Search for Heavy Neutral Leptons at the MINERνA detector



Komninos John Plows

Mansfield College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2024

Statement of originality

This thesis and the work presented in it are my own, and were produced by me as a result of my own original research. Results and figures from published works by others have been clearly attributed, and this work has not been submitted for another qualification at this or any other university.

This thesis begins by introducing the Standard Model neutrino, and neutrino oscillations which necessitate the existence of non-zero neutrino mass, in Chapters 1 and 2. Chapter 2 contains a broad overview of the theoretical foundations of HNL, and of the effective field theory developed by Pilar Coloma *et al* used in the simulation of the HNL signal. Chapter 3 and 4 detail the NuMI beamline and the MINERvA experiment respectively, and summarise the work done to understand the neutrino flux prediction and the detector response calibrations. Work on the MINERνA flux prediction was primarily carried out by Leo Aliaga and Ben Messerly, using the `dk2nu` flux driver developed for NuMI by Robert Hatcher; all their contributions have been noted, and a more technical explanation of `dk2nu` can be found in Appendix B. Chapter 4, in particular, summarises and expands upon the MINERνA detector calibration paper published in *Nuclear Instruments and Methods* ***A***, with event displays produced using Arachne. Arachne was developed by Nathaniel Tagg and Erin Cochran.

Chapter 5 introduces the `BeamHNL` simulation module for `GENIE`, and presents the full signal simulation procedure. `BeamHNL` was developed by myself and my supervisor, Dr. Xianguo Lu, in collaboration with Prof. Costas Andreopoulos and Dr. Marco Roda as part of the GENIE Collaboration's Incubator programme. The principles and configuration of `BeamHNL` have been published in *Physical Review* ***D***, and can also be found in the latest `GENIE` Physics and User's Manual; a more technical look at the module is in Appendix A. Chapter 6 details the MC signal production and discusses the organisation of the signal samples, as well as the main background-enhanced samples; these background-enhanced samples, as well as the simulation and computing infrastructure required to translate the event generator output to full reconstructed signal, were developed by the MINERνA collaboration.

Chapters 7 and 8 detail the signal selection and background estimation, respectively. The analysis reconstruction tuples, apart from those single-track specific elements developed by myself, are based on the MINERνA "data preservation" tuple Gaudi code used in MINERνA, and contributed to by all members of the collaboration. The MAT analysis framework was developed originally by Andrew Olivier, Heidi Schellman, Ben Messerly, Rob Fine, and many others; I adapted this to the current analysis.

Finally, in Chapter 9 a fake-data study based on the analysis outputs from Chapters 7 and 8 is performed to showcase what MINERνA's search results post-unblinding will look like. These results and interpretations are entirely my own work.

Acknowledgements

There are multiple people who have contributed, in ways large or small, to this thesis. An acknowledgment section that properly recognises everyone would be far longer than advisable under the page limit, so a very much non-exhaustive section will have to do instead.

First of all, huge thanks go out to my supervisors, Dr. Xianguo Lu and Prof. Giles Barr.
Xianguo has skilfully and patiently guided me throughout my D. Phil. career, providing clear and constructive advice and criticism and gladly discussing any ideas, thoughts or worries that I had. He has always been available if I need to reach out to him, in-person or online, and has helped me massively to locate and capitalise on opportunities to present and extend my work, leading to many fruitful collaborations over the years. For that, I thank him warmly.
Giles, too, has been amazingly patient and helpful, always providing his invaluable perspective on my analysis and writing procedure. He has been more than happy to sit down and honestly discuss the physics of HNL, but also of neutrino interactions and experiments in general, with me, and I owe him deep gratitude for his reassuring presence and unwavering belief in me over the course of my D. Phil.

This analysis is a result of many scientists', engineers', computer and lab technicians', and students' efforts, at Fermilab and MINERνA. I am incredibly grateful to all the members of the MINERνA Collaboration for their support in my conducting this search, and providing me with useful feedback on my analysis. Specifically, I thank Profs. Laura Fields, Kevin McFarland, and Deborah Harris for their useful discussions on modelling the HNL flux and the coherent pion production background. Also, thank you to Drs. Daniel Ruterbories, Abbey Waldron, and Alejandro Ramírez-Delgado. Dan was instrumental in maintaining much of the MINERνA infrastructure that had been developed much before I joined the experiment, and was extremely kind with his time to help me achieve the full simulation pipeline from GENIE predicition to MINERνA reconstructed files. Abbey, who heads the Special Samples production group, was instrumental in organising the delivery of the full HNL simulation. And Alex has always been very helpful with discussing the details of the CC Coherent sample, and of the physics of charged-current coherent pion production in general. Also, thank you to the many developers of the MINERνA Analysis Toolkit and of the standard analysis systematics, without which this thesis would probably have taken another year or two to complete. Another thank you goes to Ryan Howell and Oscar Moreno, with whom we had many funny and useful chats during our time doing MINERνA production.

In addition, recognition must go out to the members of the GENIE Collaboration, who helped me to develop the `BeamHNL` simulation package. Particular thanks must go to Prof. Constantinos Andreopoulos, who guided me through the procedure of development and integration with GENIE, and to Dr. Marco Roda for helping me adopt good coding practices.

There are many colleagues in the office and friends who deserve my gratitude, as well. By name, thank you to Drs. Daniel Barrow, Patrick Green, Clarence Wret, and Luke Pickering, all of whom have had useful comments and suggestions for my analysis as it progressed. Thank you to the not-as-yet-mentioned faculty of the group, Profs. David Wark and Morgan Wascko, and Dr. Kirsty Duffy, who humorously and helpfully provided insight and experience. Thank you to the other students in the office; the now-graduated Drs. Ciaran Hasnip and Kang Yang, soon-to-graduate Soniya Samani, and everyone else; thanks Tom, Federico, Weijun, Lars, Yiwen, and Abbey! I appreciate your patience with my (over)use of the coffee machine in the office, and my (very) eventual contribution to the coffee-bean stock; and your good humour and friendly chats too!

I couldn't have finished this thesis without the steady support of my friends, who are too numerous to enumerate in detail. But in particular, I thank my mates from back at the Polytechneio, and from Deree. To Thanassis, soon-to-be Dr. Karakassis, your taste in music is pure metal. To Manolis, who would have thought we'd both be able to get into Oxford? There are many lovely housemates I've had the pleasure of meeting over the years, and in particular a big thank you to Jennifer (soon-to-be Dr. Redmond!) for many helpful discussions about D. Phils and swimming. And also Reddit!

But most of all, I am deeply beholden to my family, who gave me the tools to succeed even at great personal cost, and have suffered me every step of the way. To my aunts Helen and Cathy, my uncles Ken and Marcus, and my cousins, I always enjoy seeing you and spending time with you. To my uncle Aristos and my aunt Tassa, it's a real pleasure to be around you! (I am indeed cleverer than my own shadow). And to my mother, Cally, this is it. You've sacrificed more than I could begin to fathom to let me get here, but you're the driver behind this accomplishment. You've been my teacher and mentor all this time, from primary school to D. Phil, and I can finally understand your jokes about turning into a hermit while writing! This one's for you, and Stephen and Tom.

Abstract

Heavy Neutral Leptons (HNL) are particles hypothesised to provide a mass generation mechanism for the active (observed) neutrino species, which are known to have nonzero mass from the definitive observation of neutrino oscillations. HNL are eigenstates of mass of the order $\mathcal{O}\left(0.1 - 1\ \mathrm{GeV}/c^2\right)$, which mix into the active flavour eigenstates through the extended leptonic mixing matrix. Apart from neutrino mass, they could provide a natural dark matter candidate and a mechanism for matter-antimatter asymmetry in the early Universe, giving rise to today's matter-dominated cosmos. Searches for HNL typically attempt to either confirm the existence of HNL through an excess in data that is most compatible with an HNL hypothesis, or by setting limits on the HNL parameter space $\left(M_{\mathrm{N4}}, |U_{\alpha 4}|^2\right)$ in the case of statistically insignificant excess. Such searches are intensifying as part of a global research programme at both colliders and accelerator / atmospheric neutrino experiments. One such setting is the MINERνA experiment, located in the NuMI beamline at Fermilab. MINERνA has collected a large amount of data over seven years of operation to measure the cross-sections of neutrino-nucleus interactions necessary to drive systematic uncertainties down, in order for neutrino oscillation experiments to achieve sensitivity to CP violation in the neutrino sector. With the high-energy, high-intensity NuMI beam, and with good timing, position, and angular resolution leading to sensitivity to HNL decays, MINERνA is fertile ground for an HNL search, which is performed in this thesis. Specifically, a novel, experiment-agnostic and general simulation of HNL production and decay is presented and deployed; this simulation has been incorporated in the ubiquitous `GENIE` neutrino event generator for use with neutrino experiments in the future. The event selection and background characterisation is discussed in detail, including the primary background coming from charged-current coherent and diffractive pion production from neutrino-nucleus interactions, and constraints on the background using control regions from MINERνA data are derived. Finally, by means of a fake-data study given a background of $\mathcal{O}\left(500\right)$ events, the discovery potential and limit-setting capacity of MINERνA is demonstrated, and ways to improve this capacity are expounded upon.

Contents

List of Figures ix

List of Tables xiii

List of Abbreviations xiv

1 Introduction 1

2 Heavy Neutral Leptons 4
- 2.1 Introduction . . . 4
- 2.2 The problem of neutrino mass . . . 8
  - 2.2.1 Neutrino oscillations . . . 8
  - 2.2.2 Absolute scale of neutrino masses . . . 13
- 2.3 Mass generation mechanisms . . . 15
  - 2.3.1 Low-scale seesaws . . . 18
  - 2.3.2 The νMSM . . . 18
- 2.4 The single-HNL model . . . 19
- 2.5 Effective field theories with HNL . . . 21
  - 2.5.1 HNL production mechanisms . . . 22
  - 2.5.2 HNL decays . . . 23
- 2.6 Previous searches for HNL . . . 25

3 The NuMI neutrino beam 28
- 3.1 Neutrino beam operating principles . . . 28
- 3.2 NuMI beamline . . . 30
- 3.3 NuMI beamline simulation . . . 34
  - 3.3.1 Coordinate systems . . . 36
- 3.4 Flux uncertainties . . . 40
  - 3.4.1 Hadroproduction uncertainties . . . 41
  - 3.4.2 Focusing uncertainties . . . 43
  - 3.4.3 Constraining flux uncertainties at MINERνA . . . 44

4 The MINERνA detector 47
- 4.1 Description of the detector . . . 47
- 4.2 Data acquisition (DAQ), triggering, and readout . . . 50
- 4.3 Calibration . . . 52
  - 4.3.1 Hit energy calibration . . . 53
  - 4.3.2 Timing calibration . . . 60
- 4.4 Common reconstruction . . . 63

4.4.1 Time slicing . . . 63
4.4.2 Clustering . . . 64
4.4.3 Track formation . . . 65
4.4.4 Energy reconstruction . . . 68
4.5 Detector response simulation . . . 71
4.6 Readout simulation . . . 72

5 HNL event generation 76
5.1 Effective field theory for HNL . . . 76
5.2 HNL production in a beamline . . . 79
5.2.1 Kinematics at production . . . 79
5.2.2 Propagation . . . 82
5.3 Decay . . . 90
5.3.1 Decay kinematics . . . 90
5.3.2 Decay vertex placement . . . 93
5.4 POT estimation . . . 98
5.5 The BeamHNL simulation module . . . 101

6 Monte Carlo samples 104
6.1 MINERνA running periods . . . 104
6.2 Simulation pipeline . . . 105
6.3 Sample organisation . . . 108

7 Selection of $N_4 \to \pi^{\pm}\mu^{\mp}$ events at MINERνA 111
7.1 HNL event reconstruction . . . 112
7.1.1 Vertex finding . . . 113
7.1.2 Track merging . . . 116
7.1.3 Energy deposition reconstruction . . . 119
7.2 Event signature . . . 123
7.3 One-track classification . . . 126
7.3.1 Pre-selection cuts . . . 126
7.3.2 Cost-benefit analysis . . . 131
7.3.3 Selection cuts . . . 134
7.3.4 Efficiency calculation . . . 149

8 Background estimation and systematic uncertainties 155
8.1 Backgrounds to HNL decay . . . 155
8.1.1 Coherent pion production in tracker . . . 155
8.2 Simulation of Standard Model background sources . . . 157
8.2.1 Mainline `GENIE` simulation . . . 157

8.2.2 MINERνA tunes . . . 159
8.3 Systematic evaluation . . . 160
8.3.1 Many-universe method . . . 160
Error matrix calculation . . . 169
8.3.2 Systematic uncertainties . . . 169
8.4 Background tuning . . . 171
8.4.1 Sideband fitting . . . 171
8.4.2 Tuning the background . . . 177
8.4.3 Fit results . . . 180
Unconstrained $N = 5$ fit . . . 182
Reduced fit: $N = 1$ . . . 189
Reduced fit: $N = 2$ . . . 190
8.4.4 Conclusions . . . 193

9 Inference of signal and limit extraction 200
9.1 Introduction . . . 200
9.2 Overview of inference procedure . . . 201
9.3 Feldman-Cousins limit extraction . . . 204
9.4 Fake data study for HNL . . . 207
9.4.1 Fake data study . . . 211
9.5 Setting exclusion limits on the HNL hypothesis . . . 215

10 Conclusion and future directions 223

Appendices

A The `BeamHNL` GENIE simulation package 227
A.1 The configuration file . . . 227
A.2 Running the code . . . 228

B The `dk2nu` tuple structure 230
B.1 The main tuple `dk2nutree` . . . 230
B.2 The metadata tree `dkmetaTree` . . . 231

References 233

List of Figures

2.1 SK muon disappearance results. 10
2.2 SNO electron-disappearance results. 11
2.3 Current HNL exclusion limits and sensitivities. 27

3.1 NuMI focusing horns. 29
3.2 NuMI target. 32
3.3 NuMI baffle. 32
3.4 NuMI beam sketch. 33
3.5 LE and ME fluxes at MINERνA. 34
3.6 Parent particle spectra at NuMI, from the MINERνA flux files. . . 36
3.7 MINERνA nominal SM neutrino flux breakdown (neutrino type). . 37
3.8 MINERνA neutrino flux breakdown (parent species). 38
3.9 Parent meson focusing by horns. 39
3.10 External hadroproduction data. 43
3.11 ME beam focusing systematic uncertainties. 44

4.1 The MINERνA detector views. 49
4.2 Scintillator strips at MINERνA. 50
4.3 PMT readout at MINERνA. 51
4.4 Discriminator and dead-time illustration. 52
4.5 Simplified readout schematic. 54
4.6 ADC counts versus input charge. 56
4.7 Operation of a PMT. 57
4.8 MEU determination procedure. 59
4.9 Rock muon peak number of PE, as function of time. 59
4.10 Plane alignment for a, “X” plane at MINERνA. 61
4.11 Time slewing. 62
4.12 MINERνA data event in “XZ” view. 64
4.13 Cluster types in MINERνA. 65
4.14 Long track formation in MINERνA. 67
4.15 Angular resolutions at MINERνA. 69
4.16 Energy loss by ionisation. 70
4.17 Simulated HNL decay with data overlay in MINERνA. 73
4.18 Time-smearing calibration. 74

5.1 HNL production BR validation. 77
5.2 HNL decay BR validation. 78
5.3 The phase-space integral as series of “two-body” decays. 80

5.4 HNL collimation effect. 85
5.5 Backwards emission of HNL. 86
5.6 HNL acceptance correction contours. 88
5.7 HNL forwards versus backwards emission probability. 89
5.8 Angular dependence of various HNL two-body decays. 91
5.9 $N_4 \to \pi^+ e^-$ angular dependence, lab frame. 94
5.10 $N_4 \to \pi^+ e^-$ angular dependence, HNL rest frame. 95
5.11 SM neutrino event times from data. 96
5.12 NuMI bunch and spill time structure. 97
5.13 HNL event time distributions at MINERνA. 97
5.14 HNL arrival times at MINERνA. 99
5.15 POT estimation sketch. 100
5.16 `BeamHNL` event generation schematic. 103

6.1 MINERνA running periods. 105
6.2 HNL parameter space slicing. 106
6.3 Parent entries parsed per HNL event generated. 109

7.1 HNL vertex finding comparisons, $M_{\mathrm{N4}} = 250$ MeV/c^2. 115
7.2 HNL vertex finding comparisons, $M_{\mathrm{N4}} = 475$ MeV/c^2. 116
7.3 Vertex finding comparisons, SM neutrinos. 117
7.4 Stub tracks in MINERνA from $N_4 \to \mu\pi$ decay. 119
7.5 Autocorrelation of a Landau distribution. 122
7.6 Phase space for final state of $N_4 \to \mu\pi$ decay and CC coherent $\nu_\mu + A \to \mu^- + \pi^+ + A$. 124
7.7 HNL crossing and decaying fluxes, FHC. 126
7.8 HNL crossing and decaying fluxes, RHC. 127
7.9 Reconstructed HNL muon energy distributions. 128
7.10 Phase space of MINOS-accepted muons. 130
7.11 Event vertex positions in MINERνA, "zy" projection. 132
7.12 Event vertex positions in MINERνA, "xy" projection. 133
7.13 Illustration of cylinder for vertex energy cut. 135
7.14 N_{planes} FHC distributions. 138
7.15 N_{planes} RHC distributions. 139
7.16 N_{planes} background exclusion curves. 140
7.17 $E_{\mathrm{vtx}}^{(4\ \mathrm{planes})}$ FHC distributions. 141
7.18 $E_{\mathrm{vtx}}^{(4\ \mathrm{planes})}$ RHC distributions. 142
7.19 $E_{\mathrm{vtx}}^{(4\ \mathrm{planes})}$ background exclusion curves. 143
7.20 $E_{\mathrm{vtx}}^{(10\ \mathrm{planes})}$ FHC distributions. 144
7.21 $E_{\mathrm{vtx}}^{(10\ \mathrm{planes})}$ RHC distributions. 145

7.22 $E_{\text{vtx}}^{(\text{10 planes})}$ background exclusion curves. . . . 146
7.23 $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$ FHC distributions. . . . 147
7.24 $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$ RHC distributions. . . . 148
7.25 $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$ background exclusion curves. . . . 149
7.26 Relative cut efficiencies, SM backgrounds. . . . 151
7.27 Relative cut efficiencies, HNL signal. . . . 152
7.28 SM background discrimination capability. . . . 153
7.29 HNL signal total efficiency. . . . 154

8.1 Neutrino cross-sections. . . . 158
8.2 FSI fates. . . . 159
8.3 MINOS tracking correction. . . . 161
8.4 θ_μ FHC distributions. . . . 163
8.5 θ_μ RHC distributions. . . . 164
8.6 $p_{\text{T}}^{(\mu)}$ FHC distributions. . . . 165
8.7 $p_{\text{T}}^{(\mu)}$ RHC distributions. . . . 166
8.8 $p_z^{(\mu)}$ FHC distributions. . . . 167
8.9 $p_z^{(\mu)}$ RHC distributions. . . . 168
8.10 Fractional uncertainties of θ_μ. . . . 172
8.11 Fractional uncertainties of $p_{\text{T}}^{(\mu)}$. . . . 173
8.12 Fractional uncertainties of $p_z^{(\mu)}$. . . . 174
8.13 Schematic illustration of sidebands. . . . 177
8.14 $\Delta\chi^2$ landscape plots for the C_1 sideband. . . . 181
8.15 Tuning comparisons, FHC C_1 sideband. . . . 184
8.16 Tuning comparisons, RHC C_1 sideband. . . . 185
8.17 Tuning comparisons, FHC C_2 sideband. . . . 186
8.18 Tuning comparisons, RHC C_2 sideband. . . . 187
8.19 Correlation matrices of background fit. . . . 188
8.20 $\Delta\chi^2$ curve for single-parameter fit. . . . 191
8.21 Tuned θ_μ distributions from the single-parameter fit. . . . 192
8.22 Tuned $p_{\text{T}}^{(\mu)}$ distributions from the single-parameter fit. . . . 193
8.23 Tuned $p_z^{(\mu)}$ distributions from the single-parameter fit. . . . 194
8.24 $\Delta\chi^2$ curve for double-parameter fit. . . . 196
8.25 Tuned θ_μ distributions from the double-parameter fit. . . . 197
8.26 Tuned $p_{\text{T}}^{(\mu)}$ distributions from the double-parameter fit. . . . 198
8.27 Tuned $p_z^{(\mu)}$ distributions from the double-parameter fit. . . . 199

9.1 Expected number of $N_4 \to \mu\pi$ decays in MINERνA. . . . 201
9.2 Limit extraction procedure. . . . 204
9.3 Reference Feldman-Cousins plot for $B = 3, \alpha = 0.9$. . . . 206

9.4 Feldman-Cousins confidence intervals at various choices of α. 207
9.5 Validation of varying background indicating Poissonian pathology. . 208
9.6 Validation of varying background with pathology removed. 209
9.7 Toy fake data model. 211
9.8 Effect of scale factor uncertainty on discovery potential. 213
9.9 Discovery potential of benchmark excess scenarios. 214
9.10 Diagram showing critical points for Feldman-Cousins interval. . . . 216
9.11 Feldman-Cousins high-background extrapolation fits. 218
9.12 Example of $N = 1000$ Feldman-Cousins schema obtained using extrapolation. The nominal extrapolated value is the solid red line, and the dashed blue lines are the errors on that extrapolation. In the inset, the critical point $(b_{\mathrm{D}}, \mu_{\mathrm{D}})$ lies in the centre of the square defined by $(b_{\mathrm{D}} \pm \delta b_{\mathrm{D}}) \times (\mu_{\mathrm{D}} \pm \delta \mu_{\mathrm{D}})$ (dashed black square). 219
9.13 Feldman-Cousins 90% limits on HNL in the small-signal scenario. . 222

A.1 Euler rotations, in the extrinsic $X - Z - X$ convention. 228
A.2 `GENIE EventRecord` for an example $\mathrm{N}_4 \to \pi^+\mu^-$ decay in MINERνA. 229

List of Tables

2.1 Standard Model fermionic particle content and charges. 7
2.2 Most dominant SM neutrino production channels from decay of parent particles, taken from the "Particle Listings" section of the 2022 PDG. $\mathcal{B}_{\Sigma\nu}$ is the sum of BR of all channels with a neutrino in the final state, and $\mathcal{B}$ is the SM branching ratio for each channel. Thresholds refer to HNL production via substituting $\nu_\alpha \mapsto N_4$ in each channel. 24
2.3 HNL lepton-number conserving decay modes. 25

3.1 Beam systematic sources. 45

4.1 Tracker nucleus mass/number fractions. 49

5.1 Focused parent particle multiplicities from `GEANT4` at NuMI. 100

6.1 POT per sample processed. 110

7.1 Pre-selection cut efficiencies. 129
7.2 Pre-selection cut relative efficiencies. 129
7.3 Selection cut relative efficiencies. 137

8.1 Analysis binning choices. 169
8.2 Background contaminations in sidebands. 176
8.3 $\Delta\chi^2$ values for various errors and number of fit parameters. 179
8.4 Combined samples used for fits. 182
8.5 Background scale factors from `MINUIT`. 183
8.6 Single-parameter fit results. 190
8.7 Double-parameter fit results. 195
8.8 Final values of scaling factors used in Chapter 9. 195

9.1 Critical values for LLR at various CLs. 212
9.2 Parameters for Feldman-Cousins curve extraction. 217
9.3 Summary of selection efficiencies. 220
9.4 Expected number of $N_4 \to \mu\pi$ decays in MINERνA. 220

List of Abbreviations

2p2h . . . . . "Two-particle-two-hole" scattering, i.e. from correlated pairs of nucleons in nuclear initial state.

ADC . . . . . Analogue-to-digital converter.

BBN . . . . . Big Bang Nucleosynthesis.

BSM . . . . . Beyond the Standard Model.

BooNE . . . . **Boo**ster **N**eutrino **E**xperiment.

BR . . . . . Branching ratio.

CC . . . . . Charged-current.

CKM . . . . . Cabibbo-Kobayashi-Maskawa.

C. L. . . . . . Confidence level.

CMB . . . . . Cosmic microwave background.

CP . . . . . Charge-parity exchange.

CRIM . . . . . CROC interface module.

CROC . . . . . Chain read-out controller.

CV . . . . . Central Value.

DAQ . . . . . Data acquisition.

DIS . . . . . (CC) Deep Inelastic Scattering.

DM . . . . . Dark matter.

DUNE . . . . . **D**eep **U**nderground **N**eutrino **E**xperiment.

EM . . . . . Electromagnetic.

FC . . . . . Fully contained.

FEB . . . . . Front-end board.

FHC . . . . . Forward Horn Current (NuMI in neutrino mode)

FSI . . . . . Final-State Interactions.

FV . . . . . Fiducial volume.

HNL . . . . . Heavy Neutral Lepton.

IO . . . . . Inverted ordering ($m_{\nu 3} \ll m_{\nu 1} \simeq m_{\nu 2}$).

IMD . . . . . Inverse Muon Decay.

INC . . . . . Intra-Nuclear Cascade.

KamLAND . . **Kam**ioka **L**iquid Scintillator **A**nti-**N**eutrino **D**etector.

LE . . . . . Low Energy.

LED . . . . . Light Emitting Diode.

LLR . . . . . Log-Likelihood Ratio.

LSND . . . . . **L**iquid **S**cintillator **N**eutrino **D**etector.

MAT . . . . . MINERvA Analysis Toolkit.

MC . . . . . Monte Carlo.

ME . . . . . Medium Energy.

MEC . . . . . Meson Exchange Current. See "2p2h".

MEU . . . . . Muon equivalent unit.

MI . . . . . Main Injector.

MINERνA . . The **M**ain **IN**jector **E**xpe**R**iment to study neutrino-nuclear ($\boldsymbol{\nu - A}$) interactions.

MINOS . . . . The **M**ain **I**njector **N**eutrino **O**scillation **S**earch.

MIP . . . . . Minimum ionising particle.

MPV . . . . . Most-probable value of a Landau distribution.

NC . . . . . Neutral-current.

NO . . . . . Normal ordering ($m_{\nu 1} \simeq m_{\nu 2} \ll m_{\nu 3}$).

NuMI . . . . . The **Neu**trino **M**ain **I**njector.

PC . . . . . Partially contained.

PCAC . . . . . Partially-Conserved Axial Current.

PE . . . . . Photo-electrons.

PMNS . . . . . Pontecorvo-Maki-Nakagawa-Sakata.

PMT . . . . . Photo-multiplier tube.

QCD . . . . . Quantum chromodynamics.

QEL . . . . . (CC) Quasielastic scattering.

QGS . . . . . Quark-gluon string.

RES . . . . . (CC) Resonant scattering.

RHC . . . . . Reverse Horn Current (NuMI in antineutrino mode).

RPA . . . . . Random Phase Approximation.

SBN . . . . . Short Baseline Program.

SK . . . . . Super-Kamiokande.

SM . . . . . The Standard Model.

SNO . . . . . **S**udbury **N**eutrino **O**bservatory.

SSM . . . . . The Standard Solar Model.

T2K . . . . . **T**okai **To K**amioka.

TDC . . . . . Time-to-digital conversion.

WLS . . . . . Wavelength-shifting.

1 Introduction

The discovery of neutrino oscillations, awarded the Nobel Prize in 2015 [1], opened the gates for a vast neutrino physics programme producing a wealth of measurements in three-flavour neutrino oscillations and their parameters, as well as neutrino interactions, direct measurements of neutrino mass and searches for hints beyond the Standard Model (SM). The Standard Model, predicting that neutrinos must be massless, was definitively shown to be incomplete, despite being the indisputably most accurate physical theory. Neutrinos with no mass cannot oscillate; neutrino oscillations imply that the mass of neutrinos is non-zero. Puzzlingly, the mass of neutrinos has been shown both by direct measurement and by cosmological constraints to be very small, orders of magnitude below even that of the electron. A suite of experimental observations relating to the maximal parity violation in the weak interaction, the only interaction that neutrinos participate in, provided compelling evidence for neutrinos being solely "left-handed", and lacking a right-chiral counterpart, as opposed to all other particles in the Standard Model. This lack of a right-handed copy excludes neutrinos from the Higgs mechanism which is understood to endow the Standard Model particles with their mass. (In fact, though one can write right-handed neutrino terms without violating any symmetries of the SM gauge group - just as with charged leptons - experimental measurements of the weak interaction's properties highly constrain any right-handed terms to be negligibly small, if present at all). Therefore, a dual question arises: *why are neutrino masses so small, and how do they obtain their mass*?

A minimal solution to this question occurs by challenging the premise of neutrinos being left-handed only; rather than admit these particles to an exclusive, special class, one might postulate the existence of right-handed neutrinos which interact so seldom that they would have not left any statistically significant signatures in the experiments that determined the neutrino chirality. There exists a multitude of ways to incorporate such states into the particle zoo, with various pictures of the phenomenological landscape; of these, the simplest model compatible with

observations is the seesaw mechanism which explains the smallness of the left-handed, or *active*, neutrinos through the heaviness of the right-handed, or *sterile*, ones. The heavier the steriles, the lighter the actives. Precisely how heavy the sterile states ought to be is, once again, a mystery; it is possible for these to be as light as a few eV/c^2, though in many theories the floor is at the $\mathcal{O}\left(\mathrm{MeV}/c^2\right)$ level.

Such $\mathcal{O}\left(\mathrm{MeV}/c^2\right)$ states are called *Heavy Neutral Leptons*, or HNL. HNL are a tantalising target of opportunity for beam-dump and collider experiments, since their decays leave distinctive signatures and their discovery would herald unquestionable evidence for the mechanism of neutrino mass generation. A wide experimental programme for the search for HNL is underway at neutrino experiments at Fermilab, J-PARC, and CERN, as well as the IceCube detector under the South Pole. High-resolution modern detectors with formidable reconstruction capabilities are essential to definitively rule out the full parameter space available for HNL, since the lower the strength of HNL mixing, the more likely an HNL signal is to be misinterpreted as a fluctuation of SM background. Simultaneously, the development of robust background reduction analysis techniques is required, for the same reason. This effort has been hitherto been obfuscated by the lack, in the accelerator neutrino setting, of common tools to simulate the production and decay of HNL, and to estimate the rate at which visible events from HNL are expected in a detector. Gradually, the Fermilab liquid-argon experiments are converging to some iteration of a common simulation, though adoption of an experiment-agnostic framework appears to be slow. Likewise, the T2K collaboration has implemented a simulation framework for HNL at the near detector ND280, but it may not be straightforwardly compared with realisations in other detectors. Standardisation of this step is crucial for the full exploitation of global neutrino detector data, so that constraints from the various experiments may be meaningfully combined.

The MINERνA detector at Fermilab's high-energy NuMI beam is a dedicated neutrino interaction experiment, which took data from a neutrino-dominated and an antineutrino-dominated beam at high mean neutrino energy ~ 6 GeV (known as the Medium Energy, or ME, run) over a period of seven years from 2013-2019.

During this time, it accumulated an exposure of $\sim 1.2 \cdot 10^{21}$ protons on target in either mode, which allow for measurements of the neutrino-nuclear cross-section for multiple inclusive and exclusive interaction modes to unparalleled accuracy. The detector offers excellent timing resolution and good spatial and angular resolutions, as well as superior capability to measure the energy deposition per unit length of tracks ionising the detector. Based on these traits, a search for HNL was planned out and performed using MINERνA data from the ME run, exploiting the detector's response to a unique signature of an overlapping muon and pion pair. The search attempts to characterise the irreducible background deriving from the charged-current coherent pion production off a nucleus, in a region of the coherent phase space that has heretofore not been measured by any experiment.

This thesis is organised as follows. In Chapter 2, an overview of the theoretical basis for HNL is presented. Chapters 3 and 4 detail the NuMI beamline and MINERνA detector, respectively, and discuss the operation and response of the elements within. Chapter 5 lays the groundwork for the simulation of the HNL signal, in an experiment-agnostic way, by means of a systematic treatment of relativistic kinematics; considerations on event timing and polarisation effects of the HNL are also relayed, and a schematic overview of the `BeamHNL` simulation module developed over the course of this thesis is given. The production of the MINERνA HNL signal samples is described in Chapter 6, and the double-particle signal selection and characterisation of the background are detailed in Chapters 7 and 8, respectively. Finally, the analysis strategy, taking into consideration the expected background and prior constraints on the HNL parameter space, is discussed in Chapter 9, before a short outlook on future experiments is outlined in Chapter 10.

2 Heavy Neutral Leptons

Contents

2.1 Introduction 4
2.2 The problem of neutrino mass 8
2.2.1 Neutrino oscillations 8
2.2.2 Absolute scale of neutrino masses 13
2.3 Mass generation mechanisms 15
2.3.1 Low-scale seesaws 18
2.3.2 The νMSM 18
2.4 The single-HNL model 19
2.5 Effective field theories with HNL 21
2.5.1 HNL production mechanisms 22
2.5.2 HNL decays 23
2.6 Previous searches for HNL 25

2.1 Introduction

The Standard Model of particle physics (SM) is constructed from the gauge group $SU(3)_c \times SU(2)_L \times U(1)_Y$; this factorisation describes the strong nuclear force $SU(3)_c$ and the electroweak force $SU(2)_L \times U(1)_Y$ respectively. Every currently known particle is in a representation of this group (and of the Lorentz group), giving a set of quantum numbers that are unique to the particle [2]. The quantum numbers associated with this representation are termed *colour charge*, *weak isospin* and *weak hypercharge*, respectively. Apart from these three numbers, particles are characterised by their spin, which governs their behaviour under Lorentz transformations and therefore their interactions with the rest of the Standard Model.

Concisely, the SM Lagrangian may be decomposed into a series of terms [3–8]:

$$\mathcal{L} = \mathcal{L}_{\text{kin}} + \mathcal{L}_{\text{int}} + \mathcal{L}_{\text{Yukawa}} + \mathcal{L}_{\text{Higgs}}, \tag{2.1}$$

where the gauge boson propagators live in $\mathcal{L}_{\text{kin}}$, interaction terms between gauge bosons and fermions in $\mathcal{L}_{\text{int}}$, and the elusive mass terms in $\mathcal{L}_{\text{Yukawa}}$. The last term $\mathcal{L}_{\text{Higgs}}$ allows for spontaneous symmetry breaking of the electroweak $SU(2)_L \times U(1)_Y$ group to the weak and electromagnetic forces, giving rise to the masses of the

weak gauge bosons. Specifically, the terms in $\mathcal{L}_{\text{Yukawa}}$ (the mass terms in the theory) contain couplings which are not *a priori* predicted by the theory, but must be experimentally measured.

During the development of the Standard Model, experimental observations led to the formulation of several central features which the SM was constructed to incorporate.

First, *weak CP violation*: It was established in the landmark 1957 Wu experiment [9] that the weak interaction maximally violates parity. This was a crucial conclusion which served to solve the so-called "$\theta - \tau$ puzzle"; what were apparently two distinct mesons possessed nearly identical lifetimes and masses (the *theta*, decaying into two pions, and the *tau*, decaying into three pions). Among the number of proposed solutions to this problem was Lee and Yang's original 1956 proposal that parity is not a conserved quantity [10], which implies that the θ and τ are the *same* particle (now known as the kaon). Almost contemporaneously, it was proposed in 1955 by Gell-Mann and Pais [11] that another meson, the θ^0, was not a state of definite lifetime; rather, it *mixes* with its (distinct) anti-particle $\overline{\theta}^0$. Instead, the definite-lifetime states are some fields $\theta^0_{1,2}$, which Gell-Mann and Pais related to each other through the charge-conjugation C symmetry:

$$C\theta_1 C^{-1} = \theta_1, \quad C\theta_2 C^{-1} = -\theta_2. \tag{2.2}$$

Originally, Gell-Mann and Pais hypothesised C to be conserved. After the discovery of P violation, the proposal was amended to replace C in Eq. 2.2 by CP, which was taken to be conserved [12]. This symmetry too would, however, be shown to be broken; if CP were exactly conserved, the decay $K^0_L \to \pi^+\pi^-$ would be forbidden. Instead, this decay was measured by Cronin *et al* in 1964 [13] to occur with a branching ratio (BR) of $(2.0 \pm 0.4) \times 10^{-3}$. This is interpreted as the K^0_L state being *not a definite CP state*: one says that CP is *indirectly* violated in the quark sector.

Secondly, *quark mixing*: In the quark sector, mixing was pioneered by Cabibbo [14] to relate the rates of weak leptonic decays of different mesons such that the weak coupling is universal. Soon after, the three-quark model of the Cabbibo mechanism

was extended by Glashow, Iliopoulos, and Maiani [15] to include a fourth quark. This was later generalised by Kobayashi and Maskawa [16], who predicted that a third generation of quarks would account for CP-violation in the mixing of quarks directly. The prescription for quark mixing can be summarised as:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V_{\text{CKM}} \begin{pmatrix} d \\ s \\ b \end{pmatrix}, \tag{2.3}$$

where V_{CKM} is a unitary 3×3 matrix and $d(d'), s(s'), b(b')$ are the mass (weak-interaction) eigenstates of the down-type quarks. When one enumerates all the degrees of freedom of this matrix, five of the nine parameters may be absorbed as phases of the quark fields and are unphysical; this leaves four parameters - three mixing angles and a CP-violating phase. A further comparison of the ratios of the decay rates $\eta_{+-} = \frac{\Gamma\left(K_L^0 \to \pi^+\pi^-\right)}{\Gamma\left(K_S^0 \to \pi^+\pi^-\right)}$ and $\eta_{00} = \frac{\Gamma\left(K_L^0 \to \pi^0\pi^0\right)}{\Gamma\left(K_S^0 \to \pi^0\pi^0\right)}$ showed that this phase must be non-zero in the quark sector, as $\eta_{+-} \neq \eta_{00}$; this was confirmed in B meson decays as well [17]. One says that CP is *directly* violated in the quark sector.

Thirdly, *chiral structure*: A quantum field ϕ may be decomposed into *chiral components* $\phi_{L,R}$, which are eigenstates of the projection operators

$$P_{L,R} = \frac{1}{2}\left(1 \mp \gamma^5\right), \tag{2.4}$$

with $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$ the product of the Dirac gamma matrices. Left-chiral components of fields interact with the SM weak force, while right-chiral components do not [18]. Neutrinos have no colour or electromagnetic charge, which means that *right-chiral neutrinos are singlets under the SM*. Table 2.1 lists the SM fermionic particle content and their charges under the SM gauge group.

Fourthly, *a two-component neutrino*: One must notice the absence of right-chiral neutrinos ν_R from the SM particle content. This omission is due to the construction of the SM to describe parity violation in weak interactions.

The large parity violation reported in the Wu *et al* experiment constrained the helicity of the neutrino (the projection of the expectation value of its spin along the direction of its momentum) to be either $+1$ or -1. A 1957 experiment by Goldhaber, Grodzins, and Sunyar [19] was interpreted as proof that $h = -1$ for

Species	Colour	Isospin	Hypercharge	EM charge
$\begin{pmatrix} U_L \\ D_L \end{pmatrix}$	**3**	$\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$	$-1/3$	$\begin{pmatrix} +2/3 \\ -1/3 \end{pmatrix}$
U_R	**3**	0	$-4/3$	$2/3$
D_R	**3**	0	$2/3$	$-1/3$
$\begin{pmatrix} \nu_L \\ \ell_L \end{pmatrix}$	**1**	$\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}$	1	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$
ℓ_R	**1**	$-1/2$	1	-1

Table 2.1: Standard Model fermionic particle content and charges.

neutrinos. The corollary of this statement is that the neutrino field, which is a fermion, may be described by only *two components* in its wavefunction [10, 20, 21].

Simultaneously, *the Standard Model neutrino is massless*. Because helicity is frame-dependent and chirality is not, for a massless particle the quantum numbers for helicity and chirality may be identified, which means left-chiral particles always have $h = -1$, and right-chiral particles always have $h = +1$. In the Standard Model, neutrinos are left-chiral, and antineutrinos are right-chiral.

Finally, one must distinguish between *Dirac* and *Majorana* particles. The neutrino carries no electric charge, so it is possible that it is *identical* to its own antiparticle. If the neutrino is distinct from the antineutrino, then it is a Dirac particle; if, conversely, it is identical to the antineutrino, it is a Majorana particle. Naturally, in the case of a Majorana neutrino then leptonic charge of the form $L(\nu) = +1, L(\overline{\nu}) = -1$ is not conserved; a typical probe of the neutrino's nature is the search for lepton-number-violating neutrino-less double-beta decay [22, 23]. One must also be careful with the chirality of a Majorana particle; let $\psi^C = C\overline{\psi}^T = C\gamma^0\psi^*$ be the charge-conjugated partner of a complex-valued field ψ. Then it is

$$\begin{aligned} \gamma^5\psi^C_{L,R} &= \gamma^5\left(C\gamma^0\psi^*_{L,R}\right) \\ &= -C\gamma^0\gamma^5\psi^*_{L,R} \\ &= -C\gamma^0\cdot\left(\mp\psi^*_{L,R}\right) \\ &= \pm\psi^C_{L,R}, \end{aligned} \tag{2.5}$$

where the top (bottom) sign holds for the left- (right-) chiral field. In other words, ψ_L^C is a right-chiral field, and ψ_R^C is a left-chiral field. A more thorough overview of the two-component neutrino theory may be found in [24].

2.2 The problem of neutrino mass

2.2.1 Neutrino oscillations

While the Standard Model was being constructed, a different challenge emerged, pertaining to the flux of neutrinos from the Sun. The *solar neutrino problem* [25] refers to the apparent incongruity between the theoretical expected rate of neutrino-induced charged-current (CC) interactions $\nu_e + A \to e^- + X$, based on the calculation of the emitted ν_e flux, and the measured rate of these interactions. The overall neutrino flux across all flavours α may be constrained by measuring the rate of elastic scattering off atomic electrons, $\nu_\alpha + e^- \to \nu_\alpha + e^-$ (for $\alpha = \mu, \tau$ only the neutral-current (NC) interaction is allowed in the SM, whereas for $\alpha = e$ both CC and NC channels contribute to the scattering). The observed CC neutrino-nuclear interaction rate was inconsistent with the calculation for the expected solar neutrino flux (the so-called "Standard Solar Model", or SSM) based on the observed atomic-electron elastic scattering rate. This indicated that new physics was needed to describe the neutrino.

At the same time, an "atmospheric neutrino anomaly" had been observed at the Kamiokande [26] and IMB-3 [27] detectors. Given that atmospheric neutrinos are produced primarily from the decays of pions and kaons in the atmosphere, and that these hadrons decay mostly to $\mu - \nu_\mu$ pairs (as opposed to $e - \nu_e$), and recalling that muons decay as $\mu \to e\nu_\mu\overline{\nu}_e$, one would expect the integrated flux of muon neutrinos to be about twice that of electron neutrinos from the atmosphere. However, the ratio of ν_μ-induced events over ν_e-induced ones in data was smaller than that predicted by the Monte Carlo simulations; the difference from MC was statistically significant even when systematic uncertainties due to flux, cross-section, and detector performance were considered [28]. This was another indication, in a different energy

regime, that new physics was required; both this anomaly and the solar neutrino problem could be explained in the framework of *three-flavour neutrino oscillations*.

Neutrino oscillations *in vacuo* were first proposed by Pontecorvo in a similar context to the neutral kaon mixing [29]. Pontecorvo assumed that the (anti-)neutrino $\nu, \overline{\nu}$ were mixed states, superpositions of CP eigenstates. Soon after, he pointed out that [30], in such a case, $\nu \leftrightarrow \overline{\nu}$ oscillations could alter the cross-section for the formation of argon from chlorine-37 in Davis's radiochemical experiment [31]. After the confirmation of distinct flavours of electron and muon neutrinos at Brookhaven in 1962 [32], Pontecorvo reformulated the problem of conservation of leptonic charge in terms of the *flavour* content of the neutrino, predicting that oscillations of the type $\nu_e \leftrightarrow \nu_\mu$ might be detectable in solar neutrinos. Specifically, he speculated that, if the neutrino oscillation length were small enough, then at the surface of the Earth "the flux of observable solar neutrinos would be half as large as the total flux of solar neutrinos" [33].

Three-flavour neutrino oscillations did indeed turn out to provide a consistent explanation for the aforementioned phenomena, as was shown by the Super-Kamiokande (SK) [34] and Sudbury Neutrino Observatory (SNO) [35] collaborations, as well as KamLAND [36]. Using water Cherenkov technology, SK was able to verify that there was a deficit in muon-like events as a function of zenith angle of the event, whereas there did not seem to be a similar deficit in electron-like events (see Fig. 2.1). This meant that the atmospheric anomaly could be explained in the framework of *muon disappearance* $\nu_\mu \to \not{\nu}_\mu$, or that the measured rates of muon-induced events were lower than those expected in the absence of oscillations. The fact that there was no associated *electron appearance* $\not{\nu}_e \to \nu_e$ was taken to mean that there is an effective two-flavour oscillation $\nu_\mu \to \nu_{x \neq e}$. Almost simultaneously, SNO showed that the rates of NC reactions in heavy water, $\nu_x + \mathrm{d} \to \mathrm{p} + \mathrm{n} + \nu_x$, where the neutrinos are emitted from the decay of the unstable isotope ${}^8\mathrm{B} \to {}^8\mathrm{Be} + e^+ + \nu_e$, were consistent with the prediction of the SSM. The expectation value for the non-ν_e flux from boron-8 is zero; in the absence of neutrino oscillations, one can make predictions for the rates of the CC interaction $\nu_e + \mathrm{d} \to \mathrm{p} + \mathrm{p} + e^-$, as well as

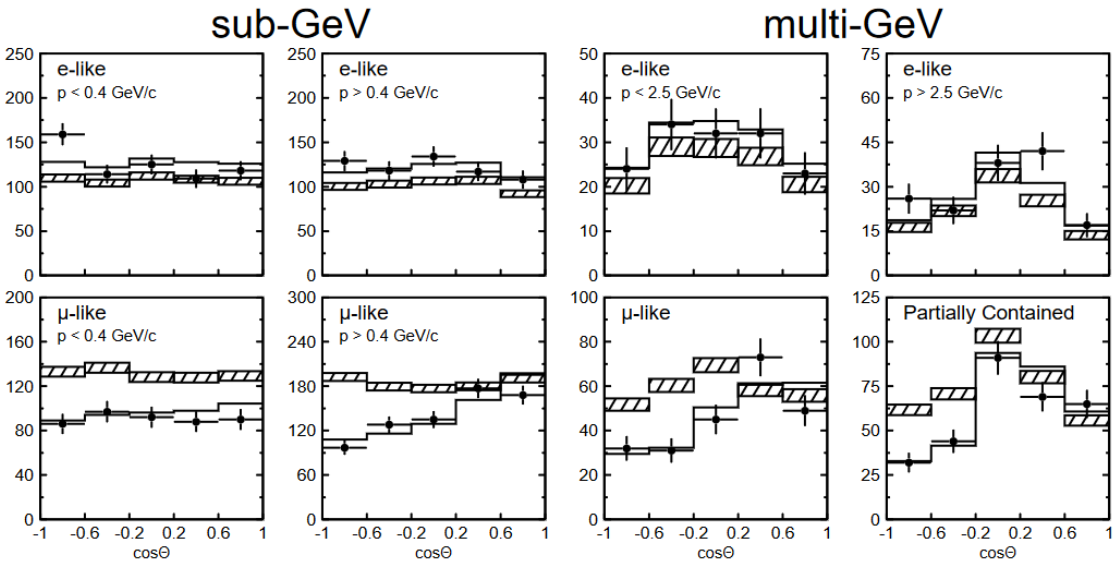


Figure 2.1: SK muon disappearance results.
Super-Kamiokande (SK) muon disappearance results: distributions of electron-like (top row) and muon-like (bottom row) events as a function of zenith angle, for single-ring fully-contained "sub-GeV" (visible energy < 1330 MeV and "multi-GeV" (visible energy ≥ 1330 MeV) samples, and for partially-contained events. The dashed regions indicate the no-oscillation expectation based on Monte Carlo, with statistical errors, normalised to the exposure of SK; the solid line is the best-fit for $\nu_\mu \to \nu_\tau$ oscillations. From [34].

the elastic scattering (ES) off of atomic electrons which proceeds both via a CC ($\nu_e + e^- \to e^- + \nu_e$) and NC ($\nu_x + e^- \to \nu_x + e^-$) component. Both the CC rate on heavy water and the ES rate were lower than one would expect based on the NC rate on heavy water, indicating that there was a deficit in electron neutrinos; this meant that there was electron neutrino disappearance $\nu_e \to \not{\nu}_e$, and that the non-ν_e flux is strongly above zero; Figure 2.2 shows how the combined rates for CC-water, NC-water, and ES establish the electron neutrino disappearance. SNO established clearly that there is *mixing* between the neutrino flavour states, but importantly, it is not neutrino oscillations at the heart of the solar electron neutrino disappearance. Instead, it is an *adiabatic conversion* of the (mostly) ν_2 state (which in the Sun, where the density is high and slowly varying, is $\simeq \nu_e$) into a mixture of ν_e and $\not{\nu}_e$ [37].

The picture was even further strengthened by the observation of $\overline{\nu}_e \to \overline{\nu}_{\mu,\tau}$ by the KamLAND collaboration [36], which measured a deficit in inverse beta-decay events $\overline{\nu}_e + p \to n + e^+$ in liquid scintillator, compared to the no-oscillation hypothesis.

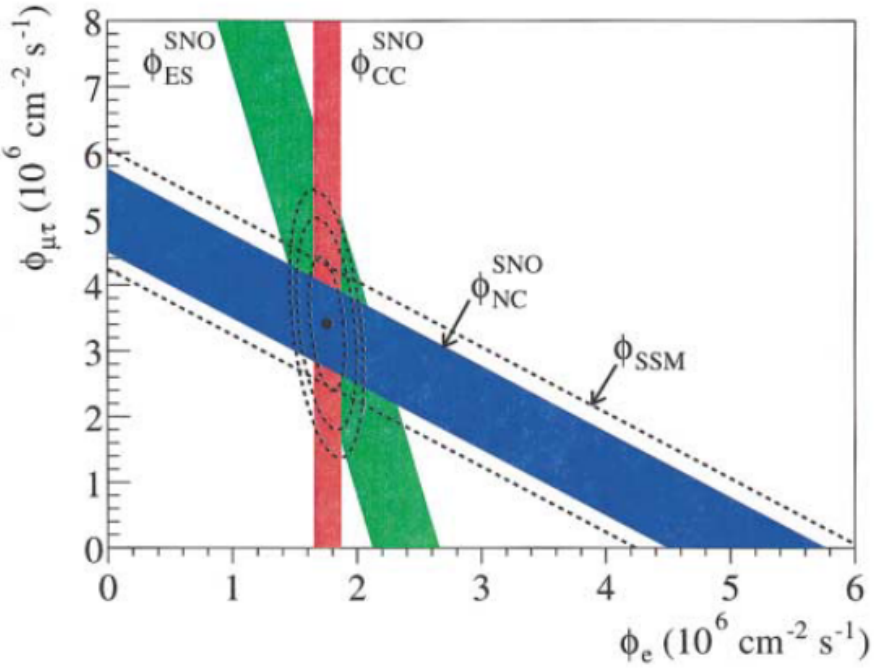


Figure 2.2: Sudbury Neutrino Observatory (SNO) electron-disappearance results: best fit of the rates, from ^{8}B neutrinos, of the CC interaction on heavy water (red), NC interaction on heavy water (blue), and elastic scattering (ES) off of atomic electrons (green) as a function of the ν_e component of the flux (horizontal axis) and the non-ν_e component (vertical axis). The SSM prediction for the overall flux normalisation is shown as the dashed lines parallel to the NC band. All three rates intersect at a point where the non-ν_e component of the flux is significantly above zero, indicating a electron neutrino disappearance. From [35].

Crucially, KamLAND showed that the probability of disappearance depended on the ratio L/E, which favoured the neutrino-oscillation explanation of the SNO electron neutrino disappearance, and provided evidence that neutrinos do indeed have non-zero mass. Since then, a plethora of experiments [38–57] have undertaken the task of measuring the neutrino oscillation parameters, which are incorporated into the standard three-neutrino mixing picture. Still more are being planned , or are in the early stages of operation [58–62]. This picture will briefly be discussed now; for more detail, the reader may consult the detailed reviews in [63–65].

In the three-flavour mixing paradigm, the neutrino flavour eigenstates ν_α are superpositions of the three mass eigenstates ν_i, with $\alpha = e, \mu, \tau$ and $i = 1, 2, 3$.

Their mixing is analogous to the quark-mixing prescription (2.3):

$$\begin{pmatrix} \nu_{e;L} \\ \nu_{\mu;L} \\ \nu_{\tau;L} \end{pmatrix} = U_{\text{PMNS}} \begin{pmatrix} \nu_{1;L} \\ \nu_{2;L} \\ \nu_{3;L} \end{pmatrix}, \tag{2.6}$$

with U_{PMNS} the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. The PMNS matrix contains four parameters of the paradigm; these are three *mixing angles* θ_{ij} and a *CP-violating (Dirac) phase* δ_{CP}. Typically, the PMNS matrix is written in the standard parametrisation:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{2.7}$$

The above parametrisation (2.7) acquires an additional diagonal matrix of phases if the neutrino is a Majorana particle; these Majorana phases do not enter oscillation probabilities. Simplifying this picture a little, the formula for two-flavour oscillations *in vacuo* depends upon the ratio of L/E:

$$P(\nu_\alpha \to \nu_\beta) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2}{4}\frac{L}{E}\right), \tag{2.8}$$

where L is the distance travelled between emission and detection points and E the neutrino energy. For the case of two flavours, $\Delta m^2 = m_2^2 - m_1^2$. In the three neutrino picture, this functional dependence on the ratio $\Delta m^2 \cdot L/E$ remains true. To first order, low-energy solar neutrinos (for which the effects of oscillations in matter can be neglected) and atmospheric neutrinos may be identified in the following way [64]:

$$\text{solar} : 1 - P(\nu_e \to \nu_e) \propto \sin^2\left(\frac{\Delta m^2_{21}}{4}\frac{L}{E}\right), \quad \text{atm} : 1 - P(\nu_\mu \to \nu_\mu) \propto \sin^2\left(\frac{\Delta m^2_{\text{atm}}}{4}\frac{L}{E}\right). \tag{2.9}$$

Current best-fit results for the mass splittings [66, 67] are $\Delta m^2_{\text{sol}} \simeq \Delta m^2_{21} = (7.4 \pm 0.2) \times 10^{-5}$ eV^2, $\Delta m^2_{\text{atm}} \simeq |\Delta m^2_{31}| = (2.50 \pm 0.03) \times 10^{-3}$ eV^2.

The fact that $\Delta m^2_{\text{atm}} \gg \Delta m^2_{\text{sol}}$ implies that there are two allowed orderings, or *hierarchies*, of neutrino masses: these are known as the normal ordering (NO, $m_{\nu 1} \simeq m_{\nu 2} \ll m_{\nu 3}$) and inverted ordering (IO, $m_{\nu 3} \ll m_{\nu 1} \simeq m_{\nu 2}$), respectively. A more detailed review on the mass hierarchy may be found in [68].

2.2.2 Absolute scale of neutrino masses

Because neutrino oscillations are only sensitive to the *difference* of neutrino masses, they cannot place constraints on the absolute mass scale, i.e. the lowest mass. Complementary measurements attempt to measure the neutrino mass scale; for a comprehensive review, see [69]. These include:

- Direct measurements of the effective "electron neutrino mass" from beta-decay spectra: These rely on the effective resolution of the emission spectrum from nuclear beta-decay from tritium to measure the "effective mass" $m_{\nu e} = \left(\sum_{i=1,2,3} U_{ei} m_i\right)^{1/2}$. A signature $m_{\nu e} \neq 0$ appears as a "kink" in the end-point of the beta spectrum, with the cutoff value being $E_{\max}^{m_{\nu e}=0} - m_{\nu e}c^2$. The current world-best limits are placed by the KATRIN experiment [70] at $m_{\nu e} < 0.45$ eV/c^2 (90% C. L.).

- Indirect measurements on the effective Majorana mass $|m_{\beta\beta}| = |\sum_i U_{ei}^2 m_i|$ via neutrinoless double-beta decay: Assuming that neutrinos are Majorana particles, they can mediate the decay $(A, Z) \to (A, Z+2) + 2e^- + 2\overline{\nu}_e$, with a corresponding line in the spectrum of the total energy of the two emitted electrons [71]. However, interpreting the limits on the rates of neutrinoless double-beta decay reactions, which is what is measured in practice, as a limit on neutrino masses is strongly dependent on the assumptions used to extract it. Most importantly, such an interpretation depends on the values of the 2 *a priori* unknown Majorana phases, which cannot be constrained by oscillation measurements as they do not enter the (dis-)appearance probabilities. Under suitable conditions, cancellations between terms containing Majorana phases can cause the Majorana mass to vanish [23]. Naturally, if neutrinos are Dirac particles, then neutrinoless double-beta decay does not occur and these limits will turn out to be meaningless.

- Cosmological constraints on the sum of neutrino masses $\sum_i m_i$ may be obtained by a number of observables, including the CMB matter power spectrum (and

its polarisation modes), the "clumpiness parameter" σ_8 of large-scale structure, and the flux of light emitted by distant quasars in the frequency range known as the "Lyman-α forest". The power, and limitations, of the cosmological probes stem from the behaviour of neutrinos created in the early Universe. For a comprehensive review, see the discussion in [72]. As the Universe grew older, it expanded and cooled, meaning the production rate of neutrinos dropped continuously; at the point where the expansion rate overtook the rate of neutrino production (about $T = 1$ MeV), the neutrinos became *decoupled* from the primordial plasma. The more massive these neutrinos, the earlier they would decouple, changing their free-streaming length. Inside length scales of this length, neutrino density fluctuations decayed rapidly [73], changing the initial conditions for large-scale structure evolution in the Universe.

Practically, the neutrino masses $\sum_i m_i$ enter as a fit parameter into multi-parameter fits to observational data, *in the context of the cosmological Standard Model* (ΛCDM). In this framework, limits on $\sum_i m_i$ are on the cusp of distinguishing between the NO and IO (assuming the ΛCDM paradigm is valid), as in the IO the sum of neutrino masses must be at least $2\Delta m_{\text{atm}} + \Delta m_{\text{sol}} \simeq 0.0982 \pm 0.0010$ eV, assuming zero mass for the lightest mass eigenstate [74]. The figure for the normal hierarchy is $\Delta m_{\text{atm}} + 2\Delta m_{\text{sol}} \simeq 0.0658 \pm 0.0005$ eV. The constraints from *Planck* 2018 data yielding $\sum_i m_i < 0.26$ eV$/c^2$ (95% C. L.) [75]. Adding in Lyman-α forest data strengthens the bound to 0.11 eV$/c^2$ at the same confidence level, and yet stronger limits (0.09 eV$/c^2$ can be placed by factoring in CMB lensing and baryon acoustic oscillation data [76]). However, these results should be viewed with caution, as modifying the underlying model assumptions could significantly change the extracted bounds [77]; furthermore, there currently exists tension between early- and late- universe measurements of the present-day Hubble expansion constant H_0, which also affects the reliability of the extracted bounds [78].

Regardless of the exact bounds on the neutrino masses, the underlying message is clear: *neutrinos posess a non-zero, albeit small, mass.* This is clearly incompatible

with the Standard Model as it was constructed, for which right-chiral neutrino (Dirac or Majorana) degrees of freedom were not included. In the next Section, there will be an overview some typical realisations of neutrino mass generation, and more detail on the mechanism understood to be valid throughout the rest of this thesis will be covered.

