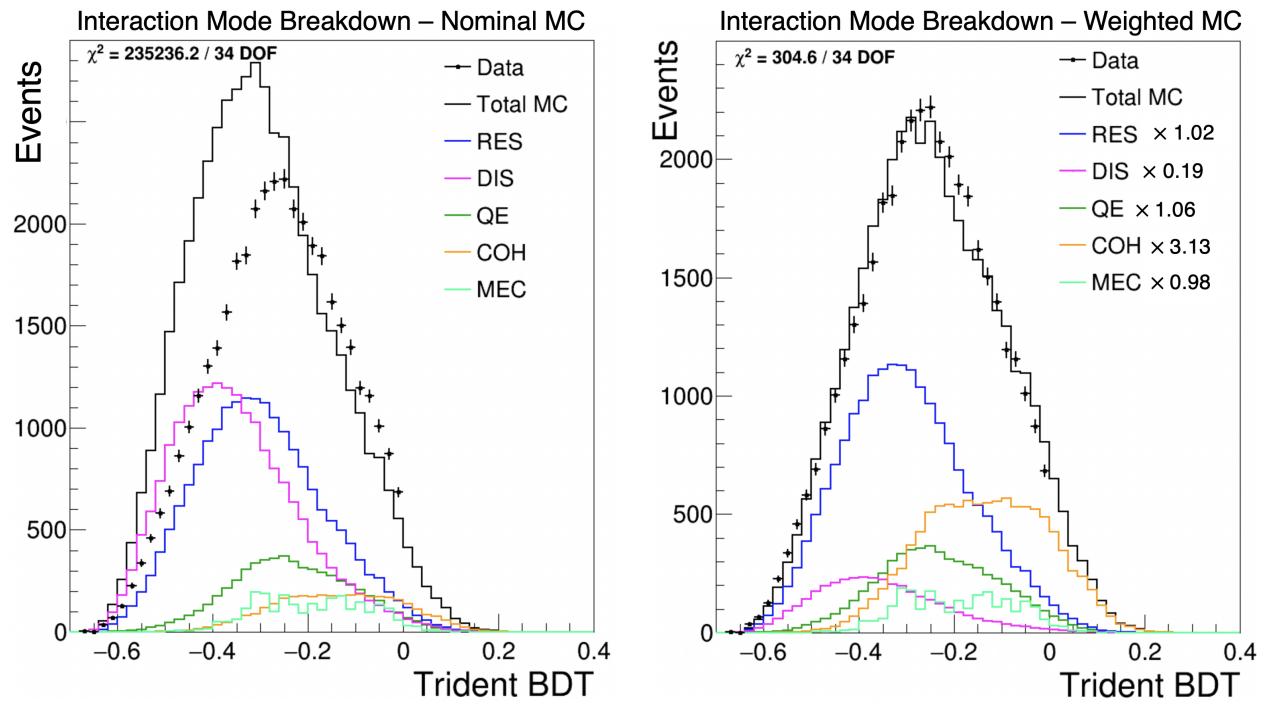


Figure 6.7: Before and after comparisons of RHC nominal MC, weighted MC, and data.

6.2 Unblinding the FHC Sideband Region

The results of unblinding the sideband region of the FHC dataset are shown in Figure 6.8. I used the scaled background interaction components to reweight the MC events in this region. The systematics do still lend a 33.5% uncertainty to the MC, and above a BDT score of about 0.15, statistical fluctuations begin to make the MC unreliable.

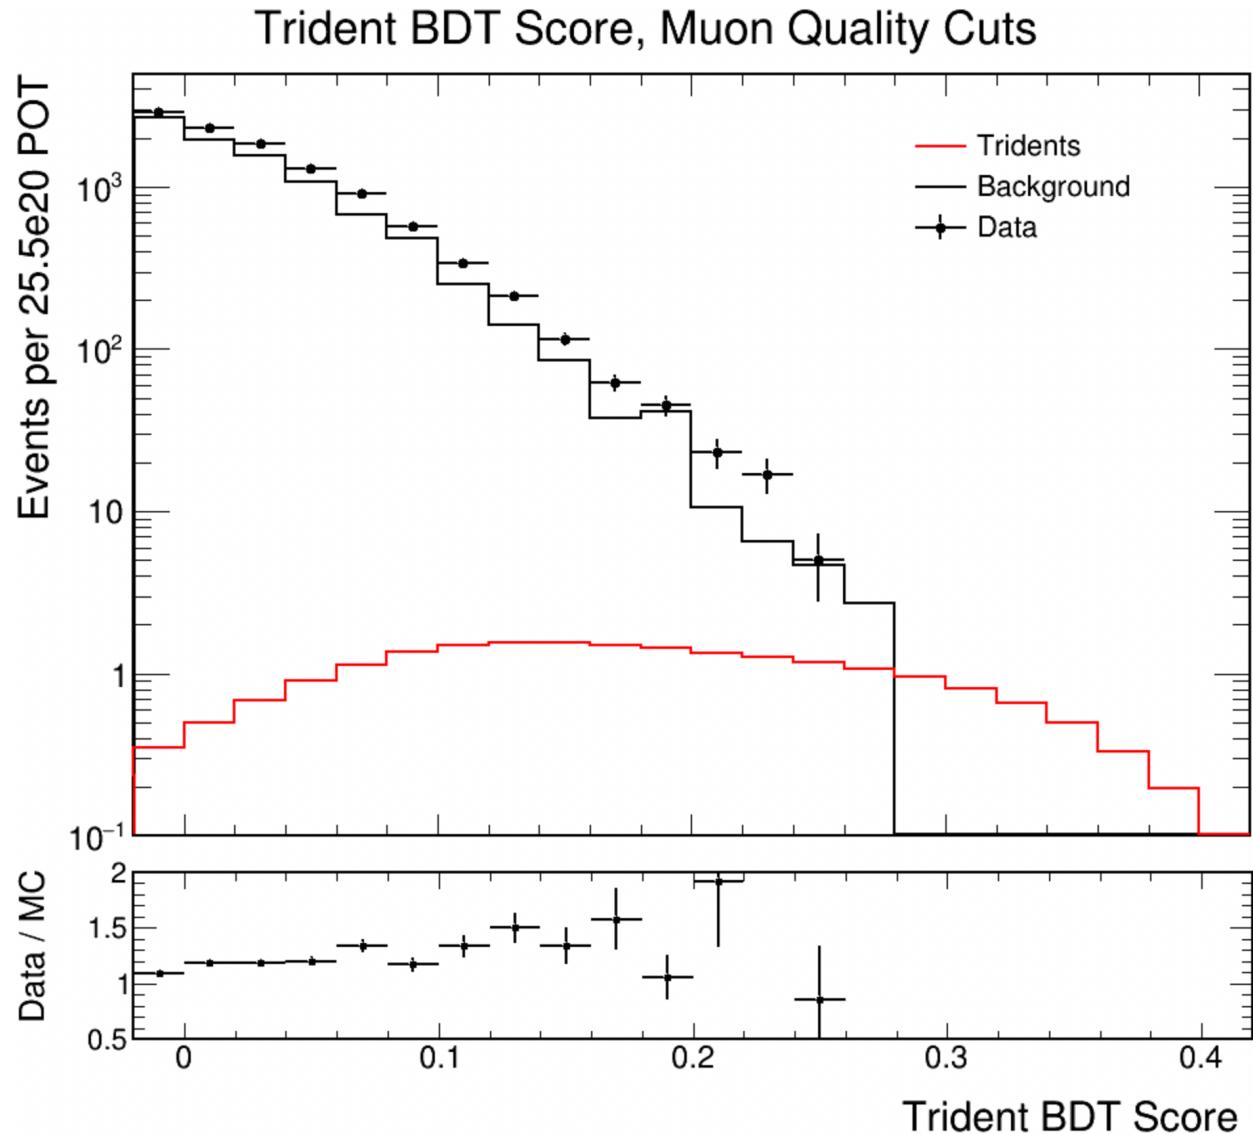


Figure 6.8: KLutz-Based Trident BDT score sideband with FHC data and MC overlayed.

6.2.1 Background Estimation

I applied the method introduced in subsection 5.6.2 to fit exponential and power law functions to the data in the BDT score range of $[0.05, 0.26]$. I used the data fits and the previously discussed MC fits to make my full background estimation envelope.

- MC Exponential: $B_{\text{exp,mc}} \pm \sigma_{\text{exp,mc}} = 5.39 \pm 1.16$ events
- MC Power Law: $B_{\text{pow,mc}} \pm \sigma_{\text{pow,mc}} = 10.22 \pm 1.51$ events
- Data Exponential: $B_{\text{exp,data}} \pm \sigma_{\text{exp,data}} = 6.05 \pm 1.39$ events
- Data Power Law: $B_{\text{pow,data}} \pm \sigma_{\text{pow,data}} = 10.64 \pm 1.68$ events

To properly consider both the MC and data-driven background estimates, we calculate an inverse-variance weighted average [189, 190] using the exponential fits as central values. This method accounts for the uncertainty on each fit by assigning weights determined by $w_i = 1/\sigma_i^2$. This weight properly accounts for how well each fit matches their respective dataset by causing estimates with smaller uncertainty to contribute more heavily to the average. Additionally the fit uncertainties are 100% uncorrelated with one another since they were fit to entirely different datasets, making these two functions great candidates to perform this weighted average with.

The weighted average and combined uncertainty of the two exponential fits are:

$$B_{\text{avg exp}} = \frac{w_{\text{mc}}B_{\text{mc}} + w_{\text{data}}B_{\text{data}}}{w_{\text{mc}} + w_{\text{data}}} = 5.66, \quad \sigma_{\text{avg exp}} = \frac{1}{\sqrt{w_{\text{mc}} + w_{\text{data}}}} = 0.89$$

\implies Average Exponential Estimate: 5.66 ± 0.89 events.

The power law fits are used as a source of relative systematic uncertainty on the average in the same way described in subsection 5.6.2. The larger of the two relative differences between exponential and power law fits is taken to keep the background estimation envelope conservative.

$$\sigma_{\text{pow, mc}} = \frac{|B_{\text{pow, mc}} - B_{\text{exp, mc}}|}{B_{\text{exp, mc}}} = 0.90, \quad \sigma_{\text{pow, data}} = \frac{|B_{\text{pow, data}} - B_{\text{exp, data}}|}{B_{\text{exp, data}}} = 0.76$$

$$\implies \sigma_{\text{abs pow}} = \sigma_{\text{frac pow}} \times B_{\text{avg exp}} = 5.07.$$

Then the overall background estimate in the signal region accounting for fit and shape uncertainty is

$$B_{\text{avg}} = 5.66 \pm \sqrt{0.89^2 + 5.07^2} = 5.66 \pm 5.15 \text{ events.}$$

6.3 Unblinding the FHC Signal Region

With a proper background estimate in hand, we proceeded to unblind the signal region of the FHC dataset, as shown in Figure 6.9. We observe 9 total events in this region, aligning with the estimated background + trident MC predictions of $5.66 + 4.66$. The extrapolated background prediction envelope and detailed fit information are presented in Figure 6.10 and Table 6.1.

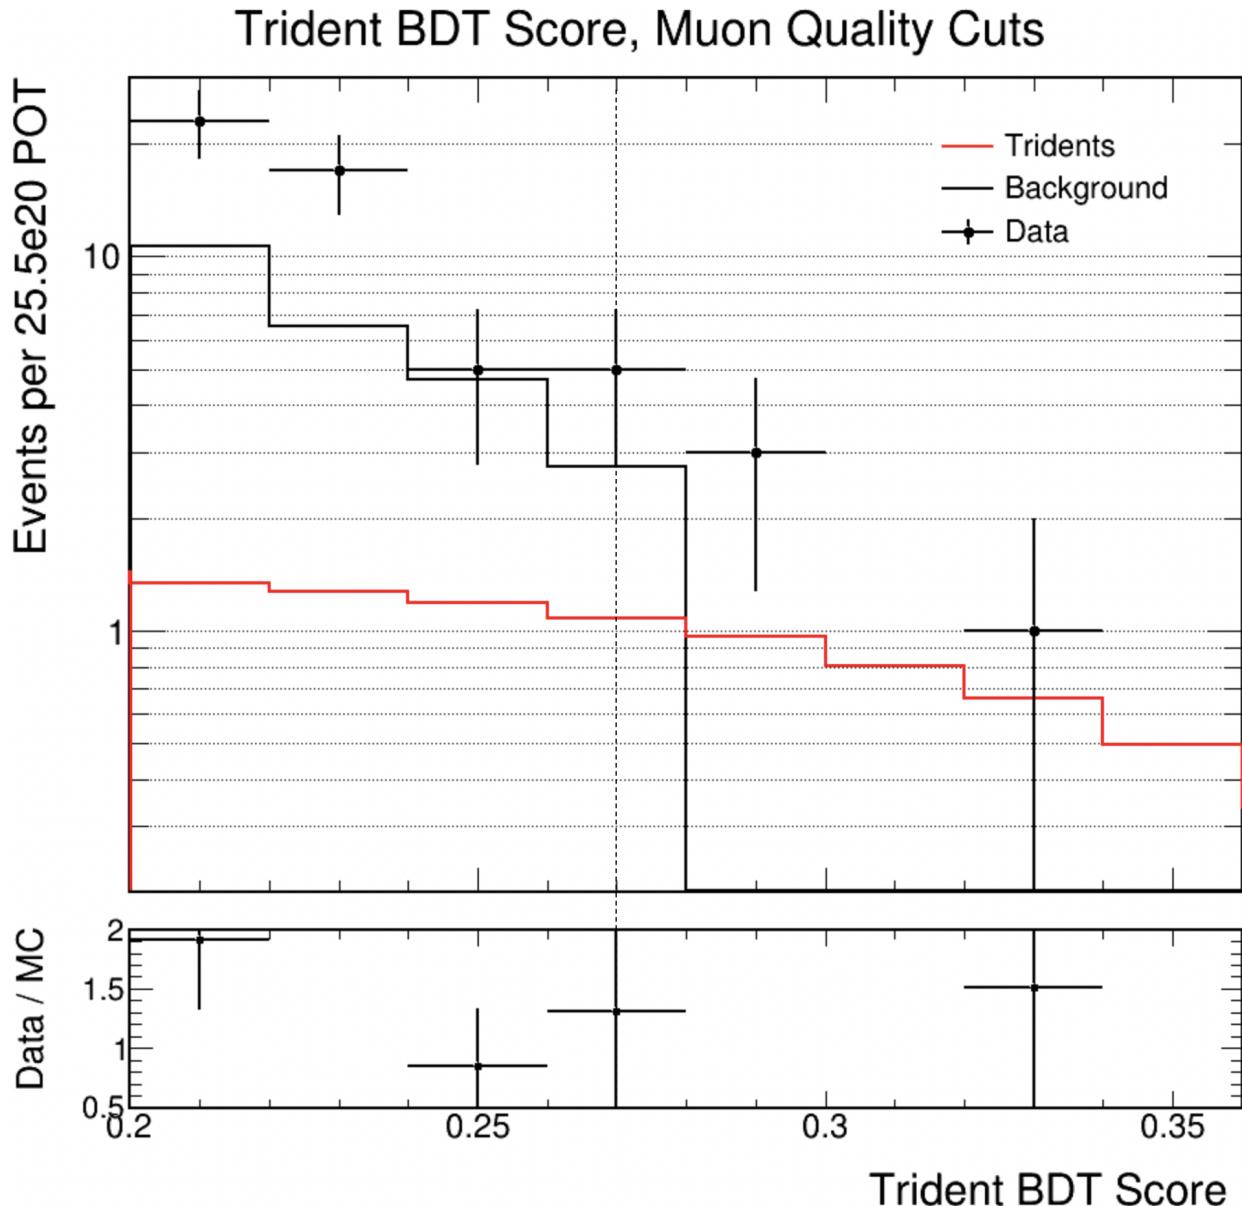


Figure 6.9: KLutz-Based Trident BDT score signal region with FHC data and MC overlayed.