2.3 Mass generation mechanisms

It is clear that the $\mathcal{L}_{\text{Yukawa}}$ term needs to be completed with additional fields in order to yield neutrino masses, but unfortunately little else is straightforward. It is not known how massive these new degrees of freedom should be, how precisely they couple to the Standard Model, or even whether they are necessarily new SM *singlet* degrees of freedom (as opposed to, for instance, a new scalar $SU(2)$ triplet that mediates neutrino masses). Generally, the most popular mechanism which generates masses for the three light (or *active*) neutrinos ν_i involves a new, unknown, physics scale which is much higher than the masses of the active neutrinos. Due to this fact, it is termed a *seesaw mechanism*: generic implementations of such seesaws seek to suppress the active neutrino mass by introducing suitably heavy new degrees of freedom, with suitable chosen Yukawa couplings such that they are consistent with their non-observation so far. It should be mentioned that seesaw mechanisms are not the only way to generate neutrino masses; more exotic alternatives are discussed in [79].

Broadly speaking, there exist three classes of seesaw, duly named Types I, II, and III: of these, only Type I seesaws actually involve new *fermionic singlets* (i.e. right-handed neutrinos). For more details on Types II and III, one may refer to [80]. In Type I seesaws, one introduces n new degrees of freedom $N_{I;R}, I = 1, \ldots, n$ that are singlets under the SM gauge group; they only couple with the Standard Model through the $\mathcal{L}_{\text{Yukawa}}$ term. There are several kinds of variations to Type I seesaws, a recap of which can be found in [81].

The most general Lagrangian that contains right-chiral degrees of freedom (without yet specifying their nature) is [82]

$$\mathcal{L} \supset i\overline{N}_R \gamma^\mu D_\mu N_R - \underbrace{\overline{\mathbb{L}_L} \cdot Y \cdot N_R \widetilde{\Phi}}_{\text{Dirac}} - \frac{1}{2} \underbrace{\overline{N_R^C} \cdot \mathcal{M} \cdot N_R}_{\text{Majorana}}, \tag{2.10}$$

where the matrices $Y, \mathcal{M}$ are Yukawa couplings and a Majorana mass term, respectively; $\mathbb{L}_L$ is short-hand for the $SU(2)_L \times U(1)_Y$ lepton doublet, and Φ is the Higgs. The second term in Eq. (2.10) gives rise after electroweak spontaneous symmetry breaking to Dirac mass terms, and the third term to Majorana mass terms [83, 84].

One note before proceeding: the careful reader may wonder if it would be possible to generate the active neutrino masses *without* invoking new degrees of freedom. Strictly speaking, this is not possible in the context of the Standard Model, *if the SM gauge group is exact*: in this case, lepton number conservation is an exact symmetry – there exists no gauge-invariant term which can combine ν_L and ν_L^C to make a neutrino mass. However, it was realised by Weinberg [85] that, if the SM were treated as an effective field theory (i.e. were not required to be renormalisable), then it is possible to generate Majorana neutrino masses involving only left-chiral neutrinos as lepton number conservation is not an exact symmetry. However, such an undertaking requires either the "integrating out" of heavy right-chiral neutrinos (shifting the problem instead to that of a suitably large new-physics scale) in the Type-I seesaw scenario [86], or the adoption of new Higgs degrees of freedom (a scalar triplet) such as in the Type-II scenario [87]. Thus, the answer is that, as seen before, it is *impossible* to generate active neutrino masses in the context of the Standard Model without invoking new degrees of freedom.

The full Lagrangian (2.10) after electroweak symmetry breaking may be written as a $(3+n) \times (3+n)$ matrix term [82, 83, 87],

$$\mathcal{L} \supset -\frac{1}{2} \begin{pmatrix} \overline{\nu_L} & \overline{\nu_R^C} \end{pmatrix} \mathcal{M} \begin{pmatrix} \nu_L^C \\ \nu_R \end{pmatrix} + \text{h. c.}, \tag{2.11a}$$

$$\mathcal{M} = \begin{pmatrix} m_L & M_D \\ M_D^T & \mathcal{M} \end{pmatrix}. \tag{2.11b}$$

The matrices $m_L, M_D, \mathcal{M}$ are 3×3, $3\times n$ and $n\times n$ respectively. Diagonalising the full mass matrix (2.11b) requires a basis change from flavour to mass eigenstates for the $3+n$ neutrinos, which means a $(3+n)\times(3+n)$ matrix $\mathcal{U}$ is required.

$\mathcal{U}$ is, of course, unitary; one wonders what its relation with the usual PMNS matrix U_{PMNS} is. Generally, one can decompose the full mixing matrix as [88]

$$\mathcal{U} = \begin{pmatrix} N & S \\ V & T \end{pmatrix}, \tag{2.12a}$$

$$N = N^{\text{NP}} U_{\text{PMNS}} = \begin{pmatrix} \alpha_{11} & 0 & 0 \\ \alpha_{21} & \alpha_{22} & 0 \\ \alpha_{31} & \alpha_{32} & \alpha_{33} \end{pmatrix} U_{\text{PMNS}}. \tag{2.12b}$$

The matrix N^{NP} gives rise to *nonunitarity* of the observed leptonic mixing matrix [89–92]. In principle, it allows for the probing of non-unitarity effects in neutrino oscillations, in a plethora of contexts [92–95]; these are also an important component of the broader family of *non-standard neutrino interactions* (see [96, 97] for an overview).

Seesaw mechanisms generally do not constrain the range of the elements of m_L, M_D, or $\mathcal{M}$: depending on the choices made for these, the phenomenology of heavy neutrinos is different. The usual choice made is the *seesaw limit* [82, 83], where all the eigenvalues of $\mathcal{M}$ are far larger than all the elements of M_D. In this limit, one may set $m_L \simeq 0$ in Eq. (2.11b), and the lepton mixing matrix $\mathcal{U}$ can be written as

$$\mathcal{U} = \begin{pmatrix} U & \Theta \\ \Theta^T & U_S \end{pmatrix}, \tag{2.13}$$

where $U \simeq U_{\text{PMNS}}$ (up to factors $\simeq \Theta\Theta^T$). As an aside, the reader should note that the eigenvalues of m_L (i.e. the *observed* active neutrino masses) are naturally small for a large range of Yukawa couplings Y, as they are suppressed by the square of the eigenvalues of $\mathcal{M}$; the seesaw mechanism solves both *existence* and *smallness* of neutrino masses! The matrix Θ is referred to as the *active-sterile mixing matrix*. However, its elements (the small active-sterile mixing angles) are not what is measured practically in experiments: rather, the experimentally relevant quantity is the mixing [82]

$$|U_\alpha|^2 = \sum_{I=1,\ldots,n} \Theta_{\alpha I}\Theta^*_{\alpha I}, \quad \alpha = e,\ \mu,\ \tau. \tag{2.14}$$

A priori, the masses of the heavy sterile neutrinos are unconstrained. Among the most popular realisations of the seesaw model is the class of *low-scale seesaws*, which will be reviewed in some brevity now.

2.3.1 Low-scale seesaws

This class of models incorporates new mass eigenstates in the mass range probable by Earth-based experiments, $\mathcal{O}\,(0.1\ \mathrm{GeV}/c^2 - 1\ \mathrm{TeV}/c^2)$. Apart from the obvious advantage that these are scenarios that can be searched for experimentally, there are a number of theoretical considerations that motivate low-scale seesaws as well [98, 99]. In particular, if neutrino masses arise at a scale similar to the electroweak scale, one has to tackle the "fine-tunning" problem of sub-eV active neutrino masses [100], if one wants to have Yukawa couplings that are not tiny. This can be done, if the sterile neutrinos can be embedded in fairly simple "left-right symmetric" extensions to the SM gauge group; the total baryon minus lepton $(B - L)$ number becomes an *approximately*, rather than *exactly*, conserved symmetry; and other considerations. Specific realisations of these include the inverse [91, 100, 101] and linear [102–104] seesaws, TeV-scale seesaws [105, 106], as well as the *neutrino Minimal Standard Model (νMSM)* [98, 107–111].

2.3.2 The νMSM

The νMSM in particular offers more than a testable theory of neutrino mass generation; it also offers a (warm) dark matter candidate, and a mechanism for generating the matter-antimatter asymmetry in the universe. In the νMSM, three new mass eigenstates are introduced; one has mass $\mathcal{O}\,(1 - 10\ \mathrm{keV}/c^2)$, and the other two are nearly degenerate, $\Delta M \ll \overline{M}$ (where $\overline{M}$ is the average of these two states). The model provides an elegant pathway to explain multiple observational shortcomings of the SM:

1. Active neutrino masses: The νMSM is a seesaw mechanism whose heavier neutrino states suppress the active neutrino masses. At least two extra

neutrino states are required to explain the two observed mass splittings $\Delta m^2_{\text{sol}}, \Delta m^2_{\text{atm}}$ [112]; the νMSM incorporates three.

2. Warm dark matter: The lightest new mass eigenstate is an $\mathcal{O}\left(1 - 10 \text{ keV}/c^2\right)$ neutrino which mixes weakly enough with the SM to decouple from the thermal plasma of the early Universe early enough to affect its evolution. Such an argument was first presented in detail by Tremaine and Gunn [113], and excludes sub-keV scale neutrinos as dark matter candidates. At the same time, upper bounds on the mass can be obtained by the analysis of power spectra from large-scale structure (see [99, 114–116] and references therein).

3. Matter-antimatter asymmetry: The heavier two mass eigenstates of the νMSM induce violation of lepton number conservation in the following way [107] - sterile states are created out of thermal equilibrium in an (almost) CP-invariant state, and mix with active states where CP violating effects lead to effective lepton (L) number generation (a process called *leptogenesis* [117]). The smallness of the active-sterile mixing implies that the density of sterile neutrinos is small in the early Universe, and they remain out of thermal equilibrium for temperatures where sphaleron processes dominate. Because sphalerons violate the baryon (B) and L numbers individually, but preserve $B - L$, a lepton number asymmetry is transferred to the baryon sector [118]. Therefore, the Sakharov conditions [119] (violation of L, out-of-equilibrium physics, and CP violation) are all met, and the νMSM naturally provides an explanation for matter-antimatter asymmetry. See [115, 120] and references therein for a detailed description of this mechanism.

2.4 The single-HNL model

The ν MSM prediction may be expressed as a "2+1" mechanism, with two nearly degenerate states contributing to the smallness of the active neutrino masses and one lighter state acting as a DM candidate. As it has been seen, it is necessary to have at least two Majorana mass eigenstates in order to source the active

neutrino masses; one must also relegate lepton number conservation to, at best, an *accidental* symmetry of the SM (as requiring lepton-number conservation means neutrinos cannot have Majorana masses) [105, 121]. Experimentally, one searches for effective couplings to the SM as in Eq. (2.14). It is worth to give an overview of the details here; crucially, *searches for HNL at experiments may assume there to be only one relevant new state.*

In order to establish successful leptogenesis, the mass splitting $\Delta M = |M_4 - M_5|$ between the two heavier HNL must be small, $\Delta M \ll \overline{M} = (M_4 + M_5)/2$ [107]. This opens up the possibility of *oscillations* between the flavour states that compose N_4 and N_5, the precise phases of which are not directly measurable at this point. However, depending on the size of ΔM and the proper time over which oscillations can occur, two heavy Majorana neutrino states may mix into an effective *quasi-Dirac pair* (or *pseudo-Dirac pair*), where lepton-number-violating effects are suppressed [122–124]. Effectively, this means experiments are sensitive to the superposition of final-state channels with practically identical kinematics, such as a muon-pion pair with a high invariant mass; the ΔM in that case would be much smaller than the peak of the invariant mass distribution [98].

In the framework of Eq. (2.14), the quantity

$$|U_\alpha|^2 \mapsto \left|U^{(\alpha)}\right|^2 = \sum_{\substack{\beta=e,\,\mu,\,\tau \\ I=1,\ldots,n}} \Theta_{\beta I}\Theta^*_{\beta I} \tag{2.15}$$

is reinterpreted as the *total* mixing with all the (*a priori* undetermined) flavours in the multiple mass eigenstates $\{N_{3+I}\}$, and α merely becomes a *label* to be used in the recasting of experimental limits. Differently put, from an experimental point of view, α carries physical meaning and the mass splitting does not: one effective mass state N_4 mixes by $\{U_\alpha\}$ with the various left-chiral, active neutrino flavour eigenstates. From a theoretical point of view, α carries no physical meaning, but the mass splitting does. Multiple interfereing mass states $\{N_{3+I}\}$ mix in some unknown way with the active flavour eigenstates, but the magnitude of the total mixing (constrained under some set of assumptions labelled by α) $\left|U^{(\alpha)}\right|^2$ is known. Re-interpretations of limits may thus be carried out *post hoc* [125–

127], so long as sufficient information about the original analyses is kept. In particular, signal selection efficiencies and full systematic uncertainty matrices are crucial for the accurate re-analysis of older searches, which may or may not be performed by researchers with access to the original simulation framework of the experiments they are recasting.

2.5 Effective field theories with HNL

The interactions of HNL with the Standard Model particles are typically feeble, due to the small mixings with the active flavour states. The precise details of the HNL interaction depend on the energy scale being considered: in the context of accelerator neutrino beamlines, at proton energies $E_p \sim \mathcal{O}\,(1 - 100\ \text{GeV})$, the scale is below the electroweak scale and heavy degrees of freedom such as the weak gauge bosons $W^{\pm}, Z$ may be "integrated out" to yield effective current operators and meson form factors [128, 129]. In practice, one uses an effective field theory to describe the interactions of low-mass SM particles with HNL, with the vertex touching an HNL leg being given by appropriate Feynman rules.

It is worth mentioning that, apart from the pure kinematic mixing, there are field theories which introduce a transition magnetic moment for HNL. These models are known as "dipole portals", because of the dominance of the dipole term $d_\alpha \cdot \overline{\nu}_{\alpha;L}\sigma_{\mu\nu}F^{\mu\nu}N_R +$ h. c. (d_α being the transition dipole moment) compared to the terms in Eq. (2.10) [130–133]. Such "dipole portals" could potentially mimic signatures causing excesses in the distribution of low-energy electron (or positron) events at LSND [46] and at MiniBooNE [134] via the radiative decay $N_4 \to \nu\gamma$ [135–137], though recent MicroBooNE analyses [138–140] disfavour (but do not yet entirely disprove, see [141]) this excess. In this work, the focus will not be on this portal, but the interested reader may refer to the above literature and references therein.

A detailed discussion of production and decay mechanisms of electroweak-scale HNL in the accelerator-neutrino context will be postponed until Chapter 5; here, a brief overview of these mechanisms will be given.

2.5.1 HNL production mechanisms

Any interaction involving a flavour eigenstate can produce HNL, provided it is kinematically accessible. In particular, one must distinguish interactions where the HNL is a real on-shell particle from those short-range interactions where it is virtual. The production mechanisms are summarised in three main categories [128]:

1. Parent particle decay [129, 142–145]: Unstable particles such as $\pi^{\pm}, K^{\pm}, D^{\pm}$ decay (semi-)leptonically. The kinematically accessible range depends on the HNL mass. Usually, this is accessible at neutrino beamlines and at atmospheric neutrino expeirments. A summary of the most relevant decay modes that produce SM neutrinos can be found in Table 2.2: one can promote a SM neutrino ν_α to an HNL N_4, if the HNL mass is below the kinematic threshold, and apply an appropriate rescaling of the SM branching ratio $\mathcal{B}$ to obtain the production rate for the HNL. This point will be revisited in Chapter 5.

2. Upscattering in the Earth / in a detector [133, 145, 146]: HNL are produced in the interaction $\nu_\alpha + A \to N_4 + X$, A, X being the struck nucleus and an *a priori* undetermined final state. Both coherent and incoherent processes can contribute to neutrino scattering and thus to HNL upscattering as well. Because HNL can be produced inside the same detector where they decay, a much different parameter space can be explored with this mechanism. Typically, this can be probed at neutrino beamlines as well as atmospheric and astrophysical neutrino telescopes.

3. Production in parton scattering processes [147–151]. This direct production typically produces either long-lived HNL such as may travel downstream to a detector and decay in flight, or short-lived HNL where displaced-vertex searches can look for the HNL decay products. These are generally subdominant for HNL with masses below $\mathcal{O}\,(5\ \mathrm{GeV}/c^2)$, making these most applicable for collider HNL searches.

4. Active-sterile neutrino oscillations [152]: These are generally sensitive to a mass-squared splitting $\Delta m^2 \sim \mathcal{O}\left(1 \text{ eV}^2/c^4\right)$, and therefore will not be considered in this work.

2.5.2 HNL decays

HNL are, unlike SM neutrinos, unstable particles. They decay to visible final states with rates that are controlled by their mixing with the active flavour eigenstates, but it is not as straightforward as with the production mechanism to write down the appropriate decay rates. There are two reasons for this.

The first is that there are already constraints on the parameter space of HNL due to the finite HNL lifetime, set by the abundances of primordial elements from Big Bang Nucleosynthesis (BBN) [153, 154]. HNL have abundances in the early Unvierse set by the scale of their mixing to the SM and their masses. Their lifetime also depends on these parameters. If the HNL are too long-lived, and there are too many of them, then their decay products (which are injected into a plasma that has reached thermal equilibrium) both heat up the plasma, accelerating the cosmic expansion rate, and mediate proton $\leftrightarrow$ neutron conversion. Both of these effects would alter the primordial abundance of light elements. Given that these abundances can be measured, upper limits on the HNL lifetime may be set; currently, the lifetime for HNL between the pion and kaon masses is set to be $\tau \lesssim \mathcal{O}\left(0.02 \text{ s}\right)$. In a realistic detector simulation, the probability of the HNL decaying inside, not downstream of, the detector must be properly estimated; it is therefore important that constraints on the lifetime of the HNL can be treated at the simulation stage.

Secondly, depending on the nature of the HNL, charge-conjugated Feynman diagrams may end up contributing to the same visible final state (in a charge-blind analysis). For many diagrams, this is tantamount to multiplying the effective decay rate by 2 for a Majorana HNL; but this is not always the case. Additionally, because of the HNL's intrinsic angular momentum, the final state decay products are expected to have an angular distribution in the HNL rest frame. Supposing HNL were detected and one knew their mass, one could then examine the angular distribution

Decay channel	$\mathcal{B}_{\Sigma\nu}$	$\mathcal{B}$	$\mathcal{B}/\mathcal{B}_{\Sigma\nu}$	Threshold (MeV/c^2)
$\mu^\mp \to \nu_\mu + \nu_e + e^\mp$	1	1	1	105.15
$\pi^\pm \to \nu_\alpha + \ell_\alpha^\pm$	1			
$\to \nu_e + e^\pm$		1.23×10^{-4}	1.23×10^{-4}	139.06
$\to \nu_\mu + \mu^\pm$		0.999877	0.999877	33.91
$K^\pm \to \nu_\alpha + X$	0.7264			
$\to \nu_e + e^\pm$		1.58×10^{-5}	2.18×10^{-5}	493.17
$\to \nu_\mu + \mu^\pm$		0.6356	0.875	388.02
$\to \nu_e + e^\pm + \pi^0$		0.0507	0.0698	358.19
$\to \nu_\mu + \mu^\pm + \pi^0$		0.03352	0.04615	253.04
$K_L^0 \to \nu_\alpha + \ell_\alpha^\pm + \pi^\mp$	0.6803			
$\to \nu_e + e^\pm + \pi^\mp$		0.4055	0.596	357.53
$\to \nu_\mu + \mu^\pm + \pi^\mp$		0.2704	0.3975	252.38
$\tau^\mp \to \nu_\alpha + X$	1			
$\nu_\tau + \pi^\pm$		0.1067	0.1067	1637.29
$\nu_\tau + K^\pm$		0.006862	0.006862	1283.18
$\nu_\tau + \nu_e + e^\mp$		0.1757	0.1757	1776.35
$\nu_\tau + \nu_\mu + \mu^\mp$		0.1714	0.1714	1671.20
other modes		0.5393	0.5393	Various
$D^\pm \to \nu_\alpha + X$	0.3830			
$\to \nu_e + e^\pm$		8.8×10^{-6}	2.3×10^{-5}	1869.15
$\to \nu_\mu + \mu^\pm$		3.74×10^{-4}	9.76×10^{-4}	1764.00
$\to \nu_\tau + \tau^\pm$		0.0012	0.00313	92.80
$\to 3-\text{body } \nu_e$		0.1496	0.3906	Various
$\to 3-\text{body } \nu_\mu$		0.09781	0.2554	Various
$\to \geq 4-\text{body } \nu_e$		0.0417	0.1089	Various
$\to \geq 4-\text{body } \nu_\mu$		0.0365	0.09529	Various
$D_s^\pm \to \nu_\alpha + X$	0.1776			
$\to \nu_e + e^\pm$		8.3×10^{-5}	4.67×10^{-4}	1967.84
$\to \nu_\mu + \mu^\pm$		0.00543	0.03058	1862.69
$\to \nu_\tau + \tau^\pm$		0.0532	0.3	191.49
$\to 3-\text{body } \nu_e$		0.06417	0.3614	Various
$\to 3-\text{body } \nu_\mu$		0.054	0.3041	Various

Table 2.2: Most dominant SM neutrino production channels from decay of parent particles, taken from the "Particle Listings" section of the 2022 PDG. $\mathcal{B}_{\Sigma\nu}$ is the sum of BR of all channels with a neutrino in the final state, and $\mathcal{B}$ is the SM branching ratio for each channel. Thresholds refer to HNL production via substituting $\nu_\alpha \mapsto N_4$ in each channel.

of their decay products as a probe between Dirac and Majorana neutrinos, as for Majorana neutrinos these charge-conjugated diagrams (with opposite angular dependencies) add coherently, cancelling out any forward-backward asymmetry; for a Dirac neutrino, this does not happen [155]. Table 2.3 enumerates a list of the most relevant HNL decay channels.

Channel	Threshold (MeV/c^2)	$\Gamma_{\text{Majorana}}/\Gamma_{\text{Dirac}}$
$N_4 \to \nu\nu\nu$	0	2
$N_4 \to \nu e^{\mp} e^{\pm}$	1.02	$2\frac{\lvert U_{e4}\rvert^2+\lvert U_{\mu 4}\rvert^2}{\lvert U_{e4}\rvert^2}$
$N_4 \to \nu e^{\mp} \mu^{\pm}$	106.17	2
$N_4 \to \pi^0 \nu$	134.98	2
$N_4 \to \pi^{\pm} e^{\mp}$	140.08	2
$N_4 \to \nu \mu^{\mp} \mu^{\pm}$	211.32	$2\frac{\lvert U_{e4}\rvert^2+\lvert U_{\mu 4}\rvert^2}{\lvert U_{\mu 4}\rvert^2}$
$N_4 \to \pi^{\pm} \mu^{\mp}$	245.23	2
$N_4 \to \nu \pi^0 \pi^0$	269.55	2
$N_4 \to \pi^{\pm} \pi^0 e^{\mp}$	275.06	2
$N_4 \to \pi^{\pm} \pi^0 \mu^{\mp}$	380.21	2

Table 2.3: HNL lepton-number conserving decay modes.

2.6 Previous searches for HNL

Accompanying the significant theoretical interest in HNL realisations of the seesaw mechanism, a vigorous experimental effort to identify HNL signatures has been underway in the last few decades. In general, the searches for HNL can be classified into categories, based on whether direct or indirect detection is being attempted and based on the length scales the experiment is sensitive to.

Generally, previous searches have set constraints on the HNL parameter space, in terms of the magnitude-squared of the mixing element $|U_\alpha|^2$ and the mass M_{N4}, for certain assumptions of the hierarchy of mixing with the various flavours $\alpha = \text{e}, \mu, \tau$. On mass scales below the kaon mass, the parameter space is quite tightly constrained with mixing elements down to $\sim 10^{-8} - 10^{-10}$ ruled out at the 90% confidence level [156]; the region of higher masses, and the hypothesis where

only $|U_{\tau 4}|^2 \neq 0$, are less stringently ruled out. Figure 2.3 shows a summarised picture of the current status of limits set by HNL searches, as well as projected sensitivities for next-generation experiments, for three different mixing assumptions where the HNL couples with only one lepton generation. For a detailed review, the reader may refer to [156] and references therein.

Decay searches [157–164] attempt to reconstruct the decay products of HNL in the detector. This of course yields model-dependent bounds, but has the most discovery potential due to their ability to severely constrain the Standard Model backgrounds, by either the topology or the kinematics of the final state. Depending on the production and decay modes analysed, different combinations of the active-sterile mixing matrix are probed; moreover, depending on the assumptions of the mixing hierarchy, the flux of HNL at the detector (and therefore the distribution of final state particle kinematics) will vary.

Peak searches [165–169], on the other hand, try to identify anomalous peaks in the spectra of leptons from the decays of hadrons to neutrinos. This has the benefit of being model-independent, as any new physics signature would show up as an anomaly in these spectra.

SM precision measurements, in general, seek to establish the level of deviation from precision SM predictions, and to interpret these as HNL signals [170]. These include electroweak observables [171, 172], lepton-flavour violating decays such as $\mu \to e\gamma, \mu \to eee$ [173–175], and deviations from unitarity of the CKM matrix [170, 171] and the PMNS matrix [92].

Apart from the completed searches, a host of upcoming analyses are being planned for the near future [129, 176–184], making these searches most opportune for exploitation.

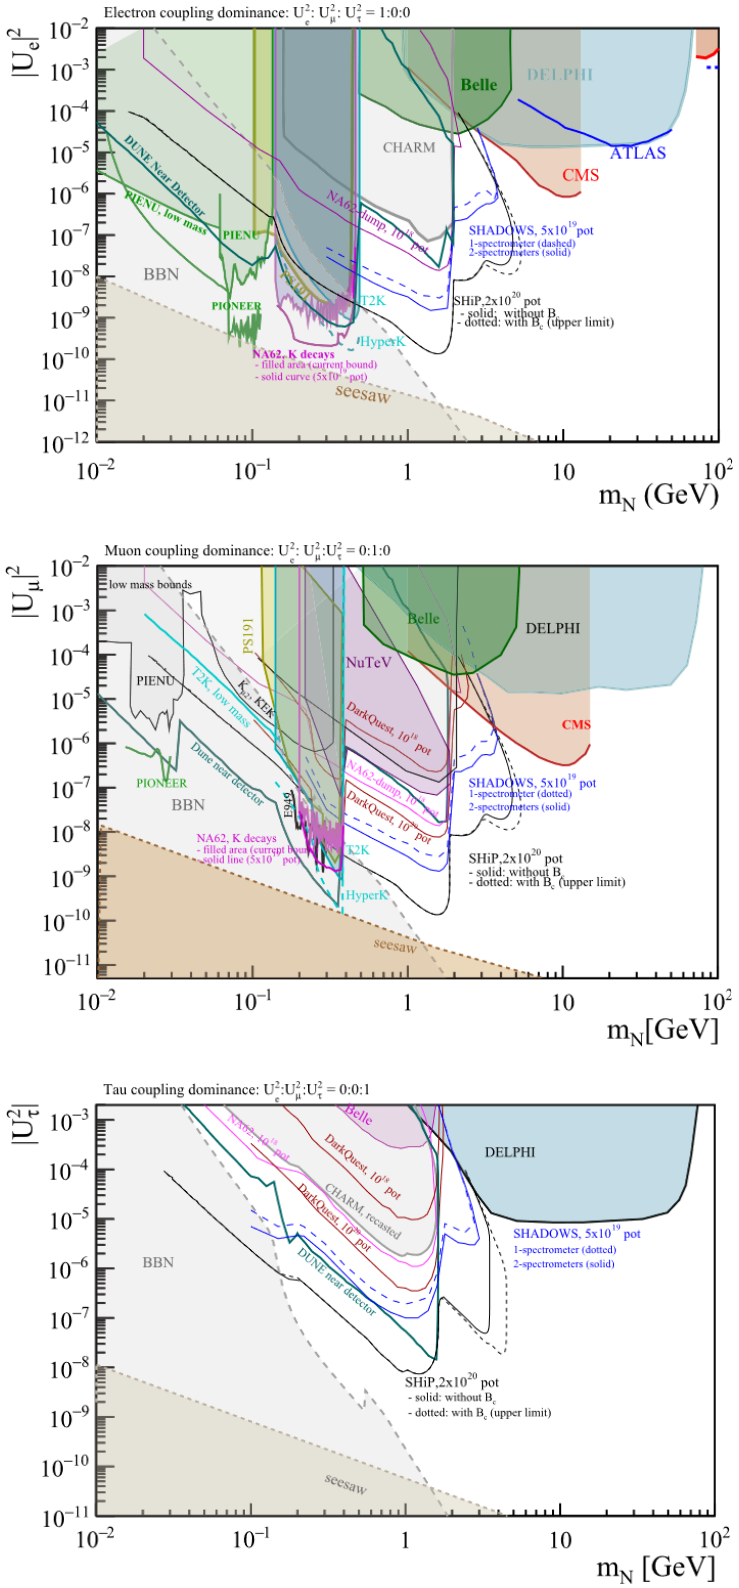


Figure 2.3: Landscape of exclusion limits (filled regions) and projected sensitivities (solid and dotted curves) in the HNL parameter space at 90% confidence level, with the mixing element $|U_{\alpha 4}|^2$ on the y-axis (top to bottom: mixing with only e, μ, τ). The BBN and seesaw curves are lower limits on HNL derived from constraints on the lifetime from Big Bang Nucleosynthesis (BBN), and from a 2-HNL seesaw scenario that can explain the observed Δm^2_{ij} from neutrino oscillations. From [156].

3 The NuMI neutrino beam

Contents

3.1 Neutrino beam operating principles 28
3.2 NuMI beamline 30
3.3 NuMI beamline simulation 34
3.3.1 Coordinate systems 36
3.4 Flux uncertainties 40
3.4.1 Hadroproduction uncertainties 41
3.4.2 Focusing uncertainties 43
3.4.3 Constraining flux uncertainties at MINERνA 44

Due to the small cross sections of neutrino interactions, oscillation and neutrino interaction experiments have need of powerful, collimated beams of neutrinos or antineutrinos to provide a high-powered flux. To this end, signficant development from the first two-neutrino experiment [32] onwards brought about the rise of secondary beam focusing using Van der Meer's magnetic horns [185], hadroproduction cross-section measurements, target optimisation and the use of the off-axis effect to create neutrino beams usable for a large range of applications; a detailed history of accelerator neutrino beams can be found in [186].

3.1 Neutrino beam operating principles

To generate neutrinos, first a beam of protons of definite momentum $\mathcal{O}\,(\mathrm{GeV}/c)$ is produced and directed onto a target that is typically a few interaction lengths long. From the primary particle interactions on the target, a chain of hadronic interactions eventually produces particles that exit the target, which decay (semi-)leptonically to neutrinos. These particles will henceforth be termed as "parents"; the entire chain of particles that led up to a neutrino being produced is called the "ancestry chain".

The target is typically constructed out of low atomic-number materials, as ancestors scatter less inside these; at the same time, the target must have sufficient thermal resistance to handle the high power of the beam. The NuMI target was

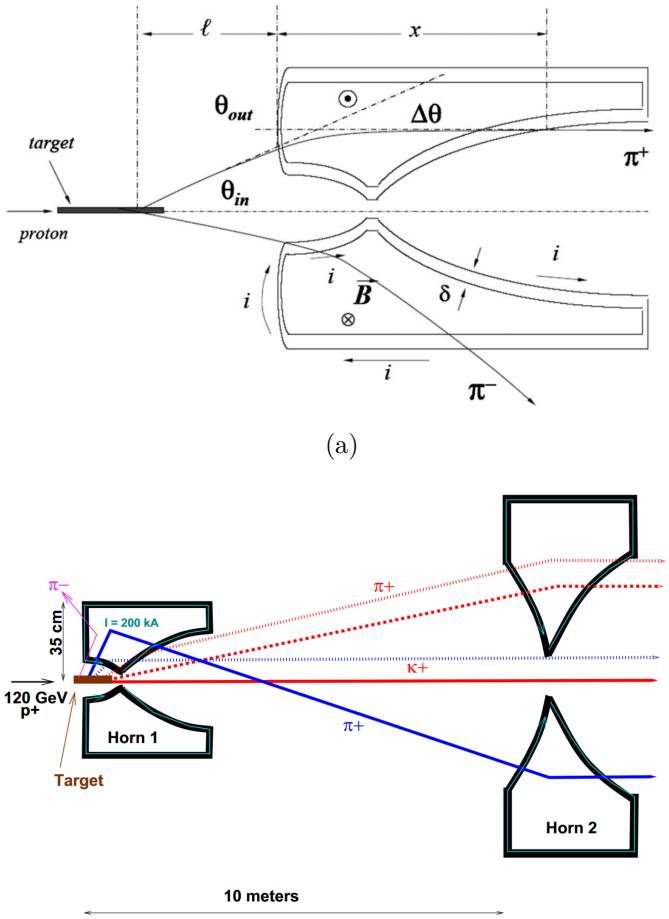


Figure 3.1: (a): Operation of a parabolic focusing horn with current i producing angular deflection of charged pions. π^+ are focused toward the proton beam axis, and π^- are focused away, producing a beam of primarily neutrinos. From [187]. (b): NuMI two-horn setup. Particles that are under- or over-focused by the first horn are corrected by the second horn. From [188].

designed so that it is separated into segments, which lets heat dissipate faster; depending on the beam power, active cooling can also be employed.

Charged hadrons that exit the target propagate forward into a series of horn magnets, which are pulsed with a current that focuses hadrons of the desired charge to be collimated with the proton beam axis, and de-focuses hadrons of the opposite charge. Because hadrons of positive charge decay to neutrinos and hadrons of negative charge decay to antineutrinos, it is possible in this way to reduce the proportion of neutrinos of the wrong kind in a beam of (anti-)neutrinos.

Figure 3.1(a) shows how a parabolic focusing horn, pulsed with a current i,

produces an angular deflection of charged pions, which in effect separates neutrino- from antineutrino- producing parents. In modern neutrino beams, additional horns are employed to correct for over- or under-focusing of particles by the first horn; an example is shown in Fig. 3.1(b). The primary neutrino-producing parents (cf. Table 2.2) that live long enough to be focused in this way are $\pi^{\pm}, K^{\pm}$. Other charged particles such as $D^{\pm}, D_s^{\pm}$ do not typically survive long enough to interact with the magnetic field and be focused.

Focused parents then propagate through an $\mathcal{O}\,(100\text{ m})$ long decay volume, which ideally is kept evacuated to prevent further scattering. Practically, a decay pipe may be filled with air or helium, which suppresses the flux by a small but tolerable amount of order $\mathcal{O}\,(\text{few }\%)$ [189]. Accelerator neutrino beamlines seek to produce $\nu_\mu(\overline{\nu}_\mu)$, but some intrinsic $\nu_e(\overline{\nu}_e)$ component is inevitable, due to either the direct production of electron (anti-)neutrinos from the decay of charged hadrons, or from the decay of secondary muons. This is minimised by a stack of material, about $\mathcal{O}\,(100\text{ m})$ long, at the end of the decay volume; the length of the decay volume may be optimised to allow higher-energy parents sufficient time to decay, at the cost of a higher electron (anti-)neutrino contamination.

3.2 NuMI beamline

The Neutrino Main Injector (NuMI) at Fermilab is a high-power beam, originally built to service the MINOS oscillation experiment and later MINOS+ [190]. Later, the off-axis NOνA oscillation experiment [53, 191] also made use of the NuMI beam. NuMI has also serviced the MINERνA neutrino interaction experiment [192], as well as the ArgoNeuT (on-axis) [193] and MicroBooNE (off-axis) [194] liquid argon experiments.

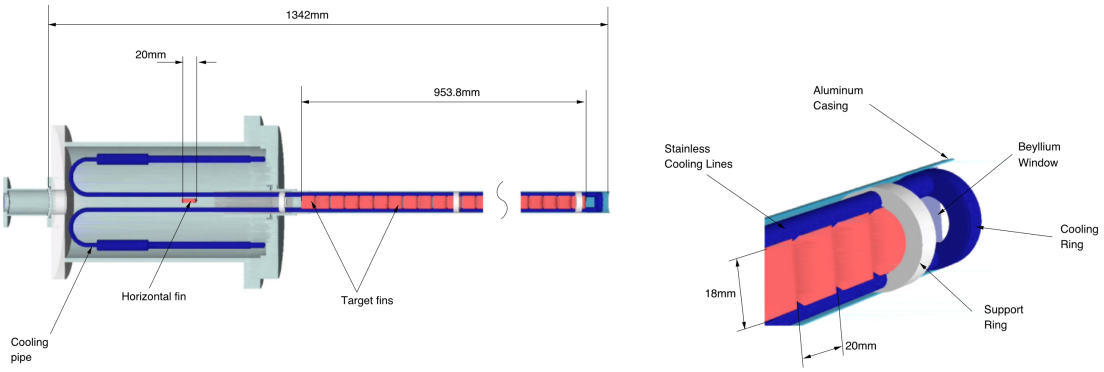
The proton beam of energy 120 GeV is produced by the Main Injector in three stages. First, H^- ions are accelerated to kinetic energy 400 MeV in the Fermilab linear accelerator. Second, they are converted into protons and accelerated to 8 GeV in the Booster, where they are placed into 1.6 μs batches with a bunch spacing of 53 MHz. Third, each batch is "injected" into the Main Injector, which

can accommodate six Booster batches (or "slots") due to its circumference being 7 times larger than the Booster's [188]. Batches are then extracted from the Main Injector and sent towards MINOS by bending them downwards by 3.343°. A set of proton beam batches accelerated together and sent towards MINOS is referred to as a "spill", which for the duration of the MINERνA run is either 5 (towards the beginning of the running period) or 6 batches making for a spill of either 8 or 10 μs.

The protons impinge on the NuMI target, which is designed to withstand a large beam power (400 kW before the NOνA upgrade, and 700 kW after), meaning large mechanical stresses, without disintegrating. It is made of graphite, and designed with small nominal dimensions so that secondary interactions of mesons from the primary p + C interactions are minimised, to maximise the neutrino yield. This necessitates its segmentation into rectangular fins, to promote fast heat dissipation. On top of that, active water cooling is employed to ensure stresses do not exceed 25.6 MPa (cf. the target's fatigue limit of 36 MPa). The target is enclosed in a helium filled container, and Budal monitors [195] upstream of the target fins scan the location of the target with respect to the beam by measuring a position-dependent signal produced mainly by knocked-out electrons from protons interacting with the monitors. The beam itself is a Gaussian profile, with the RMS horizontal and vertical radii equal at $\sigma_x = \sigma_y = 0.14$ cm.

The target, and the cooling and support lines, sit behind a baffle that protects against any missteering of the beam. It is 1.5 m long, and contains a graphite core with an 11 mm hole where the target is located, all encased within an aluminium tube. The entire setup (baffle+target) sits within the same canister, so that it moves as one object. Figure 3.2 shows the NuMI target in its canister, and Fig. 3.3 the cross-section of the baffle's inner aperture.

The NuMI horns provide focusing to the neutrino parents. Depending on the position of the target with respect to the horns, the peak of the neutrino energy spectrum shifts. During the MINERνA running period, the beam components were optimised for MINOS up to 2012, and then for NOνA from 2013 onwards. These produced neutrino fluxes at MINERνA peaked at 3.5 GeV and 5.9 GeV,



(a) (b)

Figure 3.2: The NuMI target and canister, with graphite fins and cooling, support components shown. From [188].

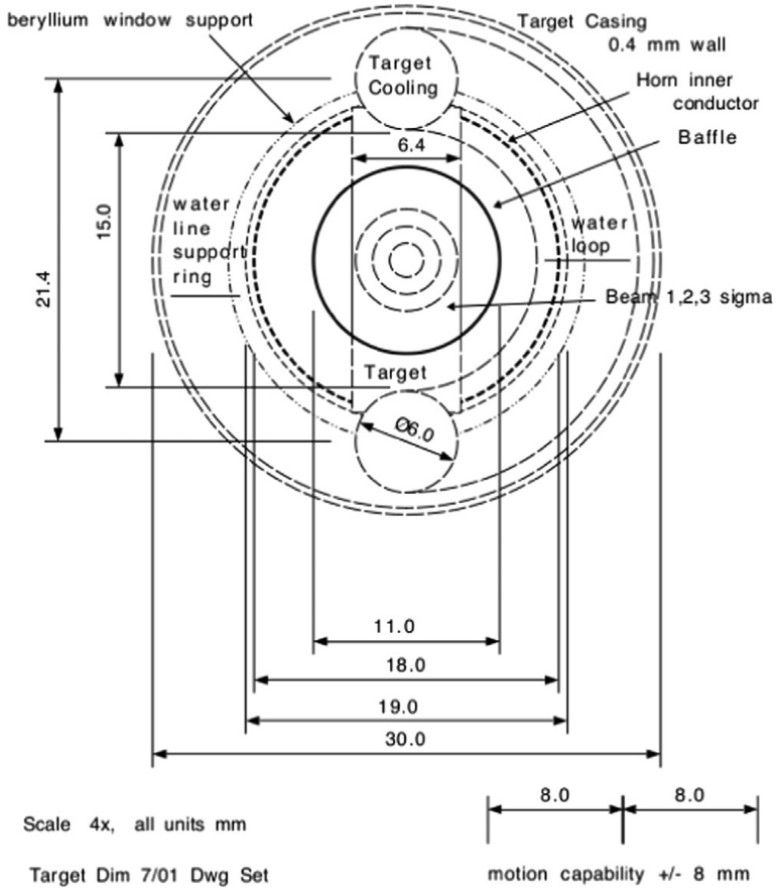


Figure 3.3: The NuMI baffle inner aperture. The beam spot 1, 2, and 3 σ are set as dashed circles in the centre of the aperture and the profile of a rectangular target fin is superimposed on top. From [188].

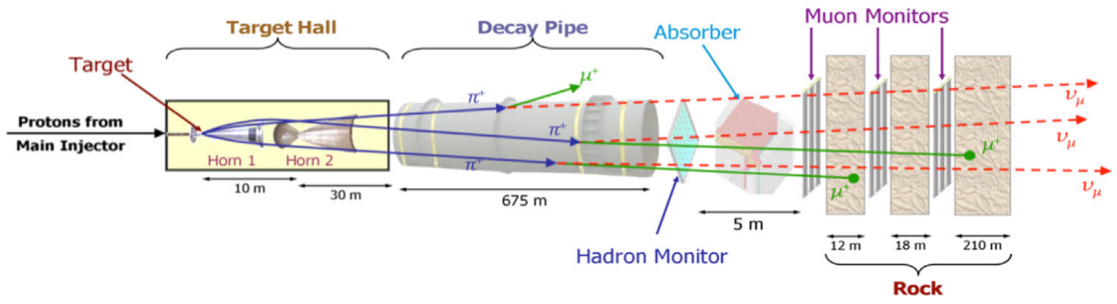


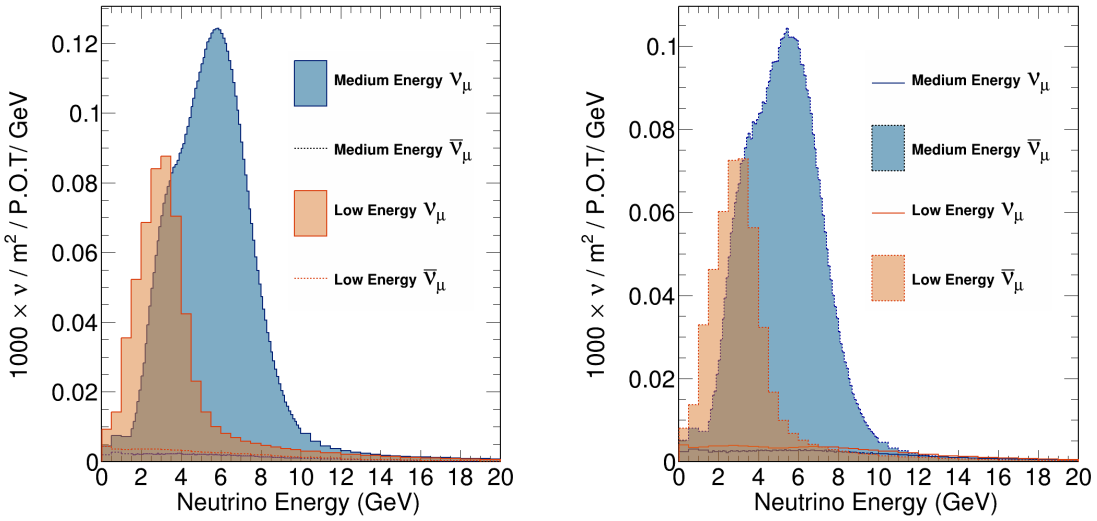
Figure 3.4: Sketch of the NuMI beam, with target hall containing the target and horns, decay pipe, and absorber / muon monitors / rock formation visible. From [188].

respectively. They are known as the Low Energy (LE) and Medium Energy (ME) runs. Fundamental beam focusing parameters are the target longitudinal position with respect to the horns, the distance between horns, and the horn current; for the selection of neutrino or antineutrino mode, the nominal horn current is pulsed at ±200 kA(185 kA) for the ME (LE) runs. The direction of current that focuses positively (negatively) charged hadrons (i.e. more (fewer) neutrino than antineutrino parents) is called the Forward (Reverse) Horn Current, or F(R)HC.

The entire target cavern is 69 m long, with about 40 m between the target (taken to be the origin in our system of coordinates) and the beginning of the decay volume. The volume is a 2 m-diameter, 675 m-long pipe which, for the duration of the ME run, was filled with low-density helium.

Downstream of the pipe, a complex of hadron absorbers, muon monitors, and surrounding dolomite rock act as a shield against almost all non-neutrino particles (muons from the interactions of neutrinos with the dolomite, or "rock muons", do in fact survive to reach MINERνA, but these are tagged and removed in the analysis). The length of the absorber is roughly 8.5 m, and the dolomite rock contributes another $\sim$ 240 m before the Near Detector cavern starts. A sketch of the NuMI beam is shown in Fig. 3.4.

Figures 3.5(a), 3.5(b) show the total flux of muon (anti)neutrinos at MINERνA for FHC and RHC. One will notice that, despite being in (anti)neutrino mode, there remains a contribution from wrong-sign neutrinos, which is caused by opposite-sign hadrons (such as underfocused pions, pions from the neck of the horns,



(a) (b)

Figure 3.5: LE and ME fluxes at MINERνA, for FHC (left) and RHC (right), after weighting flux universes. From [196].

etc.) or neutral hadrons such as K^0_L. There is also a contribution of electron (anti)neutrinos not shown in the Figure. They are, however, fully taken account of in the NuMI simulation.

3.3 NuMI beamline simulation

A detailed MINERνA GEANT4 simulation provides the backbone for the flux prediction at the detector. It is based on GEANT4, with a detailed beamline geometry developed first for MINOS and refined by MINERνA [197]. The GEANT4 version is 4.9.2.p03, with the FTFP_BERT physics list being the "nominal" flux prediction. Different physics lists are available in GEANT4, corresponding to different model parametrisations; this generates a *spread* of the predicted flux which can be constrained by external measurements. The flux uncertainties, and their constraints, will be discussed in Section 3.4.

The output of the flux simulation is a record of produced neutrinos attached to their parents, in the form of a `dk2nu` object [198]. The structure of the `dk2nu` tuple follows the ancestry of each neutrino as it is emitted in the beamline. It tracks the chain of particles (ancestors) from the primary proton interactions, with each entry in

the tuple corresponding to one neutrino produced. Each neutrino is associated with its ancestors, and information about the entire chain is stored. Crucially, the stored information includes the PDG code of the parents, the position in the beamline at which the parents produced neutrinos, the parents' momentum at the point of neutrino production, and the energy of the SM neutrino in the parent's rest frame.

Figure 3.6 shows the breakdown of saved hadron spectra as a function of hadron energy in the NuMI ME runs, FHC and RHC. The primary yield is of $\pi^+(\pi^-)$ in FHC (RHC), with a significant component of $K^+(K^-)$ at higher energies. Figure 3.9 shows the effect of the horn focusing on the parent kinematics. The top two panels show the momentum spread of $\pi^+(K^+)$ in the ME FHC run, with brighter colours indicating a higher probability density in this bin. The axes are the x and y momentum components, in the system where the beam runs parallel to the z axis. Particles of positive charge (the "correct" charge) are contained in a narrow envelope by the horn focusing, which means neutrinos produced by their decays will have higher energies due to a smaller angular deviation from their parents' momentum. The top two panels show this for π^+ and K^+. On the other hand, particles of negative charge (such as K^-) are defocused, and neutral particles (such as K^0_L) are not focused at all. The bottom two panels show the same distribution for K^- and K^0_L, which are far more spread out and therefore contribute to lower neutrino energies. Practically, though the yields of the "wrong" kind of charged particles are similar (π^+ vs. π^-, K^+ vs. K^-), focusing ensures the contamination from "wrong-kind" neutrinos is small.

From these nominal flux files, a prediction for fluxes at MINERνAcan be obtained. Figure 3.7 shows the nominal prediction broken down by neutrino flavour, and figure 3.8 shows the prediction broken down by neutrino parent (for the two neutrino flavours, and again for the antineutrino flavours, combined). It becomes apparent how the peak region (and most of the flux integrated over the entire energy range) is due to pion decay, with the high-energy tails dominated by kaon decay.

A detailed table of the `dk2nu` fields can be found in Appendix B.

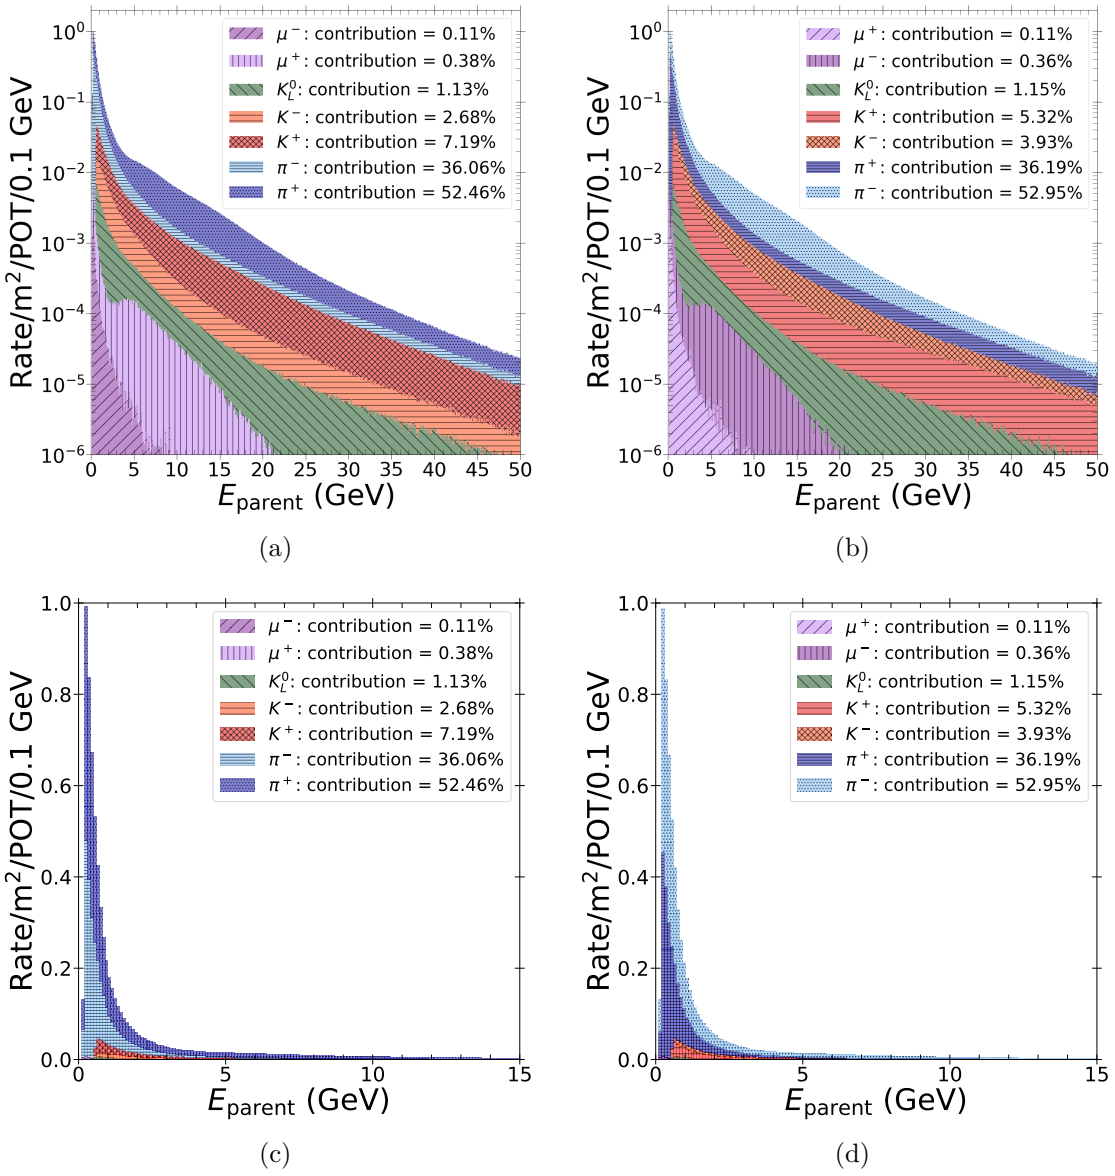


Figure 3.6: Nominal parent particle spectra at NuMI ME, FHC (a), (c) and RHC (b), (d). Top plots are log-scale, bottom are linear.

3.3.1 Coordinate systems

It is useful at this point to define some coordinate systems, and the relations between them:

- The "baseline" system of coordinates $(x_{\text{N}}, y_{\text{N}}, z_{\text{N}})$ will be referred to as "NEAR" coordinates. In this system, the origin is placed inside the target, at the same point as the beamline simulation. The z axis is perpendicular to the local gravity at the origin, and points towards the beam direction. The y axis

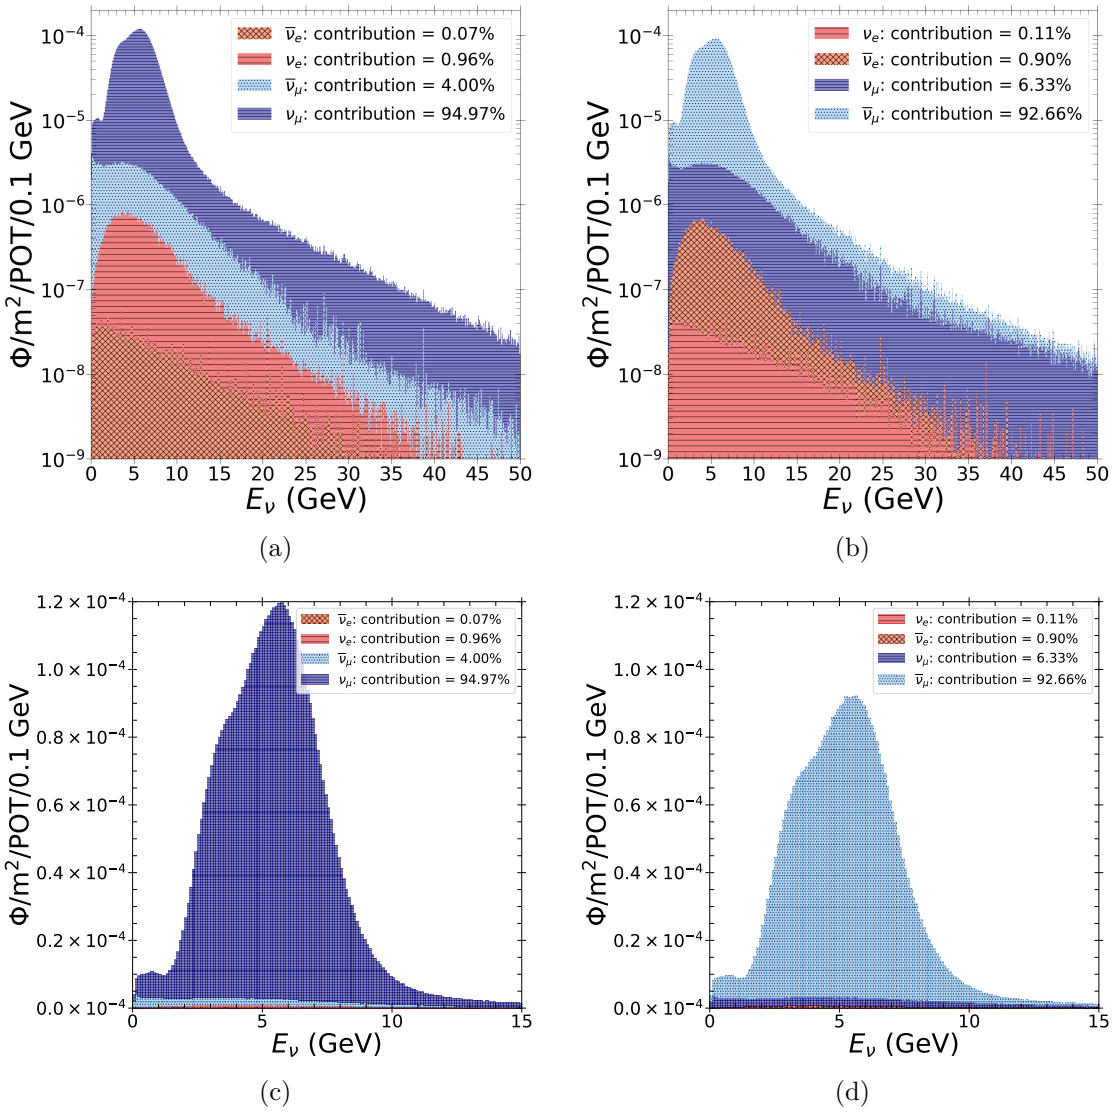


Figure 3.7: Breakdown of nominal SM neutrino flux by neutrino type at MINERνA, for FHC (a), (c) and RHC (b), (d). Top plots are log-scale, bottom are linear.

points opposite to the local gravity (that is to say, it points up), and the x axis completes the system.

- The beamline defines its own coordinate system $(x_\mathrm{B}, y_\mathrm{B}, z_\mathrm{B})$, labelled "BEAM". It shares its origin with the NEAR system, and shares the same x axis. One may obtain the BEAM coordinates from the NEAR coordinates by one rotation on the $(y_\mathrm{N}, z_\mathrm{N})$ plane by an angle $\theta = 3.343°$.

- The MINERνA coordinate system $(x_\mathrm{U}, y_\mathrm{U}, z_\mathrm{U})$ will be referred to as the

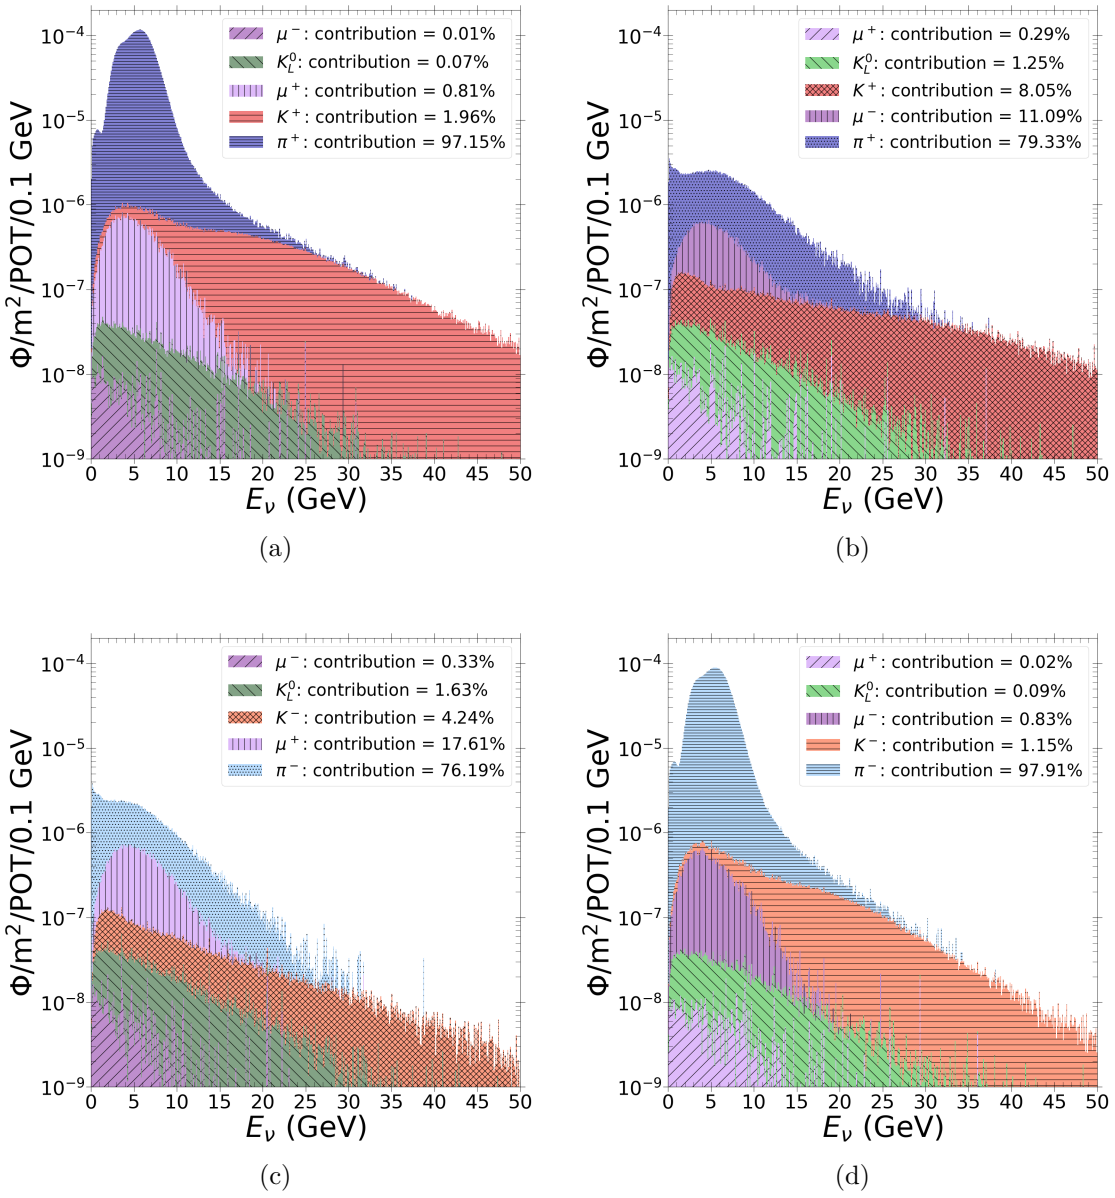


Figure 3.8: Breakdown of SM neutrino fluxes by parent at MINERνA, for $\nu_\mu + \nu_e$ (panels (a), (b)) and $\overline{\nu}_\mu + \overline{\nu}_e$ (panels (c), (d)) types. FHC on left, RHC on right.

"USER" system. It is parallel to the NEAR system, but its origin is shifted so that, in NEAR coordinates, the USER origin is at $(x_{\mathrm{U}}^{(0)}, y_{\mathrm{U}}^{(0)}, z_{\mathrm{U}}^{(0)}) = (-0.2486, -60.35, 1022.74)$ m [199, 200].

In general, one may relate two arbitrary coordinate systems between themselves by two ordered tuples: one vector of *translations*, and one tuple of *Euler angles* which specifies the relative rotation of the axes [201]. There are many conventions used in the literature; here, the *extrinsic* $X - Z - X$ *convention* will be used, in

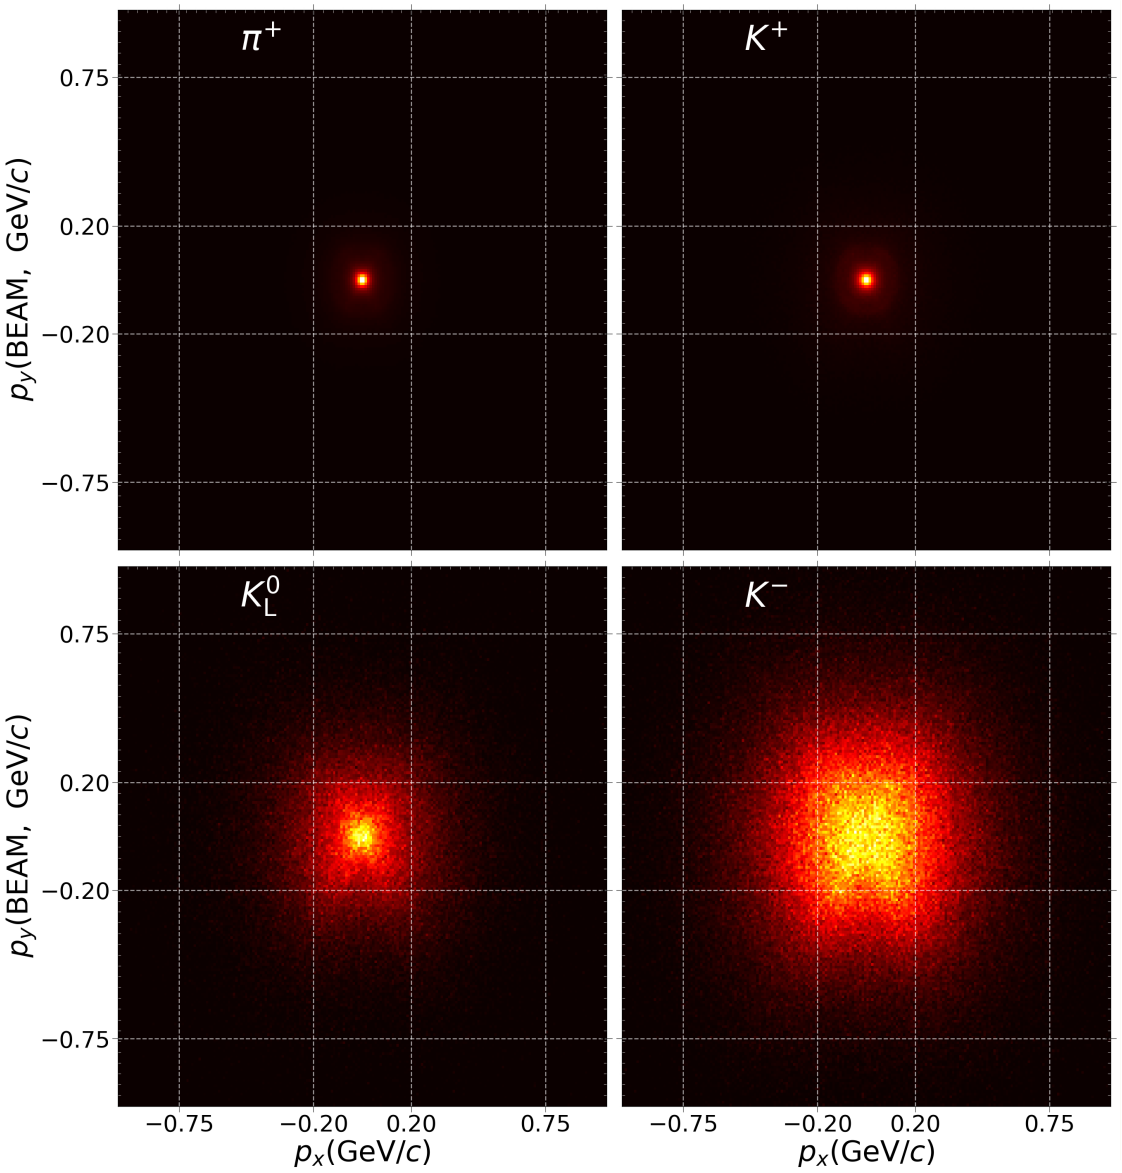


Figure 3.9: Focusing of parent mesons due to horn magnets. Clockwise from top left: π^+, K^+, K^-, K^0_L. ME, FHC run. Mesons of positive charge are focused such that $(p_x, p_y) \sim (0, 0)$ (where (x, y) are in BEAM coordinates, and z is unconstrained). Mesons of neutral or negative charge are not focused.

which one does the following operations on a system of axes: Let X, Y, Z be a global, unchanging right-handed system of axes, and let x, y, z be a local, changing right-handed system. Then one follows this series of steps:

1. Rotate (x, y, z) by an angle γ about the fixed axis X. This is equivalent to applying the rotation matrix $R_X(\gamma)$. This maps (x, y, z) to the new system (x', y', z') with $x' \equiv x$.

2. Rotate (x', y', z') by an angle β about the fixed axis Z. This is equivalent to applying the rotation matrix $R_Z(\beta)$. This maps (x', y', z') to the new system (x'', y'', z'').

3. Rotate (x'', y'', z'') by an angle α about the fixed axis X. This is equivalent to applying the rotation matrix $R_X(\alpha)$. This maps (x'', y'', z'') to the final system (x''', y''', z''').

The full rotation matrix may then be written as

$$\begin{aligned} R(\alpha,\beta,\gamma) &:= R_X(\alpha)R_Z(\beta)R_X(\gamma) = \\ &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_\alpha & -s_\alpha \\ 0 & s_\alpha & c_\alpha \end{pmatrix} \begin{pmatrix} c_\beta & -s_\beta & 0 \\ s_\beta & c_\beta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_\gamma & -s_\gamma \\ 0 & s_\gamma & c_\gamma \end{pmatrix} \\ &= \begin{pmatrix} c_\beta & -s_\beta c_\gamma & s_\beta s_\gamma \\ c_\alpha s_\beta & c_\alpha c_\beta c_\gamma - s_\alpha s_\gamma & -s_\alpha c_\gamma - c_\alpha c_\beta s_\gamma \\ s_\alpha s_\beta & s_\alpha c_\beta c_\gamma + c_\alpha s_\gamma & c_\alpha c_\gamma - s_\alpha c_\beta s_\gamma \end{pmatrix}, \end{aligned} \tag{3.1}$$

where the matrices are represented in the basis of (X, Y, Z) and $c_\theta, s_\theta \equiv \cos\theta, \sin\theta$.

3.4 Flux uncertainties

The nominal flux prediction relies on a series of steps which are subject to large systematic errors. These are encapsulated in two main groups: *hadroproduction* and *focusing*. They contain effects which must be evaluated differently, so these will be discussed separately now.

3.4.1 Hadroproduction uncertainties

Hadronic interactions in the beam lie in the regime of nonperturbative QCD, for which there exists no first-principles description. Moreover, since the hadrons from proton-target interactions must propagate in the nuclear medium before escaping, the full transport equations for particle motion are complicated (for a theoretically motivated analysis of these, see [202]). To further complicate the picture, the hadrons may also reinteract within the target before exiting, possibly producing yet more hadrons. Consequently, phenomenological models must be used in the simulation; these are then tuned to data taken by external experiments.

In the NuMI case, relevant measurements have been taken on a thick (two interaction length) NuMI replica target by MIPP at beam energy 120 GeV [203], and on a thin (few-percent interaction length) carbon target at beam energy 158 GeV by NA49 [204]. The subsequent NA61 experiment has been taking data on various targets for proton beams at 60 and 120 GeV/c, but these data have not been implemented into the MINERνA flux constraints. In the NuMI context, data from proton-target interactions with various targets at momentum 100 GeV/c is also used for the "thin-target flux" [205].

A hadron reweighting tool, `ppfx` [189, 197], was developed initially in the context of MINERνA but extended to be experiment-agnostic. It splits the external data into "thick" and "thin" target data, and considers uncertainties from the primary proton + carbon interaction modelling, pion production from nucleons produced in the primary proton + carbon interactions, meson production from nucleon interactions on non-carbon nuclei (such as the baffle), and meson interactions in the beamline itself.

The primary proton-carbon interaction is parametrised in its yields of particles as a function of transverse momentum p_{T}, and of the Feynman variable $x_{\mathrm{F}} = p_{\parallel}/p_{\parallel}^{\mathrm{max}} \simeq 2p_{\parallel}/\sqrt{s}$, where $p_{\parallel}$ is the longitudinal component of the particle's momentum, measured in the centre-of-mass frame of the proton-carbon collision and $\sqrt{s}$ is the Lorentz-invariant available energy. Feynman's scaling hypothesis [206]

was that, for inclusive high-energy hadron collisions producing a particle with mass m, the cross-section scales as

$$\frac{d^3\sigma}{d^2p_{\mathrm{T}}dx_{\mathrm{F}}} = f(p_{\mathrm{T}}, x_{\mathrm{F}}) \times \frac{\sqrt{p_{\mathrm{T}}^2 + p_{\parallel}^2 + m^2}}{E_{\mathrm{p}}^*}, \tag{3.2}$$

where $E_{\mathrm{p}}^* \simeq p_{\parallel}^{\mathrm{max}}$ is the proton's centre-of-mass energy. Manifestly, the function f is independent of $\sqrt{s}$.

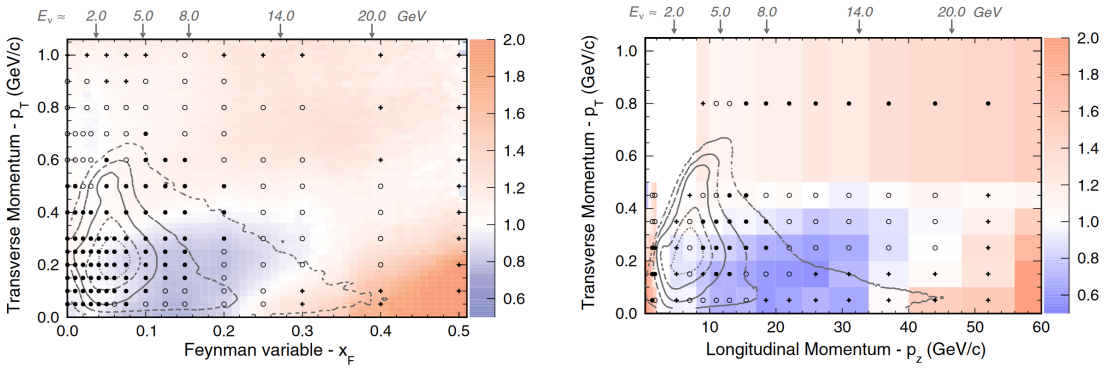
This lets one use data from different proton centre-of-mass energies such as NA49 to constrain NuMI production. In the case of NA49, run at 158 GeV, the data first need to be scaled to NuMI energies by using Feynman's x_F variable, and using FLUKA [207, 208]) to obtain this correction. Then, the Monte Carlo (MC) prediction for pion yields in each bin of $(x_{\mathrm{F}}, p_{\mathrm{T}})$ is compared to the scaled data in each bin, and a bin-by-bin weight is applied [197]. This can be written as

$$w\left(x_{\mathrm{F}}, p_{\mathrm{T}}, p\right) = \frac{F_{\mathrm{data}}\left(x_{\mathrm{F}}, p_{\mathrm{T}}\right)}{F_{\mathrm{MC}}\left(x_{\mathrm{F}}, p_{\mathrm{T}}\right)} \times \underbrace{\frac{\sigma_{\mathrm{FLUKA}}\left(x_{\mathrm{F}}, p_{\mathrm{T}}, p\right)}{\sigma_{\mathrm{FLUKA}}\left(x_{\mathrm{F}}, p_{\mathrm{T}}, p = p_0\right)}}_{g(x_{\mathrm{F}}, p_{\mathrm{T}}, p)}, \tag{3.3}$$

where $p_0 = 158$ GeV/c for the thin target data. For the thick target (MIPP) data, no such correction is necessary, and a straightforward data-MC comparison bin-by-bin is used. Figure 3.10 shows the corrective factor w as a function of p_{T} and x_{F} (or equivalently as $p_{\parallel} \equiv p_{\mathrm{z}}$).

Other corrections regarding the attenuation of the primary beam inside the target material, of the produced hadrons exiting the target, secondary interactions, or interactions with non-carbon nuclei are weighted with appropriate factors, in a similar fashion.

As for the regions where there is no relevant data, and where existing data cannot be extended usign the theoretical guidance from isospin symmetry arguments (appropriate for the isoscalar ^{12}C nucleus) or for the "material scaling" introduced in [205], the Feynman variable x_{F} is broken up into four equally sized bins, and each interaction is categorised by the combination of incident particle and produced particle. Then each of these x_{F} bins is assigned a 40% uncertainty, and is taken to be uncorrelated with the other bins, for each combination. This is motivated by the



(a) (b)

Figure 3.10: External hadroproduction data (a), from NA49 (thin target, 158 GeV/c protons), and from MIPP (b) (NuMI replica thick target, 120 GeV/c protons). Markers show the locations (bin centres) of measurements from the thin (thick) targets. Marker styles (filled, open, cross) indicate uncertainties of $< 2.5\%, < 5\%, > 5\%$ in each region. The colour scale shows the corrective ratio $F_{\text{data}}/F_{\text{MC}}$ in Eq. 3.3. Contours indicate number of $p + \text{C} \to \pi^+ X$ interactions leading to ν_μ in MINERvA LE. From [197].

observation that the underlying FTFP model is a microphysical model motivated by first principles, and which matches existing hadron production data sets to better than 40% across a broad range of kinematic ranges; thus, the "worst-case" uncertainty is taken to be 40%. This is the uncertainty applicable primarily for quasielastic $p + \text{C} \to p + X$ interactions with $x_\text{F} > 0.75$ and $\pi^+ + \text{A} \to \pi^- X$ with $x_\text{F} \in (0.25, 0.75)$ [189].

All these uncertainties are passed to "flux universes" for error propagation using the "many-universe" technique [209]. This technique will be returned to in more detail in Chapter 8. In effect, the simulation is redone multiple times in a number of "universes", where the hadroproduction parameters are each simultaneously varied by sampling from an underlying probability distribution, with its mean at the nominal value and the $1 - \sigma$ variance being the systematic uncertainty from that parameter.

3.4.2 Focusing uncertainties

Due to the high sensitivity of the flux to the beamline geometry, it is necessary to evaluate the effect of uncertainty in the beamline elements on the final flux. Like with the hadroproduction uncertainties described above, these are evaluated using the many-universe method. The various beamline uncertainties and the

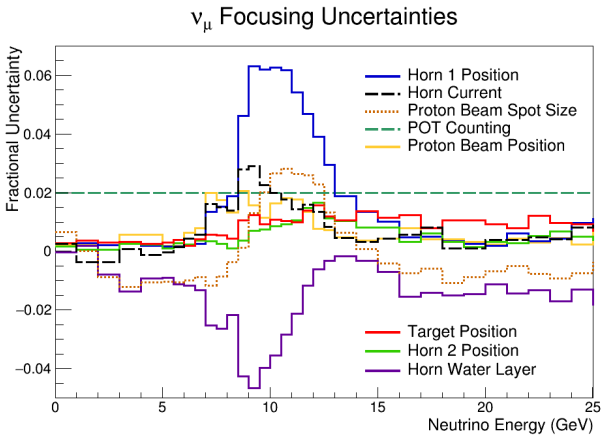


Figure 3.11: Effect of the various focusing systematic 1 σ shifts on the simulated ν_μ MINERνA ME flux. The effect of baffle scraping (0.25%) is not included here. From [210].

effect, on the flux of ν_μ at MINERνA ME, of shifting each one by 1 σ on the fractional uncertainty for each bin of neutrino energy is shown in Fig. 3.11. A list of the MINERνA ME analysis focusing parameters and their nominal values / tolerances is reproduced from [210] in Table 3.1.

In brief, they are the positions of the target and the magnetic horns, the magnitude of the horn current, the thickness of the residual water layer from the active cooling (which absorbs particles), the beam position and spot size. In addition, two flat uncertainties due to POT counting at 2% and to the beam scraping the baffle (established by measuring the baffle temperature to be 0.25%) are considered.

3.4.3 Constraining flux uncertainties at MINERνA

The *a priori* flux prediction can be further constrained by *in situ* interaction data, in a number of ways. Firstly, the *low-ν method* (where ν refers to hadronic recoil energy) exploits charged-current neutrino interactions on nuclei. The cross section of such interactions may be written approximately as [211]

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\nu} \propto \text{const} + \frac{\nu}{E_\nu} \cdot A(x_B) + \mathcal{O}\left(E_\nu^{-1}\nu/E_\nu, (\nu/E_\nu)^2\right), \tag{3.4}$$

where x_B is the Bjorken scaling variable. In the limit of low recoil $\nu/E_\nu \to 0$, the cross section is constant across ν; this relates the shape of the flux to the shape of the reconstructed event rates in the detector. MINERνA demonstrated the

Parameter	Nominal	Variance
Beam x-position (mm)	0	0.4
Beam y-position (mm)	0	0.4
Beam spot size (mm)	1.4	0.2
Horn 1 x-position (mm)	0	1
Horn 1 y-position (mm)	0	1
Horn 1 z-position (mm)	30	2
Horn 2 x-position (mm)	0	1
Horn 2 y-position (mm)	0	1
Horn water layer (mm)	1	0.5
Horn current (kA)	200	1
Target x-position (mm)	0	1
Target y-position (mm)	0	1
Target z-position (mm)	-1433	1
POT counting		2% of POT
Baffle scraping		0.25% of POT

Table 3.1: Sources of systematic uncertainties from the beamline simulation. From [210].

usefulness of this technique by identifying a "wiggle" in the ratio of flux shape to event rate, and identified its cause as originating from the muon energy scale obtained from MINOS measurements. A correction in the muon energy scale from MINOS into the MINERνA analyses by 3.6% has been incorporated into the MINERνA analyses as a result [211].

Secondly, constraints from purely leptonic interactions involving neutrinos scattering off of atomic electrons [212, 213], and from the Inverse Muon Decay (IMD) interaction $\nu_\mu + e^- \to \mu^- + \nu_e$ available over the threshold $E_{\nu_\mu} \geq (m_\mu^2 - m_e^2)/(2m_e) \simeq$ 10.9 GeV [214]. Since the matrix element involves purely leptonic tensors, the cross section is calculable and therefore the predicted rate of events straightforwardly relates to the predicted flux, and by using a Bayesian method to assign likelihoods to each flux universe (i.e. each prediction of rates at MINERνA) one can bring the total flux uncertainty down.

As can be seen in Fig. 3.8, the majority of neutrinos (especially close to the flux peak at 5.9 GeV) are produced by the decay of pions in the beamline. This presents

a unique challenge for HNL searches, since there is a very narrow mass range where HNL can be produced on-shell by pion decays. The MINERνA sensitivity, however, is sourced from the decays of kaons in the beamline, making the *in situ* flux constraints inapplicable for an HNL analysis. In this thesis, therefore, HNL flux uncertainties will not be constrained with *in situ* measurements.

4 The MINERνA detector

Contents

4.1 Description of the detector 47
4.2 Data acquisition (DAQ), triggering, and readout . . . 50
4.3 Calibration 52
4.3.1 Hit energy calibration 53
4.3.2 Timing calibration 60
4.4 Common reconstruction 63
4.4.1 Time slicing 63
4.4.2 Clustering 64
4.4.3 Track formation 65
4.4.4 Energy reconstruction 68
4.5 Detector response simulation 71
4.6 Readout simulation 72

The MINERνA experiment's primary purpose is to drive down the systematic uncertainties of oscillation analyses, by constraining the as yet poorly understood neutrino interaction models with data. At the time of its commissioning, neutrino interactions had little data to constrain the various interaction types, and on few nuclei. Since the neutrino interactions take place inside the nuclear environment, there is an obfuscation between the underlying weak interaction type and the measurable final state [215–217]; for a comprehensive review, see [218]. MINERνA has produced a wealth of physics results with the aim to tune neutrino interaction generators used to simulate interactions in oscillation experiments [192, 219–223].

In this chapter, we will always refer to the USER coordinate system unless otherwise specified. For a more detailed description of the MINERνA detector, its calibrations and design, the reader is referred to [224].

4.1 Description of the detector

The detector is placed 2.1 m upstream of the MINOS near detector in the NuMI beamline, which acts as a muon spectrometer for MINERνA analyses. Its underground position significantly reduces the rate of cosmic ray background muons

(estimated to be 18 Hz). It is composed essentially of three parts: the *inner detector* (ID), the *outer detector* (OD), and the upstream *veto wall*.

The ID is shaped like a regular hexagon, and in the longitudinal beam direction is segmented into various regions. First, a region of solid nuclear targets, containing elemental Pb, Fe, and C target planes, intersperesed with tracking planes of plastic scintillator; a water target is also in this "target region". A cryogenic container acting as a helium target sits between the beginning of the solid target region and the veto wall. Next, the fully active tracker region, composed of 62 scintillator modules. Finally, downstream calorimetry in the form of an electromagnetic (ECAL) and hadronic (HCAL) calorimeter, of 0.35 and 3 interaction lengths respectively. The OD is also shaped like a regular hexagon, and essentially consists of a steel frame with embedded scintillator, wrapped around the outside of the ID in the transverse plane. Completing this description, the veto wall is a thick steel plate, followed by a first scintillator plane, then a second steel plate - scintillator plane pair, which shields the detector from low-energy particles made in the upstream rock and allows tagging of "rock muons" that survive propagation through this region.

In the longitudinal (z) direction, the position along the detector is oft parametrised in terms of the *module number*. A module is simply a pair of scintillator planes, each as big as the ID and composed of extruded scintillator strips. The strips are triangular. In each module, one plane has its strips running parallel to the y direction (the "X" plane, as strip number on this plane gives information on the x position) and the other is rotated by $\mp 60°$ with respect to the "X" plane. Rotations of $\mp 60°$ yield "U" and "V" planes, respectively. Combining the information of module number (proxy for z) with the (x, u or v) pair allows one to place points in the (x, y, z) coordinate system, enabling 3D reconstruction.

The essential features of MINERνA can be seen in Figure 4.1.

In each tracking plane, there are 127 triangular scintillator strips, of height 1.7 cm and base 3.3 cm. The tracker material is polystyrene, doped with wavelength-shifting organic scintillator - the same that had been used for the MINOS scintillator strips [224]. At the geometrical centre a 2.6 mm hole is bored into the prism, with

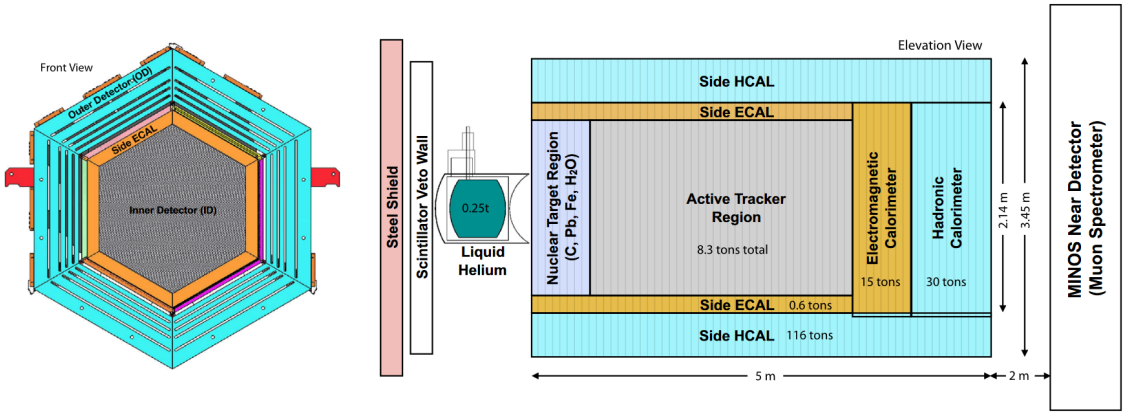


Figure 4.1: Frontal and side views of the MINERνA detector. From [224].

a 1.2 mm wavelength-shifting (WLS) fibre placed to read out scintillation light. The basic strip geometry is shown in Fig. 4.2 Because of the interleaved geometry of the strips, and exploiting the fact that minimum-ionising particles deposit more energy as they pass closer to the centre of the strip, a position resolution of 3 mm is obtained [225]. The ID tracker also contains other nuclei, primarily from the epoxy used to glue together the strips, the paint that coats either end of the strips, and the layer of white reflective paint (which contains 15% TiO_2 by weight) that each strip is coated with: Table 4.1 breaks down the mass and number fractions of the various nuclei inside the active tracker.

Species	Mass fraction	Number fraction
^{12}C	87.6%	49.0%
^{1}H	7.4%	49.4%
^{16}O	3.2%	1.33%
^{48}Ti	0.69%	0.0967%
^{35}Cl	0.55%	0.112%
^{28}Si	0.27%	0.0645%
^{27}Al	0.26%	0.0647%

Table 4.1: Mass and number fractions of the nuclei in the tracker. From [226].

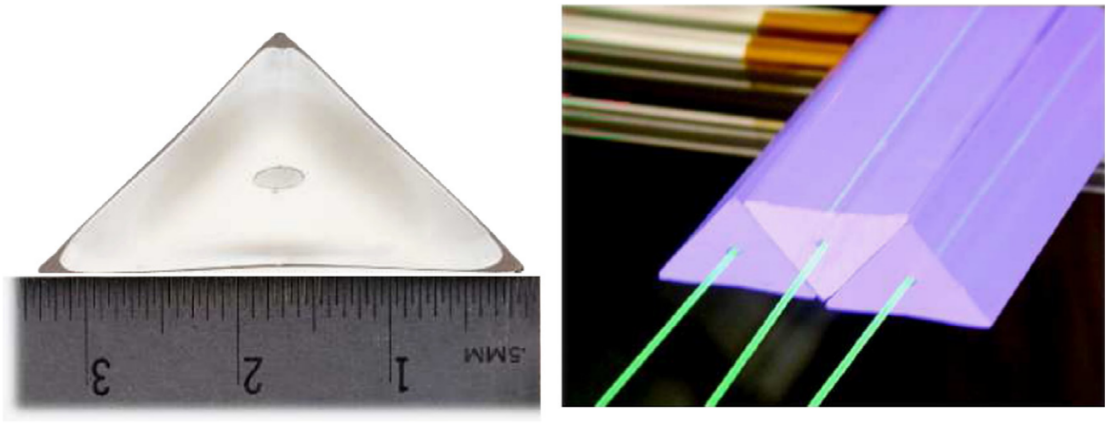


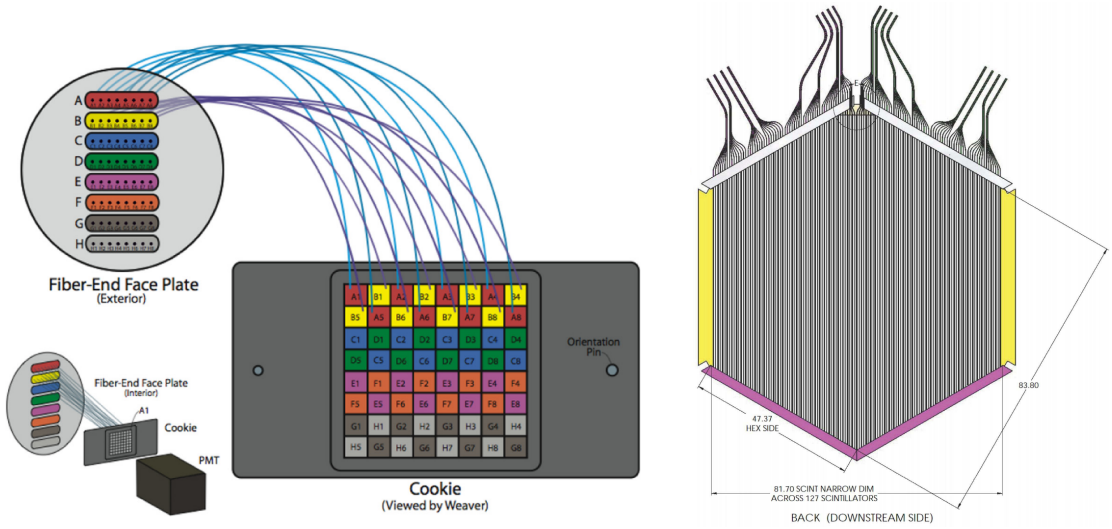
Figure 4.2: One scintillator strip (left), and three strips with fibres running through them (right). From [224].

4.2 Data acquisition (DAQ), triggering, and read-out

Only one end of the WLS fibres is read out, with the other end covered by a thick reflective coating designed to mirror light; the mirror reflectivity was measured to be $(83 \pm 7)\%$(by comparing light yields along the fibre with its one end first coated in reflective tubing, and then with absorbent black epoxy). At the other end of the fibres, they are bundled up in an 8×8 grid into clear cables, which are surrounded by light-tight tubes and an opaque sheath filled with black epoxy [224], to ensure maximal light yield. The attenuation of these cables is then obtained by fitting measurements from prepared cables of specified lengths, and fit to an exponential.

The cables are then connected to Hamamatsu multi-anode photo-multiplier tubes (PMTs), each having a grid of 64 pixels in an 8×8 array, and housed in a steel light-tight cylinder. This "cookie" array is made such that light from physically adjacent strips does not get sent to physically adjacent pixels on the cookie, to reduce the impact of signal cross-talk from adjacent pixels (see Fig. 4.3(a)). Once a particle deposits energy in the scintillator and light is sent through the fibre to the PMT, photo-electrons (PE) are produced; the PE signal is converted by an analogue-to-digital circuit (ADC) into digitised charge at one of the PMT anodes.

An "optical box" is assembled with the PMT, mounted on a circuit board that routes the signal to a front-end board (FEB) on the outside of the box. The



(a) (b)

Figure 4.3: (a): Arrangement of fibres onto a PMT readout "cookie". Notice the arrangement that separates physically adjacent fibres from each other, to reduce the impact of signal cross-talk [224]. (b): Engineer's drawing of readout fibres mounted on an ID plane [227].

optical cables enter the box on the other end to the FEBs. Steps have been taken to shield the boxes, to the best possible extent, from the residual magnetic field from MINOS, including placing a "mirror plane" between MINOS and MINERνA, orienting the boxes perpendicular to that field, and adding extra shielding to the boxes closest to MINOS.

The FEBs operate on a 9.4 ns clock tick, with a quadrature circuit enabling storage of quarter-ticks of 2.35 ns; each FEB mounts six TriP-t chips that are grouped together for storing hits [228]. Each chip records a value proportional to the electronic gain times the charge from the anode it services. They are grouped by threes, each group servicing half of one PMT: two chips operate 16 of 32 channels at low gain and 16 of 32 channels in medium gain, and the other chip operates all of its 32 channels in high gain. Each of these channels (low, medium, and high gain) is independent of each other. A discriminator is supported at the high gain channel; when the charge from an anode causes the discriminator to fire, the charge from all three channels from that time is integrated over 16 ticks, and the information is "pushed" up through the readout pipeline; afterwards, the channels require 20

ticks to reset. During this "dead time", the channels cannot time hits, nor is charge recorded during this time accurate, as it is greatly attenuated. Figure 4.4 illustrates the recording on one-half of a PMT.

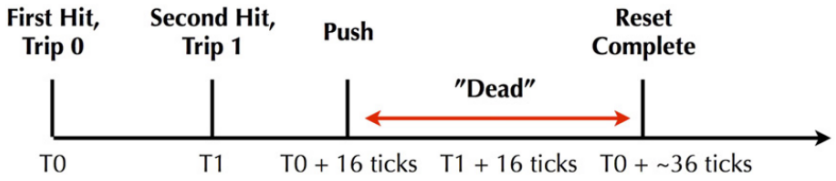


Figure 4.4: Illustration of a hit firing the discriminator of one of the high-gain channels at time $T0$. TriP-t chips 0, 1, and 5 service one half of the PMT. The high-gain "Trip 0" chip's discriminator fires: for 16 ticks, charge is collected and then pushed. For 20 ticks after that, the group of chips is "dead" and cannot read hits. From [228].

FEBs are themselves grouped into "chains", which are connected to a custom chain read-out controller, or CROC. CROCs communicate with CROC interface modules (CRIMs) which receive information from the MINOS timing system, as well as the Main Injector (MI) timing system. The MI delivers spills to MINERνA every 2.2 (2.06)μs [228]: when the signal from the MI is processed by the CRIMs, the FEBs are set to open 500 ns before the spill arrives. The readout gate stays open for 5.5 μs after the spill ends, to capture late activity such as Michel electrons from muon decay at rest.

4.3 Calibration

Raw detector readout must be calibrated in order to be translated to useful information about the event recorded. This implies calibrations to the ADC current recorded by the FEBs as well as to the timing information from the time-to-digital converter (TDC) implemented on the FEBs. Generally, calibrations may be divided into *in situ* and *ex situ*; the latter were measurements made before the detector assembly, to ensure individual components performed as required and to measure critical quantities needed. The former describe measurements of quantities that vary over the course of the experiment's long lifetime, made using suitable calibration sources.

4.3.1 Hit energy calibration

To obtain the hit energy E_i from a strip i, one must first correct the ADC readout v_i, and then apply conversion constants to translate that to an energy [224]. The corrections follow the procedure by which photons, created from the de-excitation of ionised scinitllator, eventually create a digital charge signal read out from the FEBs. This procedure may be summarised in four steps:

1. First, a light signal travelling along the WLS fibre through the length of the strip is expected to be attenuated; using calibrations that will be described later in this Chapter, a strip-dependent factor η_i is measured to correct for this;

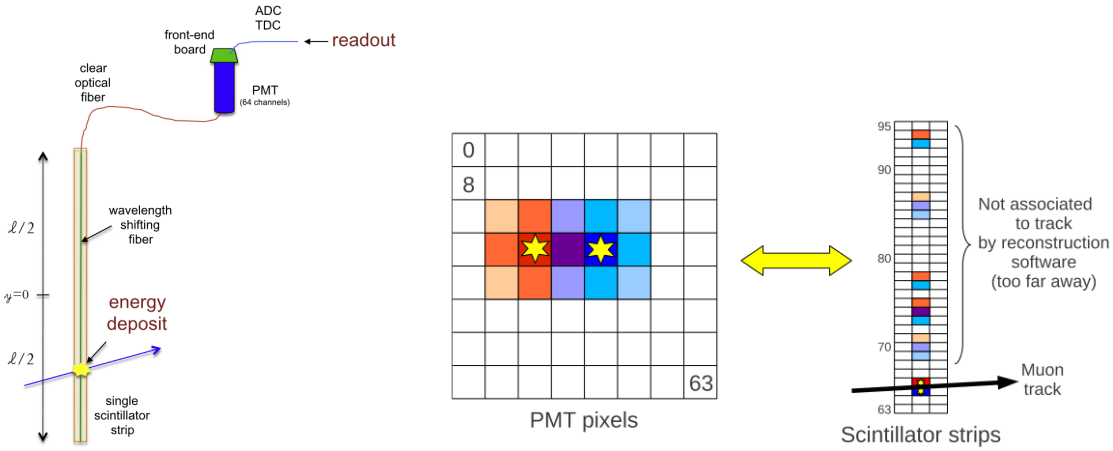
2. Second, there is attenuation of the light signal while it travels a length ℓ_i along the clear cables to reach the optical box where the PMT is located. The attenuation length λ_{clear} is measured for the cables, and a correction $\exp\left(\ell_i/\lambda_{\text{clear}}\right)$ is applied;

3. Third, a time-dependent electronics gain $G_i(t)$ is applied to account for the conversion of photons to photoelectrons;

4. Fourth, to account for the conversion of ADC to charge on the FEB, a conversion factor $Q_i^{(\text{ADC})}$ is applied.

Finally, one must bridge the gap between energy deposition in the detector and number of photons emitted from the scintillator de-excitation. In MINERνA, this is accomplished by determining an "absolute energy scale" for the detector, which is essentially a time-dependent response function $C(t)$. This response function is obtained by tuning a GEANT4 simulation of the energy deposited by rock muons, which are minimum ionising; the procedure will be explained in more detail later this Chapter. The "global" response $C(t)$ then has a strip-to-strip correction $S_i(t)$ applied to it, to account for varying responses in the strips.

$$\begin{aligned} E_i &= \mathcal{S}_i(t) \times v_i, \\ \mathcal{S}_i(t) &= \left(C(t) \cdot S_i(t)\right) \times \left(\eta_i \cdot \exp\left[\ell_i/\lambda_{\text{clear}}\right]\right) \times \left(G_i(t) \cdot Q_i^{(\text{ADC})}\right). \end{aligned} \tag{4.1}$$



(a) (b)

Figure 4.5: Schematic of the readout from hit energy deposited on a strip (a); Schematic of cross-talk from a muon track depositing energy on strips 65 and 66 on a MINERνA module (b). Pixels are mapped as in Fig. 4.3(a) (left), which leads to cross-talk on nearest neighbours (pixel space) being physically separate, which lets reconstruction algorithm reject cross-talk clusters as part of the track. From [224].

Essentially, the product $G_i(t) \cdot Q_i^{\mathrm{ADC}}$ translates ADC counts to number of PE, the factor $\eta_i \cdot \exp\left[\ell_i/\lambda_{\mathrm{clear}}\right]$ translates number of PE to number of photons, and the factor $C(t) \cdot S_i(t)$ translates number of photons to energy deposited on the strip i. A sketch of the optical readout procedure can be seen in Fig. 4.5(a).

Ex situ calibration measurements were performed on the scintillator strips, the PMTs, and the FEBs before assembly of the detector. In addition, measurements of the WLS fibre attenuation, reflection, and transmission of the clear optical cables were made. These allowed for the precise determination of the optical properties of the detector, and for the identification of anomalies that needed to be corrected for in the ADC to hit energy calibration of Eq. 4.1.

The WLS fibre attenuation was measured for each fibre, by constructing the ratio of two readouts from a test stand. The test stand consisted of an LED diode, calibrated such that it shone with the same intensity, inside a box. A "direct measurement" connecting the LED box to a readout box via optical jumper cables was made first, as a control; then, a "cable measurement" introducing clear optical cables (as described in Section 4.2) in between the jumper cables. The transmission of each cable is then the ratio cable-measurement / direct-measurement. Measurements

were made on cables with lengths 0.5, 1.08, 1.38, and 3.13 m; the transmissions were then fit to an exponential, yielding $\lambda_{\text{clear}} = 7.83$ m. Simultaneously, the transmission across the optical cable connectors was obtained by subtracting the clear fibre attenuation from the cable measurement.

Following assembly of the detector modules, a "module mapper" was deployed to extract the position dependence of optical attenuation, along the length of each strip, for both the ID and OD. The mapper performed scans over each module using two radioactive ^{137}Cs sources, moving along a predetermined path over the module. The scintillator response was recorded by PMTs and processed with custom DAQ software. This response was analysed as a function of two positions: longitudinal along the length of each fibre, and transverse along the module. The pedestal value for each fibre was recorded before and after illumination from the sources, and the mean of these two values was used for pedestal subtraction. This is due to an electronics artefact, where the pedestal value in an optical channel shifts to lower ADC counts when the channel is illuminated. For each scan, and for each longitudinal position, the distribution of the pedestal shift as a function of transverse position in the strip was fit against a Lorentzian distribution, from which the position of the fibre's centre along the transverse direction and the maximum fibre response were extracted; simultaneously, dead channels were found to be 15 out of over 32 thousand. By plotting the maximum responses as a function of longitudinal position for each strip, a response curve was constructed, which is then applied to calculate η_i for a hit whose longitudinal position is obtained via later-stage tracking.

The electronic components of the detector are characterised before installation in the detector. Each PMT is characterised for individual variations, and once installed in boxes additional tests are carried out to identify boxes with large amounts of cross-talk and repair them before installation. Figure 4.5(b) shows how cross-talk from pixels recording energy of muon track is mapped in physical space (right half). Moreover, the FEBs are measured using replica MINERνA DAQ, to determine the gain $G_i(0)$ of the high, medium, and low gain channels. A typical response curve showing the gains is shown in Fig. 4.6.

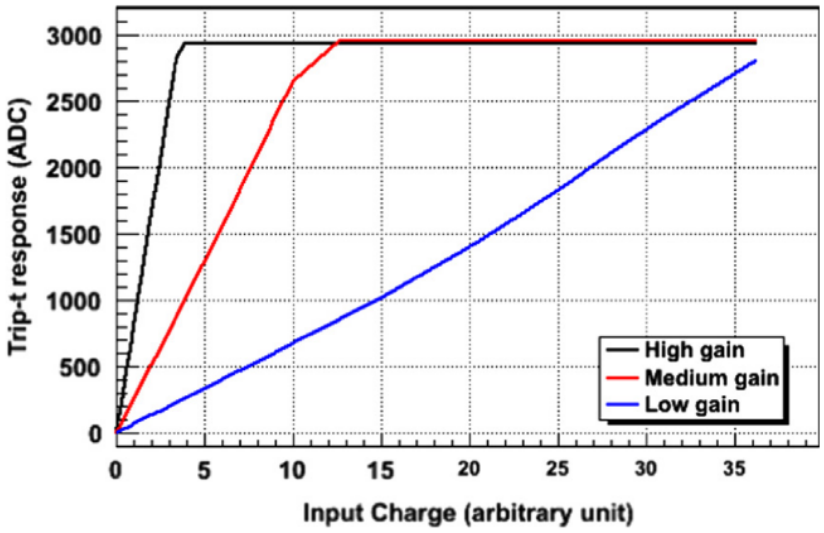
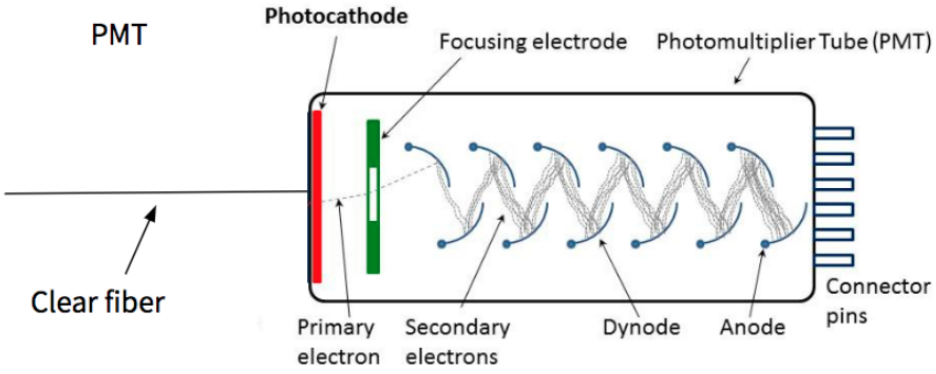


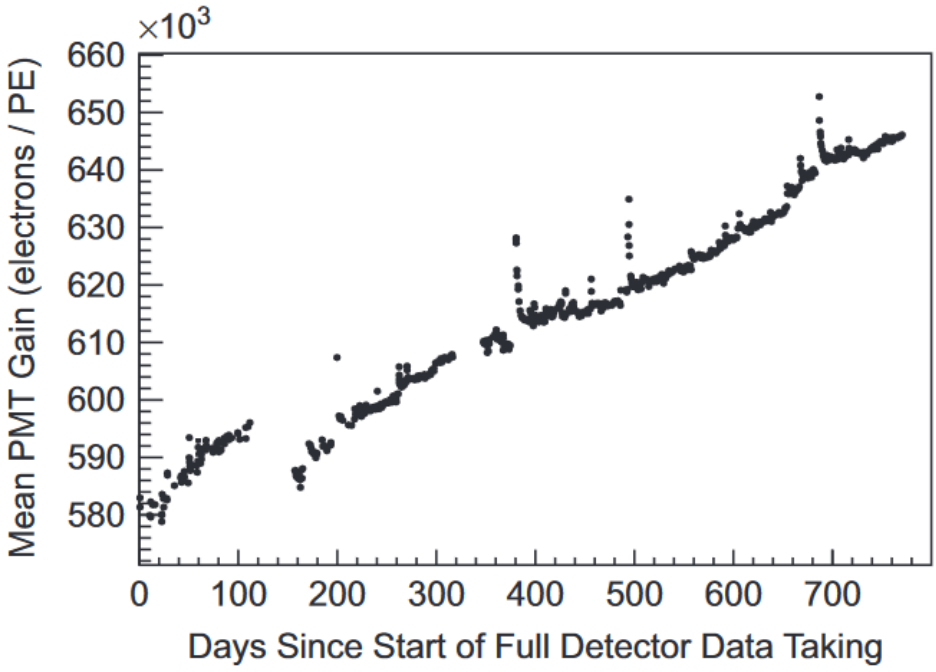
Figure 4.6: ADC counts versus input charge for high, medium, and low gain channels from the TriP-t chips. ADC counts are saturated at about 2500. From [224].

Once the detector was assembled, a suite of *in situ* measurements to correct for time-dependent response and anomalies was carried out. By monitoring the ADC counts for mixed readout gates with beam spill on and off, the pedestal value for each channel was determined (high ADC counts occurring from entry of cosmic rays are removed during DAQ); ADC counts within 3 times the RMS of the pedestal distribution are suppressed [229]. The RMS of the ADC counts in each channel is much larger than the deviation across time for each channel, meaning the pedestal for each channel is, to very good accuracy, invariant with time.

A light injection system connected to each optical box allows roughly uniform illumination of the PMT pixels, from which the spectrum of the PMT output can be constructed. Concurrently, modelling the PMT as a system of n dynodes (each multiplying the number of PE before reaching the anode and being recorded) with known underlying distribution, and known voltage differences $V_i^{(a)}, a = 1, \ldots, n$. From the light-injection spectrum and the modelling, one can obtain a value for the *PMT gain* $g = \overline{Q}/(\lambda \cdot e)$ of each pixel, where $\overline{Q}$ is the mean of the pixel's ADC distribution post pedestal suppression, λ is the mean number of PE arriving at the first dynode, and e the electron charge. The associated systematic uncertainty is less than 1% [224]. Note the PMT gain g (PE count$\rightarrow$ electrons) is different



(a)



(b)

Figure 4.7: (a): Sketch of a PMT amplifying PE received at the photocathode into a signal at the anode. From [225]. (b): PMT gain $g(t)$. Peaks are due to PMTs not stabilising immediately after high voltage is switched back on. The dip at day ~ 160 is due to nominal high-voltage being adjusted. From [224].

from the electronic gain $G_i(t)$ of each channel. g is, however, also time-dependent, as the high voltage to PMTs is switched on and off, and as the PMTs age. This dependence is explicitly shown in Fig. 4.7(b).

The strip-to-strip correction function $S_i(t)$ accounts for variations in light level between ID strips; this "smooths out" the response so that it is uniform across the

detector. Using through-going rock muon tracks, the peak energy divided by path length is montitored. Because muons are minimum-ionising particles (MIPs), the peak energy deposition along a path length is roughly uniform, but slowly varies with time. Given that the fits of these responses require thousands of rock muons, it is not possible to extract this slow variation for each strip (which would require much longer exposure to rock muons than the time in between changes to the detector hardware). Therefore, it is possible to use the peak energy to normalise the different responses of individual *planes* in the detector (where the scintillator was all manufactured in the same batch), but not for *strips* in a plane.

For strips, the truncated mean energy [230] is used as a proxy for the peak energy, which has the advantage of requiring fewer rock muons to be used as it is computed iteratively (each successive iteration using events close enough to the mean of the previous iteration). From this, a constant $s_i^{(j)}$ is calculated for strip i in plane j is obtained; it is used in a first iteration (assuming truncated mean is a perfect proxy for peak energy) to produce the energy scale for plane j. Then, these scales are used to process rock muon data, meaning the peak energy distribution can be fit plane-to-plane. Applying the fit described above to normalise the peak energy across all planes introduces an additional correction s_j, which accounts for the truncated mean $\rightarrow$ peak energy translation. These constants are then multiplied to give the final answer for the strip-to-strip correction, $S_i = s_i^{(j)} \cdot s_j$. Because this relationship is sensitive to scintillator aging (as scintillators in different planes age differently), the calibration is repeated numerous times and the strip-to-strip correction acquires a time dependence, $S_i(t)$.

Completing the chain of ADC $\rightarrow$ hit energy, the response function $C(t)$ exploits muon energy loss in the tracker to calibrate the PE-to-energy conversion. It is reported in "muon equivalent units" (MEU), and is once again obtained from rock muon data. Specifically, MINERνA rock muon tracks which are matched to tracks in MINOS have their MINOS-reconstructed momentum corrected for energy loss in MINERνA, providing an estimate for the muon momentum on entry into MINERνA. This data is analysed with a trial MEU factor C_0, and then the

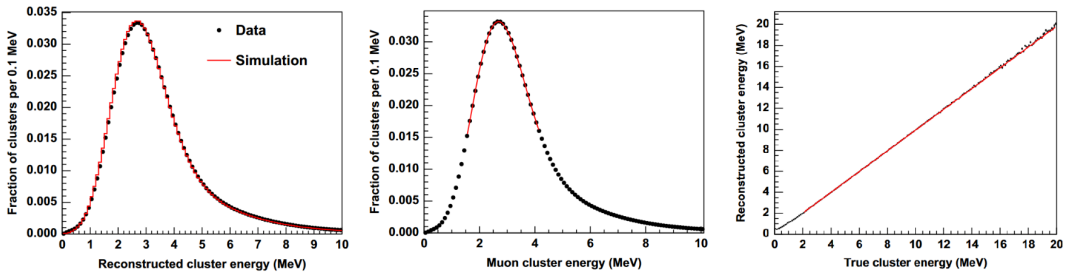


Figure 4.8: Procedure to determine the MEU factor C: tuned simulation to extract $E_{\mathrm{MC}}^{(\mathrm{reco})}$ (left), fit to extract E_{data} (middle), and slope $\int$ by comparing $E_{\mathrm{MC}}^{(true)}$ to $E_{\mathrm{MC}}^{(reco)}$ (right). From [224].

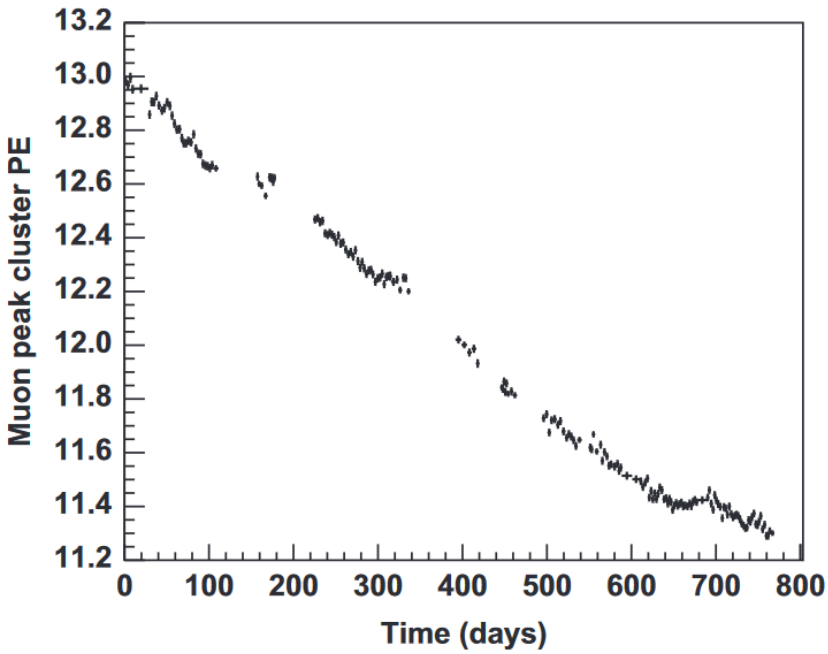


Figure 4.9: Number of PE at peak energy E_{data} from rock muons, as a function of time. From [224].

region around the peak energy deposition is fit to a fifth-order polynomial to extract the peak energy, E_{data}. At the same time, a GEANT4 simulation of the energy deposition, for which the true peak of the $\mathrm{d}E/\mathrm{d}x$ distribution $E_{\mathrm{MC}}^{(\mathrm{true})}$ is known, is used to derive the reconstructed peak energy $E_{\mathrm{MC}}^{(\mathrm{reco})}$ from the simulation (reconstructed cluster energy is obtained by smearing the true cluster energy). The true and reconstructed peaks are then plotted and fit to a linear function, extracting a slope $\int$. The factor C_0 is then corrected as $C_0 \mapsto C_0 \cdot E_{\mathrm{MC}}^{(\mathrm{reco})}/(\int \cdot E_{\mathrm{data}})$. Figure 4.8 shows the determination of the MEU according to the above steps. Primarily because of scintillator aging and environmental conditions of the detector hall, the

light yield in the strips typically decreases with time, meaning $C = C(t)$: this means the calibration must be performed at multiple times throughout MINERνA's lifetime. Figure 4.9 shows the number of PE at the peak E_{data}, which proxies for the light yield as discussed above.

Lastly, the alignment of the module positions and rotations of planes in the modules is monitored and corrected for in offline analysis, by plotting the energy from through-going rock muons as a function of the position of the track along the face of the strip (more accurately, the ordinate of the intersection between track and strip base), and along the longitude of the strip. The entire plane is treated as a rigid unit, without accounting for small variations in the strip rotations within a plane. First, a fit on the transverse position to the shape of the strip is performed to get the distance between the ordinate at which the energy deposition is highest and the centre of the strip, where the energy deposition is expected to be highest if the strip were at zero transverse offset. Performing this fit for six points along the strip, one can extract the slope of the maxima versus longitudinal position, obtaining the rotation of each strip. Figure 4.10 shows this procedure for one "X" plane. As the detector was not disturbed after installation, this procedure was carried out once.

After alignment correction, dead strips (where more than 30 % of tracks traversing at least 2 mm produce no signal) are identified, and excluded from further analysis. Simultaneously, the "plex map", which maps channels in electronic space to strips in physical space, is checked for errors (typically from when optical fibres are flipped during detector installation, leading to pairs of strips that are swapped with respect to the position they should be in). Once errors are identified, the plex map is updated and used in the data analysis [231].

4.3.2 Timing calibration

Hit timing is essential to deal with the pile-up associated with the high-intensity beam. The TDC circuit on the FEBs provides a "raw time" for each hit that fires the discriminator circuit (i.e. a hit that produces more than ~ 0.7 PE). The "raw time" T_{raw} is then corrected for the following effects:

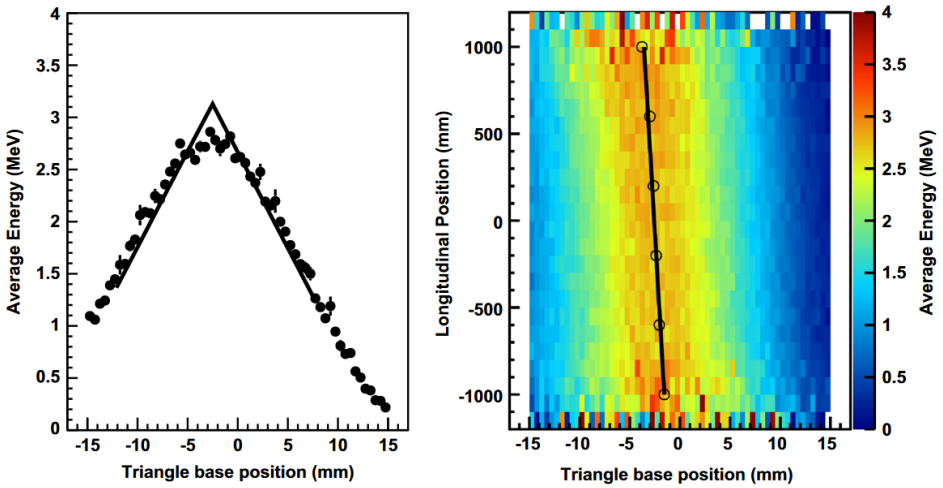


Figure 4.10: Plane alignment for a MINERνA "X" plane; on left, the energy deposited by through-going rock muons as a function of the transverse position is fit to the shape of the strip, and the position of the peak is extracted. On right, the energy is plotted with respect to transverse (x) and longitudinal (y) position, and using the result from six fits at six locations the slope (black line) is extracted. From [224].

- Transport time of photon in fibre, T_{fibre};
- Time slewing T_{slew}, i.e. the apparent delay of the output from the discriminator as a function of the accumulating input signal [232]. Since every photon at the PMT is produced by a two-stage process [231] (scintillator fluor decay, then absorption of the decay photon and re-emission of a photon that reaches the PMT) and the photon can either travel directly along the fibre to the PMT, or go backwards, be reflected at the end of the fibre and then reach the PMT, there is a delay of the discriminator firing as a function of number of PE produced in a strip. The slewing becomes larger as the number of PE is smaller, since the rising input is less steep (and it is more likely that a PE travelled a longer path to reach the PMT);
- Channel-to-channel offsets, T_{chan}, caused by time offsets between readout chains and offsets between FEBs in a single chain, but not by offsets between groups of channels on the same FEB (which are negligible for groups measured by the same TriP-t chips).