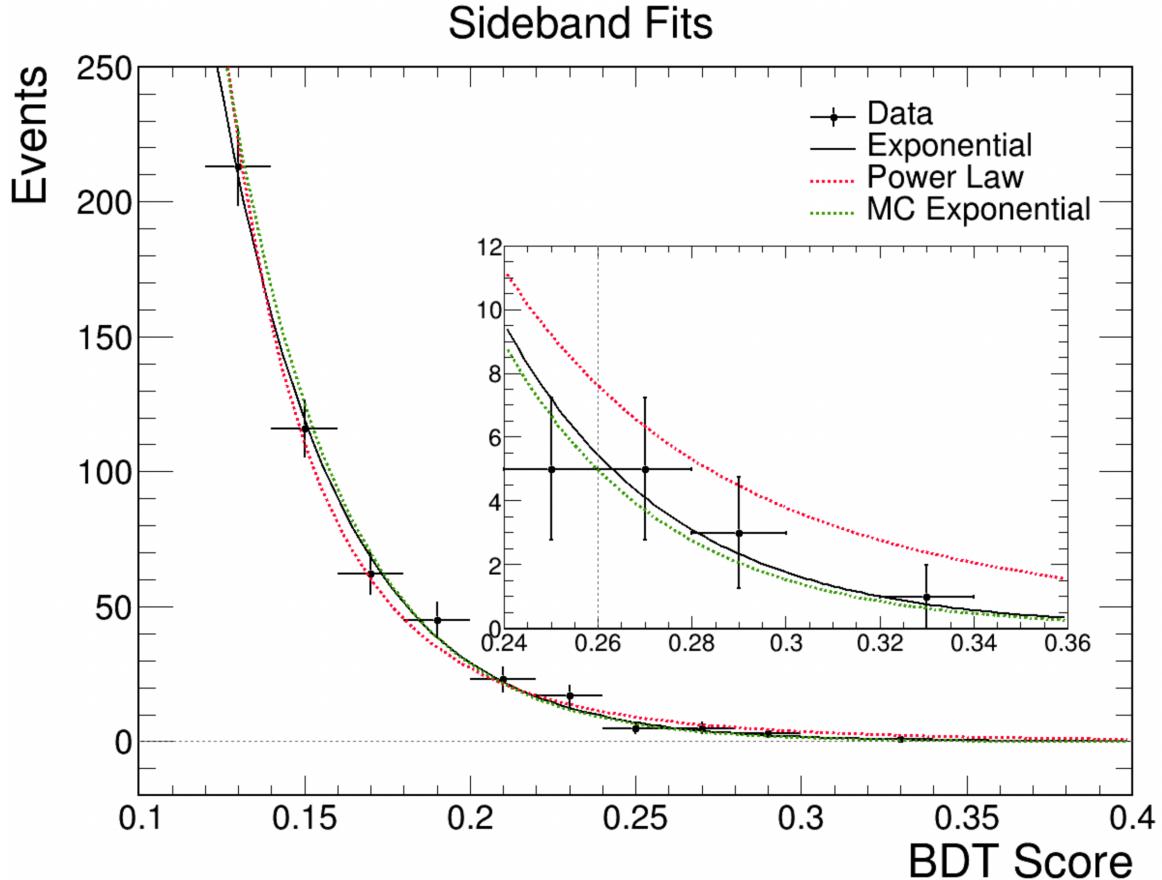


Figure 6.10: Extrapolated function fits used for the background estimation overlayed on data.

Background Estimation Fit Summaries					
Source	Fit Range	Function	χ^2/N_{DOF}	Parameters (A,B)	Estimate
MC	[0.05, 0.27]	$A \exp(-Bx)$	16.3/5	10183.2, 29.3	5.39 ± 1.16
Data	[0.05, 0.26]	$A \exp(-Bx)$	3.7/5	8049.0, 28.1	6.05 ± 1.39
MC	[0.05, 0.27]	$A x^{-B}$	22.9/6	0.0083, 5.03	10.22 ± 1.51
Data	[0.05, 0.26]	$A x^{-B}$	7.5/6	00108, 4.87	10.64 ± 1.68

Table 6.1: Details of the fits applied to the MC and real datasets to form the background estimates.

6.3.1 Trident Candidate Properties and Event Displays

This section shows the properties and event displays of the 9 candidate events. In the plots, black data points show data from the sideband region $0.0 \leq \text{BDT score} \leq 0.26$, and red data points show data from the signal region, $\text{BDT score} \geq 0.27$. This visual approach is due to the limited MC statistics in the signal region yielding poor representations of the background distributions, as seen at the bottom of Figure 6.11.

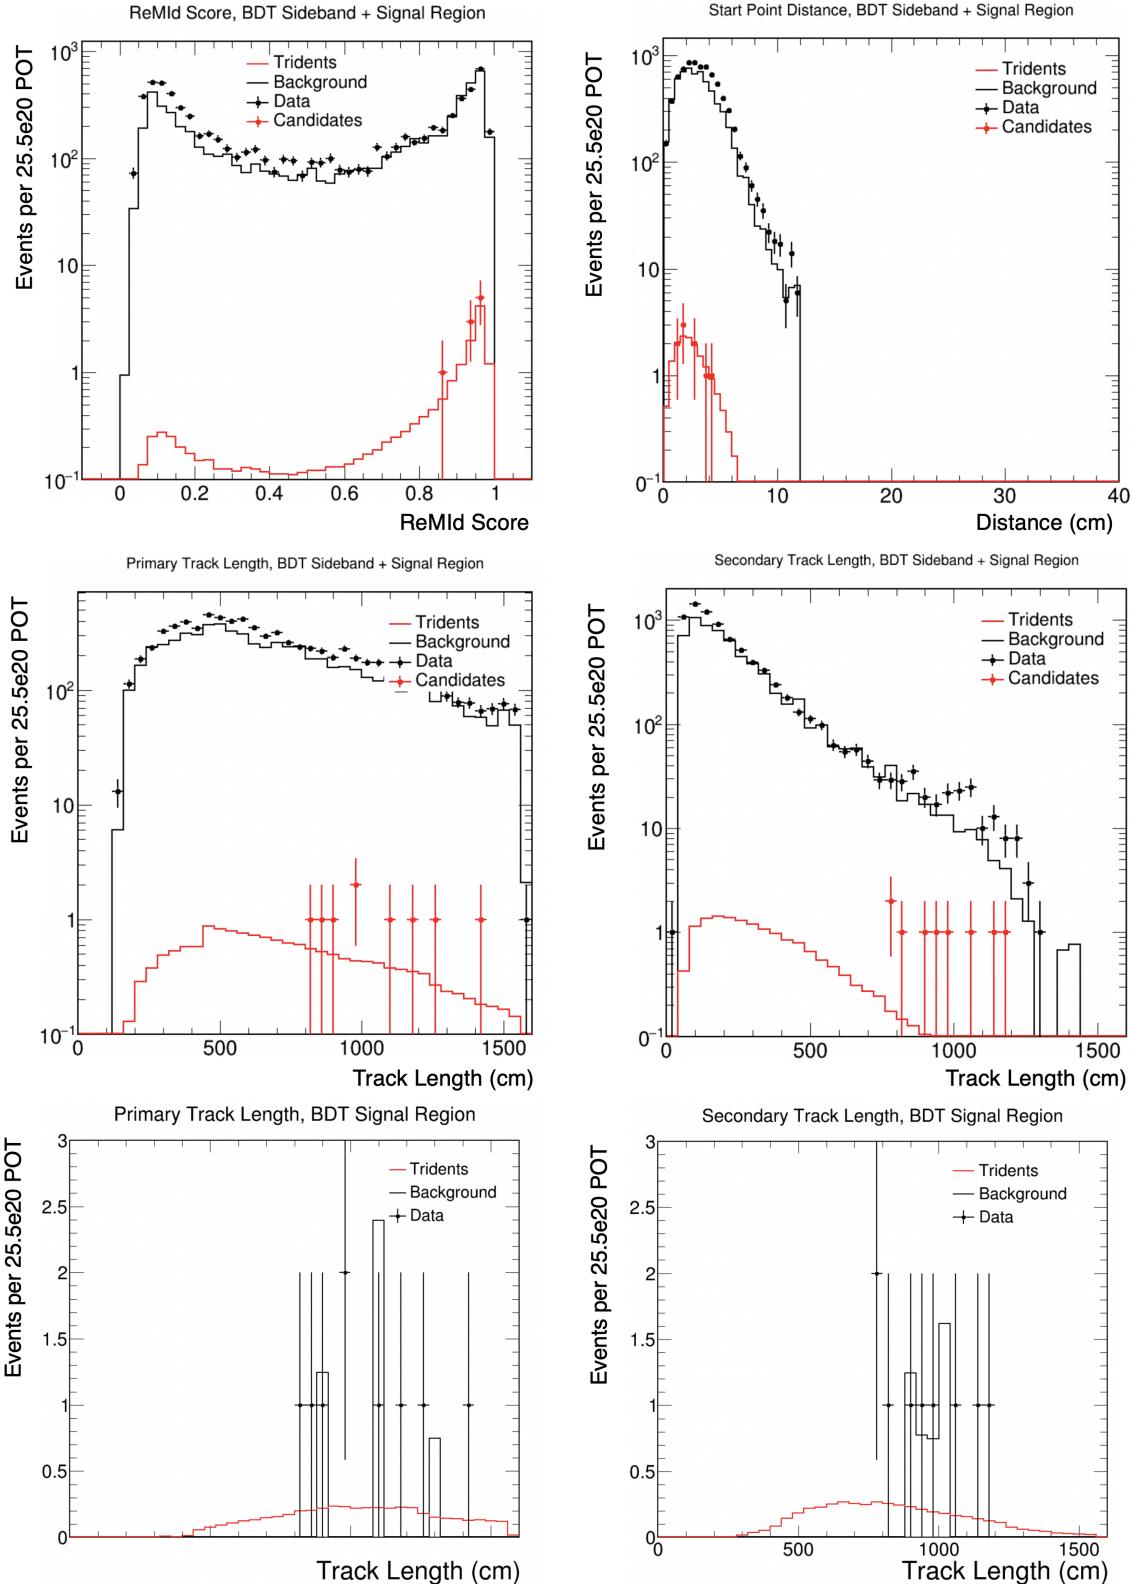


Figure 6.11: Distributions of candidate event properties (red data points). The track lengths appear longer than predicted due to inclusion of sideband events. Signal region distributions included in the bottom two plots for a more accurate representation.