The impact of N_{PE} on time slewing can be seen in Fig. 4.11. The first effect is

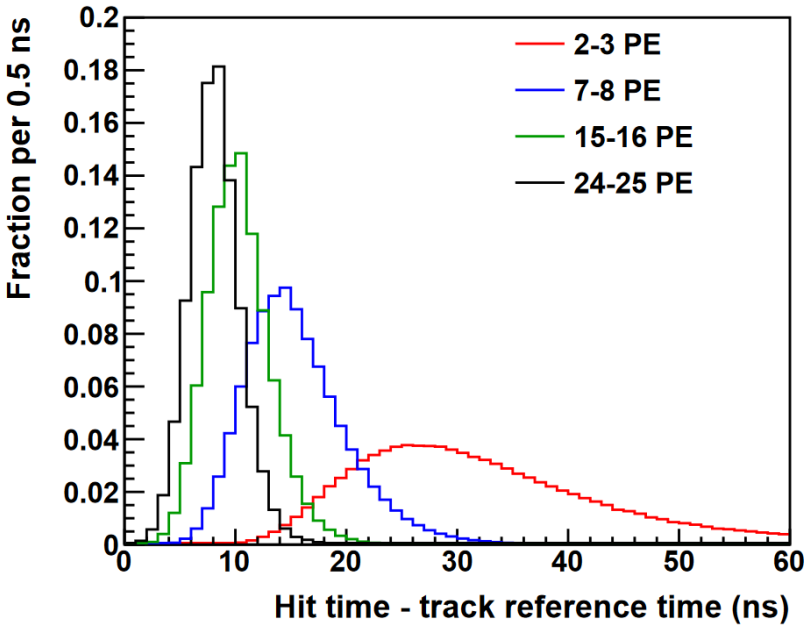


Figure 4.11: Illustration of time slewing: x axis is $T_{\text{cal}} + T_{\text{slew}} - T_{\text{avg}}$ (cf. Eqs. 4.2, 4.3). Histograms are area normalised. Time slewing is more severe for lower input yields (N_{PE}). From [231].

obtained straightforwardly by using the traversed fibre length and the speed of light in the fibre, 156 mm/ns. The other two effects are determined by an iterative procedure meant to decouple one from the other, by measuring hits from rock muon tracks. For each hit strip, a correction for the muon time of flight, T_{TOF}, is applied; then the hit time $T_{\text{raw}}^{(i+1)}$ for is corrected based on the $T_{\text{slew}}^{(i)}, T_{\text{chan}}^{(i)}$ as

$$T_{\text{cal}}^{(i+1)} = T_{\text{raw}} - T_{\text{TOF}} - T_{\text{fibre}} - T_{\text{slew}}^{(i)} - T_{\text{chan}}^{(i)}, \tag{4.2}$$

where $T_{\text{slew}}^{(0)} = T_{\text{chan}}^{(0)} = 0$. For each iteration, the electronics offsets are calculated by measuring $T_{\text{cal}}^{(i+1)} - T_{\text{avg}}^{(i)}$, where $T_{\text{avg}}^{(i)}$ is obtained along the muon track:

$$T_{\text{avg}}^{(i)} = \frac{1}{N_{\text{strips}}} \sum_j T_{\text{cal}}^{(i)}(j). \tag{4.3}$$

This procedure is repeated 8 times. From this calibration, the time slewing as a function of PE is constructed and fit to a polynomial in $1/\sqrt{N_{\text{PE}}}$. Because the channel-to-channel offsets change every time hardware is swapped, this calibration is performed each time there is a swap. Fitting the residual $T_{\text{cal}} - T_{\text{avg}}$ for rock muon tracks yields a timing resolution of 3 ns (cf. Fig. 5.11(c)).

Data: All hits $\{\mathcal{H}_i\}_{i\in I}$ in a readout gate
Result: Set of time slices $\{\mathcal{S}_i\}_{i\in I} : \mathcal{S}_i = \{\mathcal{H}_i\}$
Arrange all hits in $\{\mathcal{H}_i\}_{i\in I}$ by ascending time, $\{\mathcal{H}_{(i)}\}_{(i)\in I}$ (with $t\left(\mathcal{H}_{(0)}\right) \leq t\left(\mathcal{H}_{(1)}\right)$ etc.);
Define a rolling window $w = [w_-, w_+]$;
For each hit (i), accumulate the PE from all hits (j) within the window w centred at $(i) : t\left(\mathcal{H}_{(j)}\right) \in \left[t\left(\mathcal{H}_{(i)}\right) + w_-, t\left(\mathcal{H}_{(i)}\right) + w_+\right]$;
If the accumulated PE are above a threshold, create a time slice $\mathcal{S}_i$ and assign all hits (j) to the slice. Also include hits sufficiently close in time to the first and last $\mathcal{H}_{(j)} \in \mathcal{S}_{(i)}$;
return $\{\mathcal{S}_i\}_{i\in I}$

Algorithm 1: Time slicing algorithm.

4.4 Common reconstruction

A common reconstruction step produces high-level objects, such as clusters of hits and tracks of clusters, derived from the calibrated hits described previously. Though analysis-specific reconstruction often occurs on top of this common step, the objects obtained from here are normally used as an input to subsequent reconstruction steps.

4.4.1 Time slicing

In the ME era, the NuMI beam operated at higher intensity compared to the LE era, with pile-up becoming significant. A single readout gate of 16 μs normally contains multiple overlapping neutrino events. To combat this, a *chrono-buncher* algorithm associates hits that fulfil some common requirements to a single time *slice* $\mathcal{S}_i$: a proper subset $\{\mathcal{H}_i\}$ of the family of hits $\{\mathcal{H}_i\}_{i\in I}$. This algorithm is summarised in Algorithm 1. Hits that did not fire the electronics discriminator and hits that are untimed are not included in any slice, being assigned instead to $\mathcal{S}_0$, which contains all hits not otherwise assigned. A typical data event, such as shown in Fig. 4.12, showcases the capability of time slicing to isolate individual physics events. All ensuing reconstruction steps are performed on the hits of each slice $\mathcal{S}_i$.

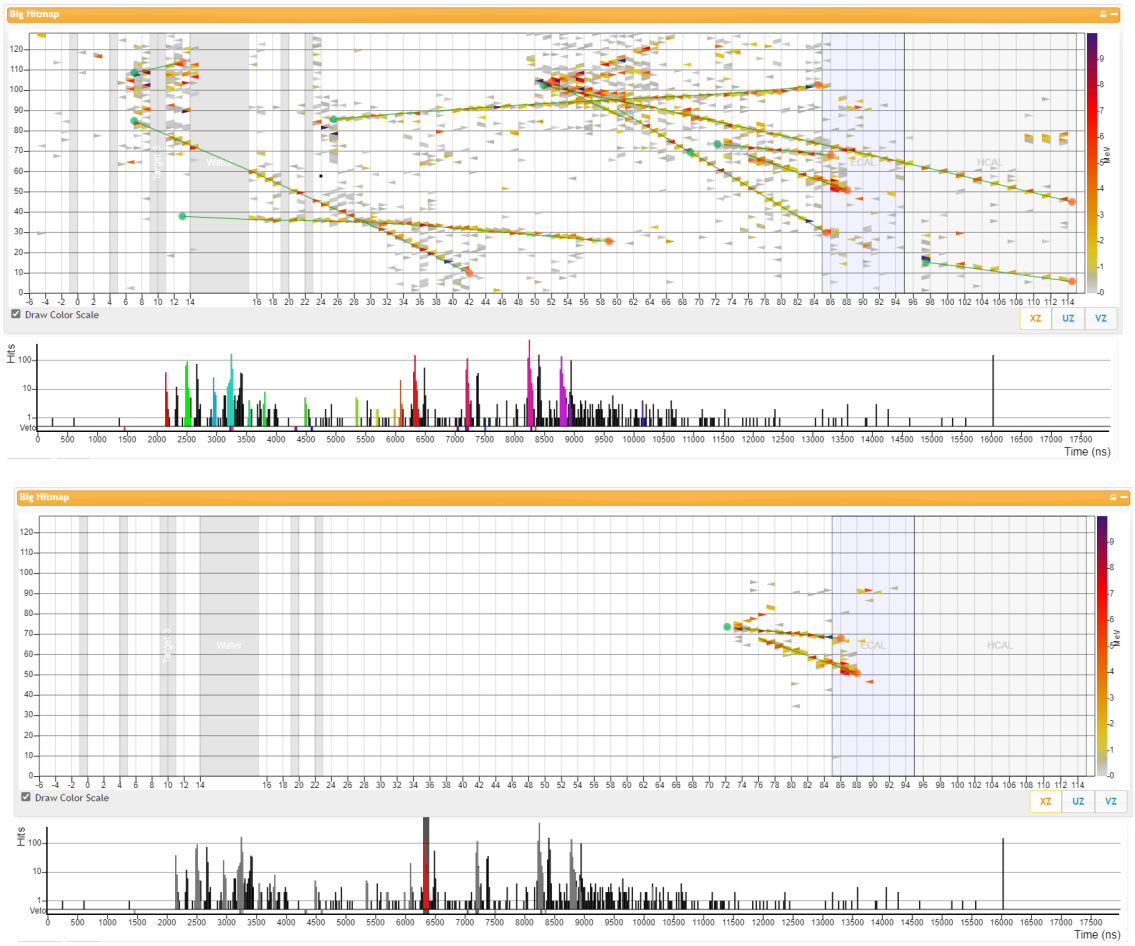


Figure 4.12: A data event from MINERνA, in the XZ view. Pile-up is evident in the top panel, with all "time slices" (coloured bands) displayed. The bottom panel shows only hits in the interval $6300 \leq t\,(\text{ns}) \leq 6400$, grouped in the same slice. This reveals one candidate neutrino interaction in the readout gate.

4.4.2 Clustering

Hits adjacent to each other and in the same plane are grouped into *clusters*, whose characteristics depend on the number and the energy of hits associated with them. The cluster position is obtained from the hit positions by taking the charge-weighted mean, and the cluster time is the time of the hit with the largest energy. There are five cluster types, cf. Fig. 4.13:

1. Trackable: Typical of minimum-ionising particles (MIPs). Require 1-3 adjacent strips, with total energy $1 \leq \Sigma E\,[\text{MeV}] \leq 12$ and each hit energy $0.5 \leq E[\text{MeV}] \leq 8$.

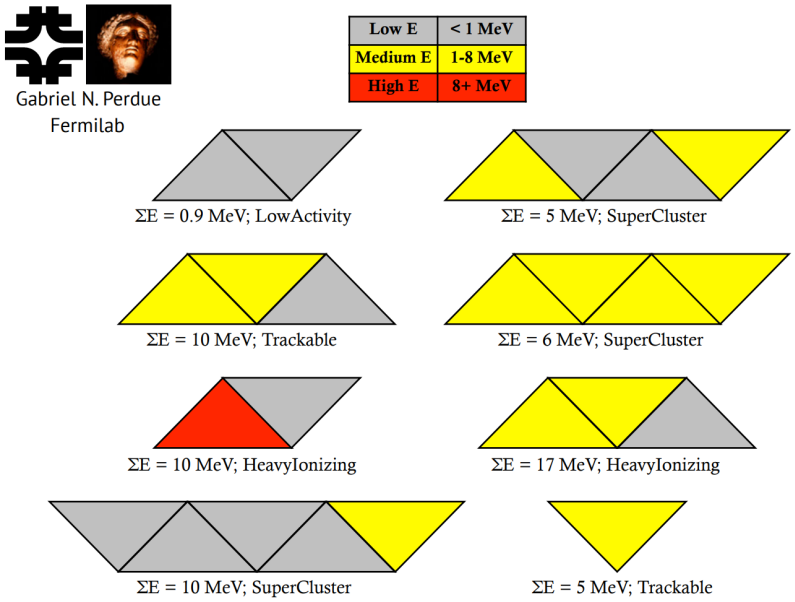


Figure 4.13: Various kinds of clusters formed in the ID from hits. Figure by Gabriel Perdue, [233].

2. Heavy Ionising: Similar to Trackable, but no upper bound on ΣE or E. These usually accompany the Bragg peak from a stopping particle.

3. Superclusters: Wide groups of hits, that satisfy the criterion $\Sigma E \geq 1$ MeV but not the other criteria. These normally indicate shower-like depositions.

4. Low Activity: $\Sigma E < 1$ MeV. Low-recoil interactions often produce these.

5. Cross-talk candidates: Single hits that are not physically adjacent to other hits, but are adjacent to other hits in electronics space. These are excluded from further reconstruction steps.

4.4.3 Track formation

The next step in the reconstruction workflow is to drive track "seeds" from each slice $\mathcal{S}_i$, to construct tracks of particles through MINERνA. A seed may be formed from a set of three clusters C_i which satisfy the following criteria [224]:

1. C_i is Trackable or Heavy-Ionising,

2. The planes P_j of clusters C_j satisfy: $P_i \neq P_j \; \forall \; i \neq j$ (no two clusters on same plane), and $P_{(j+1)} = P_{(j)} + 1$ (consecutive planes);

3. All three clusters are in the same view, X, U, or V;

4. All clusters must lie on the same 2-dimensional line.

Seeds in the same view are merged into track candidates in 2D, which are stitched together with other track candidates in other views provided their slopes and intercepts are compatible with a single 3D line. A single cluster may be part of multiple seeds, bearing in mind the possibility that two tracks deposit energy in the same cluster.

There are two algorithms that handle 3D track creation: a *long tracking* algorithm that examines all possible combinations of candidates from all three views, which cannot construct tracks that traverse fewer than 11 planes, and a *short tracking* algorithm which considers all the possible combinations of candidates across two views, which can get tracks down to 5 planes (in an $XUXV$ configuration). A seed is only associated with one track, so "used" seeds are no longer considered for further tracking attempts. All 3D tracks are fit by a Kalman filter to deal with multiple scattering in the detector. A track is allowed to be made if one plane has no clusters on it, but the subsequent plane matches the 2D track candidate: this takes care of dead channels where a cluster would not be made from the particle's passage. Track candidates are required to have their polar angle θ be less than 70° with respect to Z_{USER}, to ensure that there are enough planes to reconstruct a track. In the case of HNL two-body decays, this is true of the vast majority of events. Figure 4.14 illustrates the procedure to form a long track.

Exiting tracks are reconstructed in the MINOS near detector, matching the slope, intercept, and time of the MINERνA and MINOS tracks. In practice, this limits the muon acceptance of MINERνA to $\theta_\mu \lesssim 20°$, which is essentially all of the HNL phase space. If a match is found, the track is assumed to be a muon track, and is promoted to the primary track of the "physics event" associated with slice $\mathcal{S}_i$. Otherwise, the longest track is assumed to be the primary track, assuming

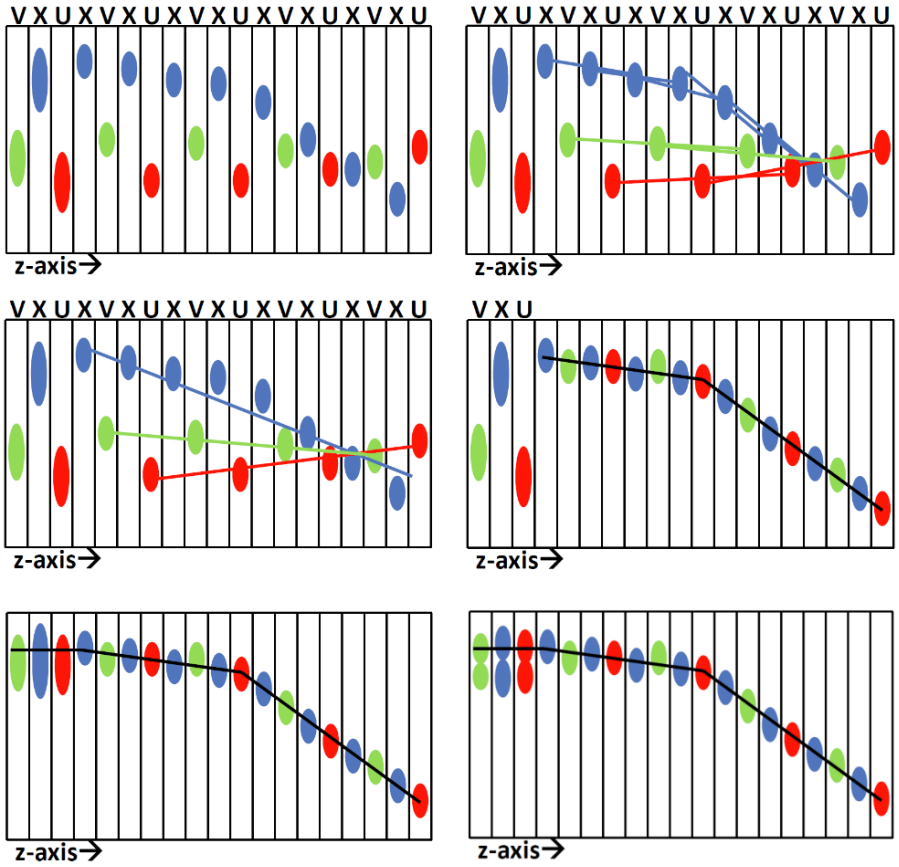


Figure 4.14: Track formation procedure for a long track. Upper left: Clusters in each view (blue: X, red: U, green: V. The first three planes contain Superclusters. Upper right: Clusters are grouped into seeds in each view. Mid left: In each view, seeds are merged into track candidates. Mid right: The X, U, and V track candidates are merged into one 3D track. Bottom left: The track is projected backwards to include the Superclusters that are not yet used. Bottom right: After cleaning, the Superclusters are broken up: three unused clusters appear. From [234].

it has at least 25 clusters. This track's starting point is dubbed the "primary vertex", and the track is called the "anchor track". After these have been found, the track-finding algorithm is run again on clusters that have not been used to make seeds, and tracks consistent with emerging from the primary vertex are constructed. The Kalman filter is run again, and the location of the primary vertex is updated. Finally, a "cleaning" procedure is carried out, to add clusters that are not likely to be associated with the tracked particle back into the pool for track finding. This improves both track pattern recognition and the vertex energy estimation.

For each track thus made, the track finding algorithm and the Kalman filter is run again, updating the locations of the vertices associated with each track.

A vertex may be classified as one of several types, depending on the number of outgoing/incoming particles:

1. StartPoint: $N_{\text{out}} \geq 1, N_{\text{in}} = 0$;
2. Kinked: $N_{\text{out}} = N_{\text{in}} = 1$ (usually indicates a particle scattering);
3. Forked: $N_{\text{out}} > 1, N_{\text{in}} = 1$;
4. StopPoint: $N_{\text{out}} = 0, N_{\text{in}} = 1$.

At this stage, clusters attached to cosmic rays and through-going rock muons (from neutrino interactions with the ~ 240 m of dolomite rock upstream of MINERνA) are tagged and excluded from reconstruction. This is done based on their characteristic signatures: cosmic rays enter through the side of the detector, whereas rock muons will deposit energy from the first plane in the target region and be accompanied by hits on the veto wall upstream of the detector.

Rock muons are notably useful for characterising the angular resolution of the detector; for muon tracks produced by front-enering rock muons, clusters were divided into an "upstream" and a "downstream" sample. Each sample was then used to construct a track, and their positions and angles were compared with each other in data and MC. The difference between upstream and downstream track angles is then a measure of the angular resolution of the detector; the MC yields an angular resolution within 1 mrad of the data-derived resolution, providing an excellent data/MC agreement [224]; Fig. 4.15 shows the data and MC angular resolutions.

4.4.4 Energy reconstruction

The mainline MINERνA analyses involve either one or no muons. As such, the reconstruction of energy generally breaks down into two steps: muon and non-muon energy reconstruction.

The former assumes that a single minimum-ionising muon traverses the detector, and potentially exits and is matched to MINOS. While in MINERνA, the muon loses energy by ionisation behaving like a MIP, as long as its momentum is $p_\mu \gtrsim 0.1$ GeV/c

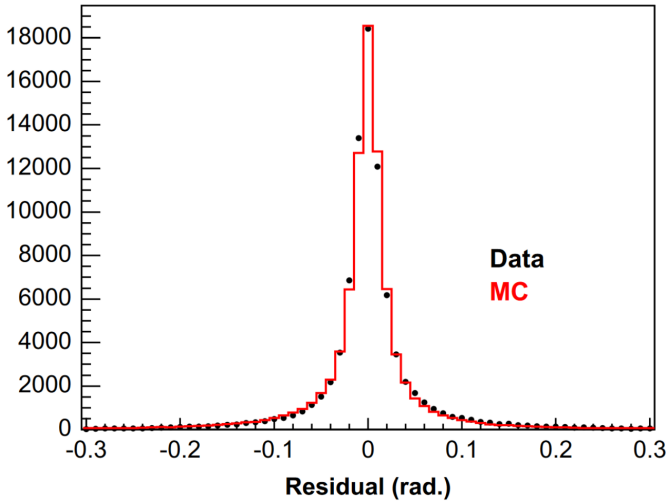


Figure 4.15: Angular resolution (difference of track angles from upstream track clusters and downstream track clusters) using rock muon tracks in data and MC at MINERνA. Adapted from [224].

(for a pion the same assumption holds for losses by ionisation for $p_\pi \gtrsim 0.15$ GeV/c). The mean energy loss by ionisation is described by the well-known Bethe-Bloch formula. It is shown in Fig. 4.16.

If the muon escapes to MINOS, then the MINOS energy reconstruction can also be utilised. For muons stopping inside MINOS, a momentum-by-range measurement gives the estimate for p_μ at the point of entry into MINOS: for escaping muons, a momentum-by-curvature measurement estimates the momentum by measuring the radius of curvature of the muon. Another benefit of exploiting the MINOS magnetic field is that charge identification of the muon may be carried out; muons contained inside MINERνA, on the other hand, cannot have that information inferred. For the mainline neutrino analyses, this information is crucial as ν_μ and $\overline{\nu}_\mu$ interactions have quite different topologies and cross-sections; for the HNL case $N_4 \to \mu^-\pi^+$, this is less of a concern, as the decay rate for both the particle and antiparticle channels is the same. (What matters for our search case is the ability to constrain a two-body decay; charge-capable detectors would be necessary in the next stage of precision physics involving HNL, post-discovery. See for example [155].) Once the MINOS momentum-at-entry is known, it is fed back into the calculation of momentum

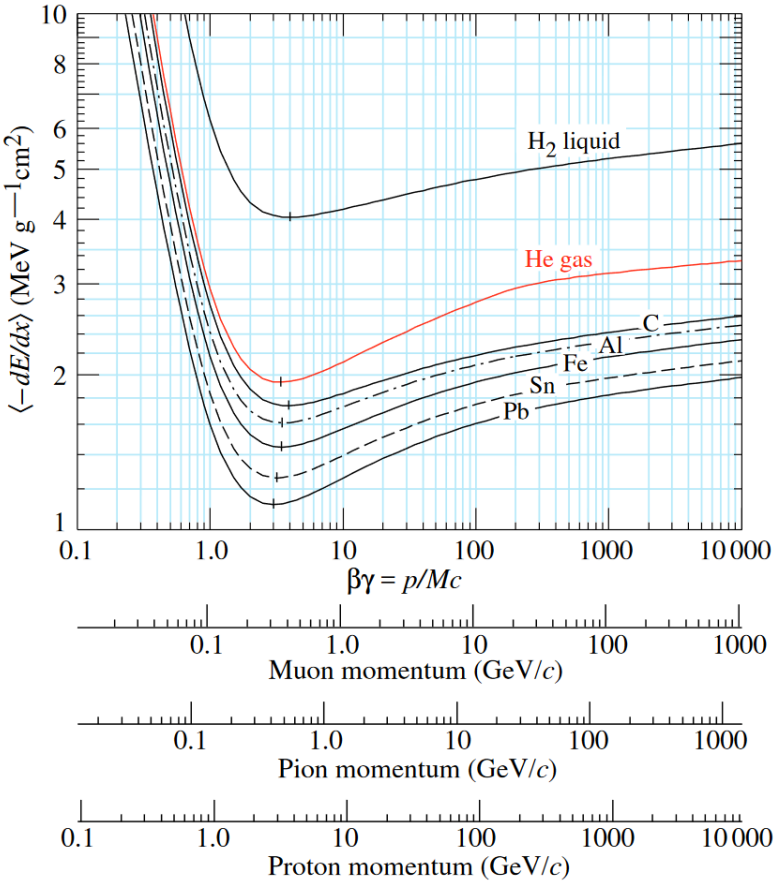


Figure 4.16: Mean energy loss by ionisation in various media. From [235].

in MINERνA , which is propagated back upstream adding in the expected energy loss for each step until the track starting position is reached.

Non-muon, or *hadronic recoil*, energy is the "other" component of the neutrino energy:

$$E_\nu^{(\mathrm{reco})} = E_\mu + E_{\mathrm{had}}. \tag{4.4}$$

This equation holds for two-body HNL decays as well, since there is no missing energy due to neutrino emission. It is not sufficient to just add up the *visible* non-muon energy, E_{vis}, to estimate the recoil energy E_{had}, though; the presence of neutral particles, or charged particles below tracking threshold, implies that energy could be "carried away" from the event without entering the above energy budget. Therefore, a *calorimetric correction* must be applied to the visible non-muon energy:

$$C^{\mathrm{SF}} = \frac{E_{\mathrm{vis}} + E_{\mathrm{passive}}}{f \cdot E_{\mathrm{vis}}}, \tag{4.5}$$

where E_{passive} is the energy deposited in passive (non-scintillating) materials ($= 0$ in the tracker), and f is the fraction of active materials in the plane ($f \simeq 0.8347$). Further calorimetric correction to clusters, based on scaling of simulated energy depositions to match the true energy depositions recorded by `GEANT4`, can be done. This will be further discussed in Chapter 7.

In the HNL analysis, this procedure affects the clusters that are not associated with the muon track, which is to say the pion clusters, *if the pion track is not overlaid on top of the muon track*. This implies that HNL decays where the muon-pion opening angle is large enough (typically $\theta_{\mu\pi} \gtrsim 7°$) lead to detectable non-muon clusters which will be calorimetrically corrected, whereas if $\theta_{\mu\pi}$ is small enough for muon and pion clusters to be counted as part of the muon track, the calorimetric energy correction will not be applied. Among other considerations, the above discussion motivates the division of the HNL sample into *two-track* and *one-track* events.

4.5 Detector response simulation

To link an output event record of final state particles generated inside the detector volume to observable parameters (hit energy and time, for each strip in which there would be a subsequent energy deposition), it is necessary to simulate the passage of particles through the matter of the detector. The procedure is as follows: propagate each final state particle through matter, deal with the effects of pile-up, and transform the result to hit ADC/TDC counts in each readout channel. Once this is done, the output can undergo the calibrations described in Section 4.3, exactly the same as data.

The readout simulation is based on `GEANT4` version `9.4.2p02`, using the `QGSP_BERT` physics list which is described in detail in [236, 237]. In summary, it applies a Bertini-style cascade model (tuned to hadron-nucleus scattering data) for hadronic inelastic reactions at energies up to 10 GeV, and switches to the quark-gluon string (QGS) hadronic model above that. `GEANT4` propagates particles in steps of 1 mm, and calculates the probability of an interaction or a decay. Each particle then

either interacts, decays, or propagates, depositing energy as determined by its particle type, charge, and momentum. For each particle in the step, the four-momentum is updated / calculated and the next step determined. The simulation does not save all the information for every particle, due to computing limitations; instead, MINERνA saves "trajectories" corresponding to particles, keeping their initial and final points. In addition, points where a particle changes material, an elastic interaction occurs, the energy loss changes by more than 5 %, or the path length is greater than 20 cm are kept [226].

In the following, the event displays shown are generated with Arachne, MINERνA's proprietary event display software [238]. *Data overlay* refers to the addition of random data gates (from the appropriate running period, determined by the simulation parameters) onto an MC event. Typically, individual interactions / decays are separated out by hit timing, as described above. Figure 4.17 shows an MC HNL decay, $N_4 \to \mu^- + \pi^+$, at $t = 5265.26$ ns after the start of the gate and module 59, overlaid with data (in this case, a through-going rock muon and a cosmic ray that exits out the side of the detector).

4.6 Readout simulation

Given the "true" Monte Carlo (MC) deposited energy in each strip (as given by `GEANT4`), a series of "de-calibration" steps are applied to translate E_{hit} to ADC counts, so that the MC hits can then be processed exactly as data. A detailed overview can be found in [231].

First, E_{hit} is passed to an optical model, which estimates the amount of photons using Birks' law [239], which describes how scintillation photons are quenched by high localised energy deposits:

$$\frac{\mathrm{d}L}{\mathrm{d}x} = A \cdot \frac{\mathrm{d}E}{\mathrm{d}x} \times \left(1 + k_B \frac{\mathrm{d}E}{\mathrm{d}x}\right)^{-1}, \tag{4.6}$$

with L being the fluorescence (amount of light), $\mathrm{d}E/\mathrm{d}x$ the energy deposited per unit length, and k_B Birks' parameter. The term in brackets is a *suppression factor* for the light produced by scintillator fluor. Birks' constant was measured using a

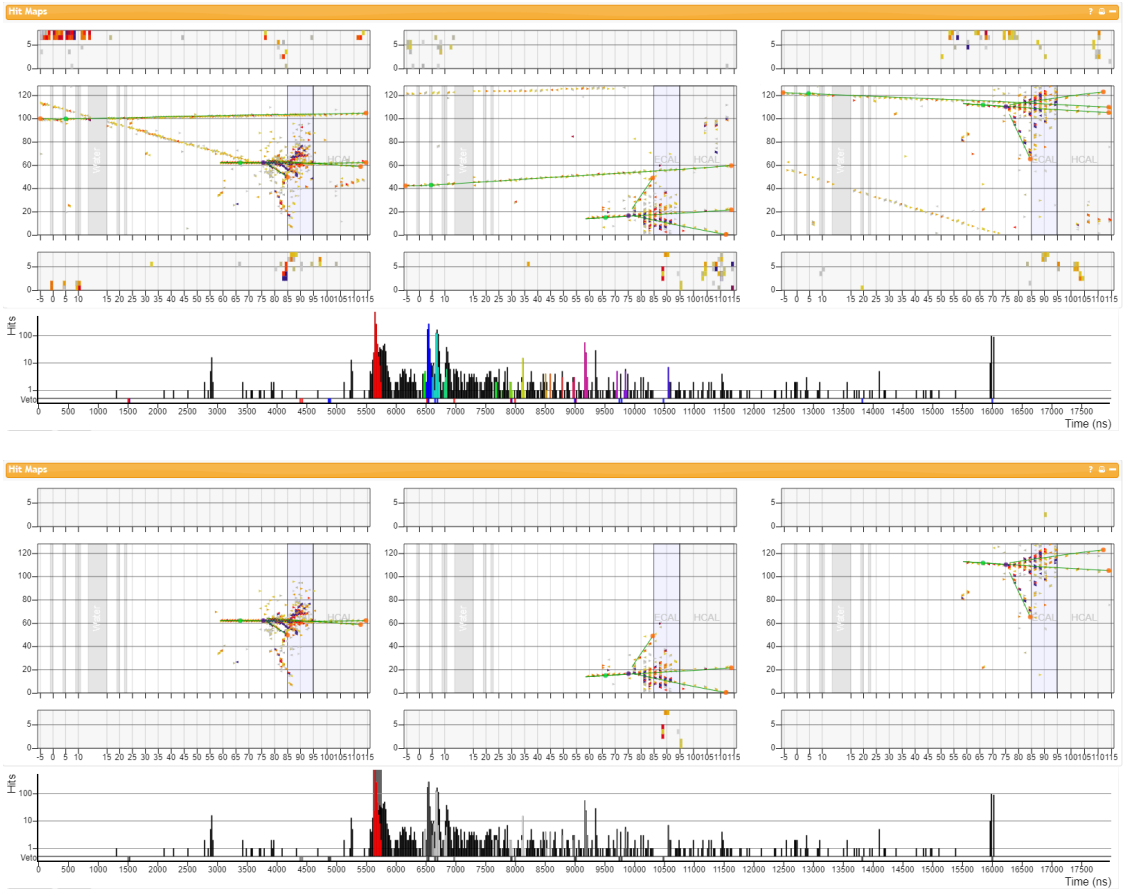


Figure 4.17: A $N_4 \to \mu^- + \pi^+$ decay overlaid with data (top panel); the HNL decay is isolated by focusing on the time slice $5620 \leq t$ (ns) ≤ 5740 (bottom panel). From left to right, the panels are the XZ, UZ, and VZ views where the X, U, and V planes of each module are shown.

test beam at Fermilab's Test Beam Facility in 2010, with a test detector composed of 40 scintillator planes in a "tracker + ECAL" configuration that mimicks the downstream tracker region and ECAL of MINERνA; a sample of protons stopping in the test detector was used to extract Birks' parameter [240]. The best-fit value is 0.0905 ± 0.0150 mm/MeV.

Afterwards, the light is translated to a mean number of PE, $\langle N_{\text{PE}} \rangle$, by applying the inverse of the PE-to-hit-energy factor in Eq. 4.1, along a channel-to-channel smearing factor. This smearing factor attempts to represent the combined uncertainty from the calibration steps; the uncertainty of the strip-to-strip correction $S_i(t)$ (estimated 1.5 %), a statistical fluctuation in the pixel gain of order $3 - 5$ %, and other sources. The aggregate effect of all these sources is determined by matching

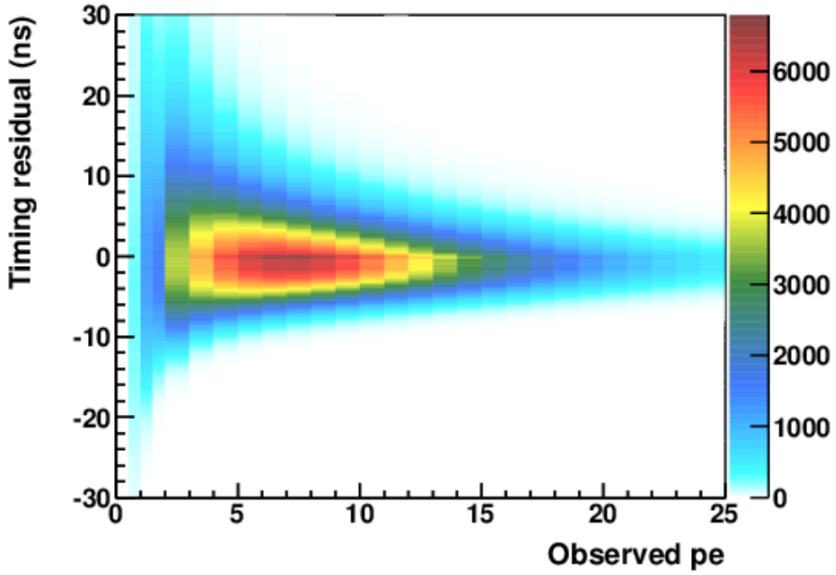


Figure 4.18: Time-smearing distribution from real data, as a function of N_{PE} and $T_{\mathrm{true}} - T_{\mathrm{sim}}$. From [231].

rock muon simulated energy depositions to data; for a collection $\{\sigma\}_i$ such that $\sigma_{\mathrm{min}} = 0, \sigma_{\mathrm{max}} = 0.1$, and $\sigma_{i+1} = \sigma_i + 0.005$, simulated rock muons have their energy smeared by a Gaussian $\mathcal{N}_i\,(\mu_i = 1, \sigma_i)$; then a chi-squared χ^2_i is calculated, and the optimal parameter $\sigma = 0.0557$ is determined from a fit.

To account for statistical fluctuations of N_{PE}, the actual number of PE is sampled from a Poisson distribution of mean $\langle N_{\mathrm{PE}} \rangle$.

Finally, the photo-electrons are converted to raw ADC counts. A PMT model adds cross-talk to the channels and amplifies by the appropriate channel gain $G_i(t)^{-1}$, and then adds random electronics noise to the event. To complete the chain, an ADC model uses the response curves of the TriP-t chips from Fig. 4.6 to go from electronic signal to ADC counts.

There are also two steps to account for the simulation of timing for each hit. First, the MC time is converted to the nearest quarter-tick that the FEBs are sensitive to, and then a random smearing is applied by polling from a data-driven timing distribution, which is a function of PE and $T_{\mathrm{true}} - T_{\mathrm{sim}}$. This histogram is constructed from data at the time the calibration procedure of Eq. 4.2 is run, and is shown in Fig. 4.18. It is least well constrained at low N_{PE}, due to the time slewing described above.

In the next Chapters, the simulation pipeline for the HNL signal, and the MINERνA dataset, will be described.

5 HNL event generation

Contents

5.1 Effective field theory for HNL 76
5.2 HNL production in a beamline 79
5.2.1 Kinematics at production 79
5.2.2 Propagation 82
5.3 Decay 90
5.3.1 Decay kinematics 90
5.3.2 Decay vertex placement 93
5.4 POT estimation 98
5.5 The BeamHNL simulation module 101

During this search for HNL, a general, experiment-agnostic simulation tool (`BeamHNL`) has been developed to describe in a generic way the production, propagation, and decay of HNL from the decays of particles in a neutrino beam. It integrates mature flux and geometry drivers with the framework of the GENIE event generator [241, 242] to produce a realistic record of HNL from their production in the beamline to their decay to visible particles. The theoretical framework for HNL has been extensively studied in the context of the mass generation mechanism and mixing with the Standard Model, as seen in Chapter 2. In particular, given the large multiplicity of low-mass hadrons produced in the NuMI beam (cf. Fig. 3.8), the framework of an effective field theory valid for low-mass HNL (with $M_{\mathrm{N4}} \lesssim m_{\mathrm{K}}$) has been implemented. Here, the workflow for modelling HNL will be discussed, broadly following the structure in [243].

5.1 Effective field theory for HNL

At low HNL masses $M_{\mathrm{N4}} \lesssim \mathcal{O}\,(1\ \mathrm{GeV}/c^2)$, one can integrate out the electroweak gauge bosons and replace them with effective operators, promoting HNL production and decay to effective vertices with appropriate Feynman rules. These vertices will depend, in general, on the kind of meson that touches the vertex; specifically,

its spin and decay constant, which will enter into the invariant matrix element describing the underlying physical process.

In the MINERνA HNL search, the effective field theory described in [129, 144] has been implemented. It describes the decays of pions, kaons, leptons, and heavier mesons to HNL, and the subsequent decays of HNL to SM particles. (Note that there exists a discrepancy on the branching ratios of HNL to vector mesons such as ρ, ω, ϕ; since all of these lie above the kaon mass, the rules in [144] and [129] are treated as equivalent). It was verified that the implementation into `BeamHNL` of the production and decay channels matches the expectation from the effective field theory, as can be seen in Figs. 5.1, 5.2. The predictions for the HNL production

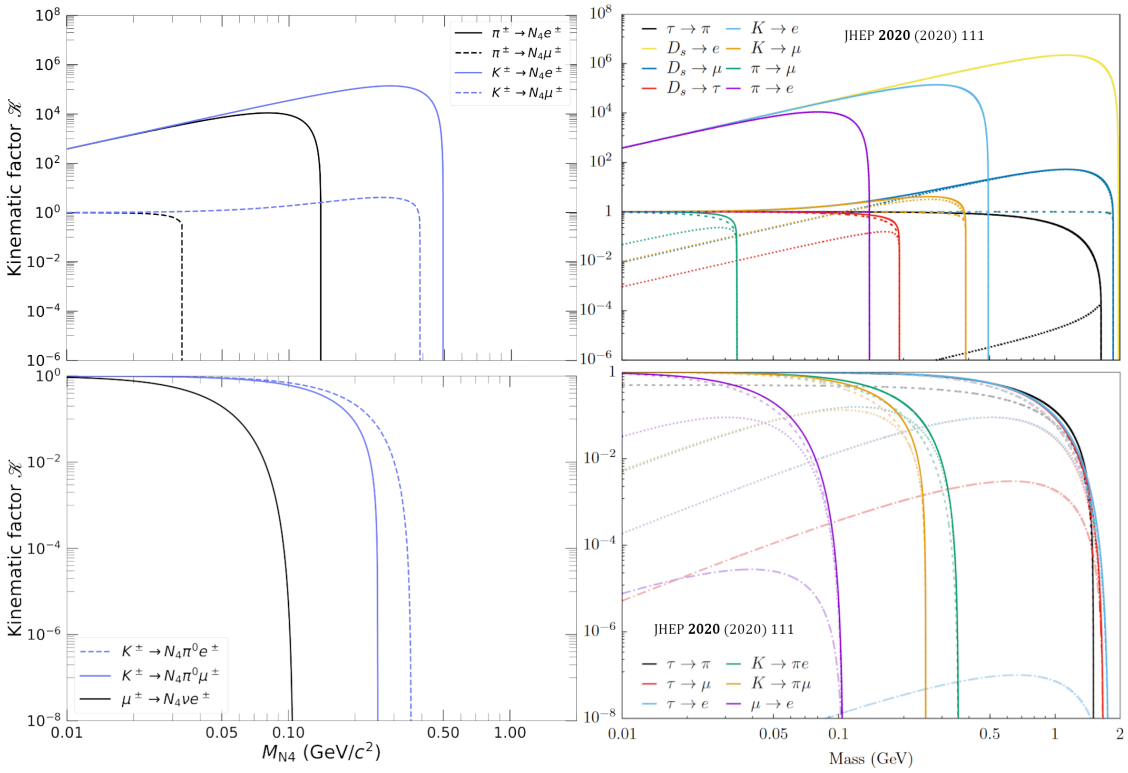
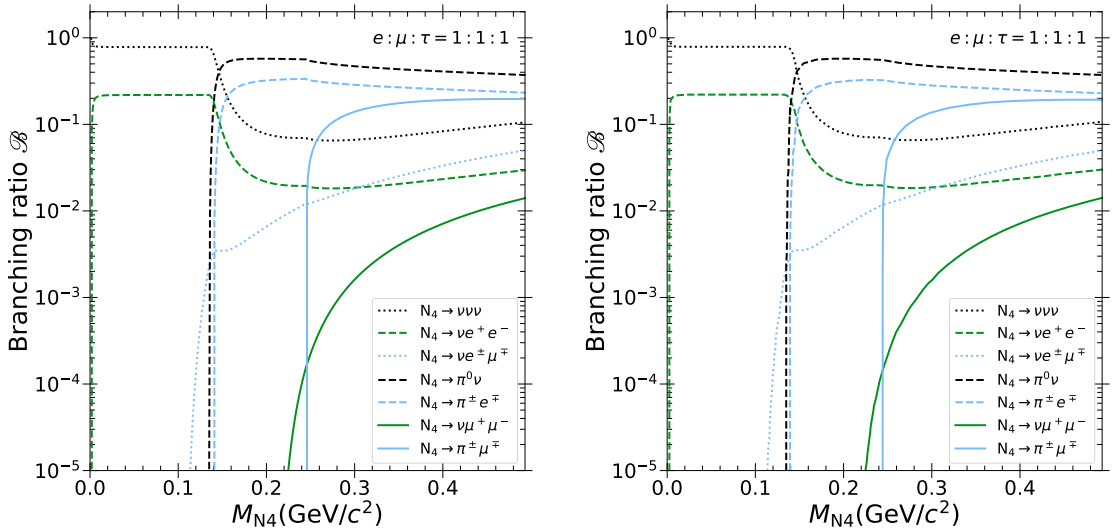


Figure 5.1: HNL production mode comparison. Left panels: `BeamHNL` implementation, right panels: from [144]. Top row: 2-body production channels, bottom row: 3-body channels. The dashed and dotted lines in the right panels are modes with helicity ± 1, the solid lines are the total kinematic factor for each channel. Line styles on the left panels are only for visualisation and refer always to the total kinematic factor.

are given in terms of the *kinematic factor*:

$$\mathcal{K}\left(P \to \mathrm{N}_4 + \dots\right) := \frac{1}{|U_{\alpha 4}|^2} \frac{\Gamma\left(P \to \mathrm{N}_4 + \dots\right)}{\Gamma\left(P \to \nu_\alpha + \dots\right)}. \tag{5.1}$$

This is readily transformed to a branching ratio for the HNL production by multiplying with the branching ratio of the SM process $P \to \nu_a + \ldots$, and the appropriate mixing element $|U_{\alpha 4}|^2$.



(a) (b)

Figure 5.2: Comparison of branching ratios for HNL decay modes; panel (a) is the `BeamHNL` implementation of the decay widths; panel (b) has been adapted from [129]. The modes $N_4 \to L\pi\pi$ are not shown as their branching ratio is negligibly small.

A note about the decay channels $N_4 \to \nu\pi^0\pi^0, \ell^\mp\pi^\pm\pi^0$; despite being *kinematically* allowed, they are heavily suppressed. The mode $N_4 \to \nu_\alpha\pi^0\pi^0$ is subleading in the matrix element $i\mathcal{M} \sim |U_{\alpha 4}|^4$ (because the intermediate fermion in the Feynman diagram is an HNL state), and the other modes are suppressed due to phase space and the pion decay constant entering the matrix element twice. This phase-space suppression was explicitly verified by using the `FeynRules` [244, 245] and `FeynCalc` [246–248] Mathematica packages. The Feynman rules for the $N_4 \to \ell^\mp_\alpha\pi^\pm\pi^0$ at $M_{N4} \lesssim m_K$ were used from the published Lagrangian in the `FeynRules` model database, and the decay width Γ was evaluated; the branching ratios for both $\alpha = e,\ \mu$ are less than 10^{-5}.

5.2 HNL production in a beamline

It is possible to predict the flux of Standard Model neutrinos directly from the simulation of parent hadrons in the beamline, given a point in space. It was also seen that, in the framework of the effective field theory, one may obtain the probability of HNL production from hadron decays. Continuing this line of thought, it will now be demonstrated how a flux of HNL at a point in space is obtained.

5.2.1 Kinematics at production

Suppose that a parent particle P has momentum $\boldsymbol{p}_P$ in the lab frame. To ensure simulation efficiency, it is constrained to decay to an HNL N_4, polling from all kinematically available channels. Fundamentally, there are two categories of HNL production channels: two- and many-body modes. With respect to the variable E^*_{N4} which denotes the HNL's energy in the rest frame of P, two-body decays produce *monochromatic* HNL as a result of four-momentum conservation; in the decay $\mathrm{A} \to \mathrm{B} + \mathrm{C}$ it is always

$$E^*_{\mathrm{B}} = \frac{m_A^2 + m_{\mathrm{B}}^2 - m_{\mathrm{C}}^2}{2m_{\mathrm{A}}}. \tag{5.2}$$

For three- or many-body decays, the decays of a scalar particle are simulated at the phase space level. The `GENIE` phase-space integrator is based on ROOT's [249] phase-space integration, which is based on the GENBOD function using the Raubold and Lynch method, described in [250]. In brief, this method factorises the full n-body phase space integral,

$$R_n = \int_{\mathrm{V}} \delta^{(4)} \left(0 - \sum_{j=1}^{n} q_j \right) \cdot \prod_{i=1}^{n} \mathrm{d}^4 q_i \delta \left(q_i^2 - m_i^2 \right), \tag{5.3}$$

into a series of sequential two-body decays. (Note that q_j represent the momenta in the rest frame of the n-body system). If the decay were, for example, $\mathrm{A} \to \mathrm{B} + \mathrm{C} + \mathrm{D}$, the sequence would be:

- $\mathrm{A} \to m_1 + M_2$, where m_1, M_2 are the invariant masses of the one-/two- body systems left after this decay (and where all the possible combinations of B, C,

D must be considered)

- $M_2 \to m_2 + m_3$, where m_2 and m_3 are the invariant masses of the other two final-state particles.

The above integral is evaluated over all the possible combinations of the final-state particles; this is to say, the invariant mass M_2 is treated as a variable. From this prescription, a weight is obtained for the decay process; from the maximum weight obtained by throwing events from the distribution a suitably large number of times the decay is accepted if its weight is smaller than the maximum weight. (In general, this procedure is referred to as *importance sampling* [250].) This succession of pseudo-"two-body" decays produces the energies of each outgoing particle in the final state.

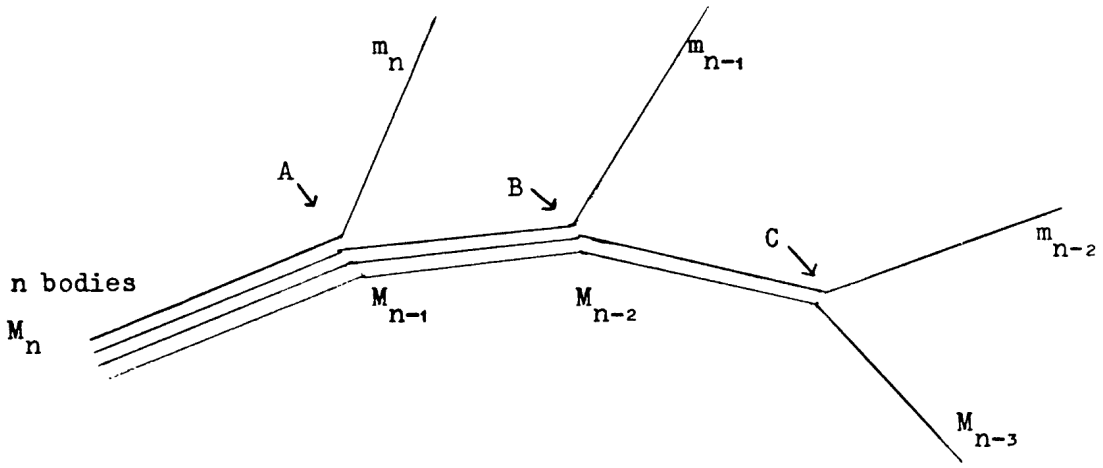


Figure 5.3: A series of successive "two-body" decays simulates the full phase-space integral R_n of Eq. 5.3. From [250].

At this stage, the kinematics are completely isotropic: the procedure to generate the kinematics similarly relies on successive two-body systems. For each particle $i = 1, \ldots, n$ in the final state, first the final state is split into three regions $\{i_- < i\}, \{i\}, \{i < i_+ \leq n\}$. In essence, $\{i_+\}$ tracks that system of particles that have not yet undergone any two-body decay. Then the kinematics of i and of the system $\{i_+\}$ are assigned from an isotropic distribution in the polar coordinates (θ, ϕ) in the rest frame of the two-body system $i, \{i_+\}$, ensuring that the momenta $\boldsymbol{q}_i$, $\boldsymbol{q}_{i_+}$ are balanced. Thirdly, all of the particles in $\{i_-\}, i$, and the stack $\{i_+\}$ are boosted by

the Lorentz-β of the stack $\{i_+\}$, and i increments: the last particle is added to the stack $\{i_-\}$, and the process repeats until all the particles have been generated and the stack $\{i_+\}$ is empty. Finally, the four-momenta $\boldsymbol{q}_i$ of all the particles are boosted into the lab frame, by using the Lorentz-β of the decaying particle with momentum $\boldsymbol{p}_P$.

For the decays of pseudoscalar mesons (the overwhelming majority of parents), this prescription is sufficient to reproduce the kinematics of the produced HNL, as there is no angular correlation between the final-state particles. For the decays of leptons such as the muon, there is an angular correlation originating from the $s = 1/2$ spin; this will produce a $1 + \alpha \cos\theta$ dependence of the final-state particles, where θ is the angle between the rest-frame momentum $\boldsymbol{q}_{N_4}$ and the direction of the polarisation vector, $\boldsymbol{\varepsilon}_P$, of P. α is an $\mathcal{O}(1)$ coefficient describing the polarisation modulus of the HNL: it tends to -1 for $M_{\mathrm{N4}} \to 0$.

However, the effect of this angular dependence in the HNL's production kinematics is diminutive. There are two reasons for this. Firstly, the direction of $\boldsymbol{\varepsilon}_P$ is determined by the conservation of angular momentum in the production of P, which in the case of muon-coupled HNL is overwhelmingly pion and kaon decay. These are either two-body decays, in which the direction of $\boldsymbol{q}_\mu$ in the pion or kaon rest frame must be collinear with $\boldsymbol{\varepsilon}_\mu$, or three-body decays, where there is no *a priori* correlation between linear and intrinsic angular momentum. Crucially, because pions and kaons are pseudoscalaR, these decays are isotropic, meaning that there is no *a priori* direction, in the HNL's own rest frame, that $\boldsymbol{\varepsilon}_\mu$ can take: it is isotropically distributed. Therefore, once the dependence $1 + \alpha \cos\theta$ is integrated over all space, the $\cos\theta$ dependence drops out. Secondly, the Lorentz boost into the lab frame further obfuscates angular correlations between final-state particles. This is less important for the HNL production than the HNL decay case, so one will return to this statement there; in brief, the more energetic the primary beam (and thus the parent particles) the less relevant rest-frame angular correlations will be.

5.2.2 Propagation

In this section, we will work in NEAR coordinates. Suppose that the HNL production vertex is at the position $\boldsymbol{r}_0$ at time $t = 0$. The question to be answered to calculate a flux is twofold: *what energy will an HNL produced at* $\boldsymbol{r}_0$ *have at a point* $\boldsymbol{r}_\mathrm{C}$? *How likely is an HNL to arrive at that point*? The workflow to follow, event by event, will broadly proceed in four steps. First, calculate the angular size of the detector from the point at which the HNL is required to be emitted (the decay of its parent), and the angular deviation ζ between the parent's momentum and the direction the HNL must travel in to be accepted (it will be seen that, for sufficiently massive HNL, there are two, one, or zero distinct regions in the parent rest frame that allow for acceptance, whereas for massless neutrinos there is only ever one single angular region). Second, obtain an energy for the HNL in the parent's rest frame, according to its production channel, and calculate the probability that that HNL may deviate by ζ from the parent's momentum. Third, calculate the rest- and lab-frame four-vectors that are consistent with the lab-frame direction of the HNL, and set the HNL momentum accordingly. Fourth, calculate the probability the HNL at this energy will be accepted.

This is an important subtlety: imposing the geometric constraint of an HNL arriving at a specific point in space (tantamount to evaluating the flux of HNL at a detector location) translates to a constraint on the four-momentum of the HNL. This is the reason behind the difference of the shapes of Figs. 3.6 and 3.8, in the case of SM neutrinos. As it will be seen, there is an additional complexity with respect to the HNL's kinematics.

One can write the relation between the rest-frame energy E^*_{N4} and lab-frame energy E_{N4} of the heavy neutrino by using the Lorentz boost: the boost matrix for a generic velocity $\boldsymbol{\beta}$ is written, in its covariant form,

$$\Lambda^{\rho}_{\ \sigma}(\boldsymbol{\beta}) = \begin{pmatrix} \gamma & -\gamma\beta_x & -\gamma\beta_y & -\gamma\beta_z \\ -\gamma\beta_x & 1+(\gamma-1)\beta_x^2/\beta^2 & (\gamma-1)\beta_x\beta_y/\beta^2 & (\gamma-1)\beta_x\beta_z/\beta^2 \\ -\gamma\beta_y & (\gamma-1)\beta_x\beta_y/\beta^2 & 1+(\gamma-1)\beta_y^2/\beta^2 & (\gamma-1)\beta_y\beta_z/\beta^2 \\ -\gamma\beta_z & (\gamma-1)\beta_x\beta_z/\beta^2 & (\gamma-1)\beta_y\beta_z/\beta^2 & 1+(\gamma-1)\beta_z^2/\beta^2 \end{pmatrix}, \tag{5.4}$$

where $\gamma^{-2} = 1-\beta^2$ and the axes x, y, z refer to the stationary observer; Greek indices indicate spacetime. The matrix Λ is invertible, with $\Lambda^{-1}(\boldsymbol{\beta}) = \Lambda(-\boldsymbol{\beta})$. This means that transformations between the rest ($\boldsymbol{q}^\sigma$) and lab ($\boldsymbol{p}^\rho$) frames can be written as

$$\boldsymbol{p}^\rho = \Lambda^\rho{}_\sigma\,(-\boldsymbol{\beta})\,\boldsymbol{q}^\sigma \Leftrightarrow \Lambda_\rho{}^\sigma\,(\boldsymbol{\beta})\,\boldsymbol{p}^\rho = \boldsymbol{q}^\sigma. \tag{5.5}$$

One can use Eqs. 5.4 and 5.5 to relate E^*_{N4} to E_{N4} as follows:

$$\begin{aligned} \boldsymbol{q}^\sigma &= \Lambda_\rho{}^\sigma(\vec{\boldsymbol{\beta}}_P)\boldsymbol{p}^\rho \\ \Leftrightarrow \begin{pmatrix} E^*_{\mathrm{N4}} \\ \vec{\boldsymbol{q}} \end{pmatrix} &= \begin{pmatrix} \gamma & -\gamma\vec{\boldsymbol{\beta}}_P \\ -\gamma\vec{\boldsymbol{\beta}}_P & \widetilde{\Lambda}_i{}^j(\vec{\boldsymbol{\beta}}_P) \end{pmatrix} \begin{pmatrix} E_{\mathrm{N4}} \\ \vec{\boldsymbol{p}} \end{pmatrix} \\ \Rightarrow E^*_{\mathrm{N4}} &= \gamma E_{\mathrm{N4}} - \gamma\vec{\boldsymbol{p}}\cdot\vec{\boldsymbol{\beta}}_P \\ &= \gamma\left(E_{\mathrm{N4}} - \|\vec{\boldsymbol{p}}\|\left\|\vec{\boldsymbol{\beta}}_P\right\|\cos\zeta\right) \\ &= \gamma E_{\mathrm{N4}}\left(1 - \beta_P\beta_{\mathrm{N4}}\cos\zeta\right), \end{aligned} \tag{5.6}$$

where $\widetilde{\Lambda}_i{}^j$ stands for the 3×3 spatial block of Eq. 5.6 and ζ is the opening angle between the neutrino momentum $\vec{\boldsymbol{p}}$ and the parent's momentum $\vec{\boldsymbol{p}}_P$, *in the lab frame*. Likewise, β_{N4} refers to the *lab-frame* momentum of the neutrino. If the neutrino were massless, one may substitute $\beta_{\mathrm{N4}} = 1$ in 5.6, recovering the closed form

$$E^*_{\mathrm{N4}} = \gamma E_{\mathrm{N4}}\left(1 - \beta_P\cos\zeta\right), \tag{5.7}$$

applicable for extraction of SM neutrino fluxes. Generally, the quantities $\beta_P, \cos\zeta$ are known from the beamline simulation, and E^*_{N4} is determined from the neutrino-generating decay, leading to a well-defined lab-frame energy.