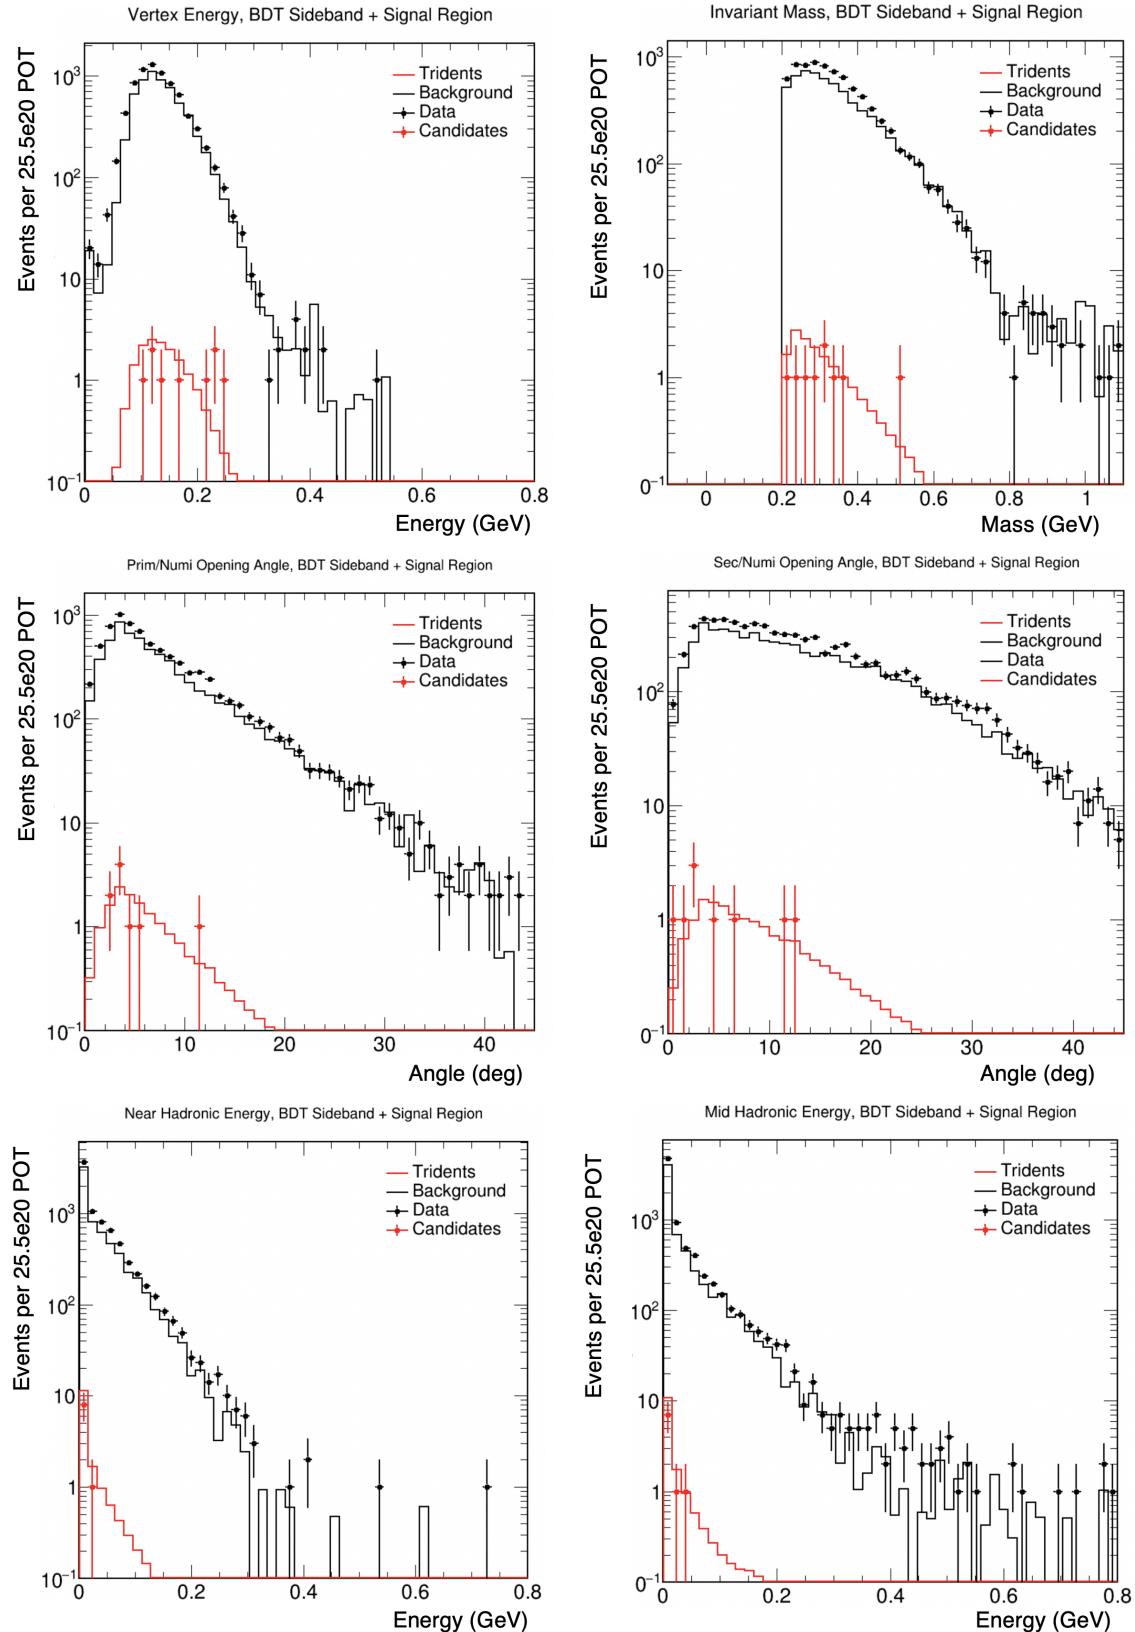


Figure 6.12: Distributions of candidate event properties (red data points).

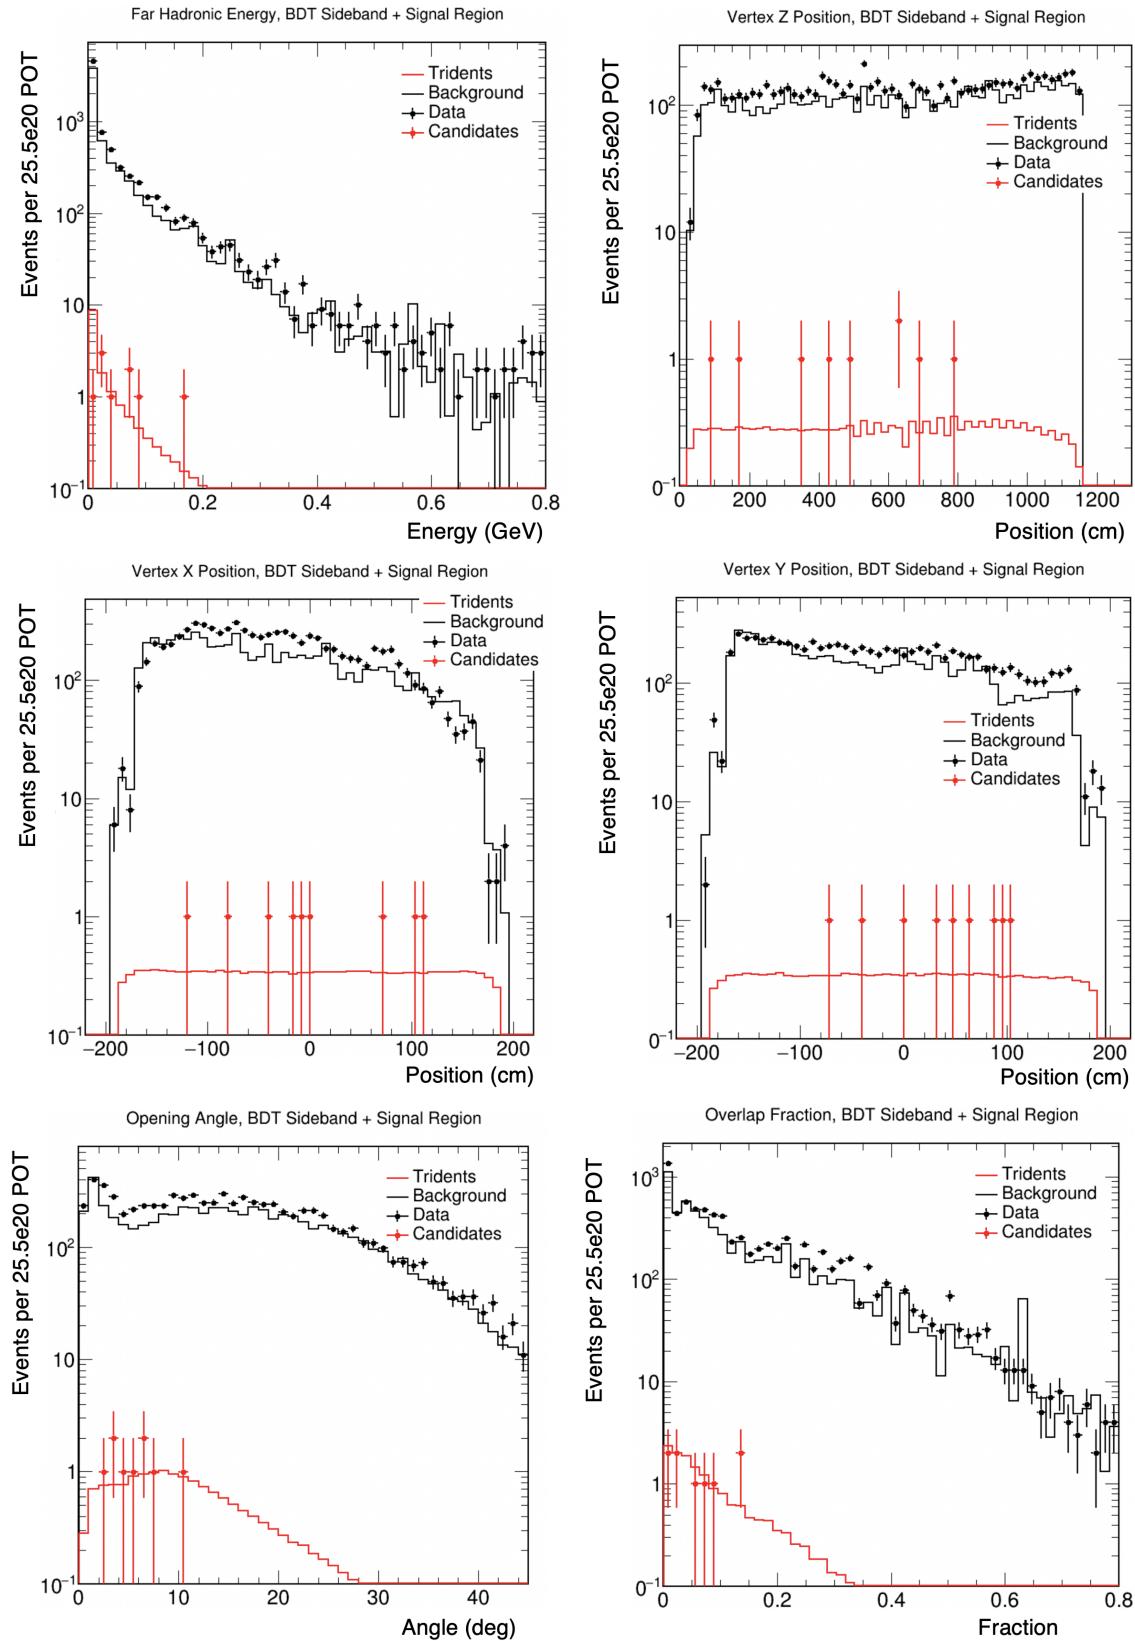


Figure 6.13: Distributions of candidate event properties (red data points).

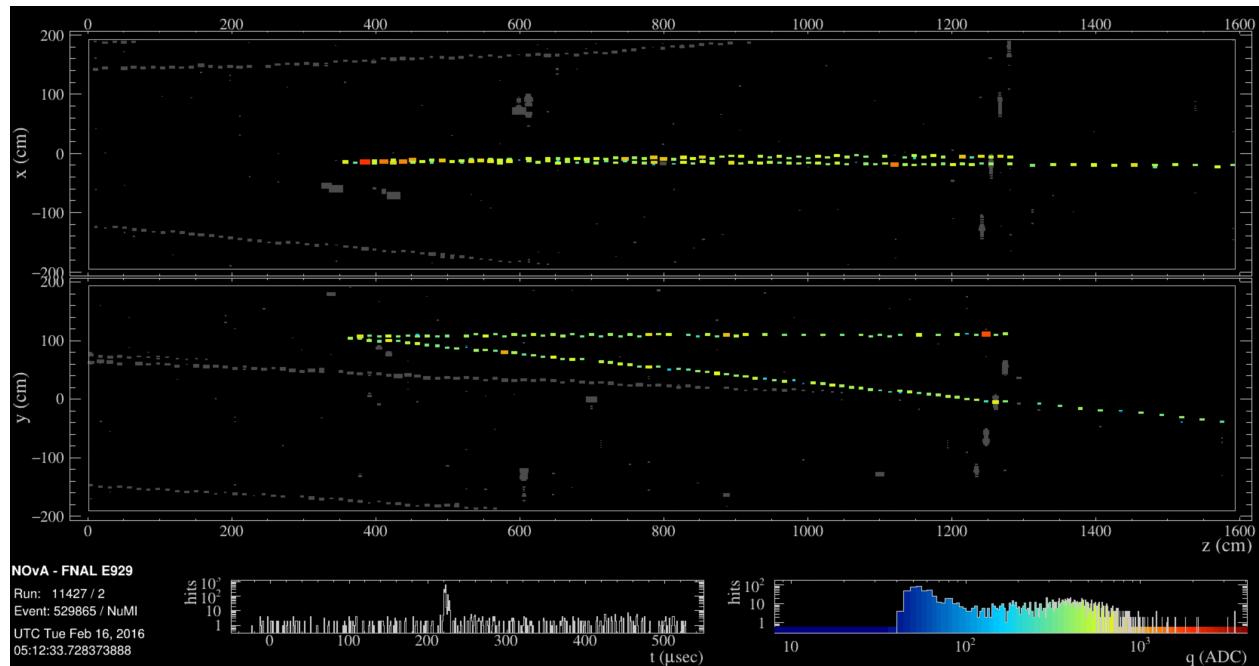


Figure 6.14: An event display featuring a candidate trident event with BDT Score 0.280.

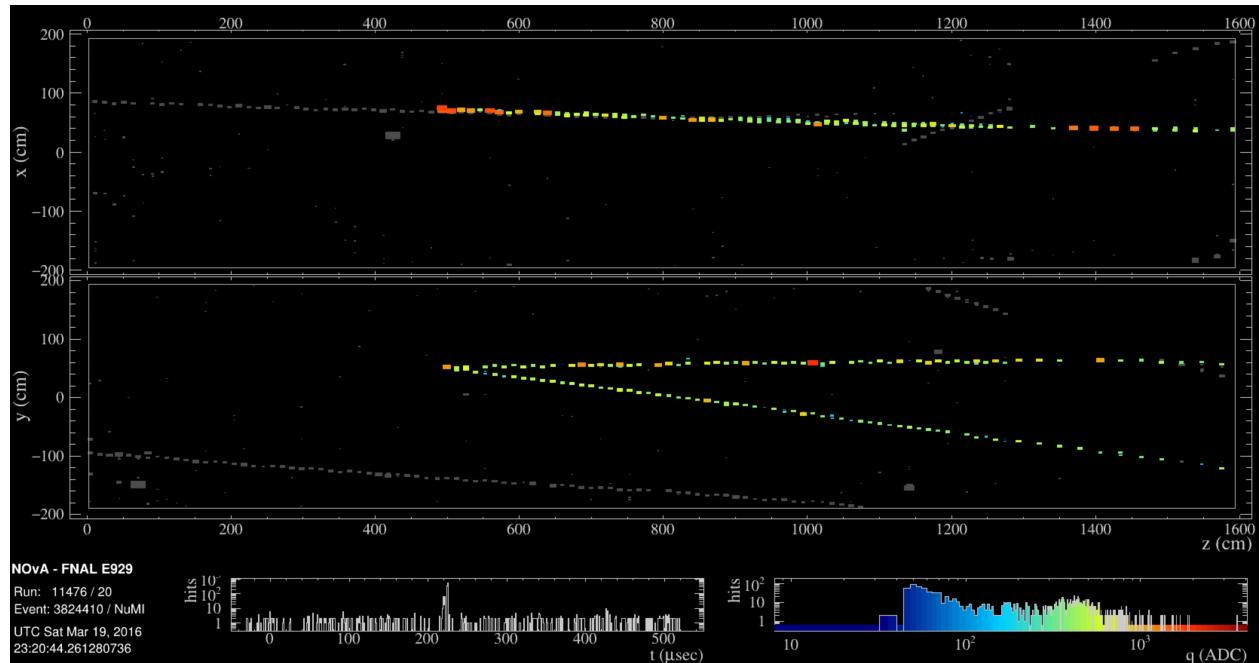


Figure 6.15: An event display featuring a candidate trident event with BDT score 0.265.