However, one will notice that for a massive neutrino the unknown *a priori* quantity β_{N4} is an extra, seemingly unconstrained, degree of freedom. This is where the subtlety of the geometric constraint enters the picture; our discussion will broadly mirror the one in [251]. Suppose that an HNL of mass M_{N4} has some momentum $\boldsymbol{q}_{\mathrm{N4}}$ in the rest frame of its parent P; equivalently, suppose that it is emitted with energy E^*_{N4} and at some angle Θ, where the case $\Theta = 0$ corresponds to the lab-frame relation $\boldsymbol{p}_{\mathrm{N4}} \parallel \boldsymbol{p}_P$. For simplicity, suppose that the rest-frame momentum $\boldsymbol{q}_{\mathrm{N4}}$ lies on the $y'z'$ plane (primes denoting rest-frame axes), and that $\boldsymbol{\beta}_P = \beta_P\hat{\boldsymbol{z}}$. Using Eq. 5.5, one finds (dropping the index N4, which will be assumed

for all quantities unless otherwise specified):

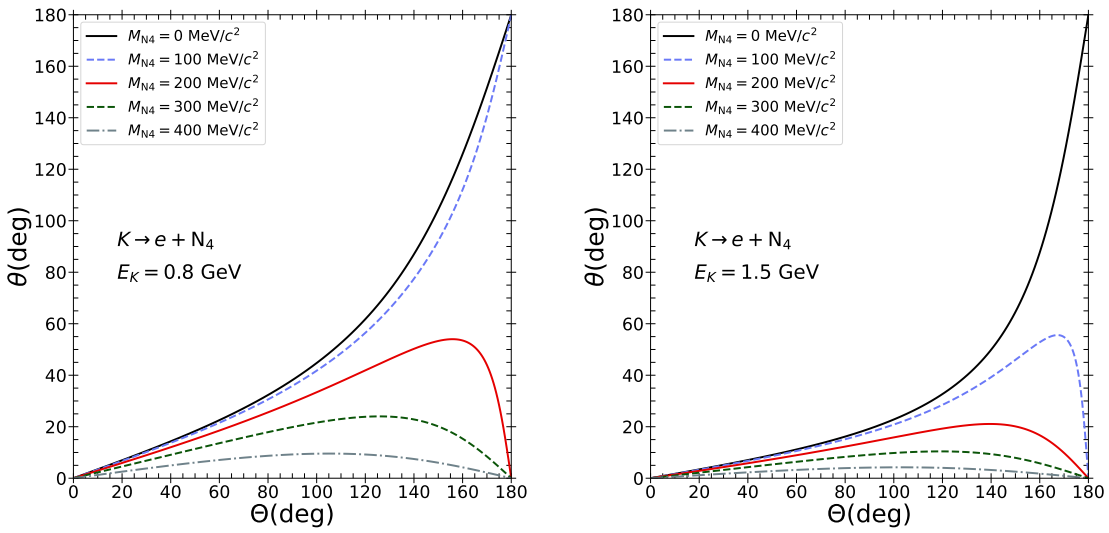
$$\begin{aligned}
\begin{pmatrix} E \\ p_x \\ p_y \\ p_z \end{pmatrix} &= \Lambda(-\boldsymbol{\beta}_P) \begin{pmatrix} E^* \\ 0 \\ q\sin\Theta \\ q\cos\Theta \end{pmatrix} \\
&= \begin{pmatrix} \gamma_P & 0 & 0 & \gamma_P\beta_P \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \gamma_P\beta_P & 0 & 0 & \gamma_P \end{pmatrix} \begin{pmatrix} E^* \\ 0 \\ q\sin\Theta \\ q\cos\Theta \end{pmatrix} \\
&= \begin{pmatrix} \gamma_P(E^* + \beta_P q\cos\Theta) \\ 0 \\ q\sin\Theta \\ \gamma_P(\beta_P E^* + q\cos\Theta) \end{pmatrix} \\
\Rightarrow \tan\theta = \frac{p_y}{p_z} &= \frac{q\sin\Theta}{\gamma_P(\beta_P E^* + q\cos\Theta)}.
\end{aligned} \tag{5.8}$$

It is evident that a suitable rotation maps the $\hat{\boldsymbol{z}}$ axis to any desired $\widehat{\boldsymbol{Z}} \parallel \boldsymbol{\beta}_P$ axis, without changing the result of Eq. 5.8; likewise, a rotation about the z' axis maps $\boldsymbol{q}_{\mathrm{N4}}$ to any combination of $q_{x'}, q_{y'}$. This means that the function $\theta(\Theta)$ can be interpreted as the polar opening angle between the parent's lab-frame momentum, $\boldsymbol{p}_P$, and the lab-frame momentum $\boldsymbol{p}_{\mathrm{N4}}$ of the HNL, if the HNL were emitted at a polar angle Θ in P's rest frame. Figure 5.4 shows the behaviour of $\theta(\Theta)$ as Θ varies from $0 \to 180$ degrees: each curve corresponds to a progressively higher HNL mass (the rest-frame HNL energy increases slightly as the mass increases, but the dominant behaviour is given by β_{N4} which scales roughly as M_{N4}^{-1}).

For massless neutrinos (black curves at the top), θ is *monotonic*, corresponding to exactly one solution of Eq. 5.11 (to be introduced shortly). However, as the mass increases beyond a critical mass M_{crit}, θ loses its monotonicity and its maximum drops below 180 degrees.

The critical mass is that mass, such that the velocity β^*_{N4} *in the parent's rest frame* is smaller than the velocity β_P of the parent *in the lab frame*. This splits the function's image $[0, 180]$ into two regions:

- $\theta \in [0, \theta_{\mathrm{max}})$: For each θ_0 in this region, there exist two function pre-images Θ_-, Θ_+ such that $\theta(\Theta_\pm) = \theta_0$. HNL emitted in the backwards direction are "swept forward" by their parent's Lorentz boost.



(a) (b)

Figure 5.4: Collimation effect for HNL of various masses, from the decays of a kaon $K \to e + \mathrm{N}_4$ with $E_K = 0.8$ GeV (a), 1.5 GeV (b). The loss of monotonicity of the lab-frame emission angle θ is evident as the mass increases.

- $\theta \in [\theta_{\max}, 180]$: There exists no Θ such that $\theta(\Theta) = \theta_0$. The HNL does not carry enough kinetic energy to sufficiently deviate from its parent's momentum.

This is the *collimation effect* of massive neutrinos. The condition for monotononicity to be lost is

$$\beta^*_{\mathrm{N4}} \leq \beta_P \Leftrightarrow \|\boldsymbol{q}_{\mathrm{N4}}\| \leq E^*_{\mathrm{N4}} \cdot \beta_P. \tag{5.9}$$

Figure 5.5 shows the first case as two regions on the rest-frame unit sphere (with the coloured centre being the HNL production vertex), where if the HNL is emitted in either of these two angular regions it will be boosted into the same angular region in the lab frame.

It is evident from Fig. 5.5 that the function $\theta(\Theta)$ is invertible, for $M_{\mathrm{N4}} \leq M_{\mathrm{crit}}$: for $M_{\mathrm{N4}} > M_{\mathrm{crit}}$, the restricted regions $\theta_+(\Theta)$ and $\theta_-(\Theta)$ (+: forward, −: backward) are also invertible. Thus Eq. 5.8 may be inverted to produce two distinct solutions

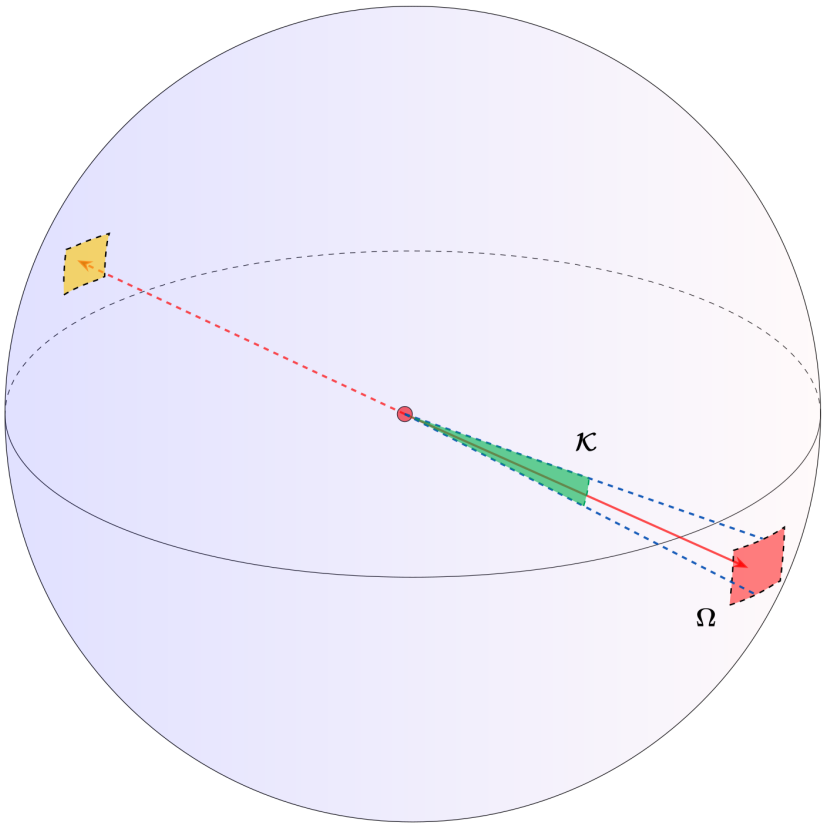


Figure 5.5: Backwards emission of HNL in the rest frame can lead to the same lab-frame momentum direction; see text.

(once again dropping the index N4):

$$\begin{aligned}
\Theta_+(t) &= \tan^{-1}\left[t\gamma_P \cdot \frac{q^2 + \beta_P E^* \sqrt{s}}{(q - t\beta_P\gamma_P E^*)(q + t\beta_P\gamma_P E^*)}\right], \\
\Theta_-(t) &= \tan^{-1}\left[t\gamma_P \cdot \frac{q^2 - \beta_P E^* \sqrt{s}}{(q - t\beta_P\gamma_P E^*)(q + t\beta_P\gamma_P E^*)}\right], \\
\sqrt{s} &= \sqrt{q^2(1 + \gamma_P^2 t^2) - t^2\beta_P^2\gamma_P^2 (E^*)^2},
\end{aligned} \tag{5.10}$$

where $t = \tan(\theta)$. If $M_{\mathrm{N4}} \leq M_{\mathrm{crit}}$, only Θ_+ survives, and Θ_- is set to 0.

Given, then, a desired lab-frame emission angle θ_0, one may find that there are zero, one, two, or (in the trivial case where $\theta_0 = 0$ and the HNL has no kinetic energy in the parent's rest frame) infinitely many solutions Θ. This is as expected; since Λ is invertible, the equation

$$\boldsymbol{\xi}^\sigma = \Lambda_\rho^{\ \sigma}(\boldsymbol{\beta})\left(\Delta t, \boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0\right)^T \tag{5.11}$$

has zero, one, two, or trivially infinite solutions, where $\boldsymbol{\xi}^\sigma$ is the rest-frame four-

vector of the neutrino's trajectory and Δt is the time taken for the neutrino of lab-frame energy E_{N4} to traverse the distance $\|\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0\|$. In the case of zero solutions, the HNL is obviously not accepted at all, and the trivial case is not interesting. For the more interesting case of one or two solutions, one may term the solutions $\Theta_0^{(+)}$ and $\Theta_0^{(-)}$, where $+(-)$ refers to a Θ in the region where the function is increasing (decreasing). The relation 5.8 may be inverted to yield

$$\cos\Theta_0^{(\pm)} = \frac{-t_0\beta_P\gamma_P^2 E^* \pm \sqrt{q^2\left(1+t_0\gamma_P^2\right) - t_0\beta_P^2\gamma_P^2\left(E^*\right)^2}}{q\left(1+t_0\gamma_P^2\right)}, \tag{5.12}$$

from which the lab-frame energy $E_0^{(\pm)}$ can be obtained as

$$E_0^{(\pm)} = \gamma_P\left(E^* + \beta_P q\cos\Theta_0^{(\pm)}\right). \tag{5.13}$$

A further subtlety arises from the observation that the derivative $\theta'(\Theta) = \mathrm{d}\theta/\mathrm{d}\Theta$ is different at $\Theta = \Theta_0$ for a SM neutrino and an HNL. Supposing one wanted to evaluate the likelihood of a neutrino of mass $m \geq 0$ being accepted in the angular region $(\theta_0 - \epsilon, \theta_0 + \epsilon), \epsilon > 0$, then the associated pre-image $\mathcal{I}_\epsilon^{(m)} = \theta_{(m)}^{-1}\left((\theta_0 - \epsilon, \theta_0 + \epsilon)\right)$ can be written as

$$\mathcal{I}_\epsilon^{(m)} = \theta_{(m)}^{-1}(\theta_0 + \epsilon) - \theta_{(m)}^{-1}(\theta_0 - \epsilon), \tag{5.14}$$

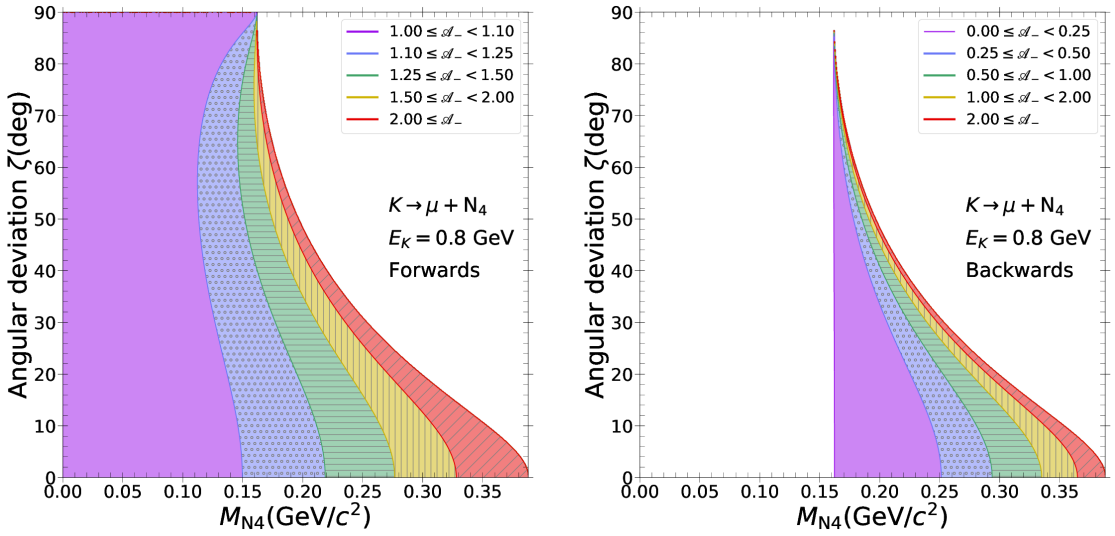
with θ^{-1} is given by Eq. 5.12. If $M_{\mathrm{N4}} \geq M_{\mathrm{crit}}$ it is understood that

$$\mathcal{I}_\epsilon^{(M_{\mathrm{N4}});\pm} = \pm\left(\theta_{(M_{\mathrm{N4}});\pm}^{-1}(\theta_0 + \epsilon) - \theta_{(M_{\mathrm{N4}});\pm}^{-1}(\theta_0 - \epsilon)\right). \tag{5.15}$$

Taking the limit of $\epsilon \to 0$ it is (dropping the index M_{N4})

$$\begin{aligned}
\lim_{\epsilon\to 0}\frac{\mathcal{I}_\epsilon^{\pm}(\theta_0)}{\mathcal{I}_\epsilon(\theta_0)} &= \pm\lim_{\epsilon\to 0}\frac{\theta_\pm^{-1}(\theta_0+\epsilon) - \theta_\pm^{-1}(\theta_0-\epsilon)}{\theta_{(\mathrm{SM})}^{-1}(\theta_0+\epsilon) - \theta_{(\mathrm{SM})}^{-1}(\theta_0-\epsilon)} \\
&= \pm\lim_{\epsilon\to 0}\frac{\left(\theta_\pm^{-1}(\theta_0+\epsilon) - \theta_\pm^{-1}(\theta_0)\right) + \left(\theta_\pm^{-1}(\theta_0) - \theta_\pm^{-1}(\theta_0-\epsilon)\right)}{\left(\theta_{(\mathrm{SM})}^{-1}(\theta_0+\epsilon) - \theta_{(\mathrm{SM})}^{-1}(\theta_0)\right) + \left(\theta_{(\mathrm{SM})}^{-1}(\theta_0) - \theta_{(\mathrm{SM})}^{-1}(\theta_0-\epsilon)\right)} \\
&= \pm\frac{2\left(\theta_\pm^{-1}(\theta_0)\right)'}{2\left(\theta_{(\mathrm{SM})}^{-1}(\theta_0)\right)'} = \frac{\left|\theta'_\pm(\Theta_0)\right|^{-1}}{\left|\theta'_{(\mathrm{SM})}(\Theta_0)\right|^{-1}}.
\end{aligned} \tag{5.16}$$

This ratio shows that an HNL at energy $E_\pm^{\mathrm{N4}}$ has a different *inherent* probability to



(a) (b)

Figure 5.6: Acceptance correction, $\mathcal{A}$, contours for forwards (a) and backwards (b) neutrino emission. Note that $\mathcal{A}_+ \geq 1$ always, due to the flatter collimation function in the forward angular region. The enveloping curve is the maximum angular deviation, $\zeta_{\max} = \tan^{-1}\left\{q\gamma_P^{-1}\left[(\beta_P E^*)^2 - q^2\right]^{-1/2}\right\}$.

be accepted in a small angular region around θ_0 than a SM neutrino, because the allowed angular range in the parent's rest frame is larger. This is introduced in the simulation as an explicit *acceptance correction*, $\mathcal{A}$. Figure 5.6 shows contours of the acceptance corrections $\mathcal{A}_\pm$ for forwards and (for masses $M_{\mathrm{N4}} \geq M_{\mathrm{crit}}$) backwards neutrino emission, as a function of mass and deviation angle ζ.

One may then straightforwardly choose between the forwards and backwards solutions of Eq. 5.11, of energies $E_0^\pm$, for any deviation angle θ_0. By applying Equation 5.16, the conditional probabilities for an accepted neutino to be forwards (backwards) emitted are just $\mathcal{A}_\pm\left(M_{\mathrm{N4}};\theta_0\right)\cdot\mathcal{P}\left(\mathrm{SM}\right)\left(\theta_0\right)$, which means the conditional probability for an HNL to be forwards (backwards) emitted with energy $E_\pm^0$ is

$$\begin{cases} \mathcal{P}_0\left(\text{forwards}\right) = \frac{\mathcal{A}_+\cdot\mathcal{P}(\mathrm{SM})(\theta_0)}{(\mathcal{A}_+ + \mathcal{A}_-)\cdot\mathcal{P}(\mathrm{SM})(\theta_0)} = \frac{\mathcal{A}_+}{\mathcal{A}_+ + \mathcal{A}_-}, \\ \mathcal{P}_0\left(\text{backwards}\right) = 1 - \mathcal{P}_0\left(\text{forwards}\right) \end{cases} . \tag{5.17}$$

Figure 5.7 shows the conditional probability $\mathcal{P}_0\left(\text{forwards}\right)$. It is $\mathcal{P}_0\left(\text{forwards}\right) \geq 0.5$, since the collimation function is always flatter in the forwards angular region than in the backwards one.

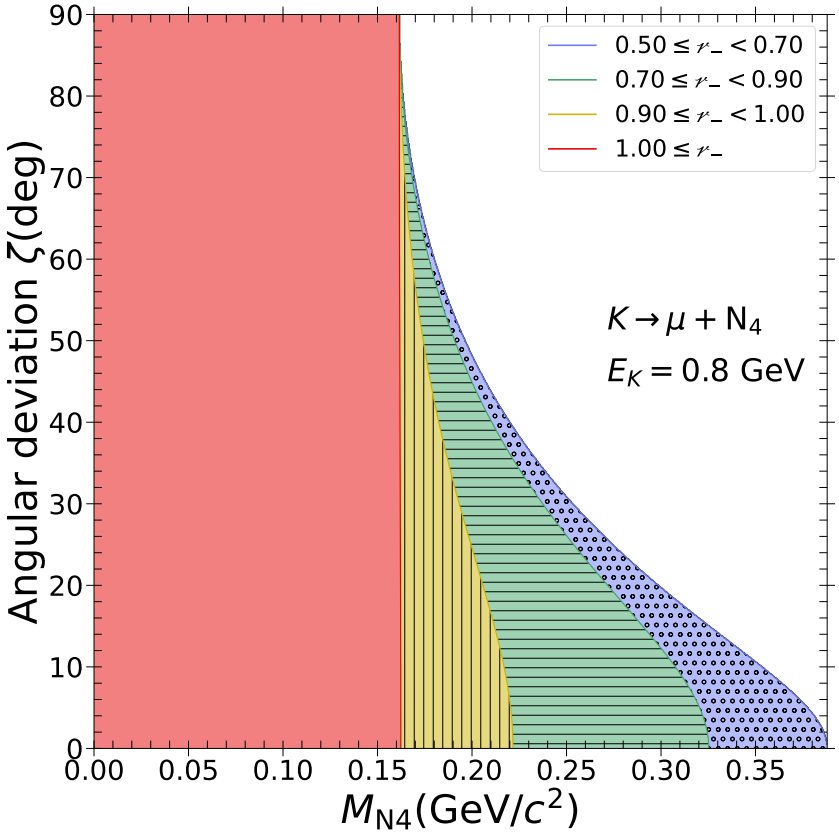


Figure 5.7: Conditional probability $r_+ = \mathcal{P}_0$ (forwards)for forwards HNL emission, as a function of M_{N4} and emission angle $\zeta = \theta_0$.

The calculation of the "base" angular size in the rest frame, Ω, of the detector is given by

$$\Omega = \omega\left(\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0\right) \cdot \mathcal{B}^2\left(\boldsymbol{\beta}_P\right), \tag{5.18}$$

where the lab-frame angular size ω is

$$\begin{aligned} \omega\left(\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0\right) &= 4\pi \times \left[\frac{1}{2}\left(1 - \cos\kappa\right) + \mathcal{O}\left(\kappa^4\right)\right], \\ \tan\kappa &= S / \left\|\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0\right\|. \end{aligned} \tag{5.19}$$

κ is the viewing angle from $\boldsymbol{r}_0$ of a circle with radius S: for a detector of transverse size S_0 (transverse here referring to the parent's momentum), and under the assumption that $\kappa \ll 1$, it is $\tan\kappa \sim \kappa = S / \left\|\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0\right\|$ so that $\omega \mapsto \omega \times \frac{S_0}{S}$.

The "boost factor" $\mathcal{B}\left(M_{\mathrm{N4}}, \boldsymbol{\beta}_P\right)$ is given by Eq. 5.6:

$$\mathcal{B} = \frac{E_{\mathrm{N4}}}{E^*_{\mathrm{N4}}} = \frac{1}{\gamma_P\left(1 - \beta_P \beta_{\mathrm{N4}} \cos\zeta\right)}. \tag{5.20}$$

Finally, then, one can answer how likely any HNL, at energy $E_\pm$, produced at $\boldsymbol{r}_0$ is to be accepted at $\boldsymbol{r}_\mathrm{C}$ by a detector of transverse size S_0: it is simply

$$\mathcal{P}\left(\text{HNL accepted, } \mathrm{E}_\pm; M_\mathrm{N4}, \boldsymbol{\beta}_P, \boldsymbol{r}_\mathrm{C} - \boldsymbol{r}_0\right) = \frac{1}{4\pi}\Omega_\pm \mathcal{A}_\pm = \frac{1}{4\pi}\omega \mathcal{B}_\pm^2 \mathcal{A}_\pm. \tag{5.21}$$

5.3 Decay

5.3.1 Decay kinematics

Similarly to the HNL production mechanism, the choice of final-state channel to decay an HNL into is made by polling from all the kinematically available channels, and constructing the branching ratios $\mathcal{B}_i$ for each available channel i (cf. Fig. 5.2).

The kinematics of the decay are determined by the HNL momentum $\boldsymbol{p}_\mathrm{N4}$ as well as the HNL polarisation vector $\boldsymbol{\varepsilon}_\mathrm{N4} = |\mathbb{P}|\,\widehat{\boldsymbol{\varepsilon}}_\mathrm{N4}$. Because of the two helicity components of the right-chiral HNL, it is not created in a fully polarised state; rather, it corresponds to a *polarisation magnitude* $|\mathbb{P}|$ which is normally < 1. The magnitude may be written, in the case of a two-body decay $P \to \mathrm{N}_4 + \ell$ [252],

$$\mathrm{sgn}\left(M_\mathrm{N4} - m_\ell\right)|\mathbb{P}| = -\frac{\left(m_\ell^2 - M_\mathrm{N4}^2\right)\sqrt{\lambda\left(m_P^2, M_\mathrm{N4}^2, m_\ell^2\right)}}{m_P^2\left(M_\mathrm{N4}^2 + m_\ell^2\right) - \left(m_\ell^2 - M_\mathrm{N4}^2\right)^2}, \tag{5.22}$$

with sgn being the sign function and $\lambda(x, y, z) = x^2 + y^2 + z^2 - 2(xy + yz + zx)$ the Källén (or "triangle" function). The left-hand side of Eq. 5.22 can be understood as the expectation value for the helicity operator, over an ensemble of HNL produced in the same reaction $P \to \mathrm{N}_4 + \ell$ (as each individual HNL must, by momentum conservation, be either in the $h = +1$ or $h = -1$ state).

Similarly, if the HNL then decays to two bodies as $\mathrm{N}_4 \to D + \ell'$, the above factor is multiplied by a factor

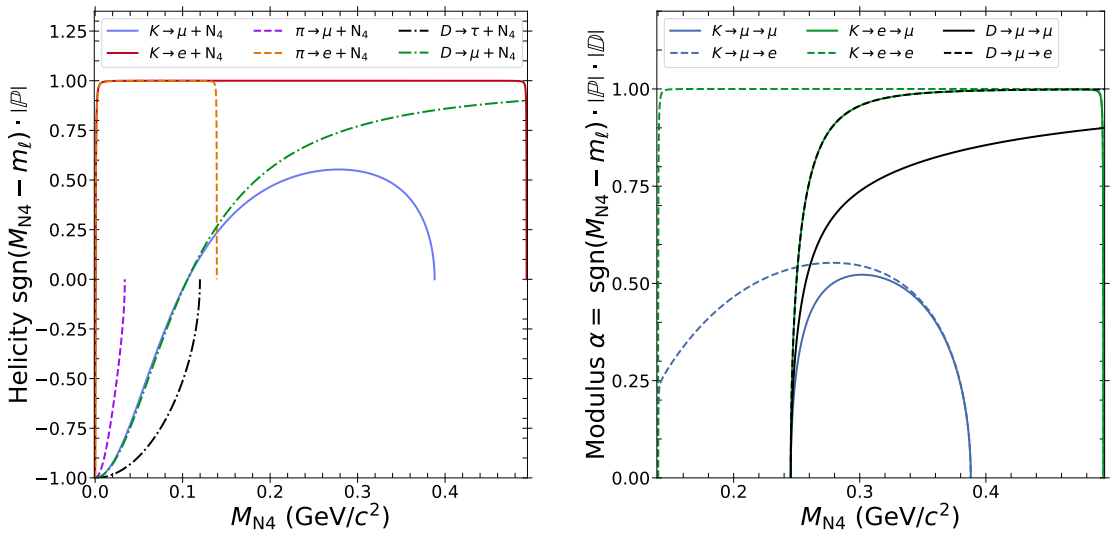
$$|\mathbb{D}| = \frac{\left(M_\mathrm{N4}^2 - m_D^2\right)\sqrt{\lambda\left(M_\mathrm{N4}^2, m_{\ell'}^2, m_D^2\right)}}{\left(M_\mathrm{N4}^2 - m_D^2\right)^2 - m_D^2\left(M_\mathrm{N4}^2 + m_{\ell'}^2\right)}, \tag{5.23}$$

so that for a particle (antiparticle) HNL produced by and decaying through a two body decay, the decay width can be written as

$$\frac{\mathrm{d}\Gamma}{\mathrm{d}\cos\Theta} \propto 1 \pm \alpha\cos\Theta = 1 \pm \mathrm{sgn}\left(M_\mathrm{N4} - m_\ell\right)|\mathbb{P}|\,|\mathbb{D}|\cos\Theta, \tag{5.24}$$

where plus (minus) refers to particle and antiparticle solutions, respectively, and Θ is the opening angle in the HNL rest frame between particle momentum and $\varepsilon_{\mathrm{N4}}$.

Figure 5.8 shows the polarisation magnitude $|\mathbb{P}|$ (left) and modulus $|\mathbb{D}|$ (right) for various two-body decays. A "helicity flip" is evident for the decays $P \to \mathrm{N}_4 + e$, as $m_e \ll M_{\mathrm{N4}}, m_P$: effectively, the left-chiral component of the HNL (the element that mixes with the SM) behaves as a left-handed $h = -1$ neutrino for very small masses where the helicity and chirality operator coincide, and as soon as the positron mass becomes negligible with respect to the neutrino mass it is the *positron* that must be helical. Since the positron was produced by a weak decay it must be right-chiral, ergo $h = +1$, and its spin points the same direction as its momentum: by momentum conservation, it is the same for the HNL, which now has the $h = +1$ component dominate. Naturally, this picture is inverted in the case $P \to \overline{\mathrm{N}_4} + e$ since the electron must be left-chiral.



(a) (b)

Figure 5.8: Helicities (a) and angular decay width coefficient α (b) for various two-body decays.

The case of three-body reactions is more complicated, with the full matrix element being too unwieldy to yield a simple closed form such as that of Equations 5.22, 5.23; work has been done to obtain these angular correlations in [144, 155, 253, 254].

Generally, these may be parametrised with respect to three angles and two invariant masses of the three-body system, in the case of leptonic decays of the HNL.

There is a salient distinction, in principle, between Dirac and Majorana HNL decays, as for Majorana particles there is no distinction between particle and antiparticle. In two-body decays of Majorana HNL to neutral particles, i.e. $\mathrm{N}_4 \to \nu + D$, the angular distribution of final-state particles is isotropic in the rest frame of N_4; if the final-state particles are charged, i.e. $\mathrm{N}_4 \to \ell + D$, a charge-blind detector would see the same as if the distribution were isotropic. Charge identification could differentiate, in principle, the $1 + \alpha \cos\Theta$ component (from $\ell^- D^+$ decays) and the $1 - \alpha \cos\Theta$ component (from $\ell^+ D^-$ decays).

However, this detection of polarisation is obfuscated by two effects. Firstly, one has the impact of the parent's Lorentz boost; the larger this boost is, the more "washed out" the rest-frame angular dependence will be. The final-state particles' momenta get forced closer to $\boldsymbol{p}_{\mathrm{N4}}$ with larger Lorentz boosts, making their lab-frame opening angle smaller. Figure 5.9 shows the angular distribution, in the lab frame, of three different contrived cases of $\mathrm{N}_4 \to \pi^+ + e^-$ decay with three masses $M_{\mathrm{N4}} = 150,\ 275,\ 450\ \mathrm{MeV}/c^2$, for the NuMI fluxes passing through the MINERνA tracker, where the directions of the spins are aligned for all events. The difference between these probability distributions is evident only in the last few bins $\cos\theta \sim 1$, corresponding to very small angles: this is the signature of a strong Lorentz boost in the forward direction washing out rest-frame kinematics, irrespective of how strong the effect appears to be in the rest frame of the parent.

Secondly, there is the question of direction of the polarisation vector, $\widehat{\boldsymbol{\varepsilon}}_{\mathrm{N4}}$: this is determined, in the rest frame of the HNL's parent, by the momentum $\boldsymbol{q}_{\mathrm{N4}}$, for a two-body decay $P \to \mathrm{N}_4 + \ell$. This occurs because P is pseudoscalar. Owing to the collinearity of $\boldsymbol{q}_{\mathrm{N4}}$ and $\boldsymbol{q}_\ell$, $\widehat{\boldsymbol{\varepsilon}}_{\mathrm{N4}}$ points the same direction in the rest frame of N_4. Due to Eq. 5.11, the direction of the HNL in P's rest frame depends on $\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0$: this is generally a weak constraint, however, as the angle of the polarisation vector with the beam axis is normally small.

Nevertheless, the collimation effect (cf. Fig. 5.4) enables backwards-emitted HNL to reach the detector, meaning HNL whose polarisation vector is pointing backwards are also accepted, for masses large enough; at parent energies large enough, this "critical mass" is small, meaning that practically all interesting masses $M_{\mathrm{N4}} \gtrsim m_\pi + m_e$ have this confusion from backwards-emitted HNL. The effect of this confusion can clearly be seen in Figure 5.10, where for three different decay chains and masses, the angular dependence of HNL decays in its rest frame are shown.

From Eq. 5.24 and Fig. 5.8(b), one can read off the polarisation moduli α for the decay chains $K^+ \to \ell^+ + \mathrm{N}_4 \to \pi^+ + e^-$, where $\ell = \mu$ for $M_{\mathrm{N4}} = 150, 275$ MeV/c^2 and e for 450 MeV/c^2. In Figure 5.10, for each mass / decay chain, the polarised decay rate $\mathrm{d}\Gamma/\mathrm{d}\cos\Theta$ expectation (dashed-dotted line) and actual value (histogram) are shown, for events in the MINERνA scintillator tracker and using the NuMI FHC flux. The angle Θ is the opening angle between the electron's momentum $\boldsymbol{q}_e$, and the unit vector $\widehat{Z'} = (0, 0, 1)$. Though the polarisation modulus increases by almost a factor of two for each successive mass shown, the collimation effect washes out the angular dependence almost completely.

5.3.2 Decay vertex placement

Throughout this section we will work in USER coordinates.

Given a detector volume V, it is possible to determine the kinematics of an HNL passing through the point $\boldsymbol{r}_{\mathrm{C}} \in V$: one must still obtain a realistic picture of the location of HNL decays in the detector. As the HNL is unstable, the probability of its survival traversing a distance L at velocity β_{N4} is described by

$$\mathcal{P}_{\mathrm{surv}}\left(L, \beta\right) = N \cdot \exp\left(-\frac{L}{\beta_{\mathrm{N4}}\gamma_{\mathrm{N4}}\tau_{\mathrm{N4}}c}\right), \tag{5.25}$$

with τ_{N4} being the HNL lifetime. Against this exponentially decaying distribution, one must determine $L = L(V; \boldsymbol{r}_0, \boldsymbol{r}_{\mathrm{C}})$. This is calculated by constructing a line $\epsilon(s) : \boldsymbol{r}(s) = \boldsymbol{r}_0 + s \cdot (\boldsymbol{r}_{\mathrm{C}} - \boldsymbol{r}_0)$, and finding its intersections with V.

Once the decay vertex $\boldsymbol{r}_{\mathrm{D}}$ has been determined, it is possible to obtain the time of the HNL decay event. This may be parametrised with respect to the

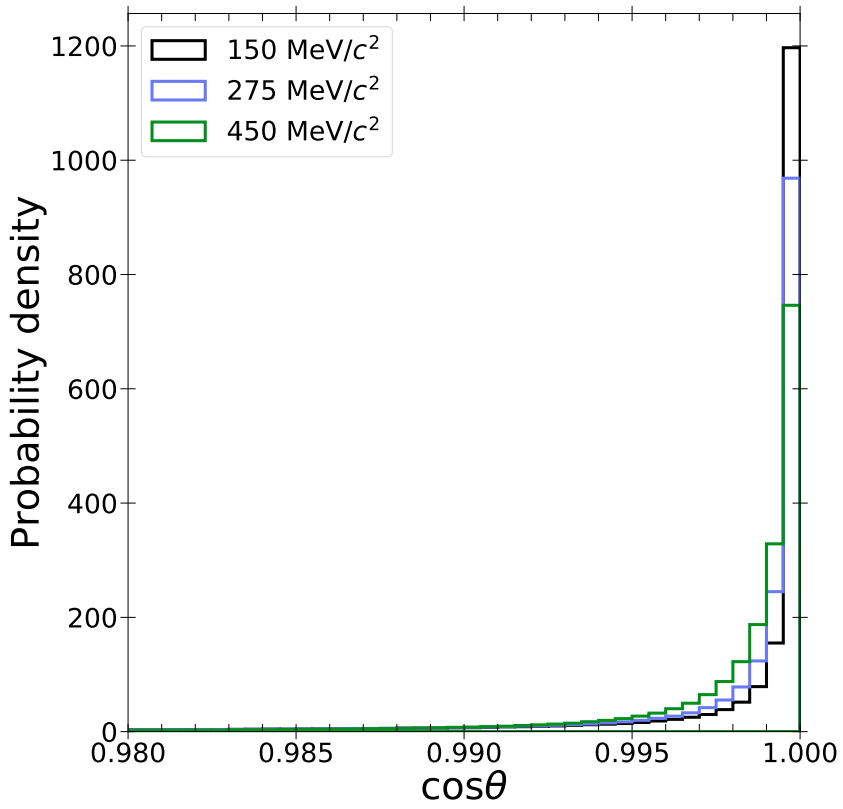


Figure 5.9: Angular dependence $\mathrm{d}\Gamma/\mathrm{d}\cos\theta$ of $\mathrm{N}_4 \to \pi^+ + e^-$ decay in lab frame.

SM neutrino arrival time, t_0, as

$$t_{\mathrm{N4}} = t_0 + \delta t = t_0 + \|\boldsymbol{r}_{\mathrm{D}} - \boldsymbol{r}_0\| \left(\frac{1}{\beta_{\mathrm{N4}} c} - 1 \right). \tag{5.26}$$

Event-by-event, the delay δt may be calculated and stored at the event vertex, which can then be overlaid on to the simulation for a SM neutrino arrival time. These arrival times can be obtained in a data-driven way from recorded neutrino interactions [224, 255, 256]. Figure 5.11 shows data-driven timing distributions from MINOS 5.11(a), MicroBooNE 5.11(b), and MINERνA 5.11(c). Each experimental derivation can recover $\mathcal{O}\,(1-5$ ns) timing resolution, which is smaller than the separation between beam bunches in NuMI; compare to Fig. 5.12(a) which shows the NuMI bunch structure (in "NuMI-only" mode [188]), with 100 ns between groups of beam bunches, and only 0.8 ns width for each beam bunch. Therefore, with sufficiently good timing resolution it is possible to develop a sensitive cut on timing; Figure 5.13 shows the shapes of HNL event times at MINERνA, modulo 18.83 ns.

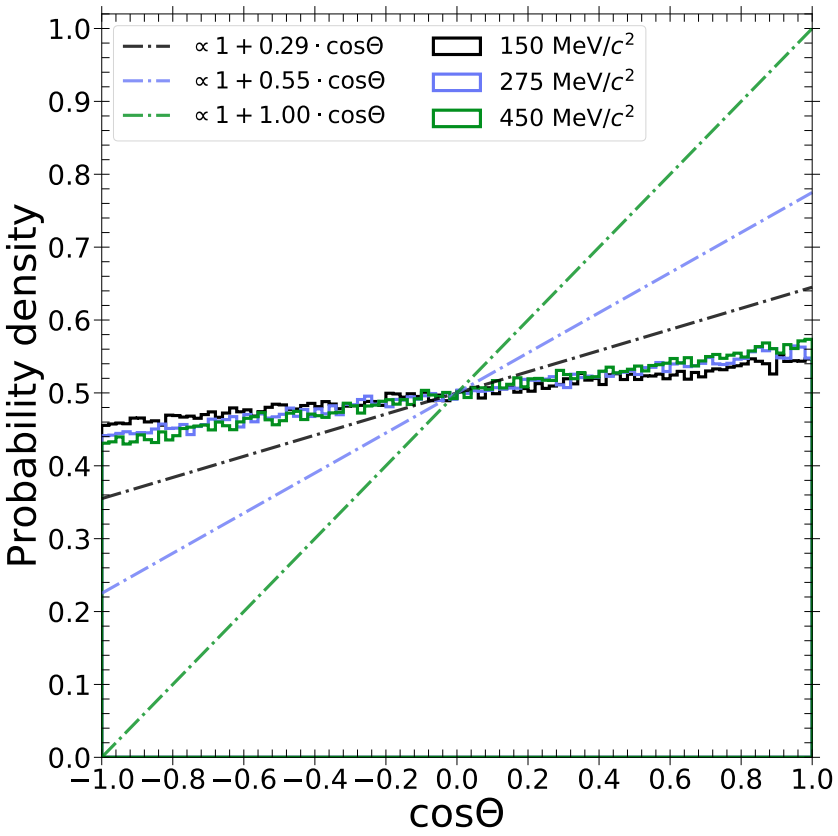
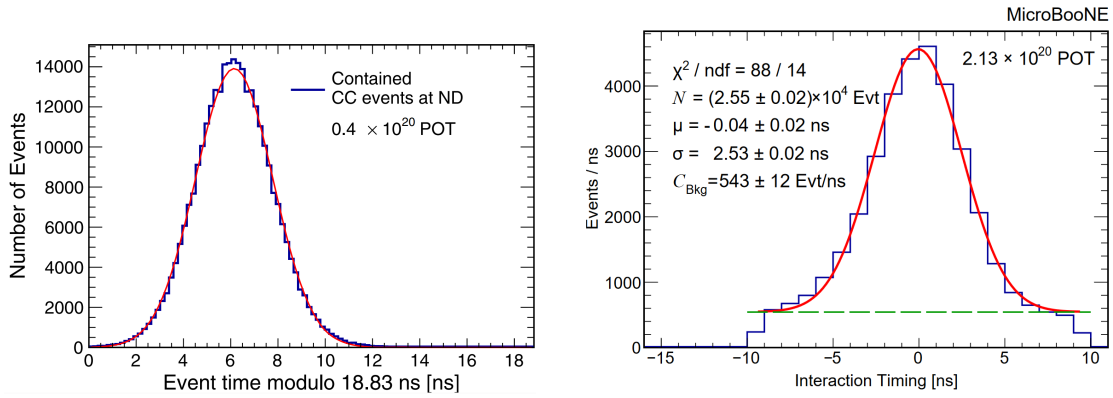


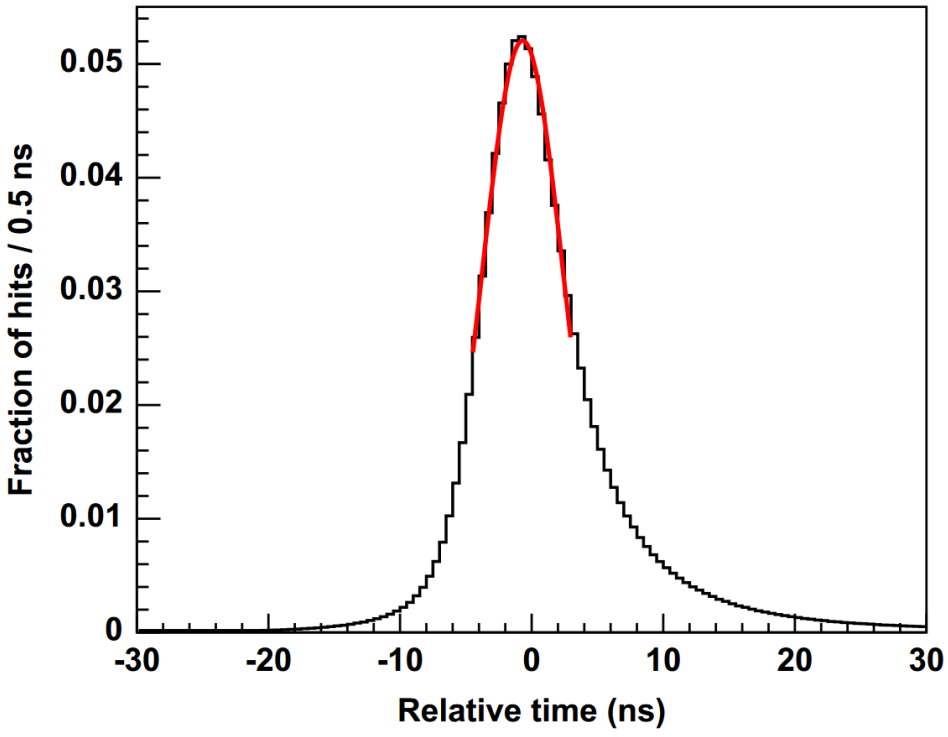
Figure 5.10: Angular dependence $d\Gamma/d\cos\Theta$ of $N_4 \to \pi^+ + e^-$ decay in HNL rest frame. Dash-dotted lines: expected distribution if polarisation vector $\widehat{\varepsilon}_{N4}$ were fixed; solid lines: obtained distribution where $\widehat{\varepsilon}_{N4}$ is determined event-by-event. See text.

The width of the SM curve is 3 ns [224], which corresponds to the resolution on the calibrated time at MINERνA; because of the relative rarity of the HNL signal, the non-zero probability to record SM interactions throughout the bunch spacing makes a bunch-based HNL trigger inefficient for MINERνA.

On the other hand, it is possible to exploit the neutrino spill structure to obtain a clear window in which the rate of SM neutrino interactions is expected to be consistent with zero. Inspecting Fig. 5.12(b), it is evident that after the NuMI beam trigger window shuts off at 11 μs after the start of the proton current, the neutrino-producing proton current is zero. Accordingly, the neutrino beam spill at MINERνA follows the same structure, with the SM interaction rate being reduced to zero by 11.5 μs. In effect, this creates a window where the only SM background is delayed detector activity (such as Michel electrons from the decays of muons at

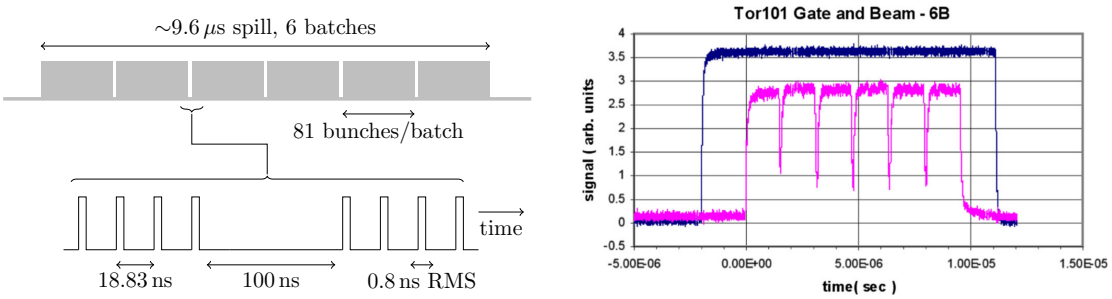


(a) (b)



(c)

Figure 5.11: Data-driven determination of SM neutrino event times for MINOS [255] (a), MicroBooNE [256] (b), and MINERνA [224] (c).



(a) (b)

Figure 5.12: NuMI bunch structure [255] (a) and spill structure [188] (b). There are six beam batches in a spill (pink lines, (b)) lasting about 9.6 μs, with 81 bunches in a batch.

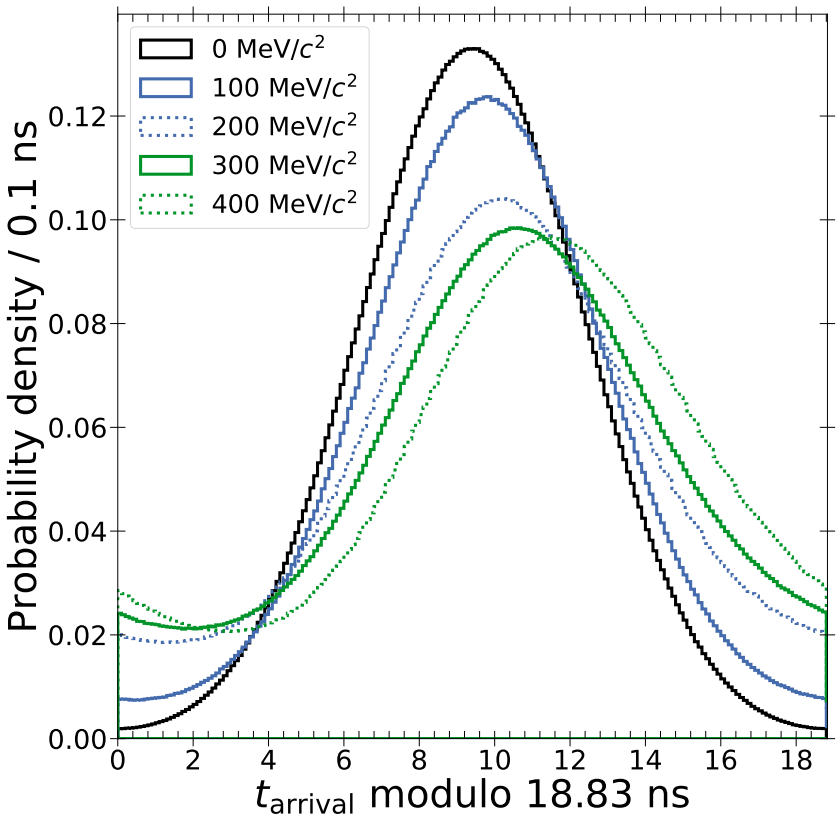


Figure 5.13: Shapes of event timing distributions at MINERνA, modulo the inter-bunch spacing 18.83 ns. SM interactions are the solid curve, and histograms correspond to different neutrino masses.

rest); if a detector reads out activity for long enough, it could majorly enhance the purity of any detected signal consistent with an HNL decay.

However, this "smoking-gun" trigger comes at the cost of a vast reduction of the signal event rate. Figure 5.14 shows the timing of HNL decay vertices at MINERνA, within the overall spill structure. The SM simulation follows the "NuMI-only" [188] spill structure in this example, ending at 9.6 μs. The region $t \geq 11.5$ μs is highlighted, as that corresponds to the shut-off of the POT current (cf. Fig. 5.12(b)) As HNL get heavier, they start to lag behind the SM neutrino more and contribute more to the "no-SM" region: the drop in signal rate is manifest, however. In our case, this trigger would not produce a significant increase in sensitivity.

5.4 POT estimation

Since the salient features of HNL detection probability depend heavily on the HNL mass and mixings to the SM, it is not a well-posed problem to derive a single number of HNL signal events for a given exposure of the detector. Nor is it sufficient to attempt this calculation setting the mixings to 1 and for a range of mass hypotheses, as the precise flavour hierarchy $|U_{e4}|^2 : |U_{\mu 4}|^2 : |U_{\tau 4}|^2$ changes the kinematics of the HNL in the detector (cf. Eq. 5.1). Instead, one may attempt the inverse problem: *given a signal event*, how many POT would it take to be produced?

Essentially, the problem can be broken down into a series of sequential steps, which depend on the event's kinematics, mixings, mass hypothesis, etc. It may be summarised in Fig. 5.15:

For the steps involving Eq. 5.25, the "decay" probability is the probability the HNL does not survive traversing the length $\ell = \|\boldsymbol{r}_\mathrm{X} - \boldsymbol{r}_\mathrm{E}\|$ of the detector, with E, X being the entry and exit points of its trajectory with the detector volume. The "survival" probability, on the other hand, refers to its survival traversing the distance between its production point D and its entry point E, $L = \|\boldsymbol{r}_\mathrm{E} - \boldsymbol{r}_\mathrm{D}\|$.

To transform the conditional probability $\mathcal{P}\left(P \to N_4 \mid \mathrm{p} + \mathrm{target} \to P\right)$ to $\mathcal{P}\left(\mathrm{p} + \mathrm{target} \to N_4\right)$, one requires the multiplicity $n\left(P\right)$ for each parent species.

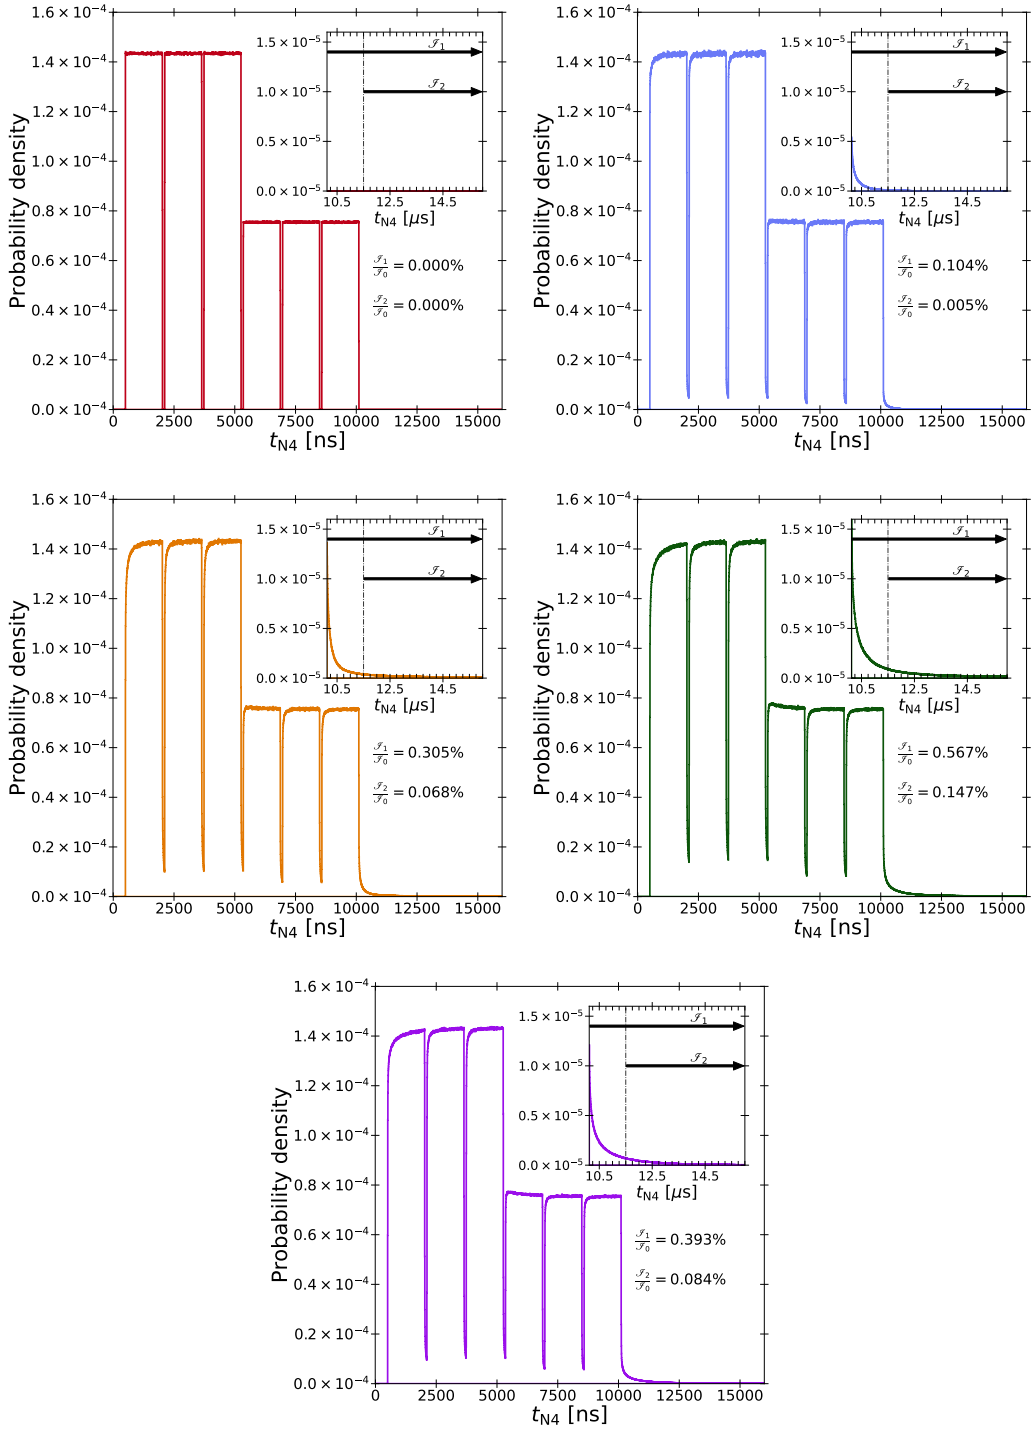


Figure 5.14: HNL arrival times at MINERνA, for different masses (0 to 400 MeV/c^2) from top left to bottom. The region $t \geq 10.1$ μs is highlighted in the inset plots.

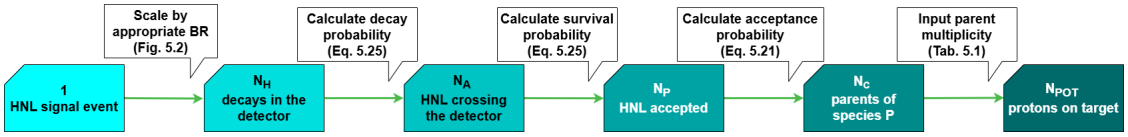


Figure 5.15: A sketch of the POT estimation process, starting from an individual signal event and arriving at the number of POT needed to produce this event. A series of sequential steps breaks up the problem into individually calculable steps.

This is obtained from the detailed `GEANT4` simulation of NuMI, `g4numi`, by executing the pseudocode in Algorithm 2:

Data: N_{POT}, n-tuple with simulated parents P
Result: $n(P)$ for all P in simulation
multiplicities = new array of zeroes;
for *entry in n-tuple* **do**
  w_{entry} = "importance weight";
  multiplicities[P] += $w_{\text{entry}}/N_{\text{POT}}$;
end
return multiplicities;

Algorithm 2: Obtaining parent particle multiplicities from `g4numi` simulation.

Species	$m_{\text{F}}^{(\text{FHC})}$	$m_{\text{Q}}^{(\text{FHC})}$	$m_{\text{F}}^{(\text{RHC})}$	$m_{\text{Q}}^{(\text{RHC})}$
π^+	1.49	1.54	0.991	0.981
π^-	1.02	0.993	1.45	1.47
K^+	0.204	0.149	0.146	9.58×10^{-2}
K^-	7.59×10^{-2}	4.67×10^{-2}	0.108	7.09×10^{-2}
K_L^0	3.20×10^{-2}	2.11×10^{-2}	3.16×10^{-2}	2.21×10^{-2}
μ^+	1.06×10^{-2}	1.32×10^{-2}	2.93×10^{-3}	3.12×10^{-3}
μ^-	3.00×10^{-3}	3.48×10^{-3}	9.73×10^{-3}	1.10×10^{-2}

Table 5.1: Parent particle multiplicities per POT at NuMI after focusing, obtained with the `g4numi` simulation. Models $m_{\text{F,Q}}$ refer to FTFP_BERT (default) and QGSP_BERT.

Supposing the true multiplicities $n(P)$ were known for all parents P, then it would be possible to apply a scalar multiplier $1/n(P)$ to each event in the last step of Fig. 5.15. One should, however, be careful to use the multiplicities corresponding to their input fluxes: if the input fluxes simulate a particular region of phase space such as in `g4numi`, only parents from that phase-space region will enter the input tuples,

and the obtained multiplicity must reflect that. Algorithm 2 ensures the correct multiplicities are obtained, introducing nonetheless the complication that different multiplicities must be used for FHC and RHC runs. This can be taken account of during the HNL sample production; we will return to this discussion in Chapter 6.

Finally, a remark; there are a number of reference physics lists in `GEANT4` ([257–259]; see [236] for documentation), of which the default FTFP_BERT (FRITIOF string model + nuclear de-excitation + Beritini cascade) is used in the default `g4numi` simulation (the official MINERνA flux files). However, different models can produce a spread of different multiplicities (cf. Tab.5.1). Therefore, one could apply a parent-particle-dependent systematic on the normalisation of the HNL, treating FTFP_BERT as the central value and the error being the largest fractional deviation for that particle. In this thesis, a simpler yet conservative uncertainty is applied, based on the overall multiplicities in Tab. 5.1. This will be discussed more in Chapter 9.

5.5 The BeamHNL simulation module

`BeamHNL` leverages the `GENIE` philosophy of separating out the neutrino event generation into monadic operations on a main event record. The essential sequence of steps is presented diagramatically in Fig. 5.16; a more detailed high-level overview is given in Algorithm 3.

The event weight is, according to the rationale of Fig. 5.15, an estimate of the N_{POT} needed to produce the one signal event; it proxies, in other words, for the *inverse probability* of the event occurring. It is straightforward to show that, for an ensemble of HNL decays $\mathcal{E}$, the probability of one individual event i conditional to the ensemble is simply the ratio of inverse weight $w_i^{-1} / \sum_{\mathcal{E}} w^{-1}$; this proof will be omitted for brevity.