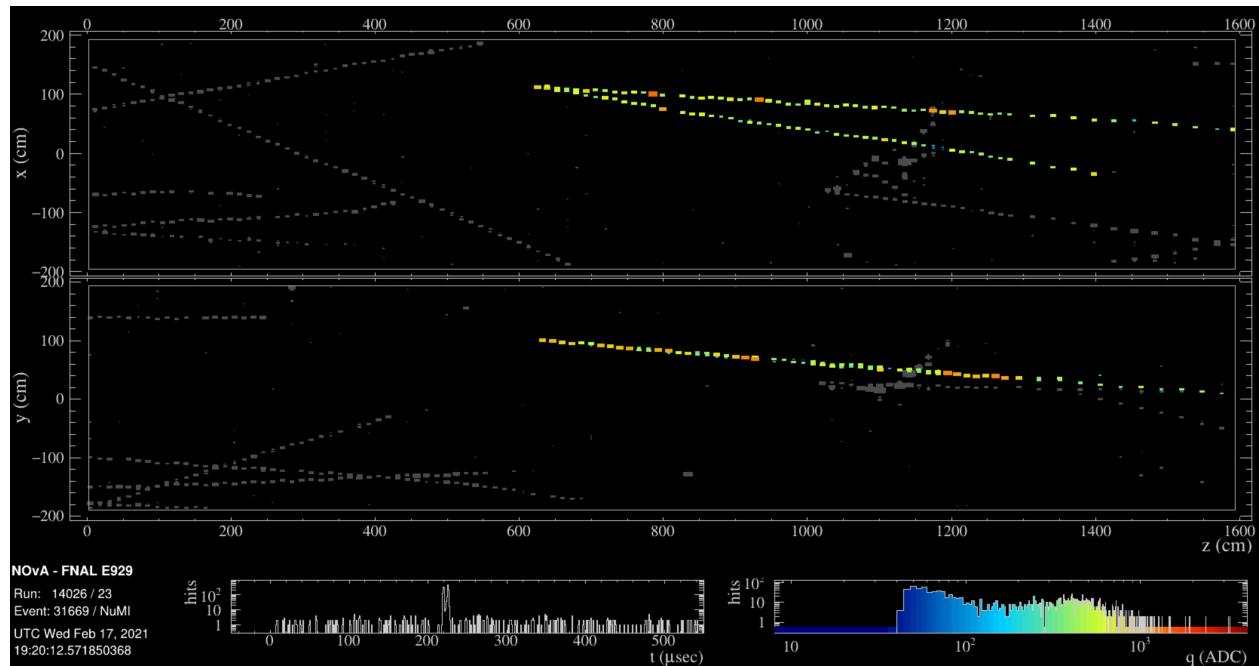


Figure 6.16: An event display featuring a candidate trident event with BDT score 0.335.

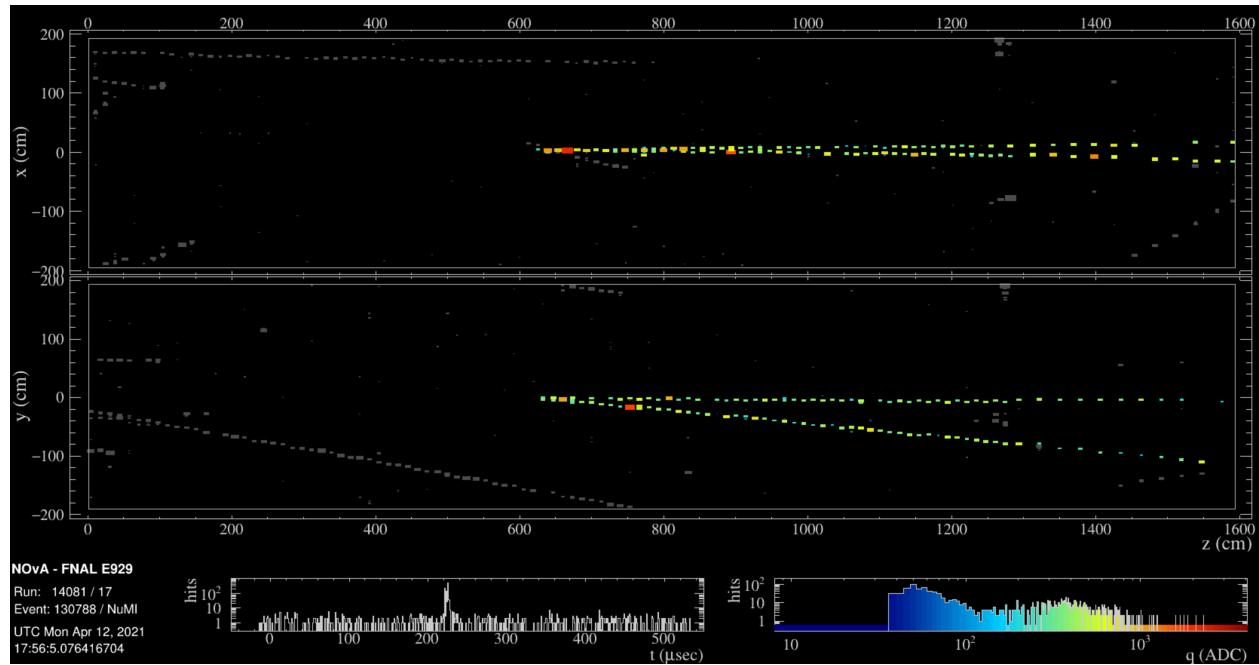


Figure 6.17: An event display featuring a candidate trident event with BDT score 0.265.

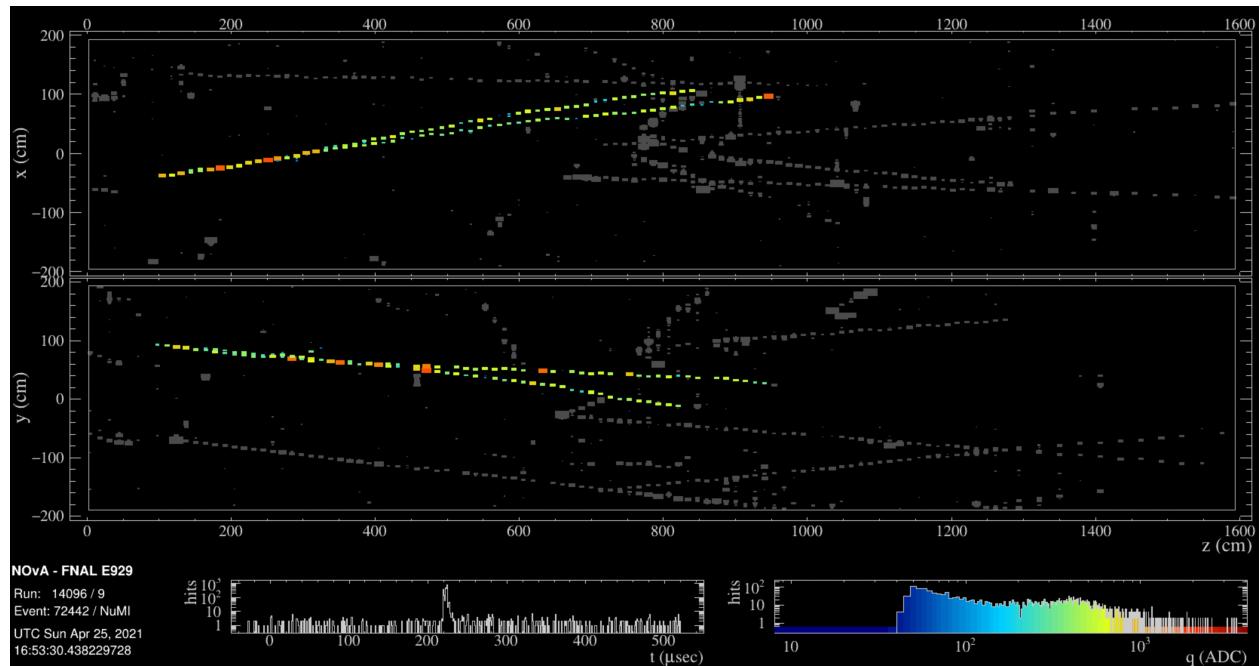


Figure 6.18: An event display featuring a candidate trident event with BDT score 0.265.

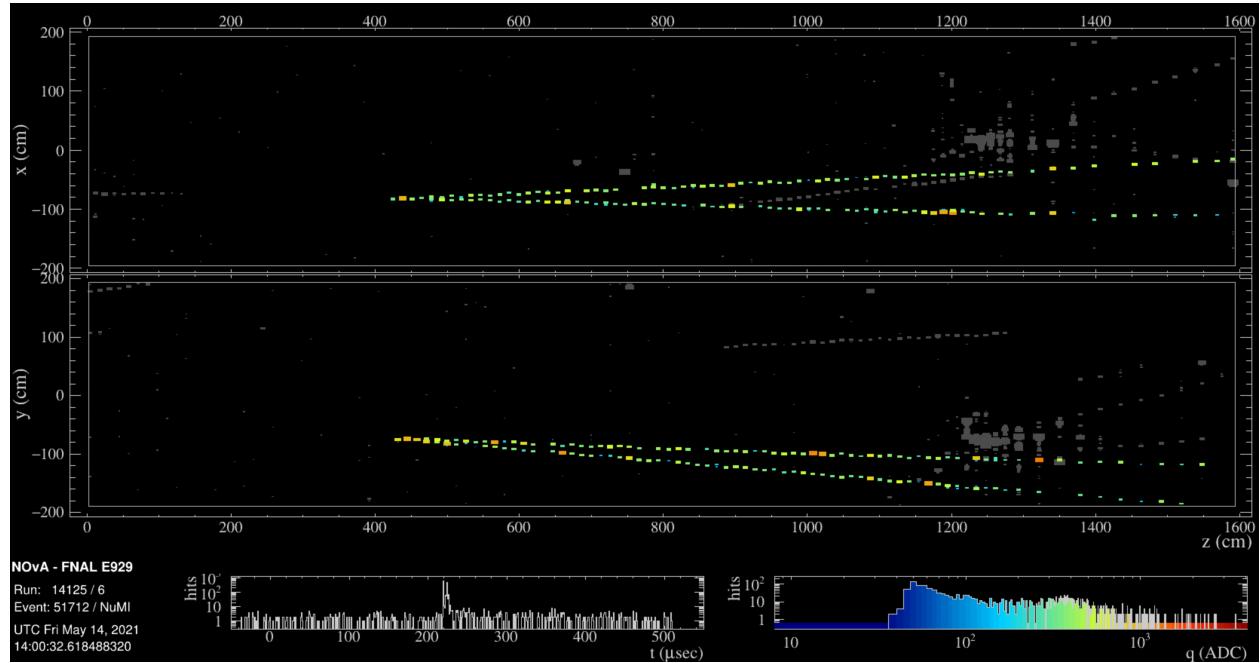


Figure 6.19: An event display featuring a candidate trident event with BDT score 0.295.

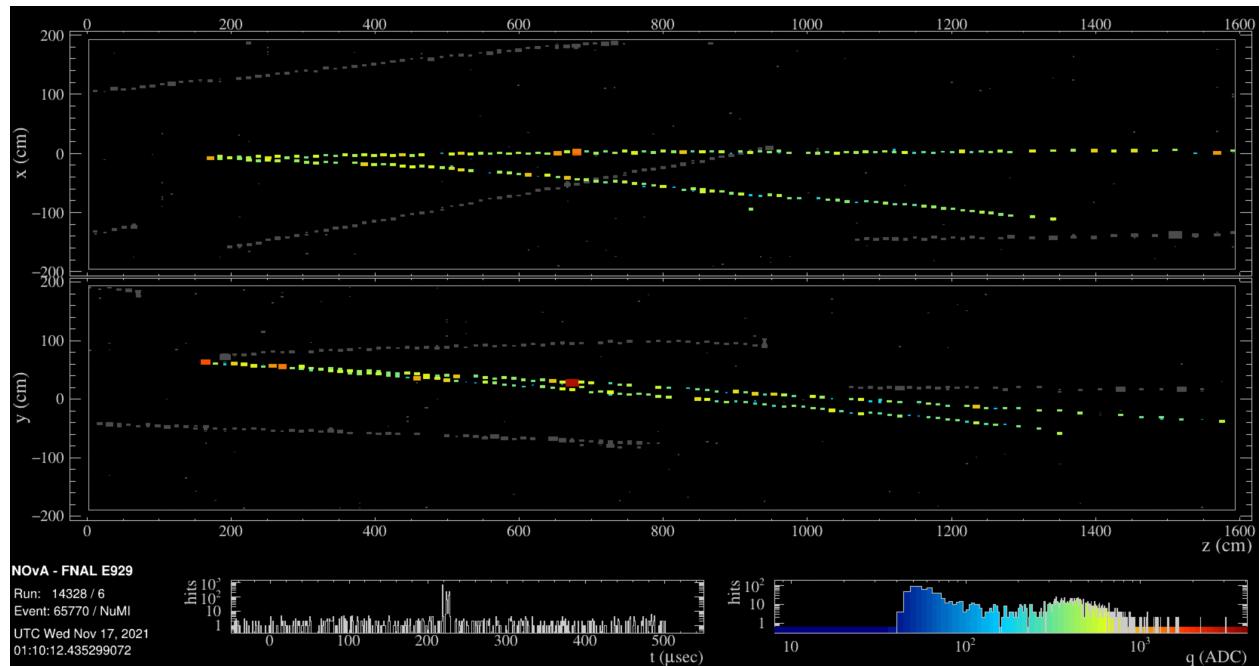


Figure 6.20: An event display featuring a candidate trident event with BDT score 0.290.

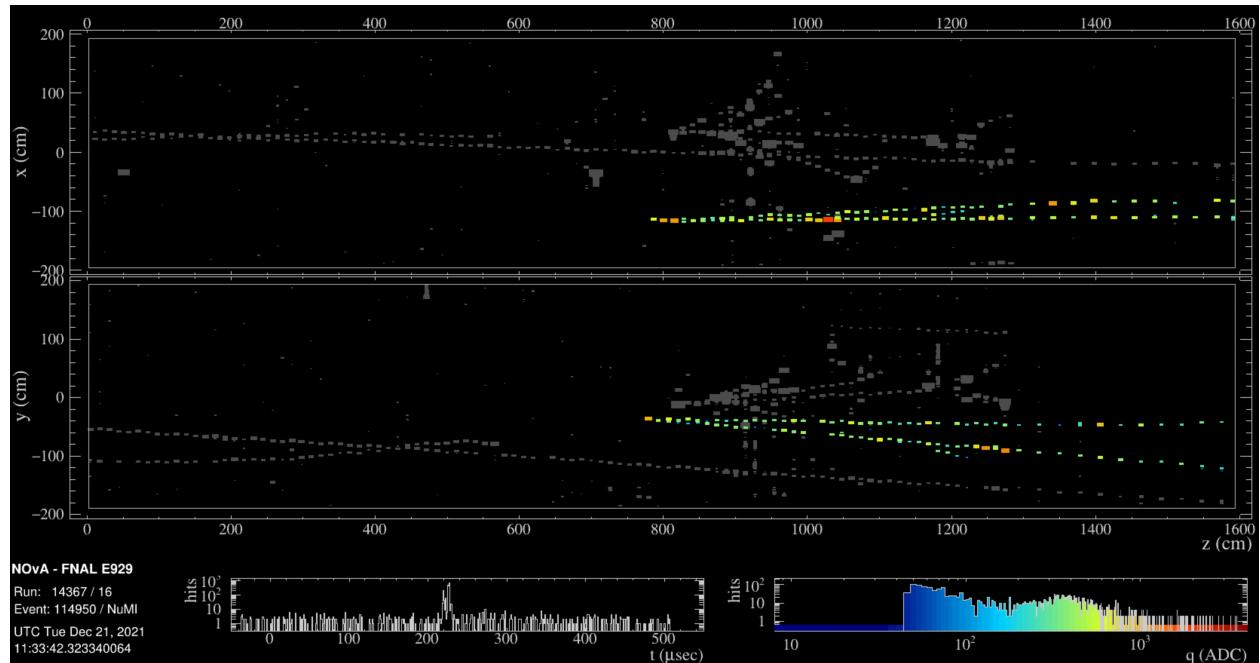


Figure 6.21: An event display featuring a candidate trident event with BDT score 0.290.