The primary user control occurrs via the configuration file `CommonHNL.xml`, which is detailed in Appendix A. It is designed so that the user may easily modify and save their configuration files and store them alongside the simulation output, so that results may be reproduced. In the Appendix, a more detailed review of the

Data: n-tuple of parent particles, geometry ROOT file, and configuration xml.
Result: n-tuple of HNL decays in the geometry.
Module `BRCalculator(`*config*`):`
  foreach *kinematically accessible channel* c_i **do**
    if c_i *: decay channel* **then**
      Calculate decay rate Γ_i
    else
      Calculate kinematic factor $\mathcal{K}_i$
    end
  end
  return ($\mathbb{M}_\Gamma$, $\mathbb{M}_\mathcal{K}$: $\mathbb{M}$ maps $c_i \mapsto \Gamma_i$ or $\mathcal{K}_i$)
for *entry in n-tuple* **do**
  eventRec = **new** EventRecord;
  Module `FluxCreator(`*entry, geom, config,* $\mathbb{M}_\mathcal{K}$`):`
    if *parent can produce HNL* **then**
      Calculate HNL 4-momentum $\boldsymbol{p}_{\mathrm{N4}}$;
      Calculate acceptance probability $\mathcal{P}_{\mathrm{acc}}\left(\text{entry}\right)$ and parent scalar multiplier s_P;
      Set weight $w_{\mathrm{entry}} = \left(\mathcal{P}_{\mathrm{acc}}\left(\text{entry}\right) \cdot s_P\right)^{-1}$;
    else
      return nonsense;
      continue;
    end
    return eventRec;
  Module `Decayer(`*entry, config,* $\mathbb{M}_\Gamma$`):`
    Select channel from $\{c_i\}$;
    Generate kinematics from c_i and populate EventRecord;
    Update weight $w_{\mathrm{entry}} \mapsto w_{\mathrm{entry}} * \left(\Gamma_i/\Gamma_{\mathrm{tot}}\right)^{-1}$;
    return eventRec;
  Module `VertexGenerator(`*entry, geom, config*`):`
    Find entry and exit points E, X with geometry;
    Assign decay 4-vertex and write to EventRecord;
    Update weight
    $w_{\mathrm{entry}} \mapsto w_{\mathrm{entry}} * \left(\mathcal{P}_{\mathrm{surv}}\left(\left\|\boldsymbol{r}_{\mathrm{E}} - \boldsymbol{r}_{\mathrm{D}}\right\|, \beta_{\mathrm{N4}}\right) \cdot \mathcal{P}_{\mathrm{dec}}\left(\left\|\boldsymbol{r}_{\mathrm{X}} - \boldsymbol{r}_{\mathrm{E}}\right\|, \beta_{\mathrm{N4}}\right)\right)^{-1}$;
    return eventRec;
  Add eventRec to EventRecords n-tuple;
end
return EventRecords;

Algorithm 3: The essential function of the `BeamHNL` generator.

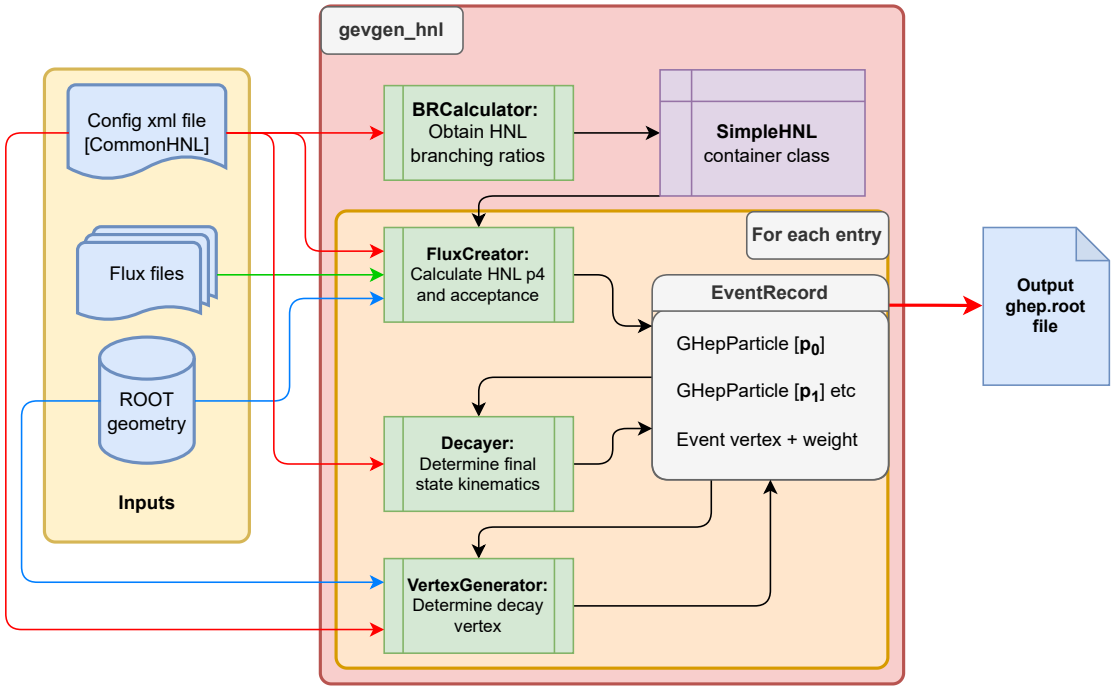


Figure 5.16: Schematic for the operation of the `BeamHNL` generation app. The essential operation takes three inputs, and operates sequentially on a `GENIE` "EventRecord" class which contains information on all the particles of the event, as well as the event vertex and weight.

code is given. For the code details including examples of how to run `BeamHNL`, the reader is encouraged to consult the most recent `GENIE` Physics and User Manual.

6 Monte Carlo samples

Contents

6.1 MINERνA running periods 104
6.2 Simulation pipeline 105
6.3 Sample organisation 108

Data organisation within MINERνA follows a scheme that allows for the tracking of the precise run period over the experiment's long lifetime, and for the unambiguous assignment of simulated Monte Carlo samples to the correct running period. This assignment is handled by the construction of a run numbering scheme for real data, and a separate run numbering scheme for Monte Carlo simulations. Additionally, simulations of particular neutrino interactions, corresponding to specific multiples of data exposure, were produced during the runtime of MINERνA; each "special sample" is categorised by the enhanced simulated signal and split into the various running periods to account for the various detector configurations and beam intensities. In accordance with prior special samples, the HNL signal was produced as a single special sample, separating the various simulated couplings, decay channels, and masses in a straightforward way. By this method, a scan of the entire HNL parameter space may be encoded in a single signal sample.

6.1 MINERνA running periods

The experiment took data in the Medium Energy configuration from September 6, 2013 to February 26, 2019, during which time the detector operated with and without filled helium and water targets, and with the NuMI beam gradually increasing its intensity to about 5×10^{13} protons per spill [260]. To deal with these changes to the beam and detector configuration, the MINERνA dataset is split into *running periods*, or *playlists*, which codify the appropriate environment in which the data was collected. The Monte Carlo simulation takes these running periods into account, and overlays random events from real data, from the proper period,

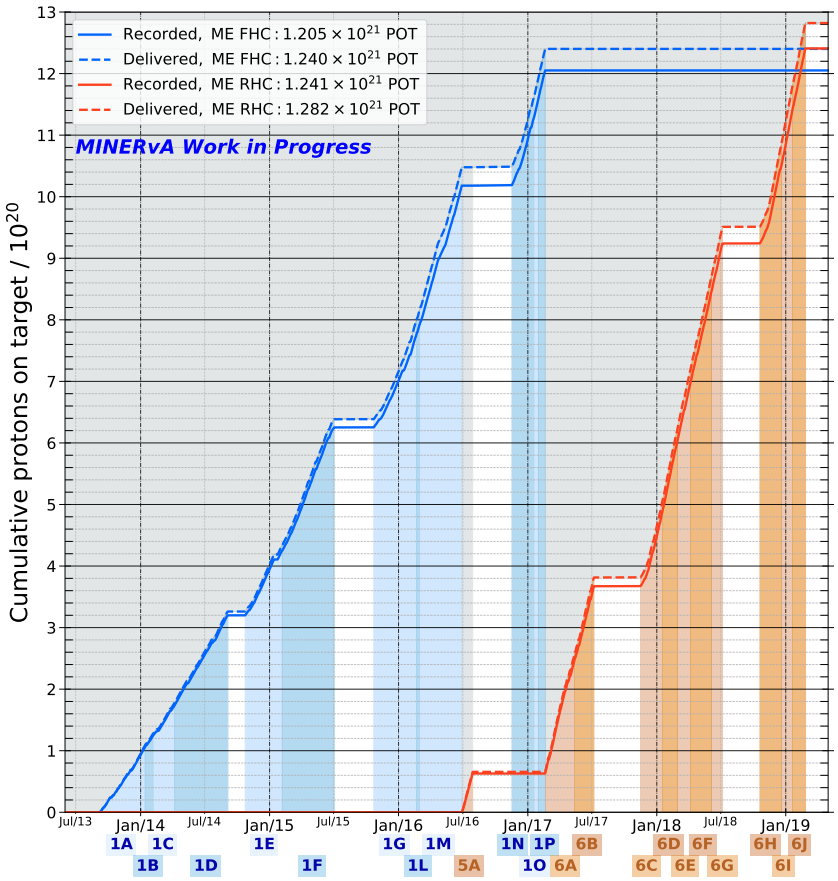


Figure 6.1: MINERνA running periods. Blue indicates FHC (neutrino) mode, and orange RHC (antineutrino) mode.

for event generation. Figure 6.1 shows the POT delivered and collected throughout the MINERνA lifetime, labelling the running periods.

6.2 Simulation pipeline

First, the GENIE prediction files running `BeamHNL` were produced, by using the nominal `dk2nu` flux files and a ROOT geometry file derived from the GDML description of the detector as input. For the full MINERνA dataset, it was decided to produce samples from all the decay channels that produce charged particles in the final state, with a branching ratio above 10^{-5}; these are the five channels $\nu_\alpha \ell_\beta^+ \ell_\beta^-, \nu_\alpha \mathrm{e}^\pm \mu^\mp$, and $\pi^\pm \ell_\beta^\mp$ with $\alpha = \mathrm{e}, \mu, \tau$ and $\beta = \mathrm{e}, \mu$.

For each channel, `BeamHNL` is set to suppress all other HNL decays. This is set for two reasons: firstly, to ensure the maximum number of statistics available

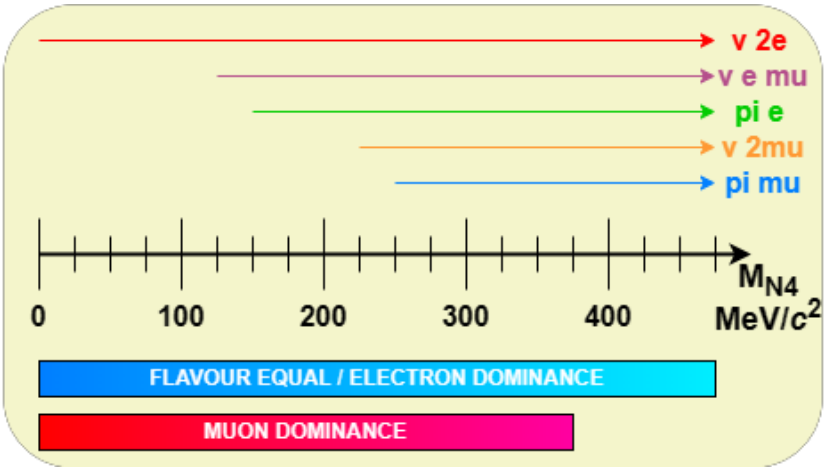


Figure 6.2: The HNL parameter space is split, for production purposes, into multiples of 25 MeV/c^2. Each of the five channels is generated at each kinematically available mass point, with thresholds depending on the final state and the mixing hierarchy assumed.

to study the decays of HNL into the desired mode. Secondly, to disentangle the different mass thresholds for the various channels. Due to the different masses in the various final states, kinematic thresholds for the HNL decays open at different M_{N4}; a strategy was developed to scan the HNL mass in increments of 25 MeV/c^2, from the smallest multiple of 25 MeV/c^2 above the appropriate threshold to the largest multiple below the HNL production threshold of $m_{\mathrm{K}^\pm} - m_{\mathrm{e}}$ (or $m_{\mathrm{K}^\pm} - m_{\mu}$, in the case where only $|U_{\mu 4}|^2 \neq 0$). A sketch of the method by which the parameter space is spanned is shown in Fig. 6.2.

Each channel is simulated in three mixing paradigms: two *flavour-dominant modes* setting $|U_{\alpha=\mathrm{e},\mu 4}|^2 = |U_{\tau 4}|^2 = 0$, and one *flavour-equal mode* with $|U_{\mathrm{e}4}|^2 = |U_{\mu 4}|^2, |U_{\tau 4}|^2 = 0$. The reason for this choice is because the available production channels can involve electrons or muons, but not taus; the helicity unsuppression mechanism from Eq. 5.1 (also see Fig. 5.1) implies that, depending on the HNL mass, SM-suppressed channels such as $K\pm \to \mathrm{N}_4\mathrm{e}^\pm$ can become significant. Therefore, a simple reweighting that linearly combines predictions from the flavour-dominant modes cannot correctly reproduce an arbitrary ratio of $|U_{\mathrm{e}4}|^2 : |U_{\mu 4}|^2$. However, one can notice that this ratio may be split into a flavour-equal and a flavour-dominant part: one can then straightforwardly scale the predictions from the appropriate modes to achieve whichever combination of $|U_{\mathrm{e}4}|^2 : |U_{\mu 4}|^2$ is desired.

The contribution of the third element $|U_{\tau 4}|^2$ only enters at the level of the BR of a channel, which may be applied as a rescaling of event weights due to the monadic operation of `BeamHNL`; see the action of the `BRCalculator` and `Decayer` modules in Algorithm 3. Consequently, any desired combination of all three mixing elements can be achieved, which enables one to deduce the expected number of signal events, for each of the above channels, for any choice of mixing elements.

The default value used for $|U_{\alpha 4}|^2$ was 10^{-8}. Due to the way the production BRs scale with one power of $|U_{\alpha 4}|^2$, and the survival and decay probabilities scale (Eq. 5.25) as the following functions of the inverse proper lifetime τ_{N4},

$$\begin{aligned} \mathcal{P}_{\mathrm{surv}}(\propto) \exp\left(-\tau_{\mathrm{N4}}^{-1}\right) &\propto \exp\left(-|U_{\alpha 4}|^2\right) \simeq 1 - |U_{\alpha 4}|^2 + \frac{|U_{\alpha 4}|^4}{2} + \ldots, \\ \mathcal{P}_{\mathrm{dec}}(\propto) 1 - \exp\left(-\tau_{\mathrm{N4}}^{-1}\right) &\propto 1 - \exp\left(-|U_{\alpha 4}|^2\right) \simeq |U_{\alpha 4}|^2 - \frac{|U_{\alpha 4}|^4}{2} + \ldots, \end{aligned} \tag{6.1}$$

the overall event rate scales as $|U_{\alpha 4}|^4 \left(1 - \frac{3|U_{\alpha 4}|^2}{2} + \ldots\right)$. Therefore, it is possible to relate the total number of POT expected to have produced a sample of HNL decays at a reference value $(\alpha_0) \leftrightarrow |U_{\alpha_0 4}|^2$ to a desired new value (α), as

$$N^{(\alpha)} = N^{(\alpha_0)} \times \left(\frac{|U_{\alpha 4}|^2}{|U_{\alpha_0 4}|^2}\right)^2 \times \frac{1 - 3\frac{|U_{\alpha 4}|^2}{2} + \ldots}{1 - 3\frac{|U_{\alpha_0 4}|^2}{2} + \ldots} \simeq N^{(\alpha_0)} \times \left(\frac{|U_{\alpha 4}|^2}{|U_{\alpha_0 4}|^2}\right)^2. \tag{6.2}$$

The prediction files of 1000 signal events each are then passed into the MINERνA simulation, which applies the de-calibration procedure described in the previous Chapter to produce a record of ADC and TDC counts per electronics channel, and the appropriate electronics-space to physical-space map. These counts are passed into the MINOS near detector's simulation, which simulates outgoing particles from MINERνA in the MINOS ND. Finally, a reconstruction step applies the general reconstruction as detailed previously to produce the primary analysis input: *reco files* containing information about reconstructed particles, trajectories, time slices, and clusters in MINERνA.

Each of these four stages (GENIE prediction, de-calibration, MINOS, and reconstruction) is declared to Fermilab's tape storage as an individual data tier, to be handled by Fermilab's SAM system [261].

6.3 Sample organisation

Because the de-calibration step contains constants which depend on the running period corresponding to the simulation, it was decided early on to split the HNL signal sample into sub-samples, one for each of the MINERνA running periods. The split into running periods imposes constraints on the MC run numbering for the HNL signal files. Each run in MINERνA is labelled as $001CRRRR.0XYZ$, with $C = 1, 2$ the current specifier (FHC or RHC), $RRRR$ run numbers depending on the detector configuration, and subrun numbers XYZ normally are just file indices. For the HNL, there exist 12(11) running periods in FHC (RHC) to simulate: each must contain the files from between $10 + 0 + 6$ mass points (in the case of $\mathrm{N}_4 \to \pi^{\pm} + \mu^{\mp}$) and $19 + 19 + 15$ points (for $\mathrm{N}_4 \to \nu + e^{\pm} + e^{\mp}$). To accommodate this requirement, the subrun numbers carry information about the signal sample characteristics: the last two digits YZ can be mapped to a unique combination of $\left(M_{\mathrm{N4}}, \{|U_{\alpha 4}|^2\}\right)$:

$$\begin{cases} 00 < YZ < 20: & \text{equal couplings,} \\ 21 < YZ < 40: & \text{electron dominance,} \\ 41 < YZ < 60: & \text{muon dominance,} \\ M_{\mathrm{N4}} = 25 * (YZ \bmod 20) & \mathrm{MeV}/c^2. \end{cases}$$

Out of the full MINERνA flux file set, a total of 10000 signal events is distributed for each parameter-space point, per running period. This was a compromise between adequate statistics for each parameter-space point, available storage space at Fermilab, and the extent of the flux files. The latter is a particularly constraining consideration close to the kinematic threshold from kaon production, with the number of flux entries parsed per generated event sharply climbing as the collimation effect becomes dominant; see Fig. 6.3.

From now on, each of the parameter-space points and channels will be treated as an independent sub-sample. The signal process for this thesis is defined to be the two-body decay $\mathrm{N}_4 \to \pi^{\pm} + \mu^{\mp}$. The Monte Carlo samples of the background processes were processed using GENIE version `2.12.6` (for the "General" (MC-General), un-enhanced, sample - a review can be found in [262]). There is also modelling of meson exchange currents (or "two-particle two-hole", 2p2h, scattering where the neutrino

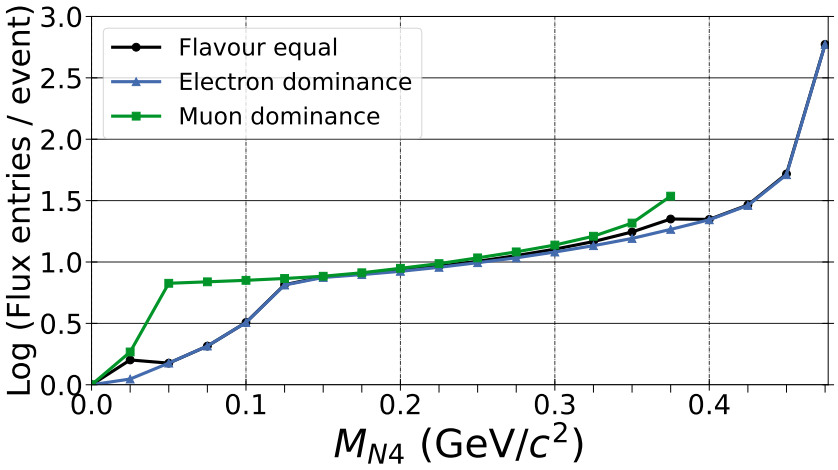


Figure 6.3: Parents parsed per generated HNL events. The sharp increase of entries close to the thresholds of $m_K - m_\ell \simeq 388,493$ MeV/c^2 is due to the rapid onset of the collimation effect. The electron- (muon-) dominant mode is the least (most) constraining, as the available rest-frame energy for the HNL is maximised (minimised).

scatters off of a pair of correlated nucleons). Coherent pion production using the Rein-Sehgal model [263] with lepton mass effects introduced [264] are included in this simulation. In addition, there are non-resonant "diffractive" pion-producing interactions from neutrino interactions on hydrogen nuclei (so named because the process is analogous to classical diffraction from coherent electromagnetic waves [265]), which are included in the simulation. Furthermore, the interactions from neutrinos interacting with atomic electrons are modelled, both for neutrino-electron elastic scattering and for the inverse muon decay (IMD, $\nu_\mu + e^- \to \mu^- + \nu_e$) interaction types (see for example [213]).

Final-state interactions (or FSI) are a collection of physical processes which complicates the characterisation of event topologies in a detector. The products of the neutrino interaction are in a nuclear environment (a many-body potential that can interact strongly, weakly, and electromagnetically). These particles must exit the nuclear environment before starting to deposit energy in the detector and be detected; this process can alter their momenta, or lead to the production of new particles, or to the absorption of particles from the intrinsic interaction. Generally, there are two classes of treatments to deal with this problem in event generators. `GENIE` employs a "semi-classical", intranuclear cascade (INC) approach.

In the `2.12.6` version used by MINERνA, FSI is handled by a fully reweightable data-driven code tuned primarily to bubble chamber data from hydrogen and deuterium targets, using the `INTRANUKE` [266] intranuclear cascade simulation; essentially, final-state hadrons are propagated inside the nucleus in steps, with each step having a certain probability to rescatter or interact. This "hA" empirical INC was complemented in `GENIE` v3 by the "hN" full INC [267], similar to the `NEUT` [268] and `NuWro` [215] generators, and by the "INCL++" cascade [269]. The other class of treatments is a dynamical, transport-theoretical approach based on the full many-body phyiscs; for details the reader is referred to [202, 270, 271].

Various signal-enhanced samples for processes of interest were also produced, and are used to characterise the principal backgrounds to this thesis. These are as follows:

- A CC Coherent (CC-COH) sample [272, 273], implementing the Rein-Sehgal partially-conserved axial current formalism (PCAC) cross-section [263]. Coherent interactions refer to the interaction of a neutrino with an entire atomic nucleus coherently, with low four-momentum transfer.

- A CC Diffractive (CC-DIF) sample [265], implementing the Rein PCAC cross-section [274]. The name refers to the analogy with classical electromagnetism, where coherent superposition of waves produces classical diffraction.

- An IMD + neutrino-electron elastic scattering (IMD) sample [212, 214].

The MC-General sample does include a CC-COH contribution, but lacks the statistics for a significant coherent event count; in this thesis, CC-COH, CC-DIF and IMD (which is not included in MC-General at all) are treated as independent samples.

Sample	FHC POT	RHC POT
Data	9.33	10.5
General	37.2	44.8
CC Coherent	783	772
CC Diffractive	312	208
IMD	2.69×10^3	N/A

Table 6.1: POT per sample processed, in units of 10^{20} POT.

7 Selection of $N_4 \to \pi^{\pm}\mu^{\mp}$ events at MINERνA

Contents

7.1 HNL event reconstruction 112
7.1.1 Vertex finding 113
7.1.2 Track merging 116
7.1.3 Energy deposition reconstruction 119
7.2 Event signature 123
7.3 One-track classification 126
7.3.1 Pre-selection cuts 126
7.3.2 Cost-benefit analysis 131
7.3.3 Selection cuts 134
7.3.4 Efficiency calculation 149

The main task of this analysis is to identify HNL candidates, and reject background processes that mimic their signatures. To accomplish this task, it is neccesary to specify a *signal definition*, and a strategy to isolate events matching this signal. The particles that MINERνA is best at reconstructing, among the channels available below the kaon mass, are muons and pions. Both of these are minimum ionising at GeV energies, and the MINOS near detector provides a clear separation between muons and pions.

Given these observations, the signal channel is defined as:

$$\mathrm{N}_4 \to \pi^{\pm} + \mu^{\mp},$$

where the HNL decay must happen inside a fiducial volume (FV) located in the tracker. One makes no distinction here between the cases $\pi^+\mu^-$ and $\pi^-\mu^+$, for the following reason. The sign of the particle does, in principle, cause a difference in the stopping power of materials for that particle, due to Coulomb interactions of the particle with atomic electrons. This effect is known as a "shell correction", and occurs when the velocities of atomic electrons are not low with respect to the propagating particle. For muons, the shell correction is small (sub-percent in iron) at MeV energies [275]. For pions, a calculation by Fermi for the Mott

scattering off atomic electrons yielded the shell correction as a modification of the effective pion mass in the Bethe-Bloch formula, of relative magnitude 9×10^{-4} for momenta of order few MeV [276]. Therefore, it is safe to assume that, for MIP pions and muons, there is no meaningful difference in the energy deposition of positive or negative particles in the detector.

The NuMI beam's high energy implies that the HNL have high energy as well. Because muons and pions have similar mass, they inherit roughly equal energy in the HNL's rest frame, and the opening angle between their trajectories will be dominated by the Lorentz boost from the HNL. This means that, generally speaking, the collimated muon and pion will traverse the first part of the detector as a pair. Each is minimum ionising, which offers a powerful discriminatory variable in the *energy deposition per unit length*, $\langle \mathrm{d}E/\mathrm{d}x \rangle$.

Therefore, further to the signal definition, the HNL phase space may be split into a *one-track region*, where the HNL decay products form a *double-MIP pair* for a suitably long track length, and a *two-track region* where the HNL products do not form a double-MIP pair. Similarly, one can split the HNL decays into *fully contained* (FC) and *partially contained* (PC) samples, where the FC events have no exiting particles and the PC have at least one.

7.1 HNL event reconstruction

The majority of HNL events at MINERνA are produced by high-energy decays, which combined with the dimensions of individual tracker strips result in double-MIP pairs. Through the similar ionisation patterns of muons and pions in the detector, it is possible to distinguish a double-MIP track from other topologies. Because the MINERνA reconstruction algorithm was not developed to accommodate double-MIP tracks, additional reconstruction steps are applied to each event to correctly estimate the energy deposition patterns of the particles.

7.1.1 Vertex finding

Due to the high collimation of the muon-pion pair, the vertex finding algorithm sometimes assigns the interaction vertex downstream of the actual simulated vertex. This is due to the projecting backwards of a track produced by the interaction products of the pion with the detector, causing its deflection. As a consequence, the upstream portion of the track is lost, and the associated energy along with it. To ensure that the tracked energy matches the expected deposition along the trajectory of a HNL's decay products, Algorithm 4 is applied. The performance of the vertex reconstruction can be seen, for the indicative cases of the lightest and heaviest simulated HNL at masses 250 and 475 MeV/c^2 and SM neutrino interactions simulated in MINERνA, in Figures 7.1, 7.2, and 7.3. The top two panels of each Figure show the improvement in the residual z-position of the reconstruced vertex, $\Delta z = v_z^{(\text{reco})} - v_z^{(\text{MC})}$; the tail at the high end where the algorithm positions the reconstructed vertex upstream of the true vertex is highly suppressed, and the true MC vertex is recovered. The lower end where the residual is negative corresponds primarily to backwards-emitted muons or pions arising from low-energy HNL decays, which will not enter the selected signal sample. In the bottom two panels, the performance of the cut is parametrised by taking the ratio of the residual Δz over the MC vertex z-distance from the tracker centre, $l_z = \left|v_z^{(\text{MC})} - v_0\right|$.

The distributions are fit to both a Gaussian (dashed) and Cauchy (solid) distribution, and it is readily seen that the tail of $\Delta z/l_z \gtrsim 0.1$ causes the Gaussian fit to perform worse than the Cauchy. For reference, the values of $\chi^2/N_{\text{d.o.f}}$ for the HNL at masses $250, 475$ MeV/c^2 are for the Cauchy (Gaussian) distributions 0.13(0.62) and 0.13(0.52) before vertex finding, and 0.13(0.72) and 0.12(0.69) after; for the SM neutrino interaction sample this is 0.33(2.36) and 0.35(2.41). Evidently, the Cauchy distribution describes the SM neutrino interactions well, and produces a biased but decent prediction for the HNL decay events, particularly after vertex finding corrects the tails. Additionally, the proportion of events where $\Delta z \leq 0.1 l_z$ (the "core" of the distribution) is computed, before and after vertex finding. This metric provides a fit-independent prescription to estimate how many events would deposit

Data: Set of clusters at known positions $\{\boldsymbol{r}_j\}$, reconstructed vertex at $\boldsymbol{r}_{\mathrm{V}}$, muon track slope $(\mathrm{d}x/\mathrm{d}z, \mathrm{d}y/\mathrm{d}z, 1)$
Result: Adjusted vertex position $\boldsymbol{r}'_V$
Let i be the index of the plane the algorithm is currently on, and a be the index of the plane immediately upstream;
Let $\{c_i\}$ be the ordered set of all clusters (not `LowActivity` or `XTalkCandidate`) upstream of V;
if *fewer than 3 clusters in* $\{c_i\}$ **then**
  // don't update;
  return $\boldsymbol{r}_V$;
end
// First get the sorted list of all clusters compatible with track direction;
Let $\{g_i\} = \{c_i : c_i$ compatible with track$\}$;
while *true* **do**
  no_more_clusters = ($\{c_i\}$.size == 0 and $i - a \leq 2$);
  Let $\boldsymbol{q}_a$ be the projection on to a of the track direction;
  Let $\{c_a\}$ be all the clusters on module a;
  for c *in* $\{c_a\}$ **do**
    expected_position = $\boldsymbol{q}_a$;
    threshold = 0.5 * strip_width;
    lower_position = c.lowerTransversePosition;
    upper_position = c.upperTransversePosition;
    if *lower_position - threshold* $\leq$ *expected_position* $\leq$ *upper_position + threshold* **then**
      $\{g_i\}$.Append(c_a);
    end
  end
  Remove $\{c_a\}$ from $\{c_i\}$;
  //Update plane number;
  $i = a$;
  if *no_more_clusters* **then**
    break;
  end
end
// Iterate over the compatible clusters and get the most upstream cluster position;
$\boldsymbol{r}'_V = \boldsymbol{r}_a : a$ minimum;
return $\boldsymbol{r}'_V$;

Algorithm 4: Vertex finding algorithm.

energy that might wrongly fail to be associated with the event vertex. As expected, there is marginal increase $\left(\Delta_{\text{core}} = \mathcal{I}_{\text{core}}^{\text{after}} - \mathcal{I}_{\text{core}}^{\text{before}}\right) \simeq 0.05\%$ in the core population of the SM neutrino interactions, which tend to have any backwards-moving products emitted at a high angle incompatible with the forward track; on the contrary, there is significant increase $\Delta_{\text{core}} = 1.68\%, 2.91\%$ for the HNL at $250, 475$ MeV$/c^2$.

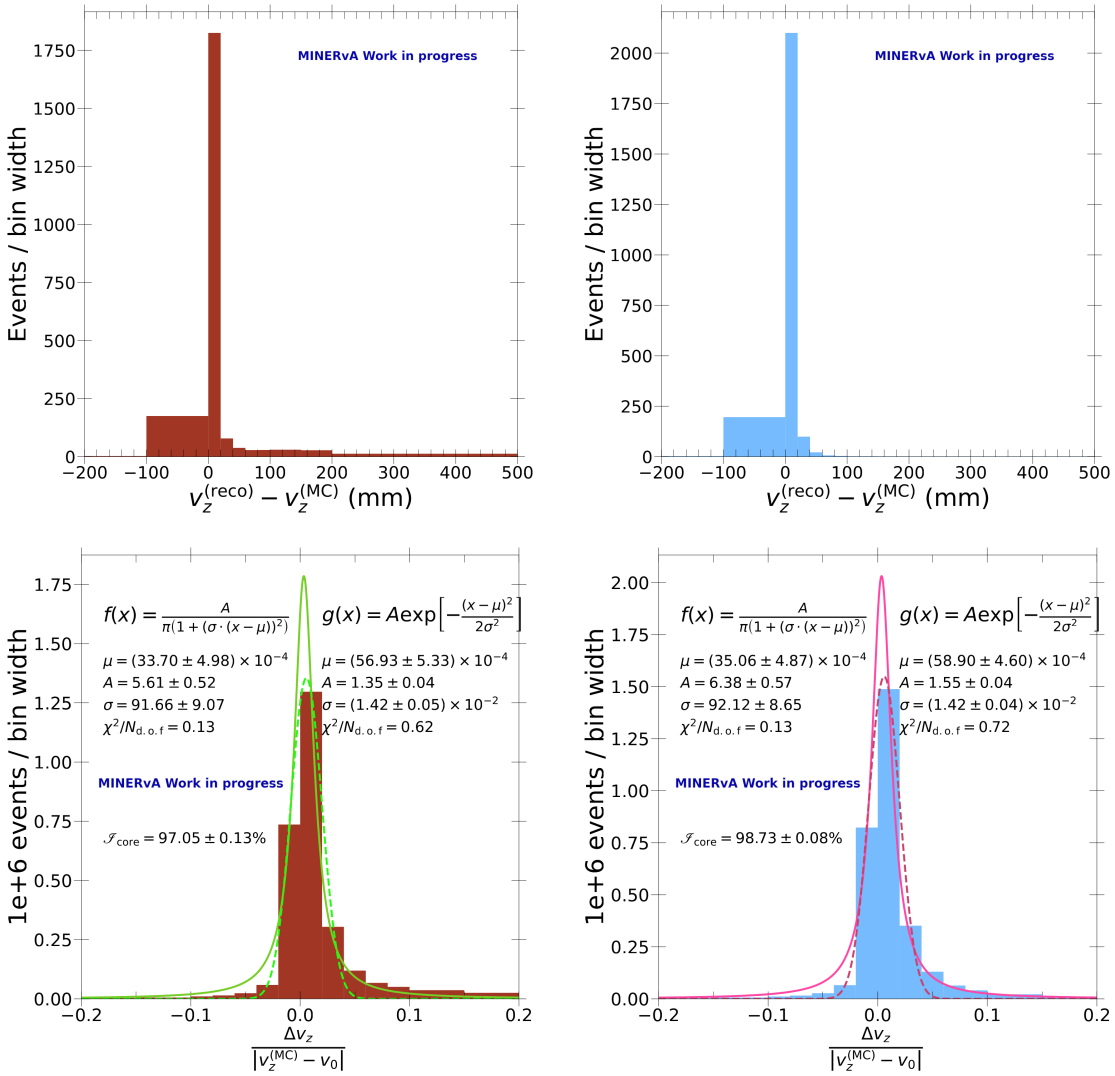


Figure 7.1: Vertex finding for HNL at $M_{\text{N4}} = 250$ MeV$/c^2$. Top panels: difference $\Delta z = z_V^{(\text{reco})} - z_V^{(\text{MC})}$ before (left) and after (right) vertex finding steps. Bottom panels: Fractional difference $\Delta z / l_z$ before (left) and after (right) vertex finding. Solid line: Cauchy fit, dashed line: Gaussian fit.

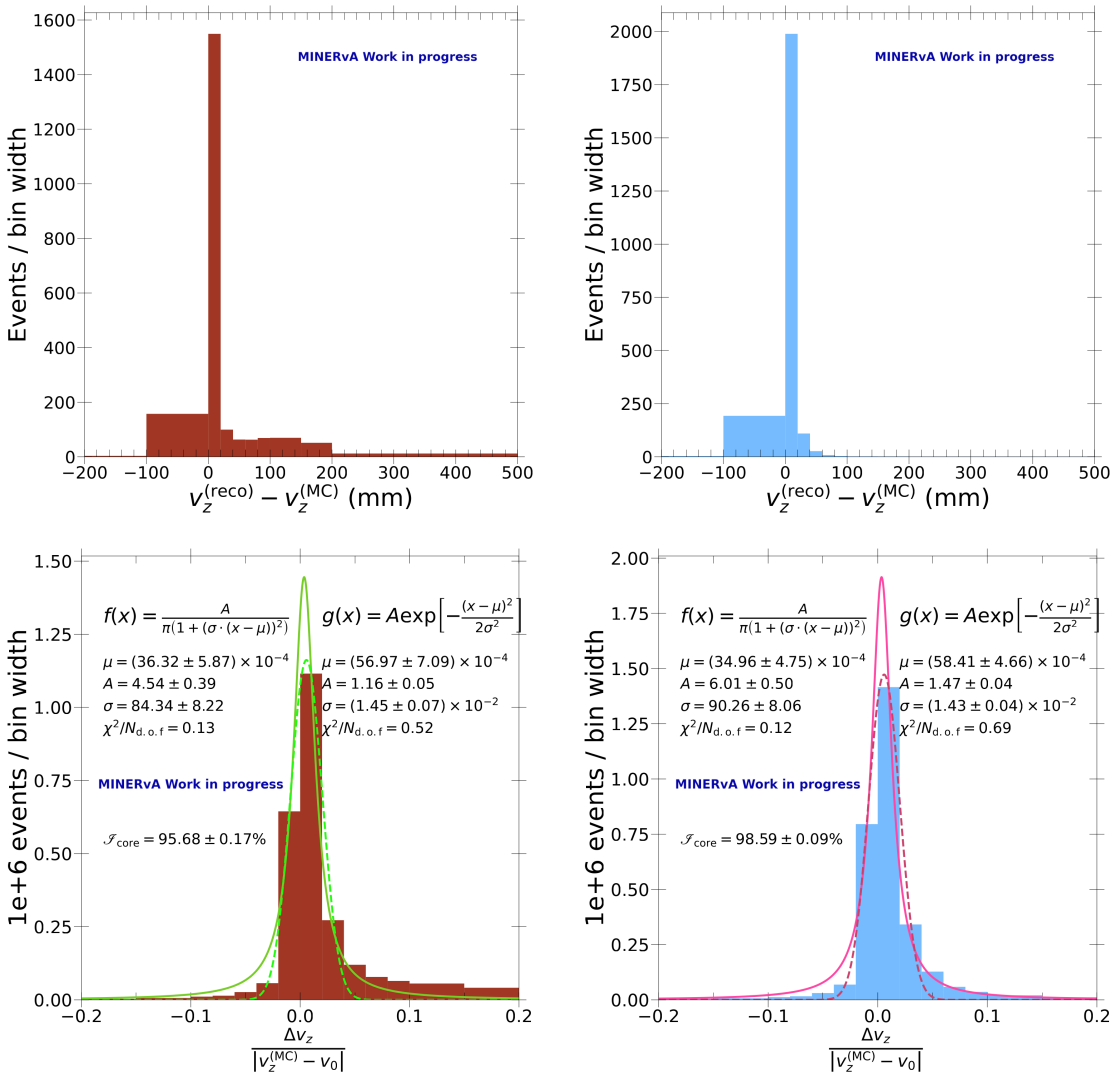


Figure 7.2: Vertex finding for HNL at $M_{\mathrm{N4}} = 475\ \mathrm{MeV}/c^2$. Top panels: difference $\Delta z = z_V^{(\mathrm{reco})} - z_V^{(\mathrm{MC})}$ before (left) and after (right) vertex finding steps. Bottom panels: Fractional difference $\Delta z/l_z$ before (left) and after (right) vertex finding. Solid line: Cauchy fit, dashed line: Gaussian fit.

7.1.2 Track merging

When the vertex gets re-positioned, it becomes necessary to assign the clusters used to find it to the assumed muon+pion track. Even if the vertex does not get re-positioned, however, the muon+pion assumption makes it necessary to perform a step of *track merging* (or more accurately, cluster merging along a known muon track), such that a direct map from muon track to cluster objects (and thus their energy depositions) gets constructed. Necessity for this step arises from the design

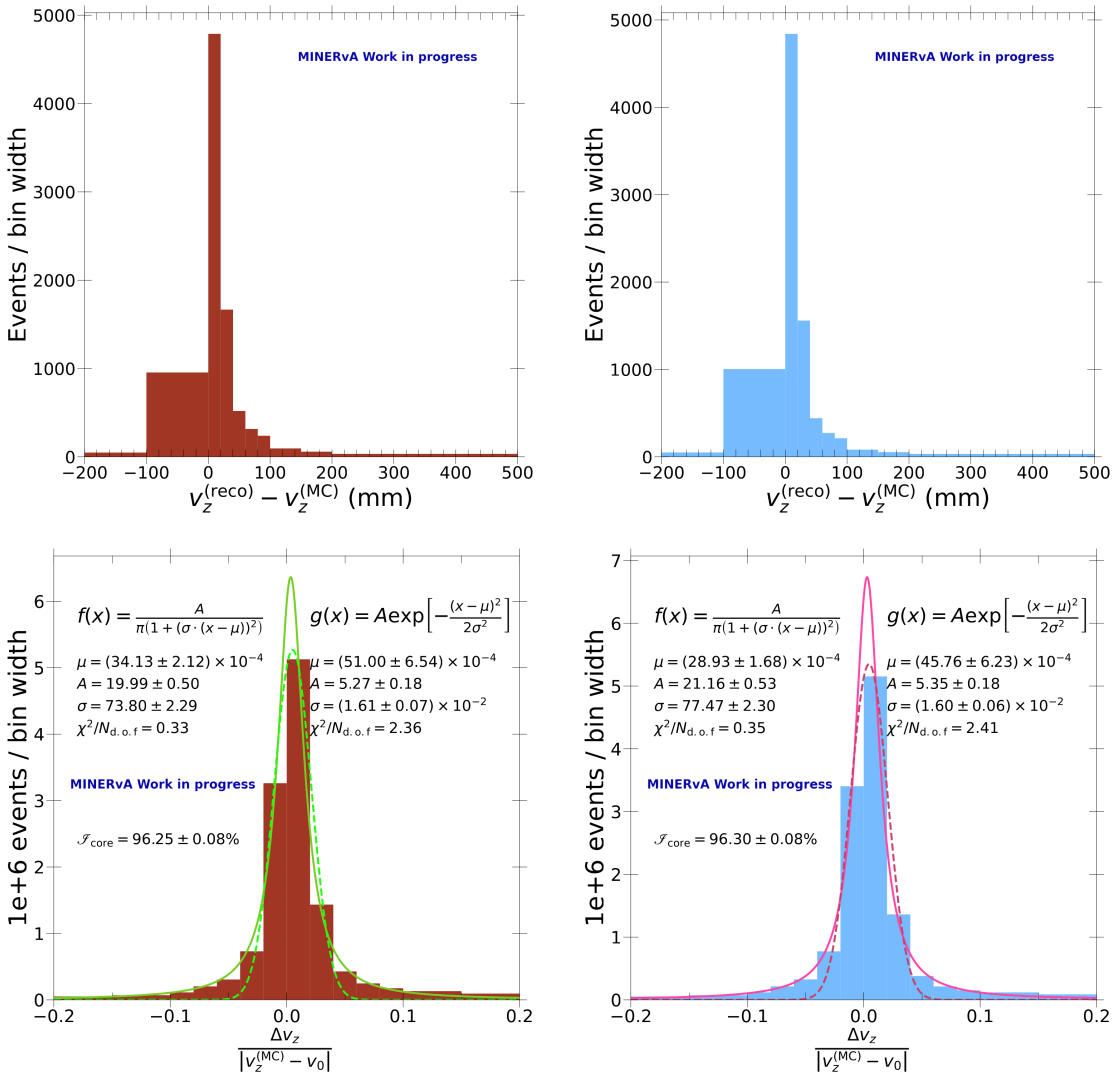


Figure 7.3: Vertex finding for SM neutrino interactions. Top panels: difference $\Delta z = z_V^{(\mathrm{reco})} - z_V^{(\mathrm{MC})}$ before (left) and after (right) vertex finding steps. Bottom panels: Fractional difference $\Delta z / l_z$ before (left) and after (right) vertex finding. Solid line: Cauchy fit, dashed line: Gaussian fit.

of the MINERνA reconstruction algorithm to only make clusters compatible with one minimum ionising particle along the length of the muon track, assuming the conditions for formation of Heavy-Ionising or Superclusters are not met. This may be seen from the event display of Figure 7.4, which shows close-up sections of an event with $\mathrm{N}_4(E = 6.151\ \mathrm{GeV}) \to \pi^+(E = 3.769\ \mathrm{GeV}) + \mu^-(E = 2.482\ \mathrm{GeV})$, which decayed in the tracker at z = 6.080 m. The left-hand panel is a close-up view of the clusters in MINERνA. From this close-up, it is obvious that no re-finding of the

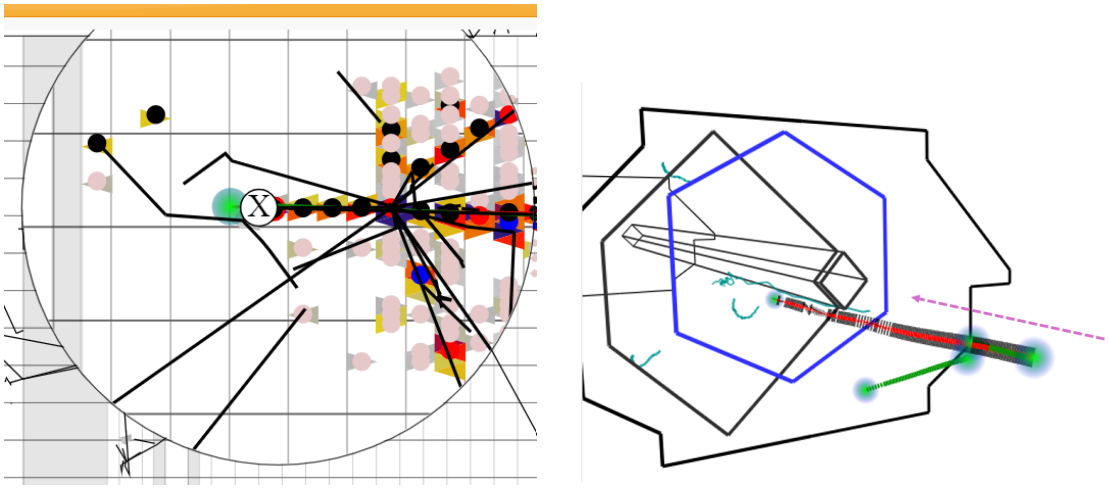
vertex was necessary, and indeed the MC vertex (white circle) is just downstream of the reconstructed vertex (green circle). However, it is not immediately apparent that for every module, the energy of the clusters visible to the eye (coloured solid circles) does not equal the energy of all the hits that made them: as an example, module 30 (the second module, just after the MC vertex) has a cluster on it with energy 2.16 MeV, but the two hits that make up the cluster each individually have energies 4.69 and 1.04 MeV, respectively. The reason for this apparent imbalance is the presence of a second cluster with identical coordinates in physical space, but which is "separated out" from the track and assigned to a track stub, visible on the right-hand panel. This grouping is appropriate for the mainline MINERνA analyses, which investigate a very different parameter space where the pion energy is typically far lower than the muon energy and where there are large opening angles between the final state products. For this thesis's analysis, however, the presence of the stub track only complicates the association of energy deposits along the path of both particles; therefore, a collection algorithm (Algorithm 5) is run after vertex finding. Finally, a pion detection algorithm is run to detect possible pion interactions along

Data: Set of clusters at known positions $\{\boldsymbol{r}_i\}$, muon (+pion) track nodes with positions $\{\boldsymbol{r}_a\}$
Result: Updated track
for *a in all planes* **do**
  Let $\{c_j\}$ be all clusters with positions $\boldsymbol{r}_j = \boldsymbol{r}_a$;
  Assign each cluster in $\{c_j\}$ to the track at point a;
end

Algorithm 5: Cluster collection algorithm.

the track routes, signalling a departure from the two-particle one-track picture. This could be because of the pion interacting strongly with a nucleus inside the tracker and causing emission of many products, or deflection of the pion, or stopping of the pion with an associated Bragg peak at the last few clusters. The algorithm, in essence, looks for possible signatures in the form of abnormally high localised energy depositions (signalling a possible pion interaction), or long departing track

signatures starting at the main track (signalling pion deflection). In brief, this function is summarised in Algorithm 6.



(a) (b)

Figure 7.4: (a): View of the clusters (black circle marks) on the path of a muon-pion pair exiting the HNL reconstructed vertex (MC in white, reconstructed in green). The region inside the circle is magnified in the event display, for ease of viewing; the view is in the XZ view where each module's "X" plane is shown (see Chapter 4, also the left-most panels of Fig. 4.17). For each cluster seen here there exists a second cluster corresponding to a second stub track, visible overlain on the red muon track in (b). In the right panel (b), the blue hexagon is the downstream end of MINERνA, with the beam direction shown as dashed line. The MINOS detector outline and central coil are also shown. Green and red tracks are MINERνA tracks, and blue tracks are MINOS tracks.

7.1.3 Energy deposition reconstruction

To obtain the metric by which one can distinguish the presence of a double-MIP pair, one must construct a robust method to estimate the energy deposition along the track. First, for each of the clusters collected onto the track, the appropriate calibration constants are applied onto the raw energy (defined as the sum of energies from all the hits that contributed to the cluster, weighted by position). This is used to construct a vector of energy depositions $V = \{\mathrm{d}E/\mathrm{d}x\}$ at each plane, which are interpreted as being representative of the plane the cluster belongs to. It is assumed that the entire plane is composed of active scintillator; by applying a "passive material correction" to the energy of each cluster, one gets the corrected $\mathrm{d}E$ in each plane. At each plane, the track slope $(\mathrm{d}x/\mathrm{d}z, \mathrm{d}y/\mathrm{d}z, 1)$ is obtained, and

Data: Muon (+pion) track and attached clusters $\{c_i\}$, and non-track clusters $\{v_i\}$
Result: Plane a with earliest pion interaction candidate
Component `EDepSearch()`:
  foreach *plane a* **do**
    Construct a window of four planes in front of a;
    if *Exists plane with* $E > 32$ MeV **then**
      break;
    else
      if *Average energy across four planes is* > 8 MeV **then**
        break;
      end
    end
  end
  $a_{\text{dep}} = a$;
  return a_{dep} *plane*
Component `GeomSearch()`:
  foreach *view* X, U, V **do**
    Separate the non-track clusters $\{v_i\}$ into three groups depending on transverse distance along each plane from the track clusters $\{c_i\}$: "fuzz" (distance ≤ 3 strips), "flag" (distance > 3 and ≤ 9 strips), and "out" (distance > 9 strips);
    "Fuzz" clusters are ignored (since these can include δ rays from the muon and are sensitive to modelling);
    Order "Fuzz" and "Out" clusters by increasing plane number;
    foreach *cluster* v_i *in "Flag"* **do**
      Check for "Fuzz" clusters in subsequent planes up to four planes ahead;
      Let $v_{j>i}$ be the most upstream cluster that matches;
      Construct the transverse (along the view) slope $s = \frac{t_j - t_i}{j - i}$;
      The error-on-slope is given by varying $t_{i,j}$ by the errors on the individual t and recalculating the slope;
      Let $\{v_k\}$ be the "Out" clusters within four planes from v_j;
      if *Exist at least two* v_k *whose position is compatible with expectation value* $t_i + (s \pm \delta s) \cdot (k - i)$ **then**
        Calculate intercept on track, $a = (t_a - t_i)/s + i$;
        return a
      end
      Let $\{v_l\}$ be the "Fuzz" clusters within four planes from v_j;
      if *Exist at least three* v_l *compatible with expectation position* **then**
        Calculate intercept on track, $a = (t_a - t_i)/s + i$;
        return a
      end
    end
  end
  return a_{geom} *plane*
return $\min(a_{\text{dep}}, a_{\text{geom}})$

Algorithm 6: Pion identification algorithm.

the length of material traversed is estimated as $L_{\text{plane}} \cdot \sqrt{(\mathrm{d}x/\mathrm{d}z)^2 + (\mathrm{d}y/\mathrm{d}z)^2 + 1}$, where L_{plane} is the longitudinal width of the plane. From this, V is constructed as an ordered list of the $\mathrm{d}E/\mathrm{d}x$ along all the planes of the track.

Because the energy deposition of a minimum ionising particle follows a Landau distribution (assuming, for now, infinitely good detector energy resolution), which has a slowly diminishing high-energy tail, it is expected that individual depositions from a Landau distribution can have values significantly higher than the most-probable value (MPV). The deposition pattern of a single MIP follows a Landau $\mathcal{L}_{\text{MIP}}$, and that of a double-MIP should be the convolution of two Landau distributions $\mathcal{L} * \mathcal{L}'$. This is true, as the probability distribution for the sum of two independent random variates X, Y distributed as $X \sim f, Y \sim g$ is just $X+Y \sim f*g$. Moreover, in the regime where $E_{\mu,\pi} \gg m_{\mu,\pi}$ and given that $m_\mu \simeq m_\pi$, one can approximate $\mathcal{L}_\mu \simeq \mathcal{L}_\pi$. Putting all this together, the deposition pattern of a muon-pion pair is expected to follow the autocorrelation *stability condition*

$$\mathcal{L}_2(E; \mu_2, \sigma_2) = \mathcal{L}_1(E; \mu_1, \sigma_1) * \mathcal{L}_1(E, \mu_1, \sigma_1) = \mathcal{L}(E; 2\mu_1, 2\sigma_1), \tag{7.1}$$

which can be verified by inspecting the distribution of $X+Y$ explicitly. In Figure 7.5, two random variates X, Y are generated with the same Landau distribution $\mu_1 = 0.5, \sigma_1 = 1$ plotted as the red dashed line; their sum is then used to generate the histogram in yellow. The blue line is a Landau distribution with $\mu_2 = 1.0, \sigma_2 = 2$, which as expected describes the histogram.

In a detector, of course, the distribution of energy depositions is convolved with a normal distribution $\mathcal{N}(\mu^{(\mathrm{G})}, \sigma^{(\mathrm{G})})$ due to the finite energy resolution of the detector. The above argument still applies, however, as the normal distribution also has an autocorrelation stability condition

$$\mathcal{N}_2(E; \mu_2^{(\mathrm{G})}, \sigma_2^{(\mathrm{G})}) = \mathcal{N}_1(E; \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})}) * \mathcal{N}_1(E; \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})}) = \mathcal{N}(E; 2\mu_1^{(\mathrm{G})}, \sqrt{2}\sigma_1^{(\mathrm{G})}). \tag{7.2}$$

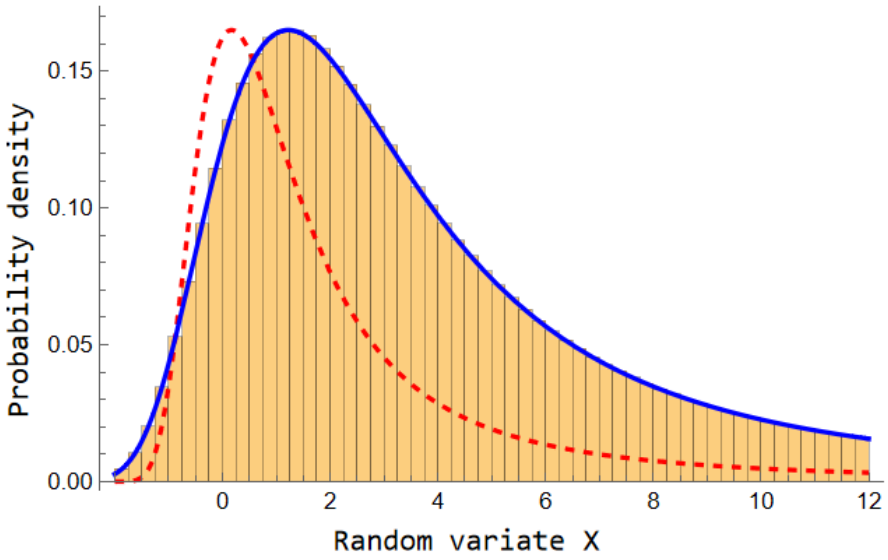


Figure 7.5: The autocorrelation of a Landau distribution with MPV μ_1 and width σ_1 (red dashed curve) is a Landau distribution with MPV $\mu_2 = 2\mu_1$ and width $\sigma_2 = 2\sigma_1$ (blue solid curve). The yellow histogram is the distribution of the sum $X + Y$, where X, Y are distributed according to the red curve.

Using the commutativity of the convolution operator, for the full distribution $(\mathcal{L} * \mathcal{N})(E)$, it is

$$\begin{aligned}
&(\mathcal{L} * \mathcal{N})_2(E; \mu_2, \sigma_2, \mu_2^{(\mathrm{G})}, \sigma_2^{(\mathrm{G})}) \\
&\quad = (\mathcal{L} * \mathcal{N})_1(E; \mu_1, \sigma_1, \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})}) * (\mathcal{L} * \mathcal{N})_1(E; \mu_1, \sigma_1, \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})}) \\
&\quad = \left(\mathcal{L}_1(E; \mu_1, \sigma_1) * \mathcal{N}_1(E; \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})})\right) * \left(\mathcal{L}_1(E; \mu_1, \sigma_1) * \mathcal{N}_1(E; \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})})\right) \\
&\quad = (\mathcal{L}_1(E; \mu_1, \sigma_1) * \mathcal{L}_1(E; \mu_1, \sigma_1)) * \left(\mathcal{N}_1(E; \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})}) * \mathcal{N}_1(E; \mu_1^{(\mathrm{G})}, \sigma_1^{(\mathrm{G})})\right) \\
&\quad = \mathcal{L}_1(E; 2\mu_1, 2\sigma_1) * \mathcal{N}(E; 2\mu_1^{(\mathrm{G})}, \sqrt{2}\sigma_1^{(\mathrm{G})}).
\end{aligned} \tag{7.3}$$

In other words, assuming the reconstructed energy distribution has no bias ($\mu_1^{(\mathrm{G})} = 0$), the observed distribution will have a most probable value at twice that of the single-MIP track, and a worse energy smearing.

To reduce the impact of these high fluctuations, the method of the *truncated mean* [277, 278] is applied: out of the first $N \leq 20$ planes over which the $\mathrm{d}E/\mathrm{d}x$ vector is constructed, those that are in the lowest fraction f are kept, while the rest are discarded, and the average is computed over the rest of the planes. The fraction f that yields the optimal behaviour for the $\langle \mathrm{d}E/\mathrm{d}x \rangle_f$ distribution (smallest

flucuations of the distribution) is between $f = 0.35$ and $f = 0.75$ [279]; in this analysis, the value $f = 0.65$ was chosen in order to maintain the largest possible pool of $\mathrm{d}E/\mathrm{d}x$ measurements to make the average, and because no substantial increase in signal / background separation was observed by switching to more aggressive values $f = 0.5$ and $f = 0.35$.

7.2 Event signature

The characteristic topology of the double-MIP, one-track signal may be isolated by a series of analysis cuts that serve to reduce the contamination of SM backgrounds. The cutoff value for two straight particle tracks to appear may be estimated as the angle that yields a separation of at least one strip (such that clusters with different positions are formed) over the minimum number of modules needed to form a track seed in a view, which is four (taking the track seed to be in the X view). To get a separation of one strip over four planes, the transverse position difference of the true pion and true muon trajectory must be at least the width of a strip, 33 ± 0.5 mm. Meanwhile, the width of a module is 45 ± 0.8 mm, meaning the minimum angle is estimated at about 10°.

To illustrate the difference in the phase space between HNL decay events and the known "irreducible" CC coherent pion background, the final-state muon energy versus muon-pion opening angle and muon energy versus pion energy is plotted in Figure 7.6. The approximate lower limit to opening angle is shown as the dashed blue line in the leftmost three panels. Immediately, the following become apparent: firstly, the HNL mass greatly influences the final state kinematics, as the Lorentz boost factor of the HNL (and hence the final state) is proportional to $1/M_{\mathrm{N4}}$; secondly, all the HNL samples look quite different from the expected SM distribution in two key ways. On the one hand, there is significantly less contamination from the SM in the low $\theta_{\mu\pi}$ region than in the high $\theta_{\mu\pi}$ region; on the other, the pion energies from HNL are significantly higher than from the coherent background.

It is worth noting here that the distribution of events in HNL energy in the detector is not the same as the spectrum of HNL that cross the detector;

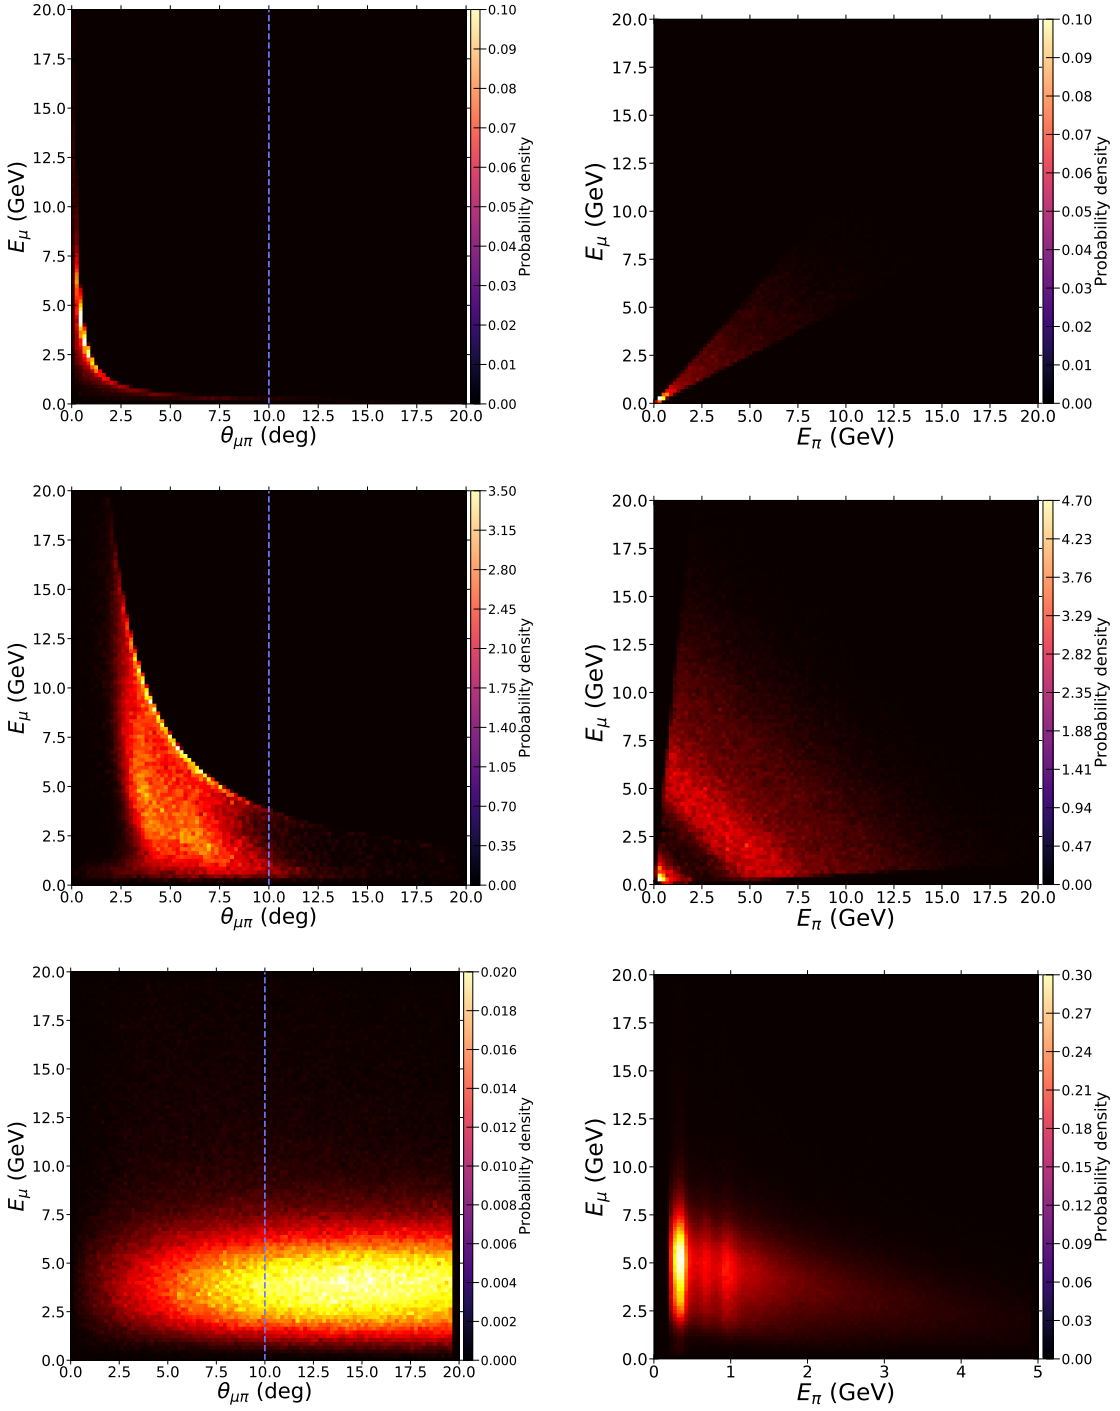


Figure 7.6: Phase space for muon+pion final states, from HNL decays at masses 250 (top) and 475 (middle) MeV/c^2, and for CC coherent pion production (bottom). `GENIE` truth-level information. Note the change in E_π axis range in the bottom right (CC coherent) panel; the pions from CC coherent are typically at much lower energy.

Figures 7.7, 7.8 illustrate the difference. The flux of HNL *crossing* the detector (top panels) is determined almost entirely by the kinematics of HNL emission, since at currently non-excluded mixing elements $|U_{\alpha 4}|^2$ the survival probability (Eq. 5.25) is $\simeq 1$. In the top right panel, one may notice how the prevalence of backwards-emitted HNL and HNL from lower-energy parents causes a "second peak" at low E_{N4}. On the other hand, the small size of the detector disproportionately enhances the lower-β HNL component, leading to a characteristic enhancement of low-energy HNL (bottom panels). This energy spectrum of HNL decays (which might also be seen as a "decaying flux", i.e. the flux of HNL that decay in MINERνA) informs the distributions of final-state particles. The effect of the finite detector size is shown in Figures 7.7 and 7.8, where the top rows are the flux of particles *crossing* the detector, and the bottom (the "decaying flux") is the energy spectrum of particles that produce a visible signature.

The regions of high and low energy muons are, nevertheless, clearly separated by their ability to exit the detector and be matched to MINOS downstream. In practice, there is a lower muon energy threshold of about 1.5 GeV and an angular acceptance 20° required for a muon to be matched in MINOS, which separates out the lower energy components. Figure 7.9 highlights this separation, for the masses $M_{\mathrm{N4}} = 250$ GeV/c^2 (left) and $M_{\mathrm{N4}} = 475$ GeV/c^2 (right); the top panels show the reconstructed distribution of muon energy for events which were MINOS matched, and the bottom panels that for events which were not MINOS matched. The high-energy tail from the lower panels corresponds to those muons which exited out the side of the detector. Figure 7.10 illustrates the difference between MINOS-matched muons for the HNL signal samples and typical backgrounds, before analysis cuts: higher-energy muons have been produced from the decays of higher-energy (hence higher-β_{N4}) HNL, and are thus more collimated with the neutrino direction. Since the HNL are constrained, for all intents, to follow the beam direction, this translates to a constraint on the angle θ_μ *with respect to the nominal beam direction* which does not follow for SM neutrino interactions, where the momentum of the muon is determined by the momentum *transfer* from the neutrino to the nucleus.

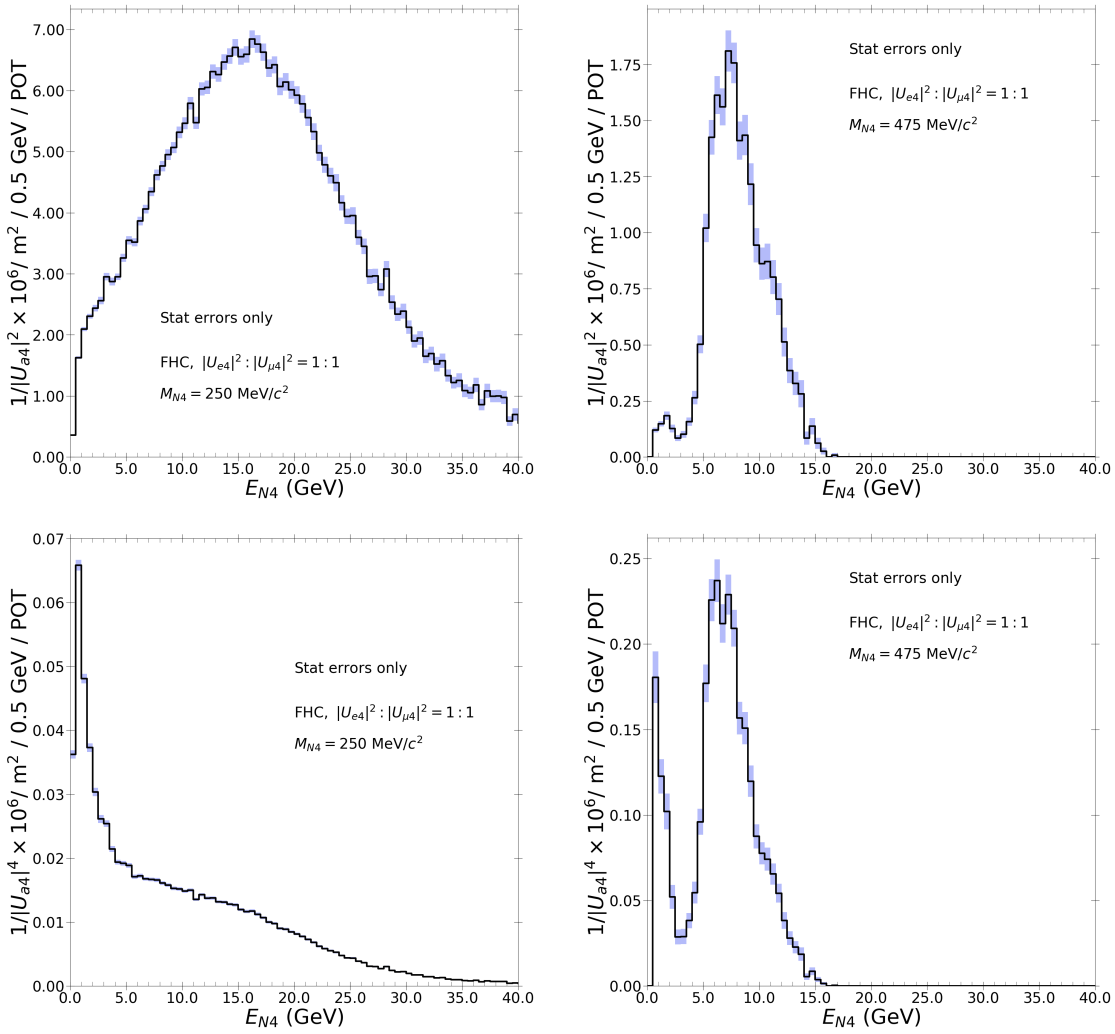


Figure 7.7: Fluxes of HNL (FHC) crossing (top panels) and decaying in (bottom panels) the MINERνA detector.

7.3 One-track classification

7.3.1 Pre-selection cuts

A sequence of pre-selection cuts are applied to the data to remove poorly reconstructed events and select for events in those parts of physical space and phase space that are best understood. These may be summarised in the following list:

- Reconstruction cuts: Removes events with objects that cause the general reconstruction algorithm to fail.

- Fiducial volume: Require reconstructed vertex to be in the tracker z range

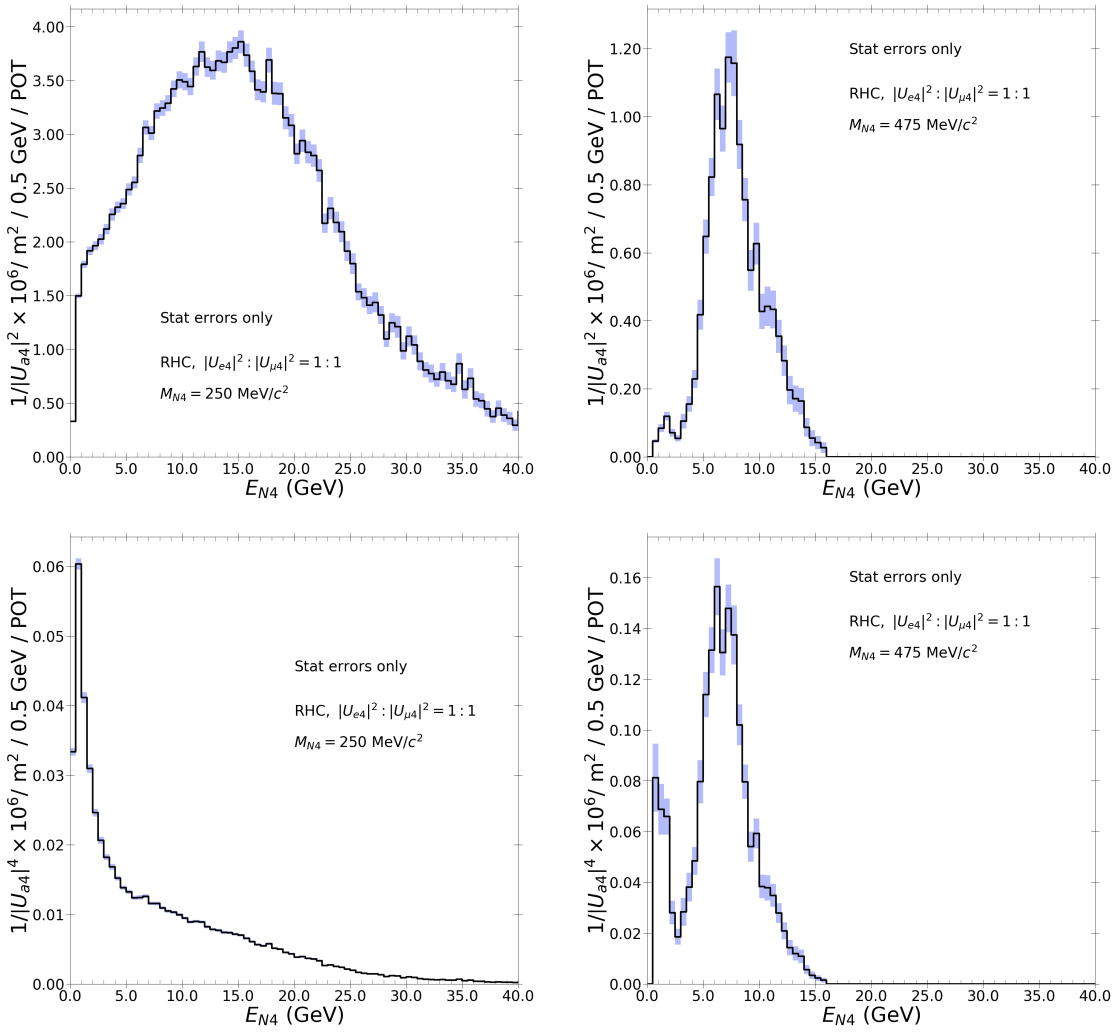


Figure 7.8: Fluxes of HNL (RHC) crossing (top panels) and decaying in (bottom panels) the MINERνA detector.

$[5980, 8422]$ mm and within a hexagonal apothem $a = 850$ mm.

- Deadtime: Remove events with reconstructed vertex within the dead time of a channel.

- Track multiplicity: Remove events with > 1 tracks after HNL reconstruction has been applied.

- Ill-defined track direction: Remove events whose track slopes form an angle less than $0.1°$ from the z axis, which is the default output of the tracker algorithm. This is an angle of $3.343°$ with respect to the nominal beam angle.

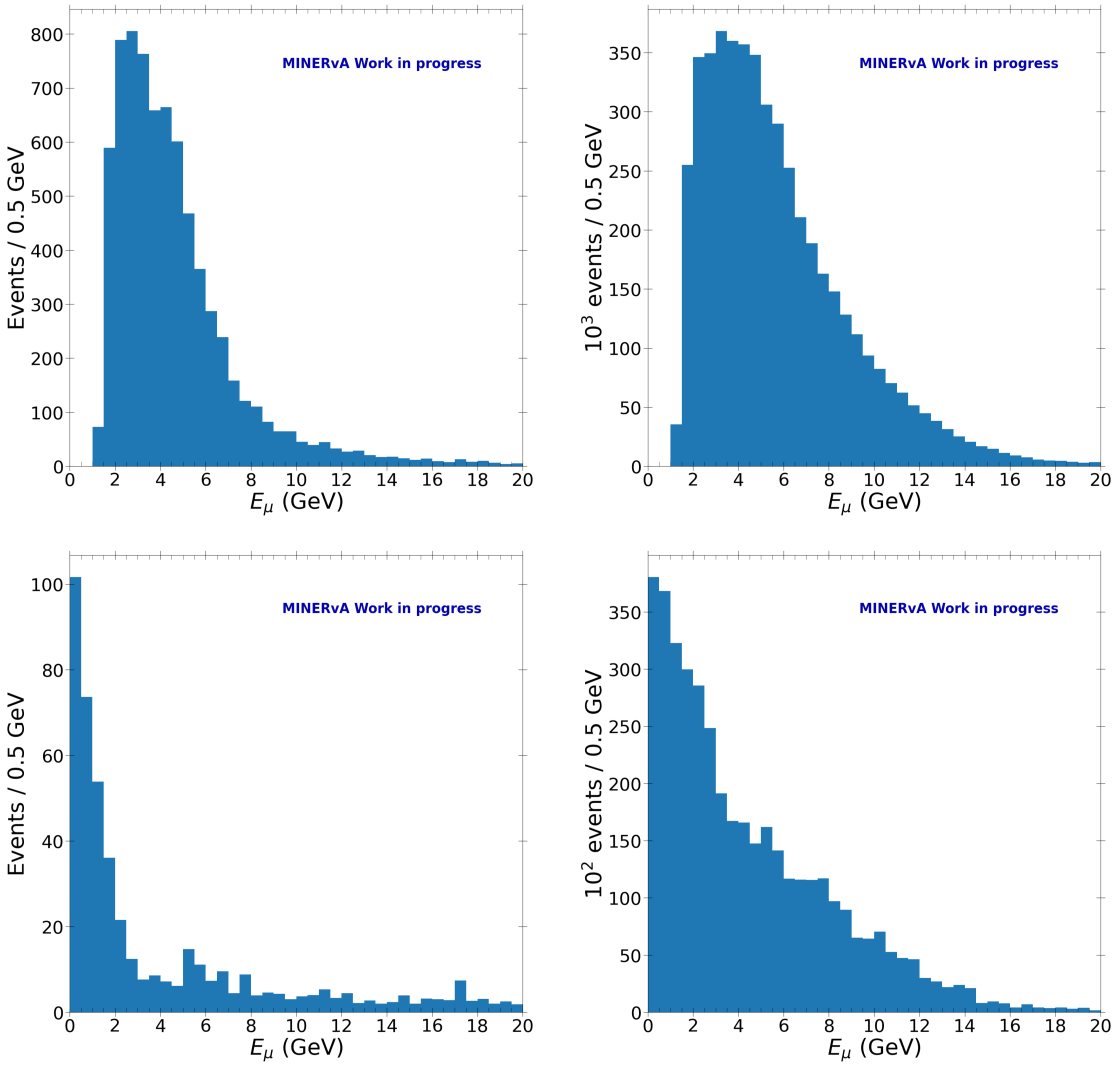


Figure 7.9: Reconstructed muon energy distributions (before analysis cuts) of HNL at masses $M_{\mathrm{N4}} = 250, 475$ MeV/c^2 (left, right) and which either were (top) or were not (bottom) accepted by MINOS.

- MINOS match: Remove events without a matching MINOS track; this ensures there exists a μ^- or μ^+.

The efficiencies in data and Monte Carlo can be summarised in Tables 7.1 (absolute efficiencies), and 7.2 (efficiency for each cut).

For most of these cuts, the efficiency in the mainline MC and the data is comparable. There are, however, a few noteworthy exceptions which bear discussion. Firstly, the reconstruction cut efficiency, which is defined as the number of events which is fully reconstructed. This is normalised to the number of events simulated,

FHC [cumulative efficiencies normalised to events in FV]							
Cut	Data	MC-General	CC-COH	CC-DIF	IMD	HNL-250	HNL-475
In FV	100%	100%	100%	100%	100%	100%	100%
One track	79.67%	75.34%	61.78%	62.12%	95.80%	94.69%	86.13%
Exclude z axis	40.79%	32.47%	33.76%	28.23%	79.58%	55.49%	46.62%
MINOS match	8.67%	10.58%	27.97%	22.73%	19.01%	35.65%	31.65%
RHC [cumulative efficiencies normalised to events in FV]							
Cut	Data	MC-General	CC-COH	CC-DIF	IMD	HNL-250	HNL-475
In FV	100%	100%	100%	100%	–	100%	100%
One track	83.59%	80.75%	63.62%	63.43%	–	94.46%	86.33%
Exclude z axis	48.79%	43.65%	36.28%	29.88%	–	59.79%	46.30%
MINOS match	17.26%	22.40%	30.71%	24.41%	–	32.13%	32.65%

Table 7.1: Pre-selection cut efficiencies for data and MC. The efficiencies are normalised to all events that pass reconstruction and are in the fiducial volume.

FHC [relative efficiencies]							
Cut	Data	MC-General	CC-COH	CC-DIF	IMD	HNL-250	HNL-475
Reconstruction	100.00%	79.19%	92.36%	92.50%	71.57%	90.46%	91.10%
Tracker	28.61%	51.95%	56.16%	64.05%	42.10%	86.22%	86.14%
Apothem	25.96%	54.57%	61.73%	65.91%	67.51%	71.97%	71.56%
Deadtime	99.40%	99.33%	98.72%	98.69%	99.75%	99.06%	98.96%
One track	79.60%	75.34%	61.78%	62.12%	95.80%	94.68%	86.13%
Exclude z axis	51.20%	43.09%	54.63%	45.45%	83.06%	58.60%	54.13%
MINOS match	21.23%	32.57%	82.85%	80.51%	23.89%	64.24%	67.89%
RHC [relative efficiencies]							
Cut	Data	MC-General	CC-COH	CC-DIF	IMD	HNL-250	HNL-475
Reconstruction	100.00%	71.99%	86.35%	87.11%	–	82.43%	84.56%
Tracker	25.59%	51.00%	54.85%	60.60%	–	86.08%	86.03%
Apothem	27.62%	57.52%	63.94%	67.78%	–	74.22%	73.66%
Deadtime	99.18%	99.14%	99.04%	99.05%	–	99.40%	99.27%
One track	83.63%	80.78%	63.62%	63.43%	–	94.47%	86.34%
Exclude z axis	58.39%	54.03%	57.02%	47.11%	–	60.12%	53.64%
MINOS match	35.41%	51.28%	84.65%	81.68%	–	56.58%	70.50%

Table 7.2: Pre-selection cut efficiencies for data and MC.

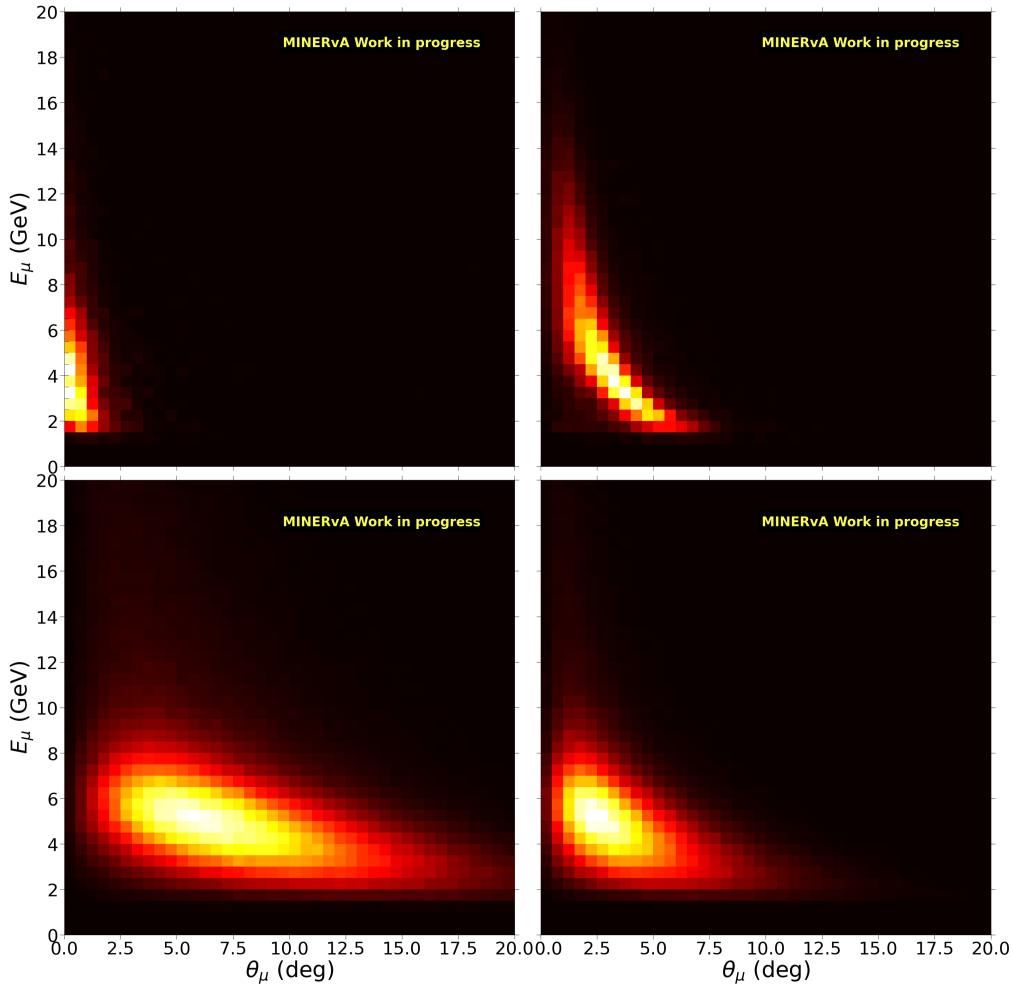


Figure 7.10: MINOS-accepted muon energy (vertical) versus angle with respect to beam nominal direction (horizontal), for the HNL masses $M_{\mathrm{N4}} = 250, 475$ MeV/c^2 (top panels), for inclusive SM CC neutrino interactions (bottom left), and for CC coherent pion production (bottom right). No analysis cuts other than successful reconstruction and MINOS acceptance have been applied. The constraint of HNL collimation is apparent by visual inspection.

for the MC samples. For data, naturally, this definition makes no sense, so it will be assumed that the reconstruction efficiency in data is the same as the efficiency in MC. Secondly, the fiducial volume cuts: for the data, the efficiency is significantly different than for either the mainline MC or for the HNL samples. This reflects the underlying simulation, which does not include the nuclear target region or the side and downstream calorimeters (apart from the IMD MC which includes everything apart from downstream calorimeters). The difference may be summarised in Figures 7.11 (Tracker) and 7.12 (Apothem). Thirdly, the MINOS

matching efficiency, which is overestimated in the MC by a factor ~ 1.5. Therefore, a correction must be made in the MC and a flat systematic uncertainty of 5.8% assigned to the MC. This uncertainty is obtained by comparing the correction in the FHC and RHC run periods, obtaining the difference, and normalising to the correction itself. The "one-track" cut performs decidedly better on the HNL and IMD samples, as expected; it is also less efficient for heavier HNL, where one would expect more events with high $\theta_{\mu\pi}$. There also appears to be a poor performance of the cut that excludes the default z axis across all samples; since this cut performs about the same in all of the samples after one-track has been applied apart from the IMD, it is likely that this performance does not reflect underlying physics.

7.3.2 Cost-benefit analysis

There is a sequence of analysis cuts made to isolate the region of events producing a collimated muon-pion pair, designed to maximise the efficiency of HNL events for all the samples that could occur in the detector. At the same time, it is of paramount importance to reduce the SM background contribution as much as possible, so that the purity of the sample is as high as possible. This is especially important in the context of a rare search, given that the level of signal dominance in a sample of selected interactions is not *a priori* determined as the mixing element is unknown. As a consequence, a "cost-benefit" approach is defined to establish cut values for each variable of interest. For each variable, a distribution is made in that variable of the signal samples for $M_{\mathrm{N4}} = 250$ MeV/c^2 and $M_{\mathrm{N4}} = 475$ MeV/c^2, as extremal points of the HNL phase space being probed. The distribution is also made for each of the samples described in Chapter 6, to understand the capability of the cuts to filter out events that do not match the criteria of double-MIP events.

Therefore, the two metrics of interest are the *signal efficiency* ϵ_{sig} and the *discrimination capability* $\kappa = -\log_{10}(\epsilon_{\mathrm{bkg}})$: for both of these, the higher the metric, the better the cut performs. These metrics can be defined cumulatively (i.e. in terms of the total selection / rejection across all cuts), and per individual cut. Their relative importance is defined, in this thesis, as a curve in the $\mathbb{R}^2$-space $(\epsilon_{\mathrm{sig}}, 1 - \epsilon_{\mathrm{bkg}})$

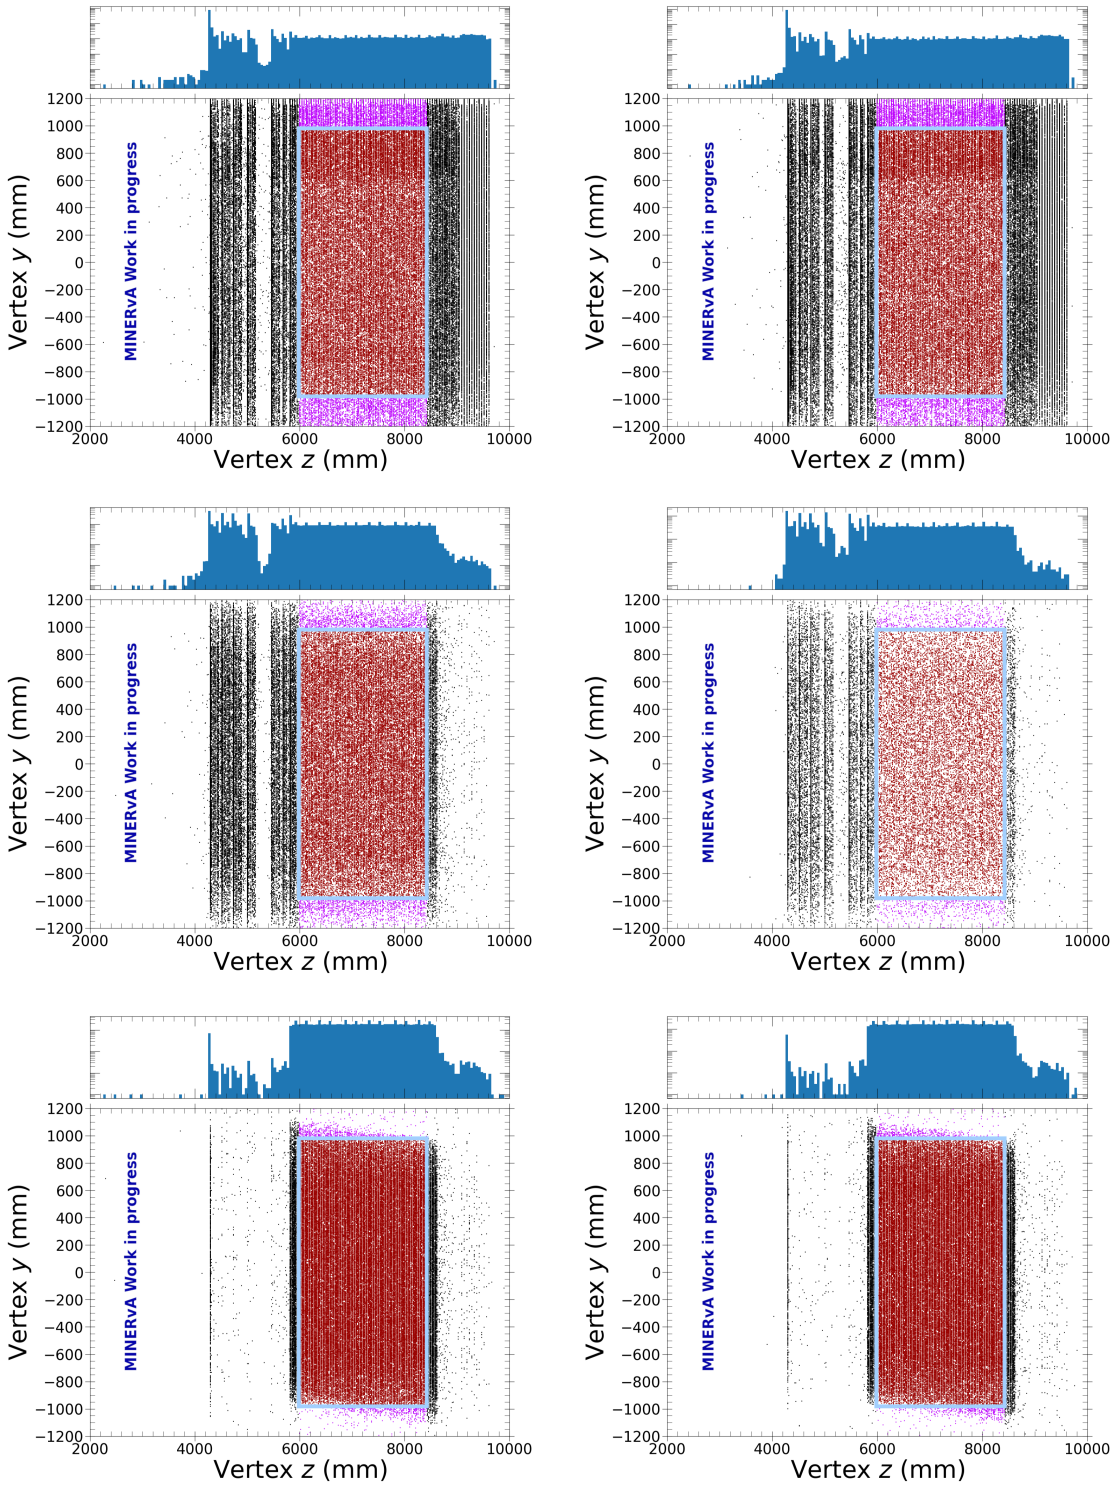


Figure 7.11: Vertices of events in data (top), mainline MC (middle), and HNL signal (bottom) for FHC (left) and RHC (right). z axis is horizontal pointing right and y is vertical pointing up. Beam axis has angle 3.343° with z. Fiducial region highlighted in blue. Red events are fully in the fiducial region (pass Apothem and Tracker cuts). Black events are outside the tracker region.

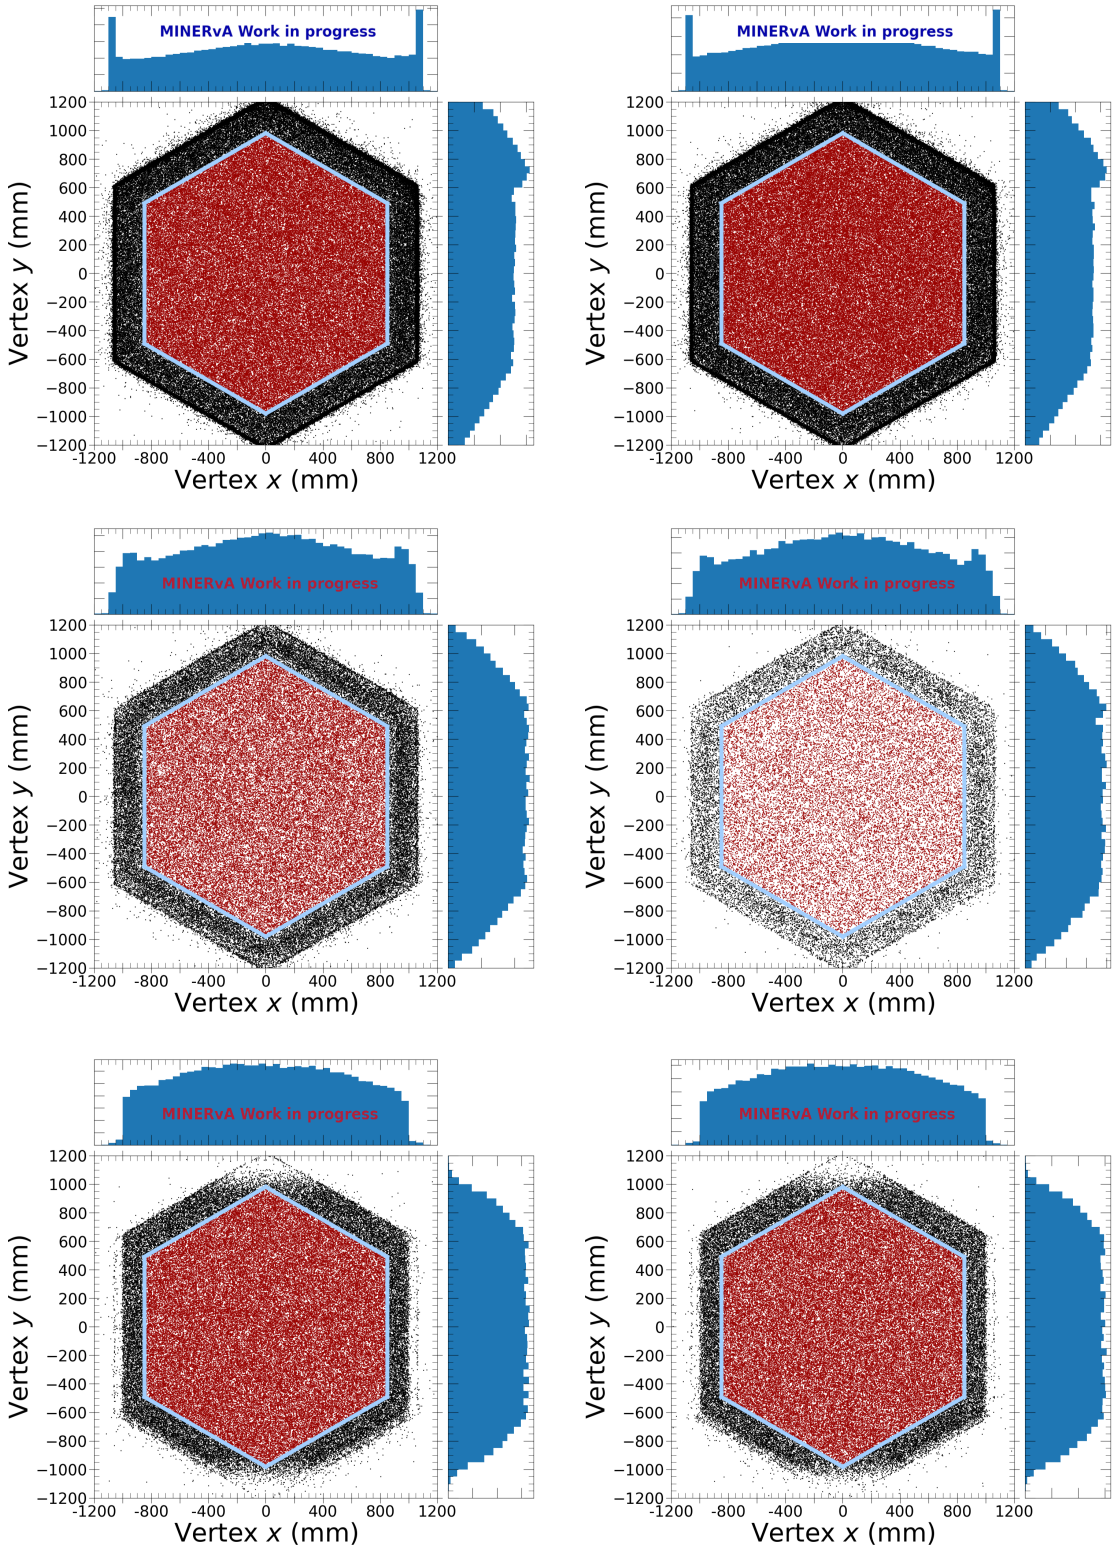


Figure 7.12: Vertices of events in data (top), mainline MC (middle), and HNL signal (bottom) for FHC (left) and RHC (right). x axis is horizontal pointing right and y is vertical pointing up. Beam axis points out of page. Fiducial region highlighted in blue. Red events are fully in the fiducial region (pass Apothem cut). Black events are outside the fiducial region.

for each cut. One may conceptualise the curve as the strictness of each cut in selecting an increasingly small subset of the signal (background) distribution. As the cut value on a variable is altered, so too do ϵ_{sig} and $1-\epsilon_{\text{bkg}}$; exceedingly strict choices for cut value will lead to almost all of signal and background distribution being thrown out, whereas unreasonably lax choices will remove almost none of either.

Suppose that for a variable X with physical bounds $[X_-, X_+]$ one makes a cut at the value X_0 such that $\forall X \leq X_0$, the event is cut out. Then the previous statment may just be written as

$$\lim_{X_0 \to X_-} (\epsilon_{\text{sig}}, 1-\epsilon_{\text{bkg}}) = (1,0)\,,$$
$$\lim_{X_0 \to X_+} (\epsilon_{\text{sig}}, 1-\epsilon_{\text{bkg}}) = (0,1)\,.$$

Of course, the explicit dependency on X_0 may just be repackaged as an implicit dependency

$$1-\epsilon_{\text{bkg}} = 1-\epsilon_{\text{bkg}}\left(\epsilon_{\text{sig}}\right). \tag{7.4}$$

From the Figures 7.16-7.25, one sees that there is a "linear region" where as one varies the cut, the background rejection $1-\epsilon_{\text{bkg}}$ increases roughly linearly with the reduction in signal efficiency ϵ_{sig}. With this in mind, one uses Eq. 7.4 to plot the points $(\epsilon_{\text{sig}}, 1-\epsilon_{\text{bkg}})$ for each value $X_0 \in [X_-, X_+]$ of the variable, and fit the first few points $\{X_{0;j}\}_j : X_{0;j} \simeq X_+$ to a line. This exercise is repeated once using $\epsilon_{\text{sig}} \equiv \epsilon_{M_{\text{N4}}=250}$ and once using $\epsilon_{\text{sig}} \equiv \epsilon_{M_{\text{N4}}=475}$ (see Figs. 7.16, 7.19, 7.22, 7.25). These points represent the tails of both signal and background distributions being accepted post-cut. Subsequently, one identifies the points j_{250}, j_{475} at which $X_{0;j}$ lies significantly below the fit line, for both HNL signal sample curves. The most aggressive of these cuts is adopted for that variable, and the next cut is optimised. Naturally, the above discussion generalises to cuts going the other direction where one is trying to find the lower bound for a cut window.

7.3.3 Selection cuts

The sequence of cuts has been optimised to select a collimated muon-pion pair, exploiting the "clean" muon signature and simultaneous energy deposition of two

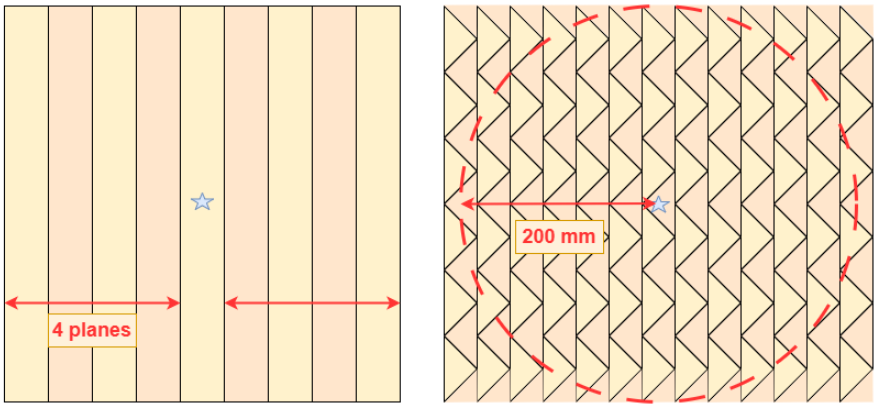


Figure 7.13: Illustration of cylinder for vertex energy cut (in this case with height = 4 planes). The energy from all clusters whose (x, y, z) position is within the cylinder is added to the "inner" vertex energy.

similar MIPs. There are four sequential selection cuts:

1. Number of clean planes, N_{planes}: This measures the length of the "clean" muon-pion track inside the tracker (i.e. no large energy depositions along the track that would signal a hadronic interaction, or a secondary track leaving the primary track). For the HNL signal samples, there is a characteristic peak in the region $15 - 30$ planes, with tails either side. The SM background, on the other hand, is strongly peaked at a small region of planes ($N < 5$), due to the emission of protons from nuclear interactions. There is also a noticeable enhancement of the distribution from $N = 12$ planes onwards, which is the number of planes in between the end of the fiducial volume and the end of the tracker region where the algorithm to calculate N_{planes} stops tracking the event; FHC is less prominent than RHC, since the proton emission from antineutrino interactions is much smaller. This enhancement is especially obvious in the IMD sample, where the track is so clean there are almost no events with $N < 12$ planes. The distributions can be seen in Figs. 7.14 and 7.15. Essentially, the stricter this cut, the more proportion of the (downstream) fiducial volume is "cut off". Figure 7.16 shows the procedure for cut optimisation: the top two panels are the plots of relative HNL efficiency (solid lines) and background rejection $\rho = 1 - \epsilon$ (dashed lines) for each of the

background samples as a function of the cut value, whereas the bottom panels show the background rejection as a function of HNL efficiency, as discussed at the end of the previous section 7.3.2. The cut value of $N = 30$ planes is driven by the CC Coherent sample.

2. "Inner" vertex energy, $E_{\text{vtx}}^{(4\text{ planes})}$: This is defined as the sum of the energies of all clusters located in a cylinder of radius in the transverse direction 200 mm and height in the longitudinal direction of 4 planes, centred on the interaction vertex. An illustration of this can be seen in Figure 7.13. From this cut, heavy ionising particles (such as protons and recoil nuclei) are picked out irrespective of their direction. The small height (180 mm) of the cylinder means that neutrons, primarily produced in the RHC mode from antineutrino-CH interactions, are not picked up by this variable, leading to a significantly different distribution for the FHC and RHC modes. Simultaneously, a degree of separation is achieved between the one-track sample with only one MIP (IMD) and the HNL signal of one track but two MIPs. Driven by the CC Diffractive mode (FHC) and inclusive sample (RHC), the cuts are set at $E_{\text{vtx}}^{(4\text{ planes})} \in [24, 31]$ MeV. The distributions and cuts can be seen in Figs. 7.17, 7.18, and 7.19.

3. "Outer" vertex energy, $E_{\text{vtx}}^{(10\text{ planes})}$: Like the previous cut, but extending the height to 10 planes. This is a height of 450 mm, which gives more time for particles to deposit energy in the detector. The HNL signal is distributed essentially like a Landau pdf, almost like the CC Coherent and Diffractive background samples which have larger high-energy tails, and the IMD distribution which is peaked at a lower MPV. The SM distribution at this point is strongly peaked at low energy as well, with a secondary bump in the HNL signal region caused by the double-MIP area of phase space. The close match of FHC and RHC distributions demonstrates the power of the "inner" cut to reduce both heavy-ionising protons and non-ionising neutrons. See Figs. 7.20, 7.21, and 7.22.

4. Truncated-mean average of energy deposition over unit length, $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$: This distribution clearly shows one single peak for both the HNL signal samples and for the IMD background sample, corresponding to either two or one particles producing ionisation deposits in the track. The CC Coherent and Diffractive samples, as well as the General sample, exhibit a more complex "double-peak" structure with a transition between two identifiable peaks around the single-MIP region at $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} \sim 1.5$ MeV/cm and the double-MIP region at $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} \sim 2.7$ MeV/cm. A cut is therefore made to the right-hand side of the double-MIP peak at $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} = 2.7$ MeV/cm to isolate the double-MIP signature, as seen in Figs. 7.23, 7.24, and 7.25. This will be elaborated on in more detail in Chapter 8.

Overall, the discrimination of SM background depends on the interaction type, which is as expected. The most "HNL-like" backgrounds are the hardest to discriminate against (which is to say, the CC Coherent and Diffractive pion production), as they overlap in phase space and topology with the HNL signal. Per sample, the relative efficiencies of the various cuts can be seen in Table 7.3.

FHC selection [relative efficiencies]							
Cut	Data	MC-General	CC-COH	CC-DIF	IMD	HNL-250	HNL-475
N planes	22.53%	22.55%	23.29%	12.60%	79.58%	52.45%	39.71%
$E_{\mathrm{vtx}}^{(4\text{ planes})}$	8.23%	8.23%	30.47%	6.75%	2.99%	49.08%	49.38%
$E_{\mathrm{vtx}}^{(10\text{ planes})}$	13.73%	14.50%	47.49%	48.51%	15.64%	65.79%	64.95%
$\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$	17.53%	15.26%	35.72%	39.78%	11.57%	64.15%	63.53%
Total	4.46×10^{-2}%	4.11×10^{-2}%	1.21%	0.164%	4.31×10^{-2}%	10.86%	8.11%
RHC selection [relative efficiencies]							
Cut	Data	MC-General	CC-COH	CC-DIF	IMD	HNL-250	HNL-475
N planes	48.16%	47.30%	27.06%	14.86%	–	53.48%	29.80%
$E_{\mathrm{vtx}}^{(4\text{ planes})}$	8.91%	8.07%	30.58%	6.70%	–	49.30%	48.92%
$E_{\mathrm{vtx}}^{(10\text{ planes})}$	14.18%	15.11%	45.53%	43.54%	–	64.28%	63.40%
$\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$	8.79%	7.61%	29.11%	33.61%	–	62.37%	53.67%
Total	5.35×10^{-2}%	4.39×10^{-2}%	1.097%	0.146%	–	10.57%	4.96%

Table 7.3: Relative efficiency of the various selection cuts, on each of the background samples and the two representative HNL signal samples.

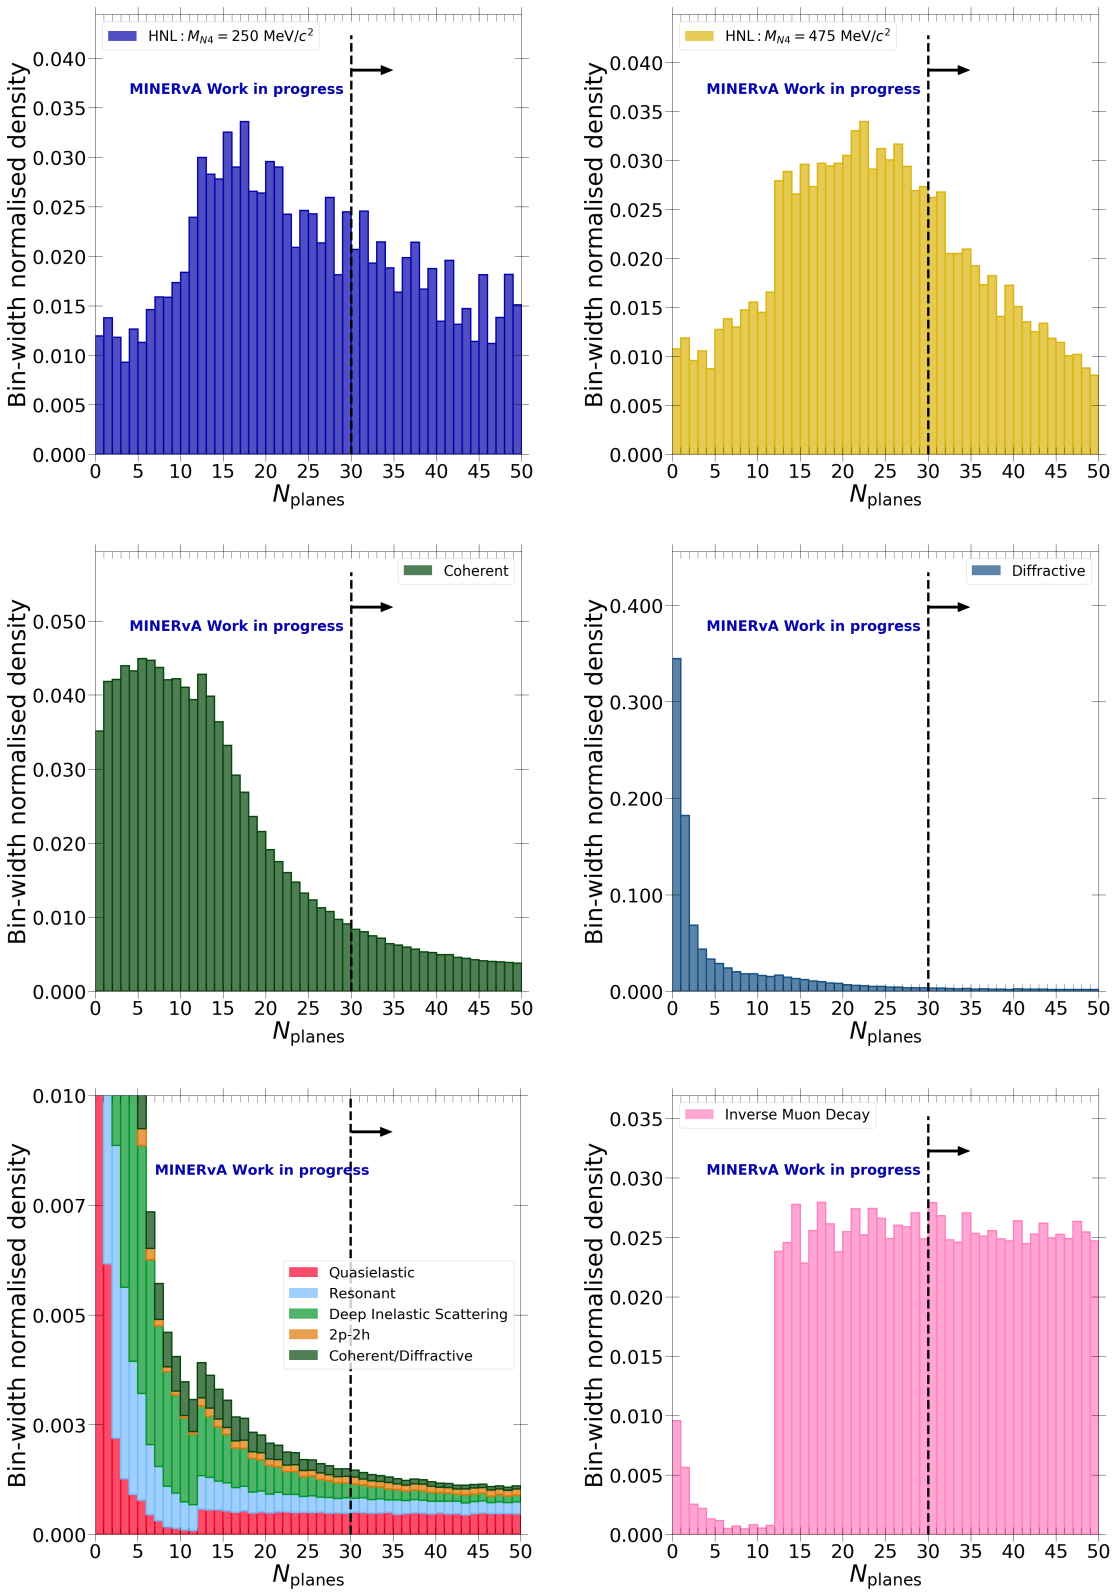


Figure 7.14: Distributions of N_{planes}, FHC. From top left: HNL ($M_{\text{N4}} = 250$ MeV/c^2), HNL ($M_{\text{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General, IMD.

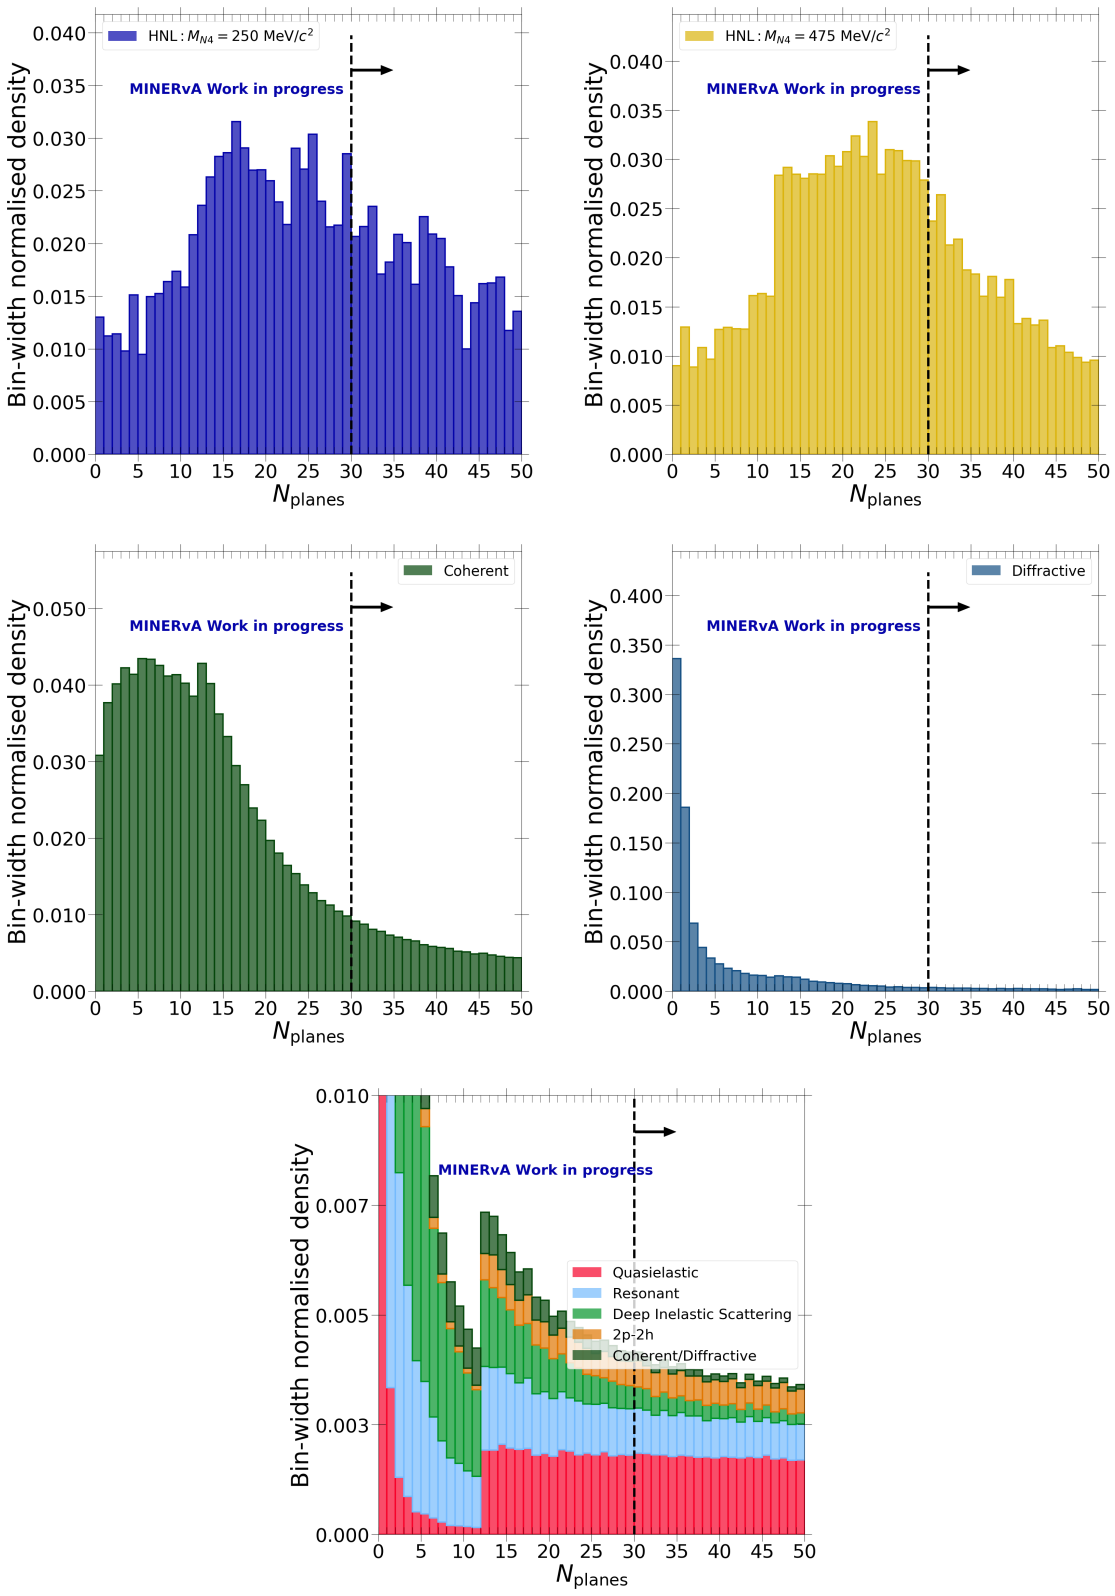


Figure 7.15: Distributions of N_{planes}, RHC. From top left: HNL ($M_{\text{N4}} = 250$ MeV/c^2), HNL ($M_{\text{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General.