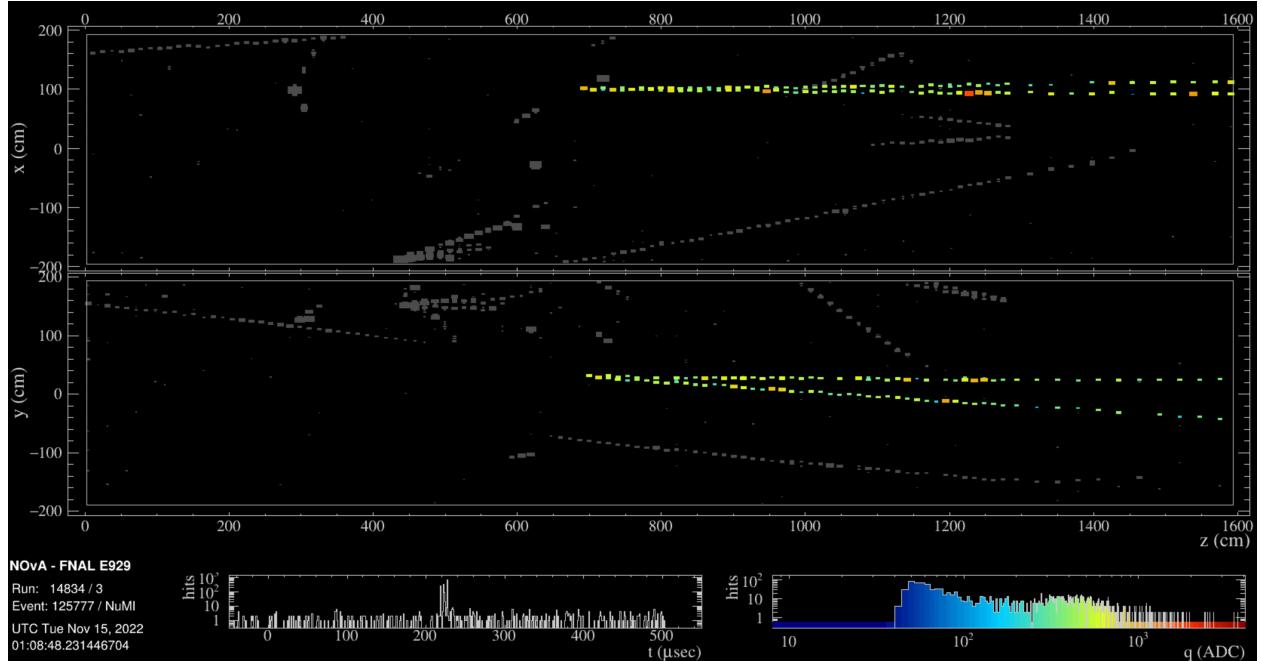


Figure 6.22: An event display featuring a candidate trident event with BDT score 0.265.

Comments on Trident Candidates

Most of these events are very reasonable trident candidates. A handful of events have vertex energy levels consistent with coherent background events, as those backgrounds tend to have slightly more hadronic energy than tridents.

The event captured in run 14096 is the only one I suspect is *not* a neutrino trident. Both tracks in this event are steeply angled with respect to the incoming NuMI beam, indicating that a neutron likely carried energy off in the opposite direction without interacting.

Additionally, a few candidate events highlight the importance of developing the KLutz dimuon tracker. Kalman tracks were poorly reconstructed in the majority of events where the muons are very close to one another or overlap in a single view. This poor reconstruction would likely have caused the candidate events to be removed from a more rigorous Kalman-based signal selection.

6.3.2 Determining Statistical Significance

With an estimate of the expected number of background events in the signal region, the next step is to interpret the number of observed events in a statistically rigorous way. Doing so requires a method of inference that accounts for both the statistical fluctuations in expecting such a low signal event count, and systematic uncertainties associated with the background prediction. I chose to follow the Feldman–Cousins method of inference, which allowed me to construct confidence intervals in the low-statistics regime.

The Standard Feldman-Cousins Method

The original Feldman-Cousins (FC) method [191] was developed for calculating confidence intervals around Poisson-distributed signal events with known background and observed event rates. I developed a standard FC calculator in ROOT which was adapted from Evan Groopman’s “FCpy” calculator [192]. I will first describe the steps involved in implementing the original FC method, then discuss changes required to make the method appropriate with unknown background events.

The standard method begins with an observation of n events ($n = 9$ for this analysis) consisting of Poisson distributed signal events with mean μ and background events with a well known mean b . The probability mass function for such an observation is

$$P(n|\mu) = \frac{(\mu + b)^n}{n!} \exp(-(\mu + b)). \quad (6.1)$$

In practice we use the Log-Likelihood for computational convenience:

$$\mathcal{L}(n|\mu) = \log(P(n|\mu)) = n \log(\mu + b) - (\mu + b) - \log(n!). \quad (6.2)$$

Since we do not know μ *a priori*, the FC method simulates observations made over many test values, μ_i , and determines a statistically valid range of μ_i that explain the data. To find the valid μ_i , we consider a “test statistic” λ as defined in the Neyman construction [193]. A test statistic is a single number quantifying how well a certain μ_i matches the observed n compared to the “best fit” μ . Here the best fit μ is defined as the number of signal expected if all values adhere to their means: $\mu_{\text{best}} = n - b$.

Then the test statistic comparing any μ_i to μ_{best} is defined as:

$$\lambda_i = -2 \ln \frac{\mathcal{L}(n|\mu_i)}{\mathcal{L}(n|\mu_{\text{best}})}. \quad (6.3)$$

In a perfect world, each Log-Likelihood would be 1, meaning we would know with 100% certainty that n observed events corresponds to μ_i signal events, yielding $\lambda = 0$. Instead, λ is non-zero, and smaller values correspond to a μ_i being more consistent with the data [194].

At this point we introduce the idea of a “critical” test statistic, λ_{crit} . These are the values actually used to define the extent of the confidence interval. For some confidence level (for example, $\text{CL} = 90\% = 0.9$) we determine the critical values using

$$\int_0^{\lambda_{\text{crit}}} P(\lambda_i) d\lambda_i = \alpha = 1 - \text{CL}. \quad (6.4)$$

To determine λ_{crit} computationally, we run Monte Carlo “psuedo-experiments” (PSE’s) that generate a probability distribution of test statistics $P(\lambda)$ that differs for each tested value of μ_i . In effect, the PSEs calculate a large number of different λ ’s, and the critical values are the largest values that show up fewer than α times. For example at a 90% CL, the largest 10% of MC generated λ values determine the critical value threshold.

Different values of μ_i are tested to find a region where any λ_i ’s generated by the test statistic fall below λ_{crit} . Any μ_i ’s that pass the test of $\lambda_i < \lambda_{\text{crit}}$ are added into the confidence interval, and μ_i ’s that do not pass the test are not included. I calculated a standard FC 90% confidence interval using values of $n = 9$, $b = 5.66$ in my calculator. I ran one million PSE’s on μ_i values ranging from 0 to 12 in increments of 0.1, which yielded a confidence interval of [0.6, 7.5] trident events.

The Profiled Feldman-Cousins Method

The original FC method does not include a means to handle “nuisance parameters” - variables that are not of interest to our final result, but are still subject to fluctuation. In this analysis, the small background MC sample requires an estimate on the number of background events b with large uncertainty. Therefore, b is an unknown nuisance parameter rather than a known quantity, as is assumed in the standard FC method. Implementing nuisance parameters is the basis of the

“profiled” FC method used to interpret one of NOvA’s recent results [195, 196].

To cover the phase-space of possible outcomes due to background fluctuation, I implemented an additional Log-Likelihood calculation to the original FC calculator. I sampled the expected backgrounds from a random Gaussian distribution with a mean of 5.66 and a standard deviation of 5.15, skipping contributions from any background counts less than 0. Utilizing this method, the 90% confidence interval broadened to $[0.0, 9.1]$. Figure 6.23 shows the full “confidence belt” calculated using the FC method. The belt was constructed by considering possible values of observed events n , and for each tested μ determining the range $[n_1, n_2]$ such that $P(n \in [n_1, n_2]) = \alpha$. The red dashed line was placed at the observed number of events ($n = 9$). The resulting confidence interval $[\mu_1, \mu_2] = [0.0, 9.1]$ is the union of all values of μ for which the corresponding acceptance interval is intersected by the vertical line in Figure 6.23.

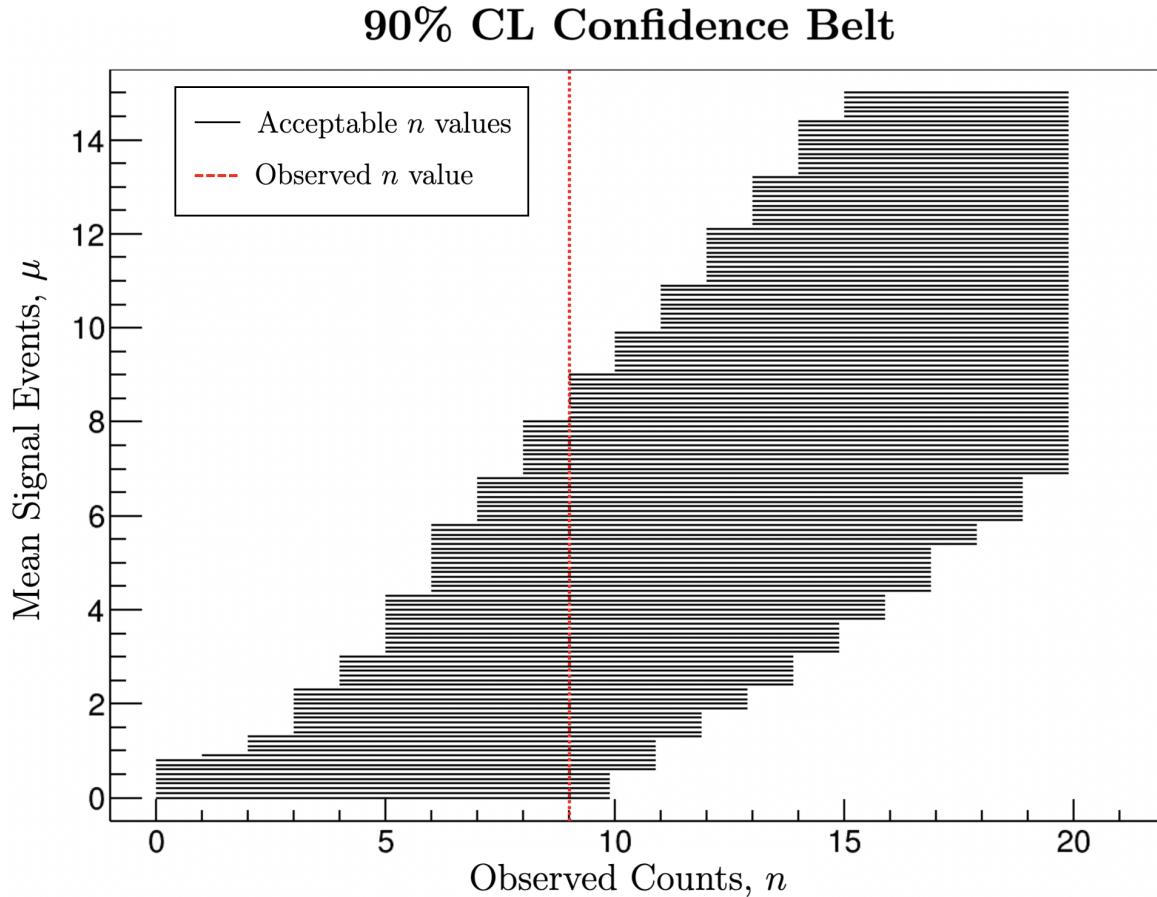


Figure 6.23: A Feldman-Cousins confidence belt representing observed events compared to mean signal event count.

The red dashed line intersects the range of acceptable n values for $\mu = 0$, meaning that the number of events observed does not provide conclusive evidence at 90% CL of trident events in our selected sample. Had the background statistics been plentiful enough as to not require the profiled FC method, our 90% CL confidence belt would have ruled out $\mu = 0$, allowing us to make a more definitive statement of trident observation.

6.4 Results

With 9 events observed and a central value background estimate of 5.66 events, we found the best fit signal value of 3.34 trident events to be consistent with the Standard Model prediction of 4.66. Within errors the measurement was consistent with zero signal events, and limited to below 9.1 signal events at the 90% confidence level. This limited the occurrence of trident events to $1.95 \times$ the Standard Model prediction. Treating the best fit signal value of 3.34 as the mean with fluctuation ranging from zero to 9.1 events, the ratio of events measured compared to the SM is

$$\frac{\sigma(\nu_\mu \rightarrow \nu_\mu \mu^+ \mu^-)_{\text{exp}}}{\sigma(\nu_\mu \rightarrow \nu_\mu \mu^+ \mu^-)_{\text{SM}}} = 0.72^{+1.23}_{-0.72}, \quad \langle E_\nu \rangle \approx 18 \text{ GeV}$$

The number of signal events having the potential to fluctuate to zero resulted in asymmetric interval bounds, representing an upper limit at 90% CL. Comparing this to the next-lowest energy measurement made by CHARM-II and overlaying the data on a plot of BSM models (originally created by Kaneta and Shimomura [197]) shows the power inherent in studying trident production with a relatively low-energy beam, shown in Figure 6.24.

The NuTeV and CCFR measurements used neutrinos with average energies of more than 100 GeV, which fall outside the bounds of the plot, but the BSM models do fall entirely within the range of their measurements. The measurement made by CHARM II has lower energy, with an average neutrino energy of 20 GeV which provides increased sensitivity to BSM physics. However, NOvA's low energy beam results in my analysis being the lowest-energy investigation of trident production to date, with an average neutrino energy of 18 GeV. This low energy consequently makes this limit the most sensitive probe of BSM physics so far, allowing us to disfavor the light Z' model exhibiting strong gauge coupling even more strongly than the CHARM II measurement.

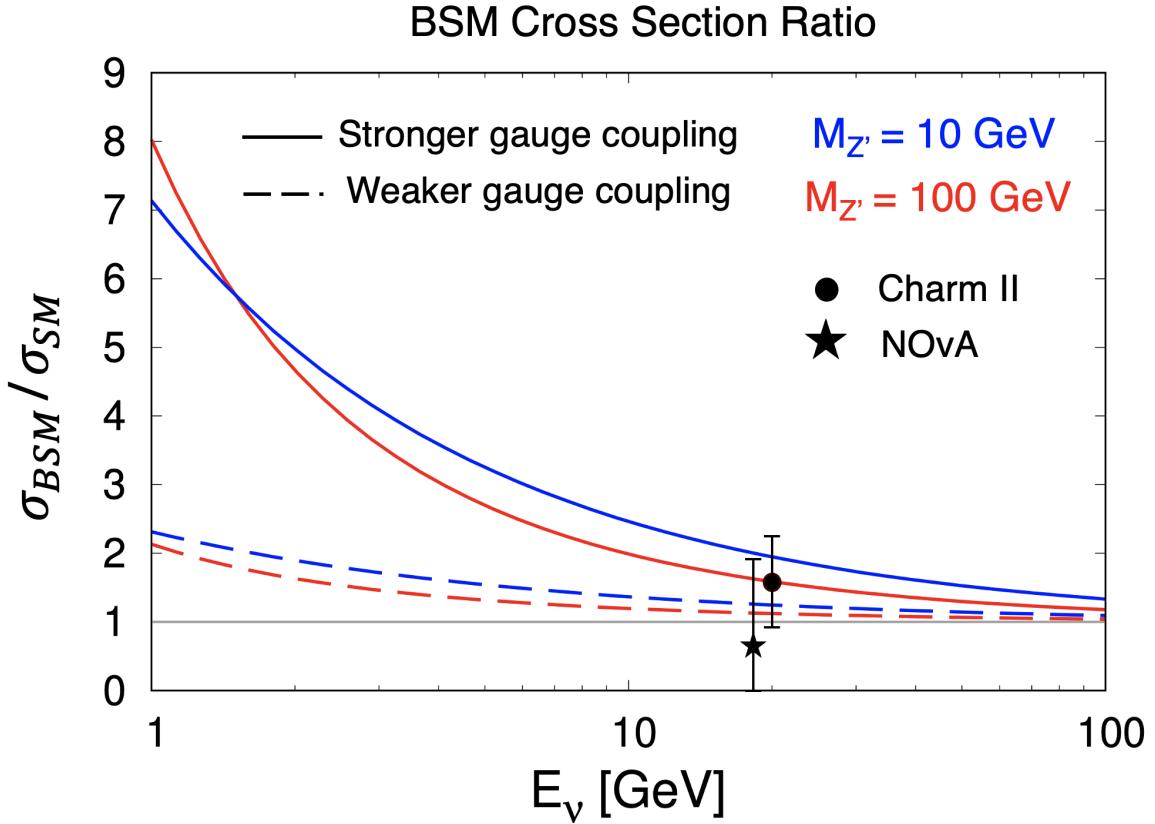


Figure 6.24: Some BSM models [197] overlayed with experimental results from the NOvA and Charm II [133] experiments.

Future Prospects

This upper limit on trident production lays a solid foundation for future observations and a potential cross section measurement of trident production in NOvA. The analysis could be improved upon in a handful of ways. The biggest improvement would come from the inclusion of the RHC dataset, which means understanding and eliminating the discrepancy between data and MC. Possible solutions this issue may be to train an RHC-specific BDT to be used in the data reduction step. The current BDT only used FHC datasets to form cut criteria, even when those cuts are applied to the RHC dataset. Inclusion of the RHC dataset would improve the available statistics by an additional 50%. Generating more background statistics would help reduce the systematic uncertainty associated with trident-like topologies significantly, and result in a measurement at a higher confidence level.

Chapter 7

Conclusion

This analysis shows sensitivity to measuring the trident production process in NOvA and places an upper limit on trident production. In this thesis I have presented a number of novel analysis techniques, including the construction and implementation of a custom event tracker and selection strategy. My result is the first progress in the study of trident production in nearly 30 years, and the energy range of trident candidate events resulting from the analysis is lower than any previous measurement, providing insight into potential physics beyond the Standard Model. I found no significant disagreements between the event rates determined experimentally and those predicted by the Standard Model. This upper limit fits in with previous trident production measurements well, with a mean value matching that observed by NuTeV, but with a lower average neutrino energy and better precision.

$$\frac{\sigma(\nu_\mu \rightarrow \nu_\mu \mu^+ \mu^-)_{\text{exp}}}{\sigma(\nu_\mu \rightarrow \nu_\mu \mu^+ \mu^-)_{\text{SM}}} = \begin{cases} 1.58 \pm 0.64, & (\text{CHARM-II}), \quad \langle E_\nu \rangle \approx 20 \text{ GeV} \\ 0.82 \pm 0.28, & (\text{CCFR}), \quad \langle E_\nu \rangle \approx 160 \text{ GeV} \\ 0.72^{+1.73}_{-0.72}, & (\text{NuTeV}), \quad \langle E_\nu \rangle \approx 50 - 300 \text{ GeV} \\ 0.72^{+1.23}_{-0.72}, & (\text{NOvA}), \quad \langle E_\nu \rangle \approx 18 \text{ GeV} \end{cases}$$

While the result presented here does not constitute a definitive measurement of the trident production cross section, it does set an upper limit at world-leading neutrino energies. I hope this analysis reignites interest in a rarely studied process with strong sensitivity to physics beyond the Standard Model, and lays a solid technical foundation for a future cross section measurement in NOvA. Adjustments to the selection procedure, increased statistics from the reverse horn current dataset, and more statistics in the sample of trident-like backgrounds offer promising ways to improve the analysis. These improvements lead to exciting prospects for future studies of the trident production process in NOvA.

Bibliography

- [1] Pierre Radvanyi and Jacques Villain. “The discovery of radioactivity”. In: *Comptes Rendus. Physique* 18.9-10 (2017), pp. 544–550.
- [2] Henri Becquerel. “On the rays emitted by phosphorescence”. In: *Compt. Rend. Hebd. Séances Acad. Sci.* 122.8 (1896), pp. 420–421.
- [3] P. Curie, M. Curie, and G. Bémont. “On a new, strongly radioactive substance, contained in pitchblende”. In: *Comptes Rendus de l’Académie des Sciences* 127 (1898). English translation available at <https://web.lemoyne.edu/giunta/curiesra.html>, pp. 1215–1217.
- [4] Ernest Rutherford. “The Scattering of α and β Particles by Matter and the Structure of the Atom”. In: *Philosophical Magazine*. Series 6 21.125 (1911), pp. 669–688.
- [5] G. H. Briggs. “The decrease in velocity of alpha particles from radium C”. In: *Proceedings of the Royal Society of London. Series A* 114 (1927), pp. 341–354.
- [6] Charles Drummond Ellis and W. A. Wooster. “The average energy of disintegration of radium E”. In: *Proceedings of the Royal Society of London. Series A* 117 (1927), pp. 109–123.
- [7] Wolfgang Pauli. “Dear radioactive ladies and gentlemen”. In: *Phys. Today* 31N9 (1978), p. 27.
- [8] Fred L. Wilson. “Fermi’s Theory of Beta Decay”. In: *American Journal of Physics* 36.12 (1968), pp. 1150–1160.
- [9] H. Bethe and R. Peierls. “The “Neutrino””. In: *Nature* 133 (1934), p. 532.
- [10] Luis W Alvarez. *A Proposed Experimental Test of the Neutrino Theory*. Tech. rep. Lawrence Berkeley National Lab. (LBNL), Berkeley, CA (United States), Apr. 1949.
- [11] Louise Lerner. “The first nuclear reactor, explained”. <https://news.uchicago.edu/explainer/first-nuclear-reactor-explained>. 2021.

[12] B. Pontecorvo. “Inverse beta process”. In: *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* 1 (1991), pp. 25–31.

[13] Frederick Reines and Jr. Clyde L. Cowan. “The Neutrino”. In: *Nature* 178 (1956), p. 446.

[14] Jr. C. L. Cowan et al. “Detection of the Free Neutrino: A Confirmation”. In: *Science* 124 (1956), p. 103.

[15] F. Reines. “The neutrino: from poltergeist to particle”. In: *Rev. Mod. Phys.* 68 (2 1996), pp. 317–327.

[16] Jr. C. L. Cowan et al. “Detection of the Free Neutrino: A Confirmation”. In: *Science* 124 (1956), pp. 103–104.

[17] T. D. Lee and C. N. Yang. “Question of Parity Conservation in Weak Interactions”. In: *Phys. Rev.* 104 (1 1956), pp. 254–258.

[18] T. D. Lee and C. N. Yang. “Parity Nonconservation and a Two-Component Theory of the Neutrino”. In: *Phys. Rev.* 105 (5 1957), pp. 1671–1675.

[19] Chien-Shiung Wu. “The Discovery of the Parity Violation in Weak Interactions and Its Recent Developments”. In: *Nishina Memorial Lectures: Creators of Modern Physics*. Tokyo: Springer Japan, 2008, pp. 43–70. ISBN: 978-4-431-77056-5.

[20] C. S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. In: *Phys. Rev.* 105 (4 1957), pp. 1413–1415.

[21] M. Goldhaber, L. Grodzins, and A. W. Sunyar. “Helicity of Neutrinos”. In: *Phys. Rev.* 109 (3 1958), pp. 1015–1017.

[22] Vidya Bhojaraju. “What is Helicity?” ScienceABC: Pure Sciences. <https://www.scienceabc.com/pure-sciences/what-is-helicity.html>. Jan. 2021.

[23] Lee Grodzins. “The Tabletop Measurement of the Helicity of the Neutrino”. In: *Nuclear Physics B - Proceedings Supplements* 229-232 (2012). Neutrino 2010, pp. 5–13.

[24] Boris Kayser. “Neutrino Mass, Mixing, and Flavor Change”. 2002. arXiv: [hep-ph/0211134](https://arxiv.org/abs/hep-ph/0211134) [hep-ph].

[25] M. Schwartz. “Feasibility of Using High-Energy Neutrinos to Study the Weak Interactions”. In: *Phys. Rev. Lett.* 4 (6 1960), pp. 306–307.

[26] T. D. Lee and C. N. Yang. “Theoretical Discussions on Possible High-Energy Neutrino Experiments”. In: *Phys. Rev. Lett.* 4 (6 1960), pp. 307–311.

[27] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1 1962), pp. 36–44.

[28] Martin. L. Perl et al. “Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation”. In: *Phys. Rev. Lett.* 35 (22 1975), pp. 1489–1492.

[29] Martin L. Perl. “The Discovery of the Tau Lepton”. In: *3rd International Symposium on the History of Particle Physics: The Rise of the Standard Model*. Sept. 1992, pp. 79–100.

[30] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (2001), pp. 218–224.

[31] Alain Blondel. “The number of neutrinos and the Z line shape”. In: *Comptes Rendus Physique* 3.9 (2002), pp. 1155–1164.

[32] S. Navas et al. “Review of particle physics”. In: *Phys. Rev. D* 110.3 (2024), p. 030001.

[33] Masayuki Nakahata. “History of solar neutrino observations”. In: *Progress of Theoretical and Experimental Physics* 2022.12 (Mar. 2022), 12B103.

[34] Bruce T. Cleveland et al. “Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector”. In: *The Astrophysical Journal* 496.1 (1998), p. 505.

[35] Raymond Davis Jr. “Nobel Lecture: A half-century with solar neutrinos”. In: *Nobel Lectures, Physics 2002*. Accessed: 2025-08-08. The Nobel Foundation, 2002.

[36] K. S. Hirata et al. “Observation of ^8B solar neutrinos in the Kamiokande-II detector”. In: *Phys. Rev. Lett.* 63 (1 1989), pp. 16–19.

[37] Atsuto Suzuki. “The start of things for Kamiokande: The Kamioka Nucleon Decay Experiment”. In: *Progress of Theoretical and Experimental Physics* 2022.12 (Feb. 2022).

[38] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Zh. Eksp. Teor. Fiz.* 53 (1967), pp. 1717–1725.

[39] V. Gribov and B. Pontecorvo. “Neutrino astronomy and lepton charge”. In: *Physics Letters B* 28.7 (1969), pp. 493–496.

[40] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. In: *Zh. Eksp. Teor. Fiz.* 34 (1957), p. 247.

[41] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the unified model of elementary particles”. In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880.

[42] Guido Fantini et al. “The formalism of neutrino oscillations: an introduction”. 2020. arXiv: 1802.05781 [hep-ph].

[43] Johanna L. Miller. “Physics Nobel Prize honors the discovery of neutrino flavor oscillations”. In: *Physics Today* 68.12 (Dec. 2015), pp. 16–21. eprint: https://pubs.aip.org/physicstoday/article-pdf/68/12/16/10100897/16_1_online.pdf.

[44] Y. Suzuki. “Solar neutrino results from Super-Kamiokande”. In: *Nuclear Physics B - Proceedings Supplements* 77.1 (1999), pp. 35–42.

[45] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (8 1998), pp. 1562–1567.

[46] A. Bellerive et al. “The Sudbury Neutrino Observatory”. In: *Nuclear Physics B* 908 (2016). Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015, pp. 30–51.

[47] Q. R. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 89 (1 2002), p. 011301.

[48] N. Allemandou et al. “The STEREO experiment”. In: *Journal of Instrumentation* 13.07 (2018), P07009.

[49] Balázs Meszéna. “Neutrino Oscillations”. Wolfram Demonstrations Project. <http://demonstrations.wolfram.com/NeutrinoOscillations/>. 2015.

[50] A. Yu. Smirnov. “The MSW effect and Solar Neutrinos”. In: (2003). arXiv: hep-ph/0305106 [hep-ph].

[51] H. Päs. “Neutrino masses and particle physics beyond the standard model”. In: *Annalen der Physik* 514.8 (Sept. 2002), 551–572.

[52] Jose Valle. “Neutrino masses: Evidences and implications”. In: *Journal of Physics: Conference Series* 485 (Feb. 2014).

[53] Sandip Pakvasa. “CP Violation in the Neutrino Sector”. In: *Journal of Physics: Conference Series* 556.1 (2014), p. 012060.

[54] W. Buchmüller, R.D. Peccei, and T. Yanagida. “LEPTOGENESIS AS THE ORIGIN OF MATTER”. In: *Annual Review of Nuclear and Particle Science* 55.1 (Dec. 2005), 311–355.

[55] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino cross sections across energy scales”. In: *Reviews of Modern Physics* 84.3 (Sept. 2012), 1307–1341.

[56] J. N. Abdurashitov et al. “Measurement of the solar neutrino capture rate with gallium metal”. In: *Phys. Rev. C* 60 (5 1999), p. 055801.

[57] W. Hampel et al. “GALLEX solar neutrino observations: Results for GALLEX IV”. In: *Phys. Lett. B* 447 (1999), pp. 127–133.

[58] Borexino. “Experimental evidence of neutrinos produced in the CNO fusion cycle in the Sun”. In: *Nature* 587.7835 (Nov. 2020), 577–582.

[59] K. Eguchi et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Physical Review Letters* 90.2 (Jan. 2003).

[60] Maury Goodman and Thierry Lasserre. “Double Chooz, A Search for the Neutrino Mixing Angle theta-13”. In: (2006). arXiv: [hep-ex/0606025](https://arxiv.org/abs/hep-ex/0606025) [hep-ex].

[61] Daya Bay Collaboration. “A Precision Measurement of the Neutrino Mixing Angle theta-13 using Reactor Antineutrinos at Daya Bay”. In: (2007). arXiv: [hep-ex/0701029](https://arxiv.org/abs/hep-ex/0701029) [hep-ex].

[62] J. K. Ahn and the RENO Collaboration. “RENO: An Experiment for Neutrino Oscillation Parameter theta-13 Using Reactor Neutrinos at Yonggwang”. In: (2010). arXiv: [1003.1391](https://arxiv.org/abs/1003.1391) [hep-ex].

[63] F.P. An and the Daya Bay Collaboration. “Legacy of the Daya Bay Reactor Neutrino Experiment”. https://indico.fnal.gov/event/45106/contributions/198092/attachments/133513/168123/LOI_DayaBay.pdf.

[64] Takaaki Kajita. “Atmospheric Neutrinos”. In: *Advances in High Energy Physics* 2012.1 (2012), p. 504715. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2012/504715>.

[65] M. G. Aartsen et al. “Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data”. In: *Phys. Rev. D* 91 (7 2015), p. 072004.

[66] Roger Wendell. “Tau Neutrino Searches at Super-Kamiokande and Hyper-Kamiokande”. Presentation at NuTau 2021, Brookhaven National Laboratory. https://indico.bnl.gov/event/10495/contributions/51760/attachments/37420/61656/wendell_nutau_20210929_up.pdf. Sept. 2021.

[67] R. Abbasi et al. “Observation of Seven Astrophysical Tau Neutrino Candidates with IceCube”. In: *Phys. Rev. Lett.* 132.15 (2024), p. 151001. arXiv: 2403.02516 [astro-ph.HE].

[68] M.G. Aartsen et al. “Measurement of atmospheric tau neutrino appearance with IceCube DeepCore”. In: *Physical Review D* 99.3 (Feb. 2019).

[69] Roshan Mammen Abraham et al. “Tau neutrinos in the next decade: from GeV to EeV”. In: *J. Phys. G* 49.11 (2022), p. 110501. arXiv: 2203.05591 [hep-ph].

[70] Daniel A. Dwyer. “The neutrino mixing angle θ_{13} : Reactor and accelerator experiments”. In: *Physics of the Dark Universe* 4 (2014). DARK TAUP2013, pp. 31–35.

[71] S.H Ahn et al. “Detection of accelerator-produced neutrinos at a distance of 250 km”. In: *Physics Letters B* 511.2–4 (July 2001), 178–184.

[72] K. Abe et al. “Indication of Electron Neutrino Appearance from an Accelerator-Produced Off-Axis Muon Neutrino Beam”. In: *Phys. Rev. Lett.* 107 (4 2011), p. 041801.

[73] P. Adamson et al. “Measurement of Neutrino and Antineutrino Oscillations Using Beam and Atmospheric Data in MINOS”. In: *Phys. Rev. Lett.* 110 (25 2013), p. 251801.

[74] Sacha E. Kopp. “The NuMI Neutrino Beam at Fermilab”. In: (2005). arXiv: [physics/0508001](https://arxiv.org/abs/physics/0508001) [physics.acc-ph].

[75] Ishwar Singh, Brajesh C. Choudhary, and Louise Suter. “Latest Three-Flavor Neutrino Oscillation Results from NOvA”. FERMILAB-CONF-24-0792-PPD. 2024.

[76] P. F. de Salas et al. “2020 global reassessment of the neutrino oscillation picture”. In: *Journal of High Energy Physics* 2021.2 (Feb. 2021).

[77] P. Grafström. *Lifetime, cross-sections and activation*. Tech. rep. CERN-AB-2008-064. Lecture notes from CERN Accelerator School, Vacuum Technology Course. CERN Accelerator School, 2008, pp. 209–221.

[78] Enrico Fermi. *Nuclear Physics: A Course Given by Enrico Fermi at the University of Chicago*. Revised edition of lectures originally delivered in 1949. University of Chicago Press, 1974.

[79] P. A. M. Dirac. “The Quantum Theory of the Emission and Absorption of Radiation”. In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 114.767 (1927), pp. 243–265.

[80] D. Griffiths. *Introduction to Elementary Particles*. New York, USA: John Wiley & Sons, 1987.

[81] Francis Halzen and Alan Martin. *Quarks & Leptons: An introductory course in modern particle physics*. New York, USA: John Wiley & Sons, 1984.

[82] Sheldon L. Glashow. “The renormalizability of vector meson interactions”. In: *Nucl. Phys.* 10 (1959), pp. 107–117.

[83] Steven Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (21 1967), pp. 1264–1266.

[84] Abdus Salam. “Weak and Electromagnetic Interactions”. In: *Conf. Proc. C* 680519 (1968), pp. 367–377.

[85] Debdatta Bhattacharya. “Neutrino and Antineutrino Inclusive Charged-Current Cross Section Measurement with the MINOS Near Detector”. Accessed: 2025-08-10. Ph.D. thesis. University of Pittsburgh, 2009.

[86] Daniel Ivan Scully. “Neutrino Induced Coherent Pion Production”. PhD thesis. University of Warwick, Department of Physics, 2013.

[87] Matthew A. Judah. “Measurement of the Inclusive Electron-Neutrino Charged-Current Cross Section in the NOvA Near Detector”. Fermilab Thesis Report FERMILAB-THESIS-2019-29. PhD thesis. Colorado State University, 2019.

[88] Biswaranjan Behera. “Measurement of the Double Differential Inclusive Muon Neutrino Charged-Current Cross-Section in the NOvA Near Detector”. Fermilab Thesis Report FERMILAB-THESIS-2018-20. PhD thesis. Indian Institute of Technology Hyderabad, 2018.

[89] Andrew T. C. Sutton. “Domain Generalization with Machine Learning in the NOvA Experiment”. Ph.D. thesis, University of Virginia, FERMILAB-THESIS-2022-01. PhD thesis. University of Virginia, 2022. ISBN: 978-3-031-43583-6.

[90] Dave Schmitz. “CTEQ Summer School 2011 Lecture”. Presentation at the CTEQ Summer School, National Center for Theoretical Sciences, Taipei, Taiwan. Lecture slides. 2011.

[91] G. Rajasekaran. “Fermi and the theory of weak interactions”. In: *Resonance* 19.1 (Jan. 2014), 18–44.

[92] F. J. Hasert et al. “Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment”. In: *Phys. Lett. B* 46 (1973), pp. 138–140.

[93] Dominique Duchesneau. “Experimental aspects of neutrino oscillation physics”. Apr. 2009.

[94] P. Vilain et al. “A precise measurement of the cross section of the inverse muon decay $\nu_\mu + e^- \rightarrow \mu^- + \nu_e$ ”. In: *Physics Letters B* 364.2 (Feb. 1995), pp. 121–126.

[95] S. R. Mishra et al. “Inverse muon decay, $\nu_\mu + e \rightarrow \mu^- + \nu_e$, at the Fermilab Tevatron”. In: *Physics Letters B* 252.1 (Dec. 1990), pp. 170–176.

[96] D. Ruterbories et al. “Constraining the NuMI neutrino flux using inverse muon decay reactions in MINERvA”. In: *Physical Review D* 104.9 (Nov. 2021).

[97] J. D. Bjorken. “Asymptotic Sum Rules at Infinite Momentum”. In: *Phys. Rev.* 179 (5 1969), pp. 1547–1553.

[98] M. Abdullah et al. “Coherent elastic neutrino-nucleus scattering: Terrestrial and astrophysical applications”. In: (2022). arXiv: 2203.07361 [hep-ph].

[99] C. H. Llewellyn Smith. “Neutrino Reactions at Accelerator Energies”. In: *Phys. Rept.* 3 (1972), pp. 261–379.

[100] Dieter Rein and Lalit M. Sehgal. “Neutrino Excitation of Baryon Resonances and Single Pion Production”. In: *Annals Phys.* 133 (1981), pp. 79–153.

[101] J. L. Hewett et al. “Fundamental Physics at the Intensity Frontier”. In: (2012). arXiv: 1205.2671 [hep-ex].

[102] J. E. Amaro et al. “Meson-Exchange Currents and Quasielastic Antineutrino Cross Sections in the Superscaling Approximation”. In: *Physical Review Letters* 108.15 (Apr. 2012).

[103] Tomasz Golan. “Modeling Nuclear Effects in the NuWro Monte Carlo Neutrino Event Generator”. PhD thesis. University of Wrocław, 2014.

[104] R. Shneor et al. “Investigation of proton-proton short-range correlations via the C-12(e, e-prime pp) reaction”. In: *Phys. Rev. Lett.* 99 (2007), p. 072501. arXiv: nucl-ex/0703023.

[105] J. Nieves, I. Ruiz Simo, and M.J. Vicente Vacas. “The nucleon axial mass and the Mini-BooNE quasielastic neutrino–nucleus scattering problem”. In: *Physics Letters B* 707.1 (Jan. 2012), 72–75.

[106] M. Martini et al. “Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei”. In: *Physical Review C* 80.6 (Dec. 2009).

[107] L. Alvarez-Ruso et al. “NuSTEC White Paper: Status and challenges of neutrino–nucleus scattering”. In: *Progress in Particle and Nuclear Physics* 100 (May 2018), 1–68.

[108] Alessandro Candido et al. “Neutrino structure functions from GeV to EeV energies”. In: *Journal of High Energy Physics* 2023.5 (May 2023).

[109] M. Diemoz, F. Ferroni, and E. Longo. “Nucleon Structure Functions From Neutrino Scattering”. In: *Phys. Rept.* 130 (1986), pp. 293–380.

[110] G. Onengut et al. “Measurement of nucleon structure functions in neutrino scattering”. In: *Phys. Lett. B* 632 (2006), pp. 65–75.

[111] P. Adamson et al. “First measurement of muon-neutrino disappearance in NOvA”. In: *Physical Review D* 93 (Mar. 2016), p. 051104.

[112] Maria Martinez-Casales. “Cross section model tuning and multiplicity studies in NOvA”. In: *FERMILAB-POSTER-19-055-ND* (2019).

[113] Gabriel Magill and Ryan Plestid. “Neutrino trident production at the intensity frontier”. In: *Physical Review D* 95.7 (Apr. 2017).

[114] Wolfgang Altmannshofer et al. “Neutrino tridents at DUNE”. In: *Physical Review D* 100.11 (Dec. 2019).

[115] W. Czyz, G. C. Sheppey, and J. D. Walecka. “Neutrino production of lepton pairs through the point four-fermion interaction”. In: *Nuovo Cim.* 34 (1964), pp. 404–435.

[116] Jørgen Løvseth and Mark Radomski. “Kinematical Distributions of Neutrino-Produced Lepton Triplets”. In: *Phys. Rev. D* 3 (11 1971), pp. 2686–2706.

[117] R W Brown et al. “INTERMEDIATE BOSON. III. VIRTUAL-BOSON EFFECTS IN NEUTRINO TRIDENT PRODUCTION.” In: *Phys. Rev., D* 6: No. 11, 3273-92. (Jan. 1972).

[118] K. Koike et al. “Neutrino production of lepton pairs. I. -”. In: *Prog. Theor. Phys.* 46 (1971), pp. 1150–1169.

[119] K Koike et al. “NEUTRINO PRODUCTION OF LEPTON PAIRS. II.” In: *Progr. Theor. Phys. (Kyoto)* 46: No. 6, 1799-1804. (Jan. 1971).

[120] Wolfgang Altmannshofer et al. “Neutrino Trident Production: A Powerful Probe of New Physics with Neutrino Beams”. In: *Physical Review Letters* 113.9 (Aug. 2014).

[121] Sebastián de Jesús Sánchez Falero. “Neutrino Trident Production Revisited for DUNE-like and MINERvA-like Scenarios”. MA thesis. Pontificia Universidad Católica del Perú, 2017.

[122] Bei Zhou and John F. Beacom. “Neutrino-nucleus cross sections for W-boson and trident production”. In: *Physical Review D* 101.3 (Feb. 2020).

[123] Hans Bethe and Walther Heitler. “On the stopping of fast particles and on the creation of positive electrons”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 146.856 (1934), pp. 83–112.

[124] K. Fujikawa. “The self-coupling of weak lepton currents in high energy neutrino and muon reactions”. In: *Annals Phys.* 75 (1973), pp. 491–502.

[125] R. Belusevic and J. Smith. “W-Z interference in ν -nucleus scattering”. In: *Phys. Rev. D* 37 (9 1988), pp. 2419–2422.

[126] Matheus Hostert. “Neutrino Trident Production: A Short Review”. In: *Prospects of Neutrino Physics Workshop*. 2019.

[127] Peter Ballett et al. “Neutrino trident scattering at near detectors”. In: *Journal of High Energy Physics* 2019.1 (Jan. 2019).

[128] Wolfgang Altmannshofer. “Probing New Physics with Neutrino Tridents”. In: *Hadronic Contributions to New Physics Searches (HC2NP 2019)*. Puerto de la Cruz, Tenerife, 2019.

[129] Wolfgang Altmannshofer et al. “Quark flavor transitions in L_μ - L_τ models”. In: 89.9, 095033 (May 2014), p. 095033. arXiv: 1403.1269 [hep-ph].

[130] Shao-Feng Ge, Manfred Lindner, and Werner Rodejohann. “Atmospheric trident production for probing new physics”. In: *Physics Letters B* 772 (Sept. 2017), 164–168.

[131] Gabriel Magill and Ryan Plestid. “Probing new charged scalars with neutrino trident production”. In: *Physical Review D* 97.5 (Mar. 2018).

[132] R. Santacesaria. “New results from the CHARM-II experiment”. In: (1990). Ed. by K. K. Phua and Y. Yamaguchi.

[133] D. Geiregat et al. “First observation of neutrino trident production”. In: *Phys. Lett. B* 245 (1990), pp. 271–275.

[134] Jeffrey D. Richman and Patricia R. Burchat. “Leptonic and semileptonic decays of charm and bottom hadrons”. In: *Rev. Mod. Phys.* 67 (1995), pp. 893–976. arXiv: hep-ph/9508250.

[135] Janet M. Conrad, Michael H. Shaevitz, and Tim Bolton. “Precision measurements with high-energy neutrino beams”. In: *Reviews of Modern Physics* 70.4 (Oct. 1998), 1341–1392.

[136] S. R. Mishra et al. “Neutrino Tridents and W Z Interference”. In: *Phys. Rev. Lett.* 66 (1991), pp. 3117–3120.

[137] T. Adams et al. “Evidence for diffractive charm production in nu-mu Fe and nubar-mu Fe scattering at the Fermilab Tevatron”. In: *Physical Review D* 61.9 (Mar. 2000).

[138] T. Adams et al. “Neutrino trident production from NuTeV”. In: *29th International Conference on High-Energy Physics*. July 1998, pp. 631–634. arXiv: hep-ex/9811012.

[139] D. S. Ayres et al. “The NOvA Technical Design Report”. FERMILAB-DESIGN-2007-01. Oct. 2007.

[140] P. Adamson et al. “The NuMI neutrino beam”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 806 (2016), pp. 279–306.

[141] Robert Miles Zwaska. “Accelerator Systems and Instrumentation for the NuMI Neutrino Beam”. PhD thesis. Texas U., 2005.

[142] Alexander Himmel. “FHC and RHC Horn Focusing Figure”. NOvA Internal Document, DocDB-20855.

[143] Fermilab Creative Services. <https://vms.fnal.gov/>.

[144] Erika Catano-Mur. “Recent results from NOvA”. arXiv.2206.03542. June 2022.

[145] Mathew Muether. “Silicone Coated APD Photos”. NOvA Internal Document, DocDB-6041.

[146] Anna Cooleybeck. “Blessing package for FD All POT plots”. NOvA Internal Document, DocDB-62786.

[147] A. Aurisano et al. “The NOvA simulation chain”. In: *Journal of Physics. Conference Series* 664.7 (Dec. 2015).

[148] S. Agostinelli et al. “GEANT4 - A Simulation Toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303.

[149] A. Fasso et al. *FLUKA: A Multi-Particle Transport Code*. Tech. rep. SLAC-R-773. SLAC National Accelerator Laboratory, 2005.

[150] L. Aliaga et al. “Neutrino flux predictions for the NuMI beam”. In: *Physical Review D* 94.9 (Nov. 2016).

[151] C. Andreopoulos et al. “The GENIE neutrino Monte Carlo generator”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 614.1 (Feb. 2010), 87–104.

[152] M. A. Acero et al. “Adjusting neutrino interaction models and evaluating uncertainties using NOvA near detector data”. In: *The European Physical Journal C* 80.12 (Dec. 2020).

[153] CERN. “The ROOT Manual”. <https://root.cern/manual/>.

[154] Kanika Sachdev. “Event Topologies in NOvA Detectors”. NOvA Internal Document, DocDB-9908.

[155] Ryan Patterson et al. *Developing novel techniques for readout, calibration and event selection in the NOvA long-baseline neutrino experiment*. Tech. rep. California Inst. of Technology (CalTech), Pasadena, CA (United States), Oct. 2016.

[156] The Attenuation and Calibration Threshold of the NOvA Detectors. “Alexander Radovic”. NOvA Internal Document, DocDB-13579.

[157] Luke Vinton. “Calorimetric Energy Scale Calibration of the NOvA Detectors”. NOvA Internal Document, DocDB-13579.

[158] Prabhjot Singh. “Attenuation Calibration of the NOvA Detectors”. NOvA Internal Document, DocDB-15472.

[159] M Baird et al. “Event Reconstruction Techniques in NOvA”. In: *Journal of Physics: Conference Series* 664 (Dec. 2015), p. 072035.

[160] Biswaranjan Behera, Gavin Davies, and Fernanda Psihas. “Event Reconstruction in the NOvA Experiment”. In: (2017). arXiv: 1710.03772 [physics.ins-det].

[161] Mattias Ohlsson and Carsten Peterson. *Track Finding with Deformable Templates: The Elastic Arms Approach*. Tech. rep. p888. CERN, 1992.

[162] Evan D. Niner. “Observation of Electron Neutrino Appearance in the NuMI Beam with the NOvA Experiment”. Fermilab-thesis FERMILAB-THESIS-2015-16. PhD thesis. Indiana University, 2015.

[163] J. C. Dunn. “A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters”. In: *Journal of Cybernetics* 3.3 (1973), pp. 32–57.

[164] Nicholas Jacob Raddatz. “Measurement of Muon Neutrino Disappearance with Non-Fiducial Interactions in the NOvA Experiment”. FERMILAB-THESIS-2016-05. PhD thesis. Fermilab, 2016.

[165] R Fruehwirth. “Application of Kalman filtering to track and vertex fitting”. In: (1987).

[166] Michael David Baird. “An Analysis of Muon Neutrino Disappearance from the NuMI Beam Using an Optimal Track Fitter”. PhD thesis. Indiana U., 2015.

[167] G. Lutz. “Optimum Track Fitting in the Presence of Multiple Scattering”. In: *Nucl. Instrum. Meth. A* 273 (1988). Ed. by O. Botner, p. 349.

[168] Kelli Michaels and Erica Smith. “Neutrino Tridents in the NOvA Near Detector”. In: Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States). July 2020.

[169] Matthew Strait. “Tech Note: 2016 geometry and composition changes”. NOvA Internal Document, DocDB-23132.

[170] NOvA Collaboration. “Header File: Exposures.h”. NOvA GitHub Repository.

[171] Reed Bowles and Erica Smith. “Trident Update for ND Working Group, June 2023”. NOvA Internal Document, DocDB-59108.

[172] NOvA Collaboration. “Header File: SpillCuts.h”. NOvA GitHub Repository.

[173] NOvA Collaboration. “Macro: BuildMetricsTreeOnMon.C”. NOvA GitHub Repository.

[174] A. Hoecker et al. “TMVA: Toolkit for Multivariate Data Analysis”. In: *PoS ACAT* (2007), p. 040.

[175] Hall Crannell et al. “Interaction Lengths of Energetic Pions and Protons in Iron”. In: *Phys. Rev. D* 7 (Feb. 1973).

[176] Giovanni Punzi. “Sensitivity of searches for new signals and its optimization”. In: (2003). arXiv: [physics/0308063 \[physics.data-an\]](https://arxiv.org/abs/physics/0308063).

[177] Alexander Booth. “Ana2024: ND Data/MC”. NOvA Internal Document, DocDB-62102.

[178] Matthew Judah. “2021 Kaon Systematic Technote and Discussion”. NOvA Internal Document, DocDB-49704.

[179] Erika Catano-Mur. “Blessing package - BEN decomposition 2017”. NOvA Internal Document, DocDB-24816.

[180] Adam Lister. “Steriles Paper: Kaon Issue”. NOvA Internal Document, DocDB-61409.

[181] M.A. Ramírez et al. “Neutrino-Induced Coherent π^+ Production in C, CH, Fe, and Pb at $\langle E_\nu \rangle \sim 6$ GeV”. In: *Physical Review Letters* 131.5 (Aug. 2023).

[182] Statistics Easily. “What is Fitting? Understanding Model Fitting in Statistics”. <https://statisticseasily.com/glossario/what-is-fitting-understanding-model-fitting-in-statistics/>.

[183] ROOT. “TFitResult Class Reference”. <https://root.cern.ch/doc/master/classTFitResult.html>.

[184] GeeksforGeeks. “Covariance Matrix”. <https://www.geeksforgeeks.org/mathematics/covariance-matrix/>.

[185] P.D. Dauncey et al. “Handling uncertainties in background shapes: the discrete profiling method”. In: *Journal of Instrumentation* 10.04 (Apr. 2015), P04015–P04015.

[186] V. Khachatryan et al. “Observation of the diphoton decay of the Higgs boson and measurement of its properties”. In: *The European Physical Journal C* 74.10 (Oct. 2014).

[187] J. R. Klein and A. Roodman. “Blind analysis in nuclear and particle physics”. In: *Ann. Rev. Nucl. Part. Sci.* 55 (2005), pp. 141–163.

[188] Fred James. “MINUIT – Function Minimization and Error Analysis”. CERN Program Library Long Writeup D506. 1994.

[189] G. Cowan. “Statistics”. In: *Review of Particle Physics* (2023).

[190] G Zech. “Comparing Statistical Data to Monte Carlo Simulation - Parameter Fitting and Unfolding”. In: (1995). <https://lib-extopc.kek.jp/preprints/PDF/1995/9508/9508210.pdf>.

[191] Gary J. Feldman and Robert D. Cousins. “Unified approach to the classical statistical analysis of small signals”. In: *Physical Review D* 57.7 (Apr. 1998), 3873–3889.

[192] Evan Groopman. “FCpy: Feldman-Cousins Confidence Interval Calculator”. National Institute of Standards and Technology. <https://catalog.data.gov/dataset/fcpy-feldman-cousins-confidence-interval-calculator>. 2022.

[193] Jerzy Neyman. “Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability”. In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 236 (1937), pp. 333–380.

- [194] Glen Cowan and Eilam Gross. “Discovery Significance with Statistical Uncertainty in the Background Estimate”. <https://www.pp.rhul.ac.uk/~cowan/stat/notes/SigCalcNote.pdf>. 2008.
- [195] M. A. Acero et al. “Monte Carlo method for constructing confidence intervals with unconstrained and constrained nuisance parameters in the NOvA experiment”. In: (2025). arXiv: 2207.14353 [hep-ex].
- [196] Andrew Dye. “The Profiled Feldman-Cousins Method for Confidence Interval Construction for the NOvA 3-flavor analysis”. Fermilab-THESIS-2025-01. PhD thesis. University of Mississippi, 2025.
- [197] Yuya Kaneta and Takashi Shimomura. “On the possibility of a search for the $L_\mu - L_\tau$ gauge boson at Belle-II and neutrino beam experiments”. In: *Progress of Theoretical and Experimental Physics* 2017.5 (May 2017).

Curriculum Vitae

Education

August 2018 - October 2025: Indiana University (IU)

- Ph.D. in Physics
 - Primary emphasis in Experimental High Energy Physics.
 - Secondary emphases in Biophysics and Electronics.
- M.S. in Physics.

August 2014 - May 2018: Wichita State University (WSU)

- B.S. in Physics and Mathematics with Honors
 - Graduated Summa Cum Laude

Teaching Experience

Spring 2023 - Fall 2025: Visiting Assistant Professor at WSU

- Taught Topics in Modern Physics.
- Taught calculus-based Physics 1 and 2 courses.
- Received outstanding reviews from students throughout multiple classes.

Fall 2018 - Spring 2019: Assistant Instructor at IU

- Taught laboratory courses in physics 1 and 2.
- Ensured students followed proper lab procedures and analysis techniques.

Spring 2016 - Spring 2022: Tutor to Undergraduate Students at WSU & IU

- Tutored students in first and second year physics and calculus courses.

Spring 2016 - Spring 2018: Physics Laboratory Instructor at WSU

- Instructed lab classes for physics 1, physics 2, and the physics of sound.

- Taught students proper data management and analysis techniques.
- Created lesson plans and exams, prepared and graded assignments and lab reports.

Research Experience

Summer 2018 - Fall 2025: Graduate Research Assistant at IU

- Performed an analysis identifying neutrino trident events with the NOvA near detector.
- Gained expertise in implementing classical and neural network-based event reconstruction techniques for vertices, tracks, and particle identification on the experiment.
- Worked in data quality for three years identifying and solving problems with detector instrumentation.
- Performed release management and software maintenance on the cluster of NOvA's data quality machines.
- Learned methods of Monte Carlo particle interaction simulation using GENIE.

Spring 2016 - Spring 2018: Undergraduate Researcher at WSU

- Performed energy resolution studies to optimize purity and efficiency of a selected sample in the NOvA near detector.
- Accepted by Department of Energy for SULI internship in 2017.
 - Traveled to Fermi National Accelerator Lab to work directly with research scientists for the summer.
- Learned methods for performing multivariate analyses with boosted decision trees.

Leadership Experience

Fall 2021 - Fall 2022: Young NOvA Institutional Board Representative.

- Participated in and took meeting minutes for NOvA Institutional Board meetings.
- Assisted with planning events, organizing journal club meetings, and creating a “buddy system” for new graduate students and postdocs on NOvA.

Fall 2020 - Fall 2021: Student Representative on IU Physics Student-Staff-Faculty Relations and Long Range Planning Committees

- Provided input on administrative moves made by the department from a student's perspective.
- Helped plan events oriented at improving interdepartmental relations and grow a sense of community.

Spring 2020 - Fall 2021: President of IU Physics Graduate Student Council (PGSC).

- Acted as the primary liaison between the department administration and students.
- Organized several committees to improve student involvement within the department, including committees on Diversity, Academic Advising, and Undergraduate Relations.

Fall 2019 - Spring 2020: Vice President of IU PGSC.

- Traveled to local elementary schools to perform physics outreach.
- Advised other graduate students about available coursework and research opportunities.

Fall 2018 - Spring 2019: Student Representative of IU PGSC.

- Helped plan social events and outreach.
- Worked as a recruiter for incoming graduate students.

Fall 2016 - Spring 2017: President of WSU Society of Physics Students (SPS).

- Organized society events and outreach.
- Performed marketing for events through electronic and physical media.
- Increased participation and attendance at monthly meetings by about 10 people.
- Planned and arranged for the building of demonstrations to be displayed at local schools and outreach events.

Fall 2015 - Spring 2016: Vice President of WSU SPS.

- Assisted in planning events.

Presentations

- **November 2022:** Physics seminar presentation given at Wichita State University: Neutrino Tridents at the NOvA Near Detector.
- **May 2022:** Poster presentation at the International Conference on Neutrino Physics and Astrophysics: Neutrino Tridents in the NOvA Experiment.
- **April 2022:** Presentation at American Physical Society April Meeting 2022: Neutrino Tridents in the NOvA Experiment.
- **February 2022:** Poster presentation for Indiana Lawmakers at Indiana Science Communication Day: Studying Particle Physics with the NOvA Experiment.
- **October 2020:** Plenary presentation at NOvA collaboration meeting: Data Quality Spill Cuts Improvements.
- **October 2017:** Created and presented a poster at FNAL, open to all faculty: Progress on a Single Prong NCPi0 Analysis in the NOvA Near Detector.

Computer Skills

- Proficient with Linux, Windows, and MacOS, including all Microsoft Office utilities.
- Experienced in creating technical documents and papers using LaTeX.
- Knowledgeable in software development and release management with SVN and GIT.
- Hands-on experience training and evaluating convolutional neural networks within the TensorFlow framework.
- Experienced in programming and performing data analysis with MATLAB, PYTHON, C++, ROOT, and FORTRAN.