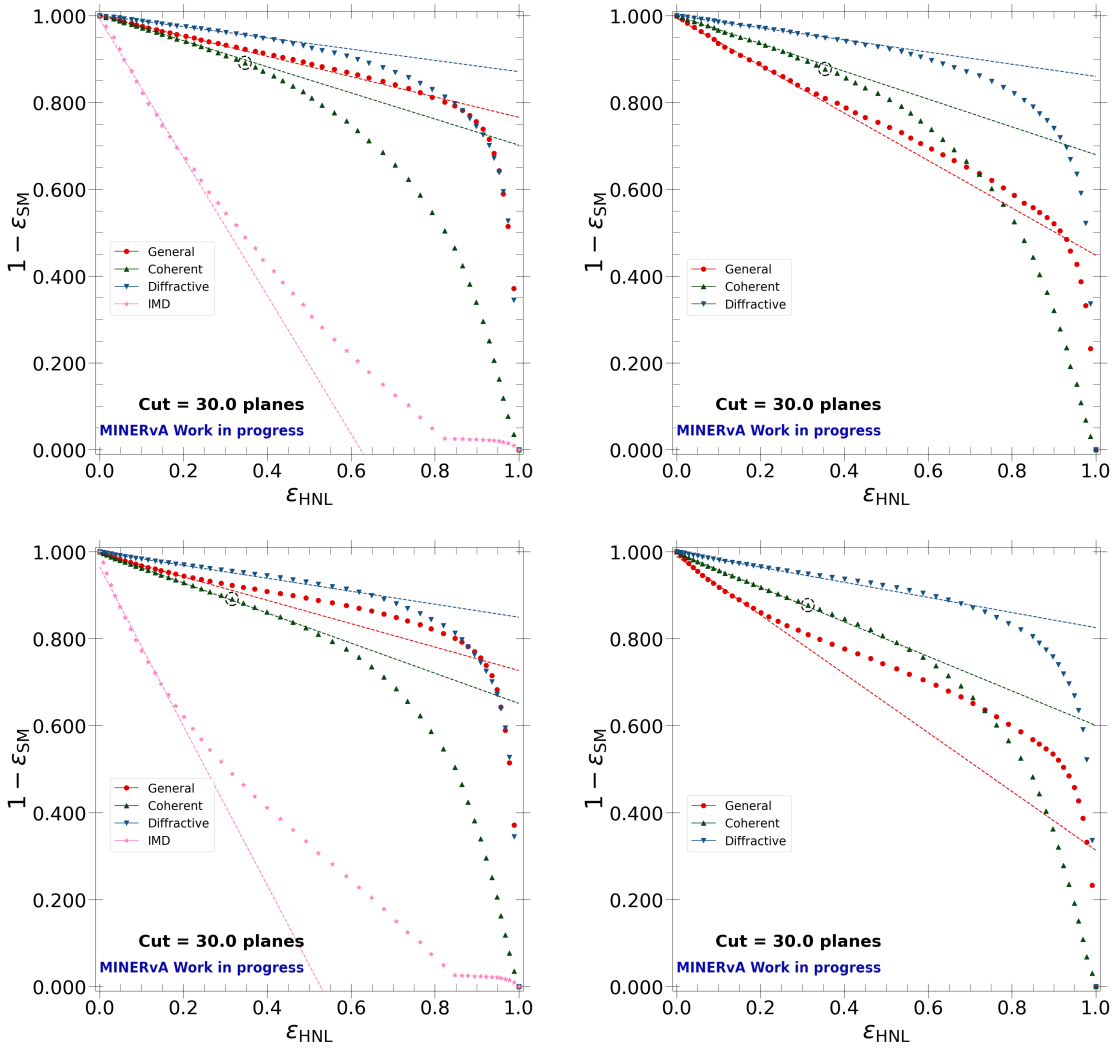


Figure 7.16: Background exclusion curves as a function of HNL signal efficiency, for various choices of the N_{planes} cut value represented as individual points. Top panels: $M_{\text{N4}} = 250$ MeV/c^2, bottom panels: $M_{\text{N4}} = 475$ MeV/c^2, left: FHC, right: RHC. On the extreme left hand side of each plot, the first few points are fitted to a line, representing the gradual inclusion of high-N_{planes} tails of both HNL signal and background efficiency. The CC Coherent (green) curve drives the choice of cut value $N_{\text{planes}} = 30$. See text for details.

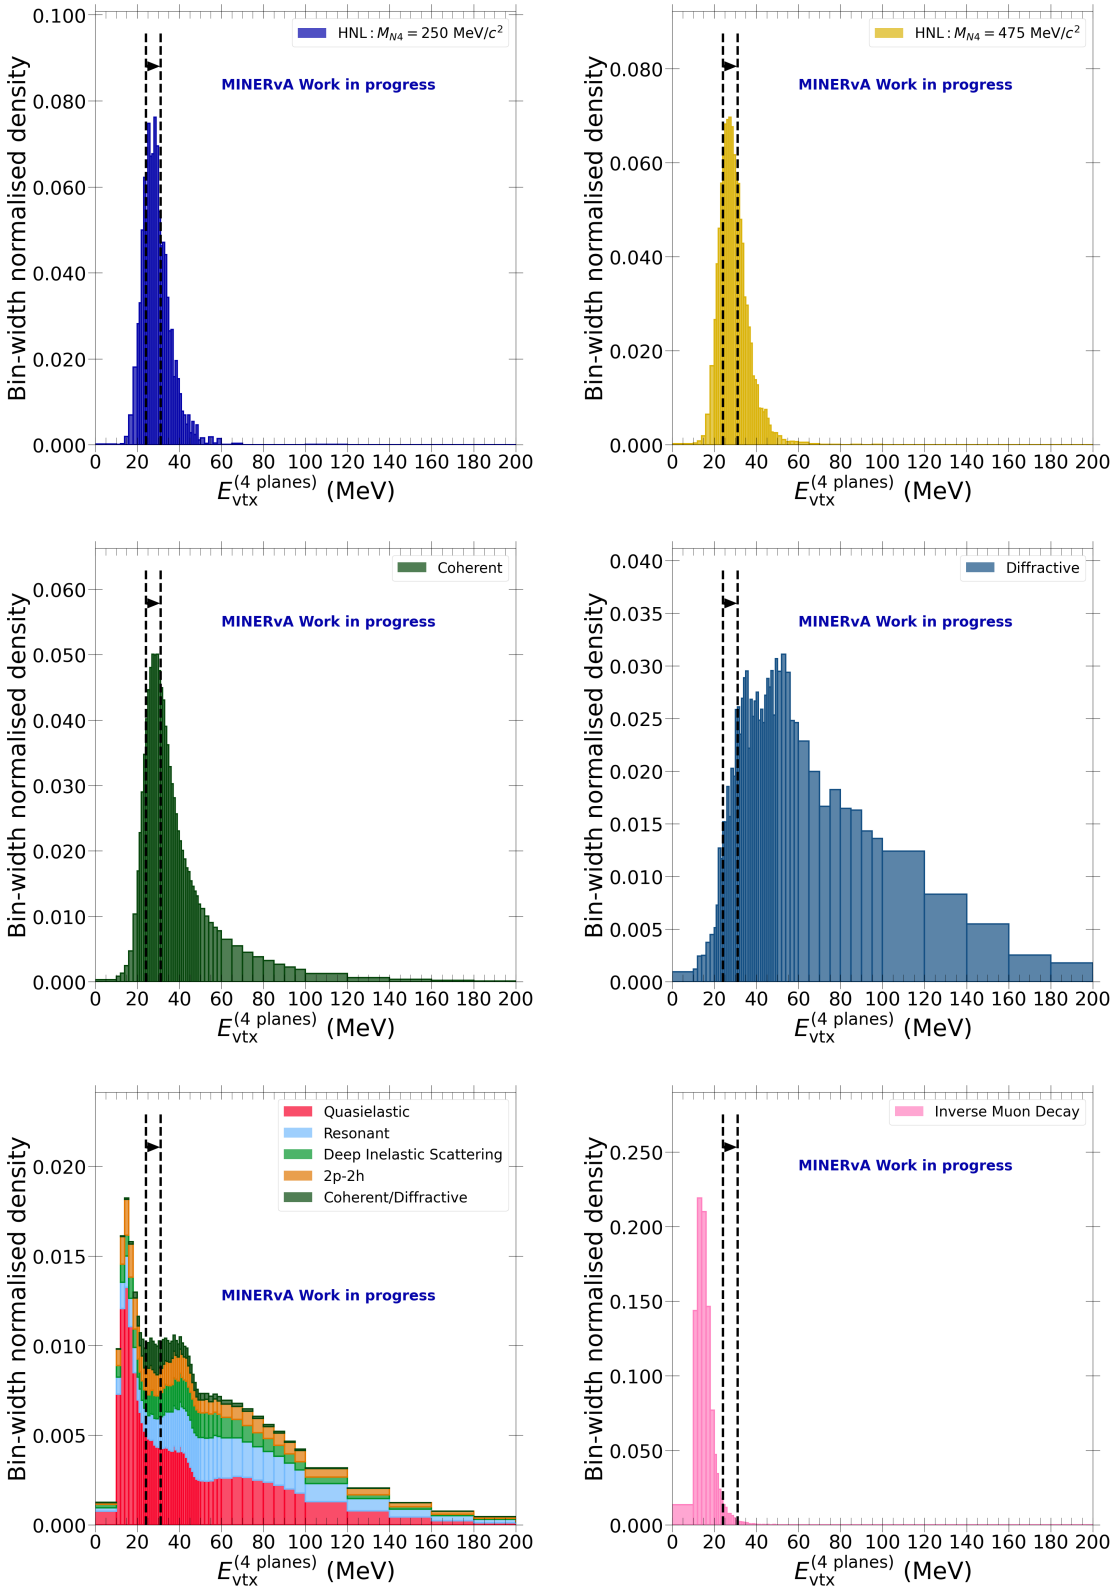


Figure 7.17: Distributions of $E_{\text{vtx}}^{(4\text{ planes})}$, FHC. From top left: HNL ($M_{\text{N4}} = 250$ MeV/c^2), HNL ($M_{\text{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General, IMD.

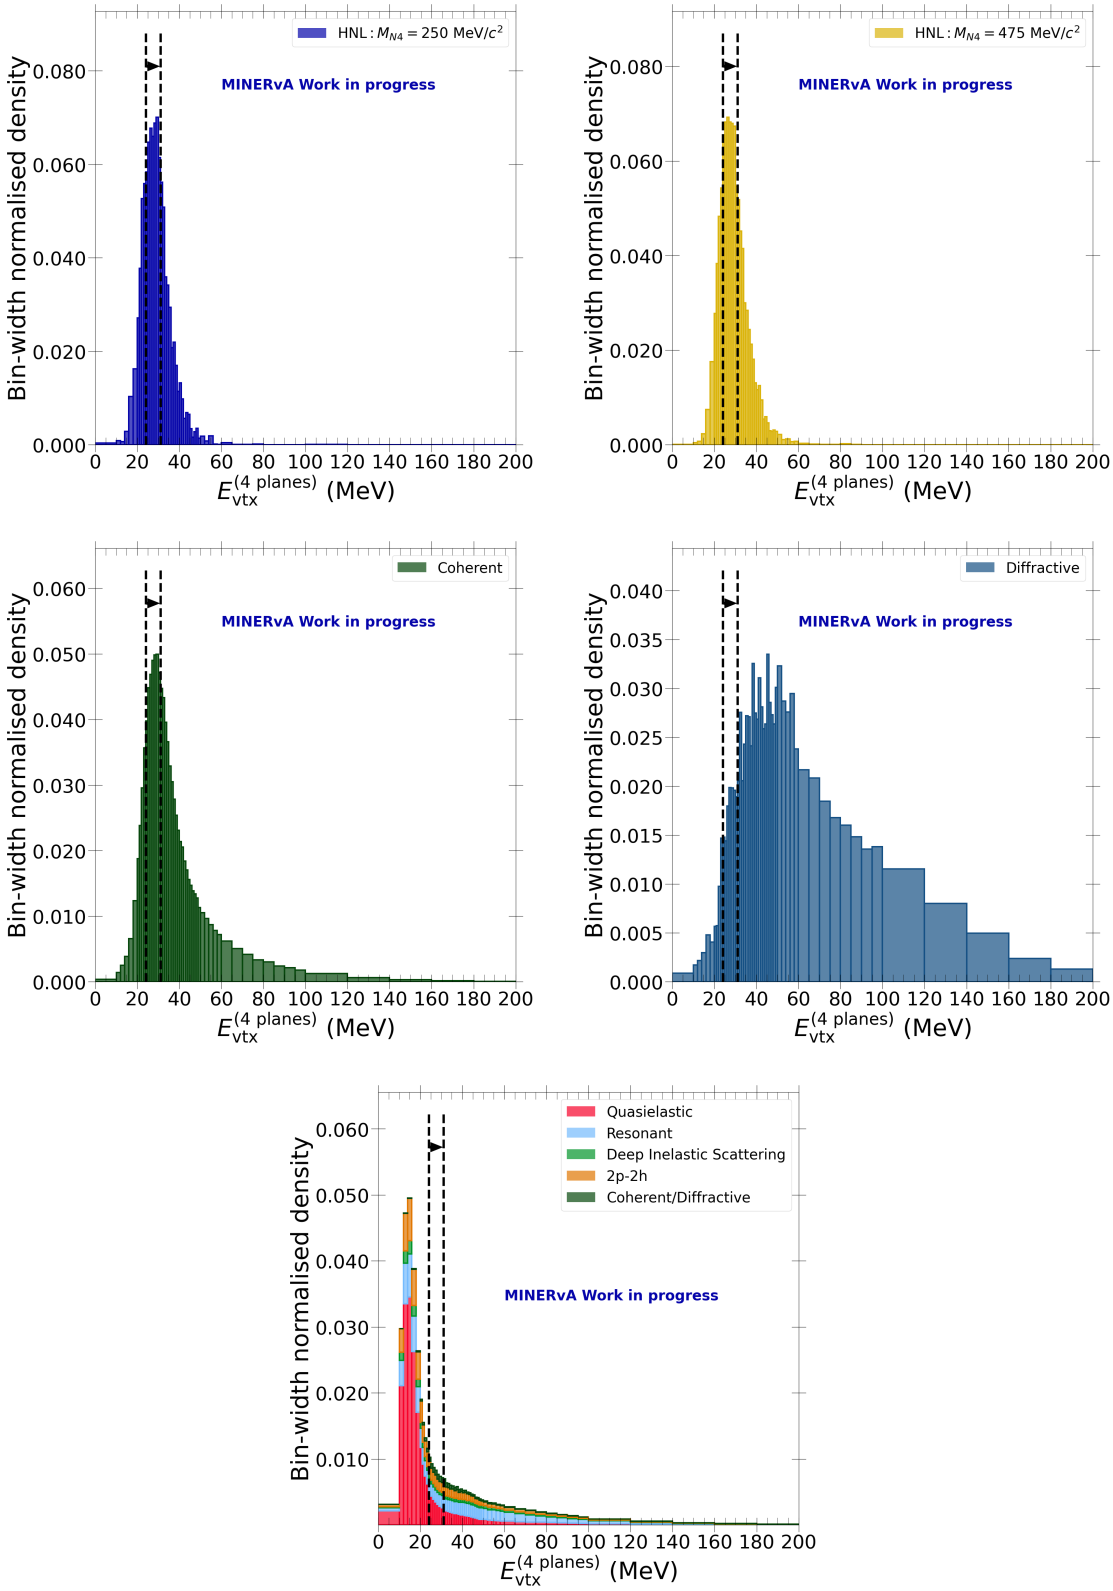


Figure 7.18: Distributions of $E_{\text{vtx}}^{(4\ \text{planes})}$, RHC. From top left: HNL ($M_{\text{N4}} = 250$ MeV/c^2), HNL ($M_{\text{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General.

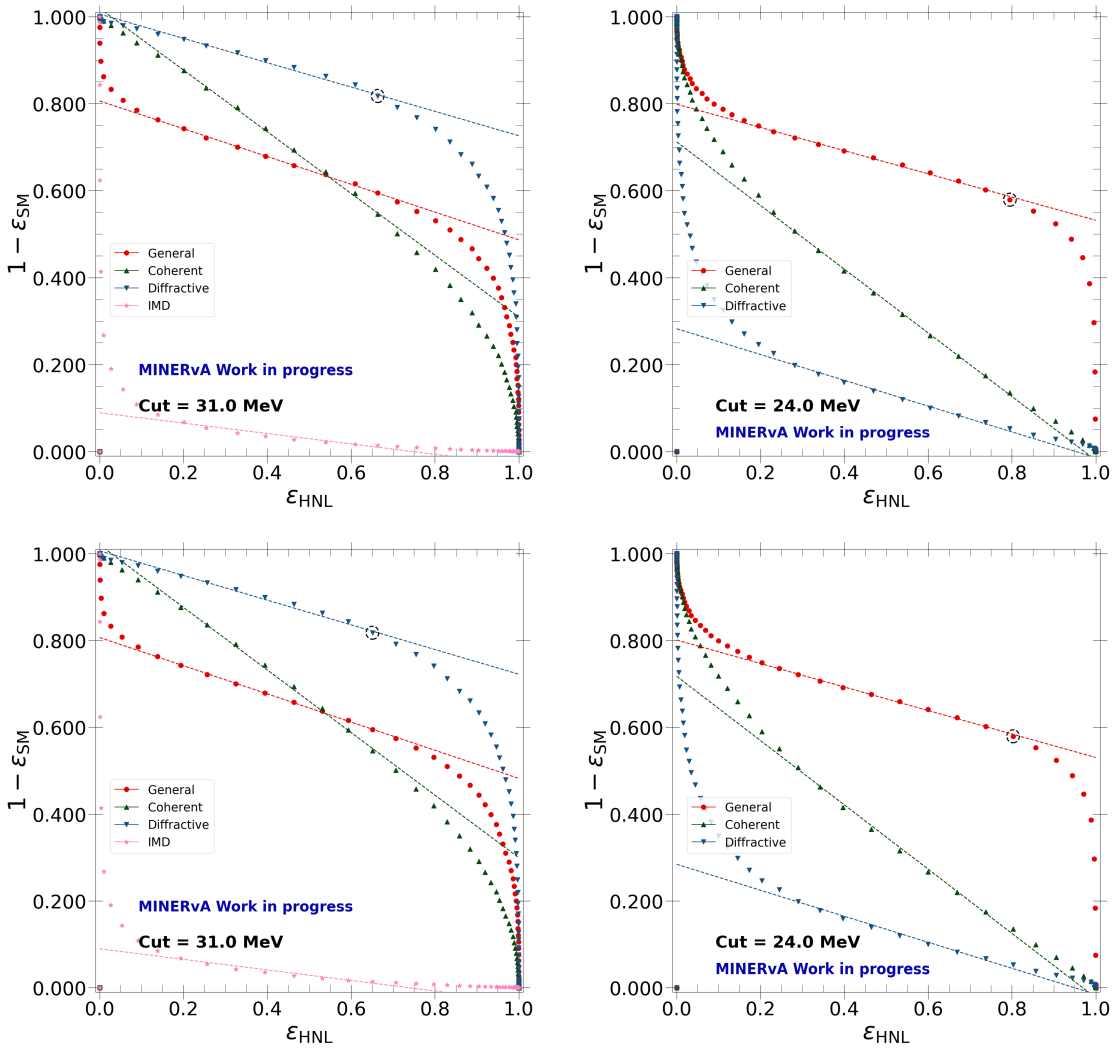


Figure 7.19: Background exclusion curves as a function of HNL signal efficiency, for various choices of the $E_{\text{vtx}}^{(4\text{ planes})}$ cut value represented as individual points. Top panels: $M_{\text{N4}} = 250$ MeV/c^2, bottom panels: $M_{\text{N4}} = 475$ MeV/c^2, left: FHC, right: RHC. FHC CC Diffractive drives the upper cut $E_{\text{vtx}}^{(4\text{ planes})} \leq 31$ MeV, and RHC General drives the lower cut $E_{\text{vtx}}^{(4\text{ planes})} \geq 24$ MeV. See text for details.

Note that the correction from the data/MC discrepancy in the MINOS matching stage does not enter Table 7.3, as it will scale the overall discrimination capability κ of each MC sample after all cuts have been applied. There are various details to note in the efficiencies of these cuts for the various samples. The data and MC-General samples perform in roughly similar ways throughout the cuts. As for the other background samples, there is significant variation: in particular, the CC Coherent and CC Diffractive samples are impacted significantly differently

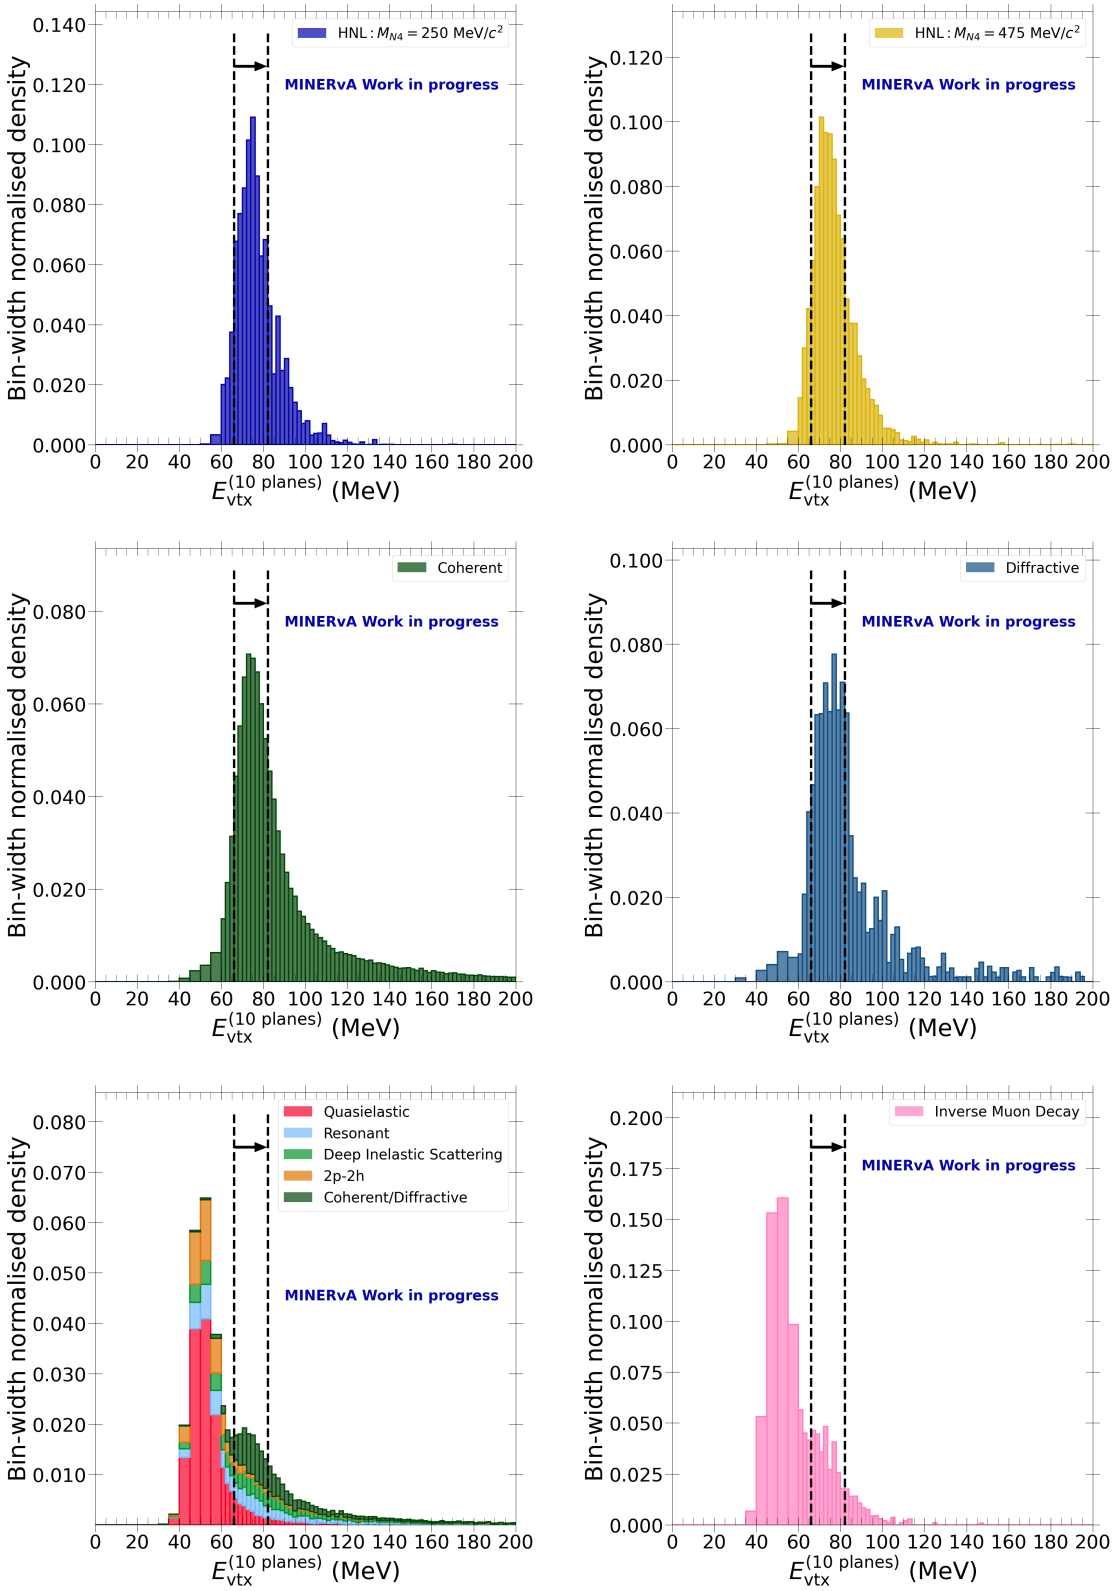


Figure 7.20: Distributions of $E_{\text{vtx}}^{(10\ \text{planes})}$, FHC. From top left: HNL ($M_{\text{N4}} = 250\ \text{MeV}/c^2$), HNL ($M_{\text{N4}} = 475\ \text{MeV}/c^2$), CC Coherent, CC Diffractive, General, IMD.

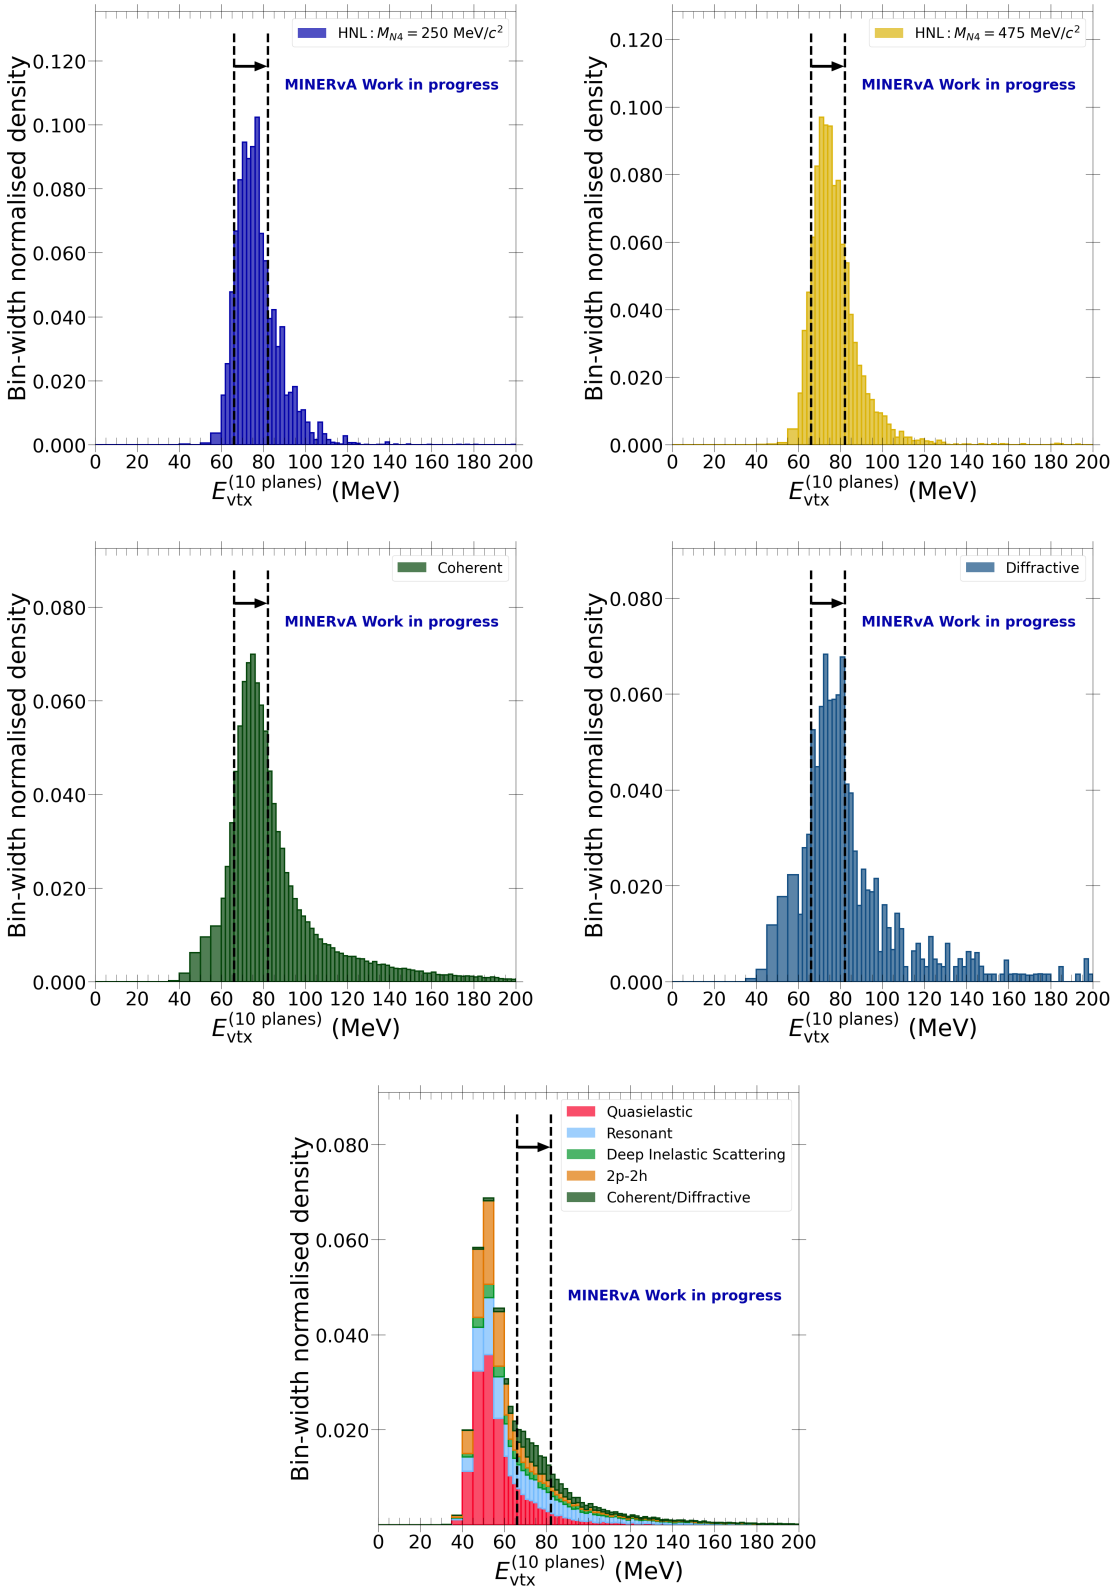


Figure 7.21: Distributions of $E_{\text{vtx}}^{(10\ \text{planes})}$, RHC. From top left: HNL ($M_{\text{N4}} = 250\ \text{MeV}/c^2$), HNL ($M_{\text{N4}} = 475\ \text{MeV}/c^2$), CC Coherent, CC Diffractive, General.

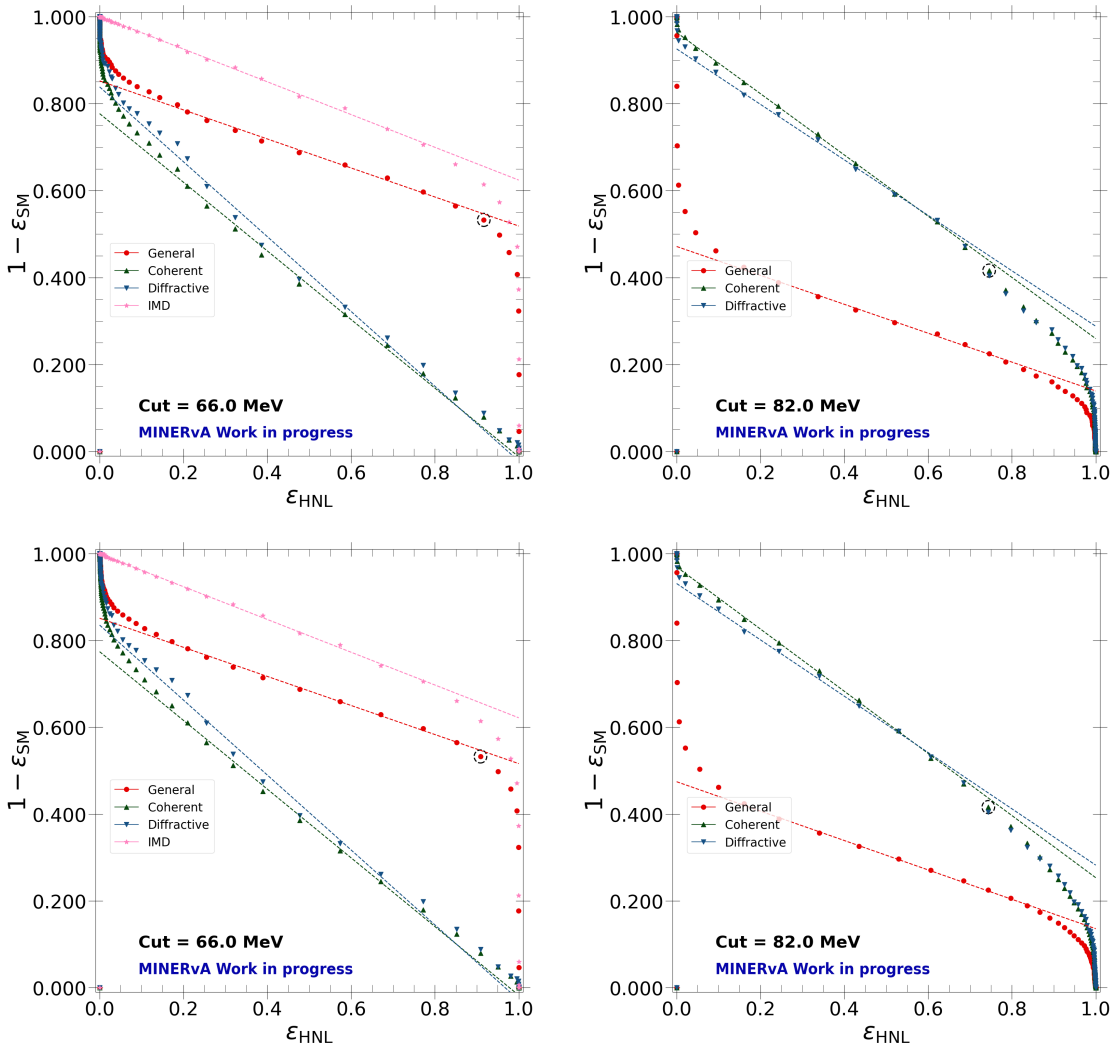


Figure 7.22: Background exclusion curves as a function of HNL signal efficiency, for various choices of the $E_{\text{vtx}}^{(10\text{ planes})}$ cut value represented as individual points. Top panels: $M_{\text{N4}} = 250$ MeV/c^2, bottom panels: $M_{\text{N4}} = 475$ MeV/c^2, left: FHC, right: RHC. RHC General drives the upper cut $E_{\text{vtx}}^{(10\text{ planes})} \leq 82$ MeV, and FHC CC Coherent drives the lower cut $E_{\text{vtx}}^{(10\text{ planes})} \geq 66$ MeV. See text for details.

by the first two selection cuts, with the Diffractive sample cut more often. This happens, because the final-state nucleus is a free proton, hence much more likely to deposit energy along the track in the detector and be cut out. After the first vertex energy cut, both CC Coherent and CC Diffractive samples behave alike, as only events with low four-momentum transfer produce events that pass the cuts. The IMD sample is harshly impacted by all cuts due to its unique topology of a single final-state muon. The HNL, as well as CC Coherent/Diffractive, samples

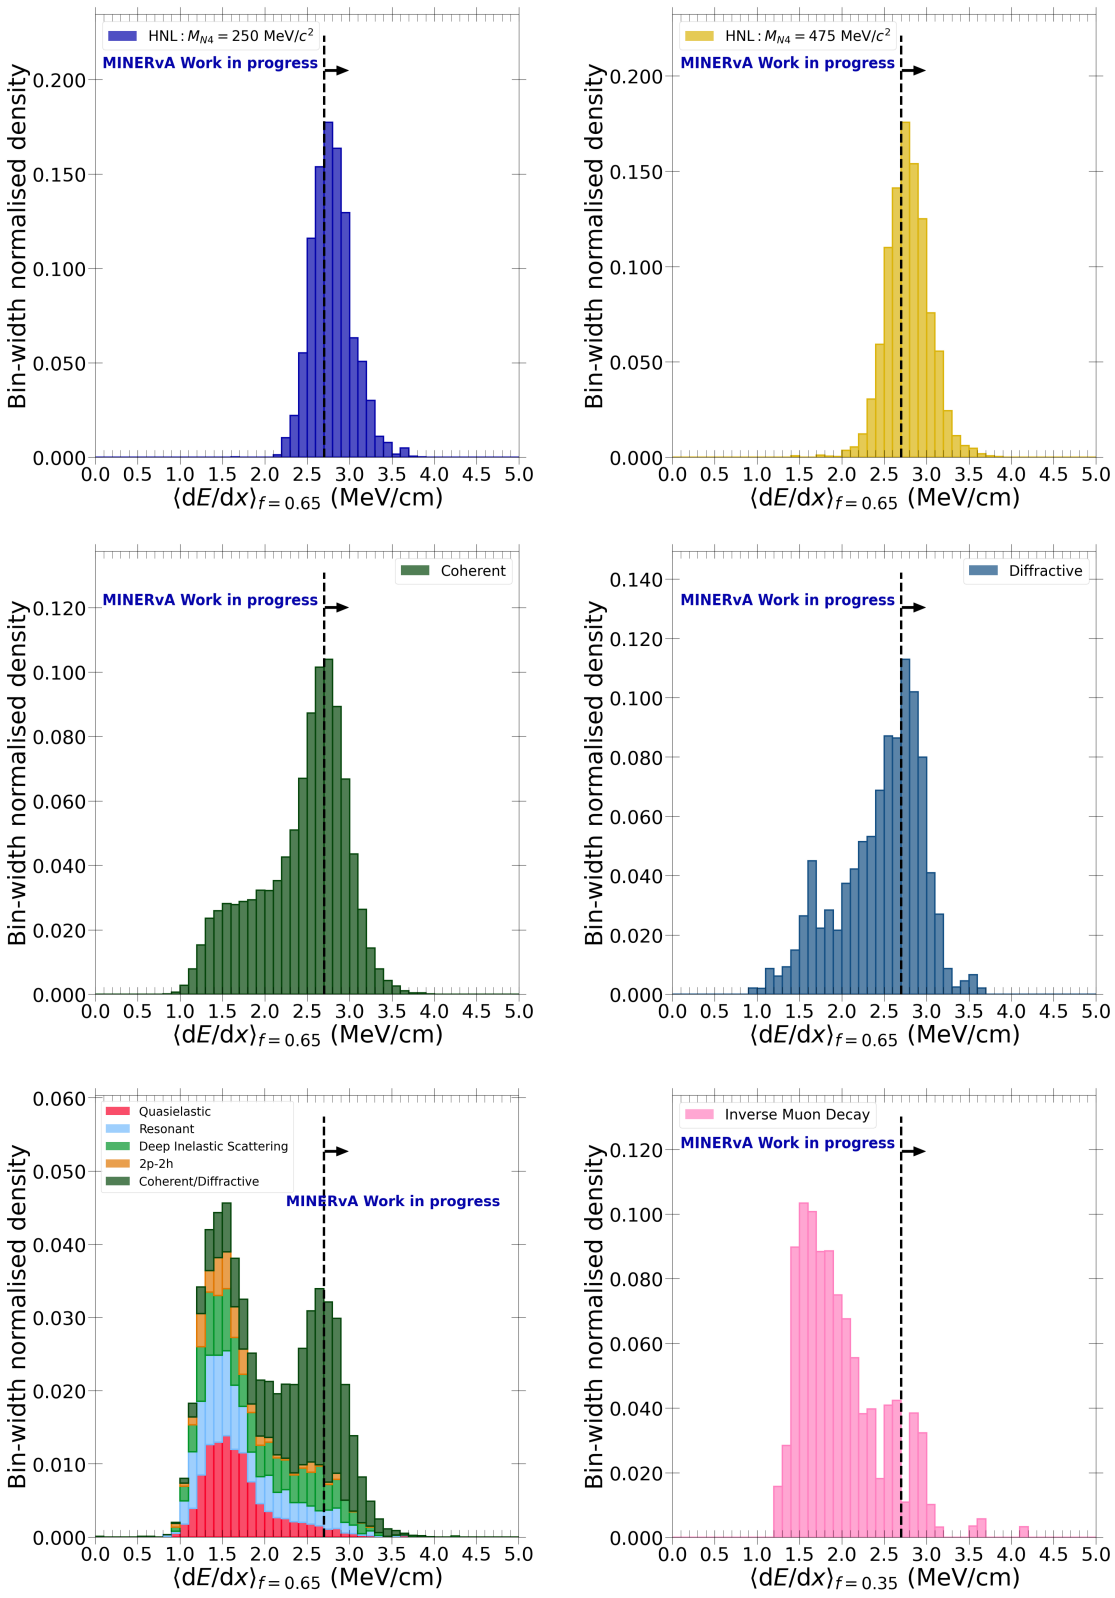


Figure 7.23: Distributions of $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$, FHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General, IMD.

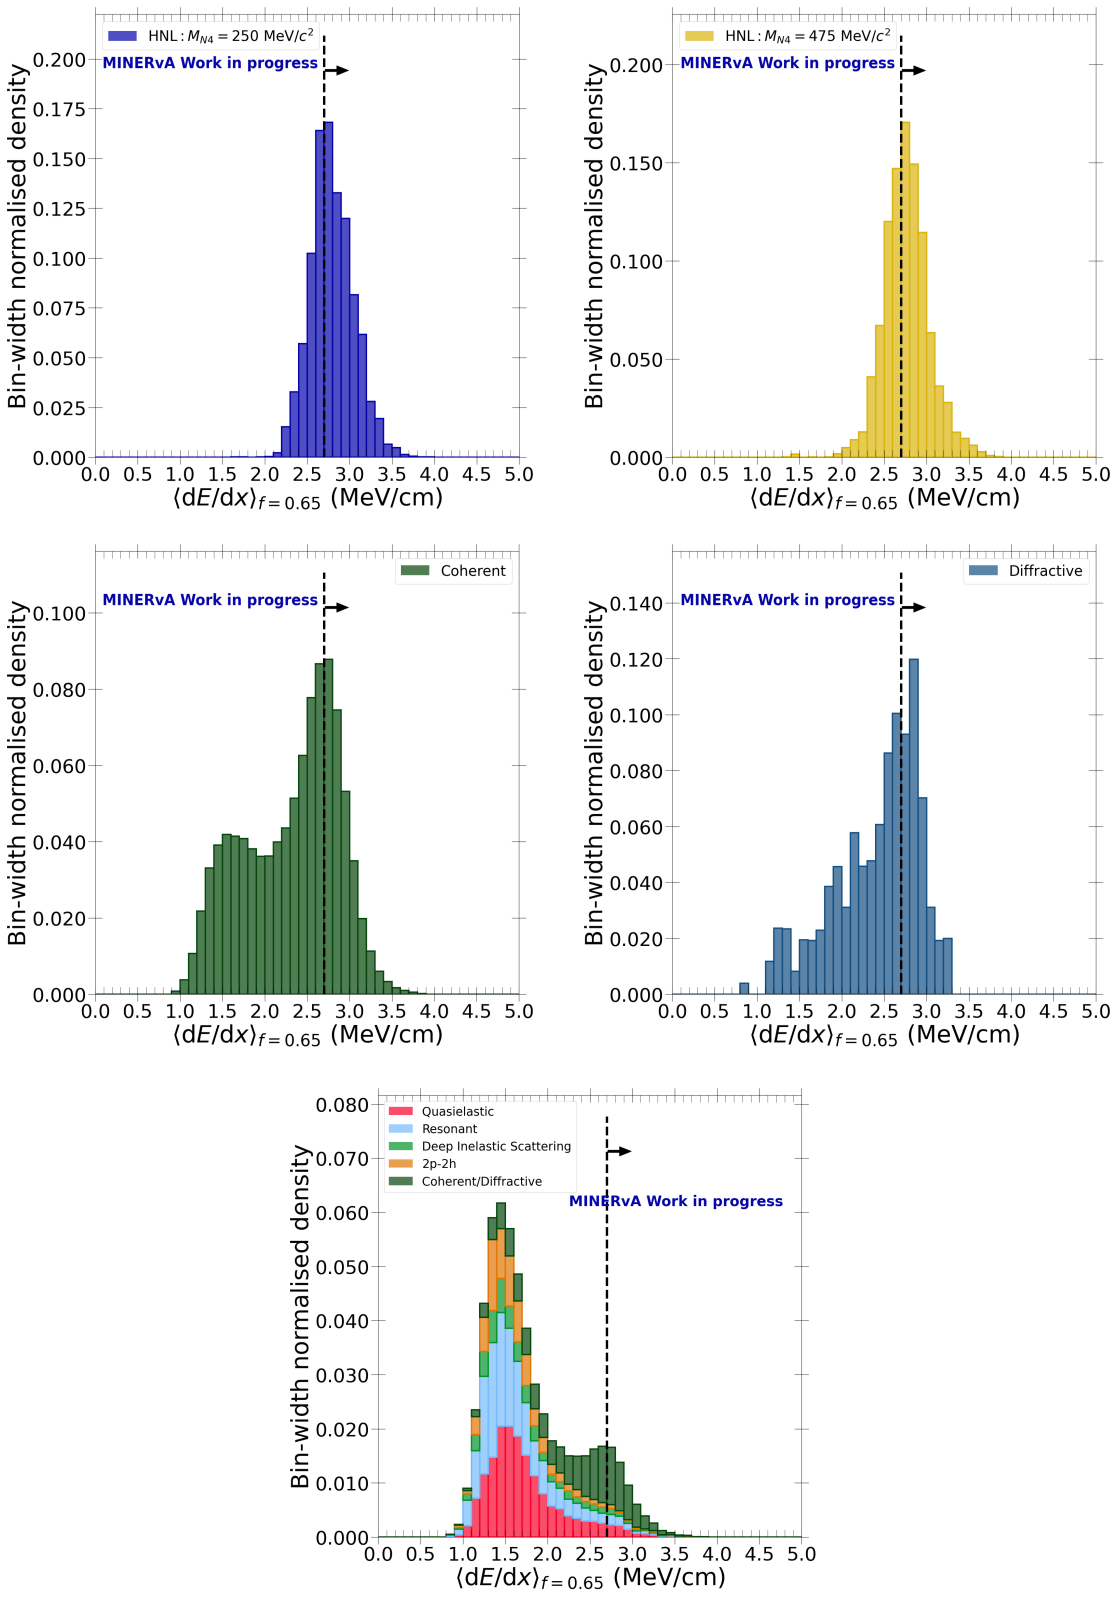


Figure 7.24: Distributions of $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$, RHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General.

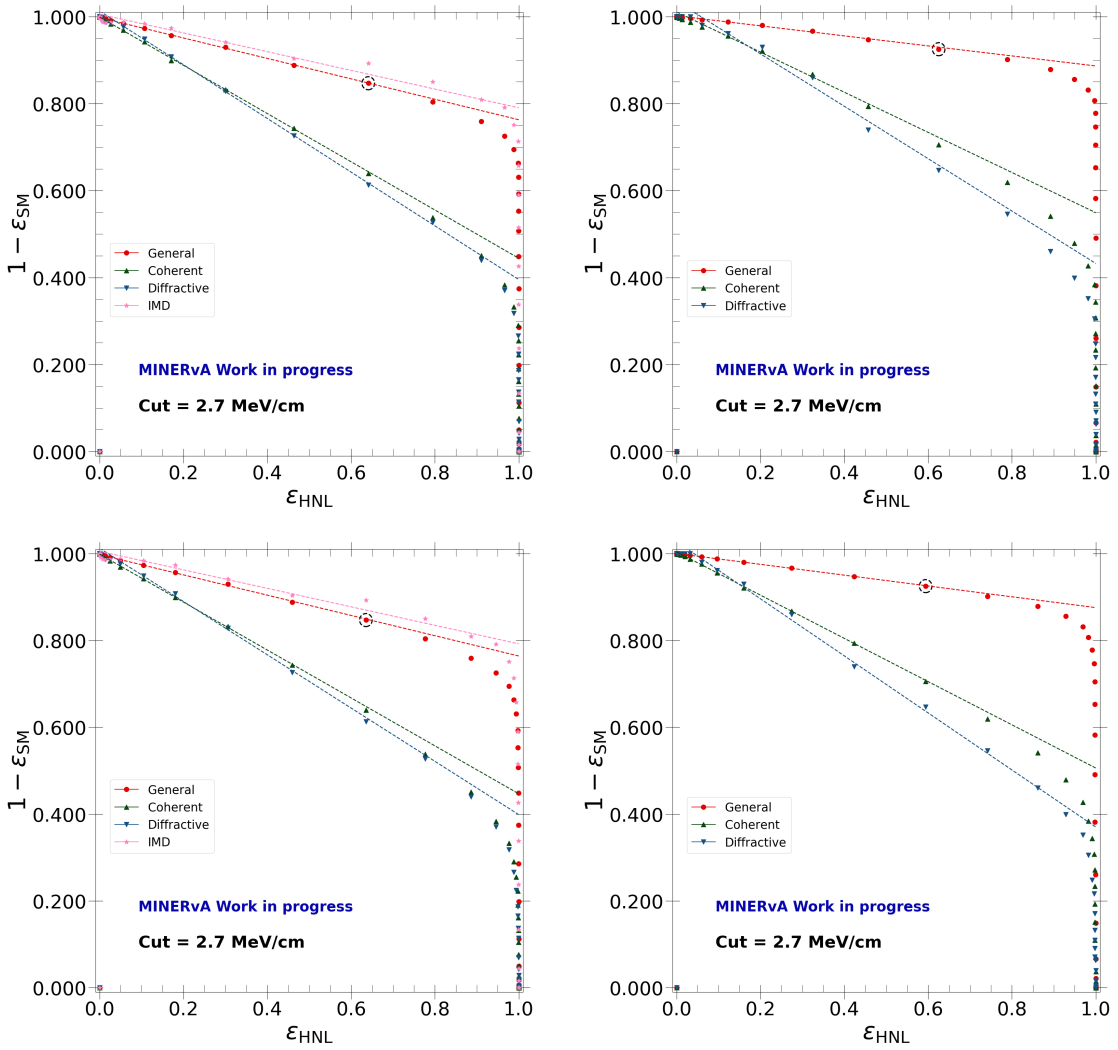


Figure 7.25: Background exclusion curves as a function of HNL signal efficiency, for various choices of the $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$ cut value represented as individual points. Top panels: $M_{\mathrm{N4}} = 250$ MeV/c^2, bottom panels: $M_{\mathrm{N4}} = 475$ MeV/c^2, left: FHC, right: RHC. See text for details.

behave similarly in FHC and in RHC, though this is not the case in the data or the MC-General sample: this is due primarily to proton emission in neutrino events causing energy to be deposited close to the vertex.

7.3.4 Efficiency calculation

The selection efficiency of each sample is, as shown, a vital quantity to estimate in order to proceed with this analysis. This quantity is defined as the integral of the MC distribution that survives selection cuts, over the integral of the entire

distribution. In practice, it is calculated as the product of the relative efficiencies for each sequential cut $1, 2, \ldots, n$, as

$$\epsilon = \frac{\int_{\text{cut 1}} \mathrm{d}V_1 f(V_1)}{\int_{\text{no cuts}} \mathrm{d}V_1 f(V_1)} \cdot \frac{\int_{\text{cut 2}} \mathrm{d}V_2 f(V_2)}{\int_{\text{cut 1}} \mathrm{d}V_2 f(V_2)} \cdot \ldots \cdot \frac{\int_{\text{cut n}} \mathrm{d}V_n f(V_n)}{\int_{\text{cut } n-1} \mathrm{d}V_n f(V_n)}. \tag{7.5}$$

This prescription also allows for a conservative estimate of effects of systematic uncertainties causing migration of events into and out of cuts. With each cut, a histogram of events including the statistical and systematic uncertainties is produced, and one can also define the errors on the relative efficiency (replacing the integrals in Eq. 7.5 with sums over all the bins in each variable and the integrands with the bin contents $b_i^{(n)}$ and their errors $\delta b_i^{(n)}$)

$$\epsilon^{\pm}_{\text{cut } n} = \frac{\sum_{\text{bins } i} \left(b_i^{(n)} \pm \delta b_i^{(n)} \right)}{\sum_{\text{bins } i} \left(b_i^{(n-1)} \mp \delta b_i^{(n-1)} \right)}. \tag{7.6}$$

This then allows one to construct a conservative estimate of the efficiency in the presence of bin migration, as $\epsilon \in \left[\prod_{\text{cuts } j} \epsilon^{-}_{\text{cut } j}, \prod_{\text{cuts } j} \epsilon^{+}_{\text{cut } j} \right]$. In other words, the product of the lower bounds for the relative efficiency is guaranteed to be a lower bound for the total efficiency, and the same is true for the upper bounds.

Figure 7.26 summarises the relative efficiencies of the background for the various cuts, with the correction due to MINOS matching applied. The uncertainty on each point is the full statistical and systematic uncertainty, including the flat MINOS systematic. This may be contrasted with the characteristic HNL signal samples at 250 and 475 MeV/c^2, at FHC and RHC, in Fig. 7.27. Multiplying all the relative efficiencies together, and taking the logarithm, one gets the discimination capability κ for each background sample as in Fig. 7.28. Once again, the MINOS correction has been applied and the asymmetric error bars encompass the statistical and systematic uncertainties.

Finally, one can compare with the HNL signal efficiencies themselves, which are fairly poorly constrained due to the large statistical errors associated with the sample production, in Fig. 7.29.

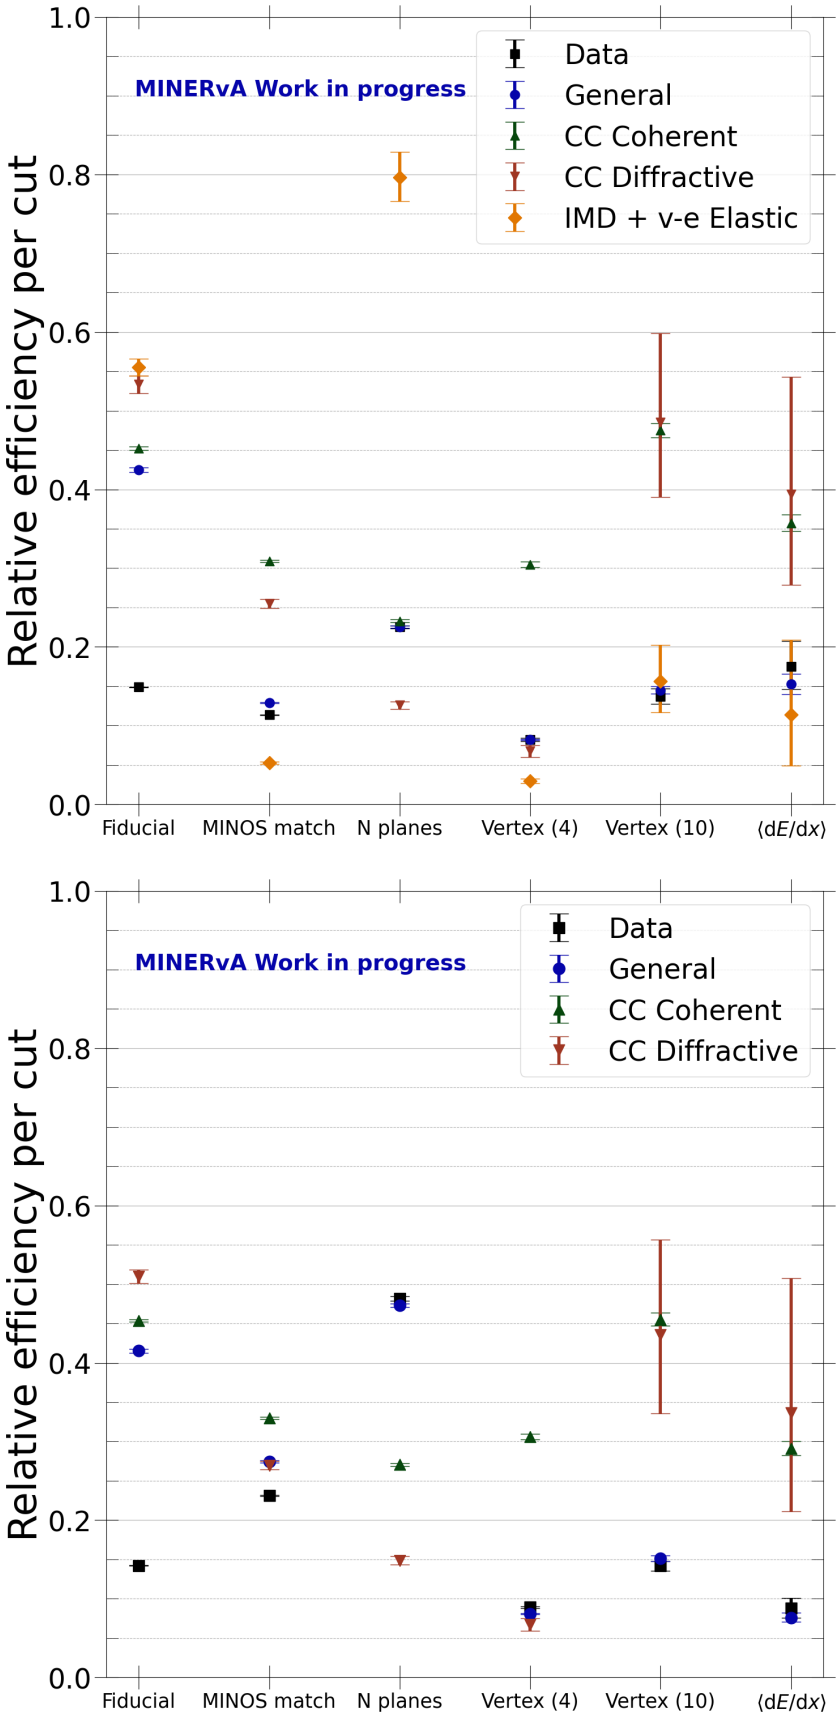


Figure 7.26: Relative efficiencies for the SM backgrounds, FHC (top) and RHC (bottom).

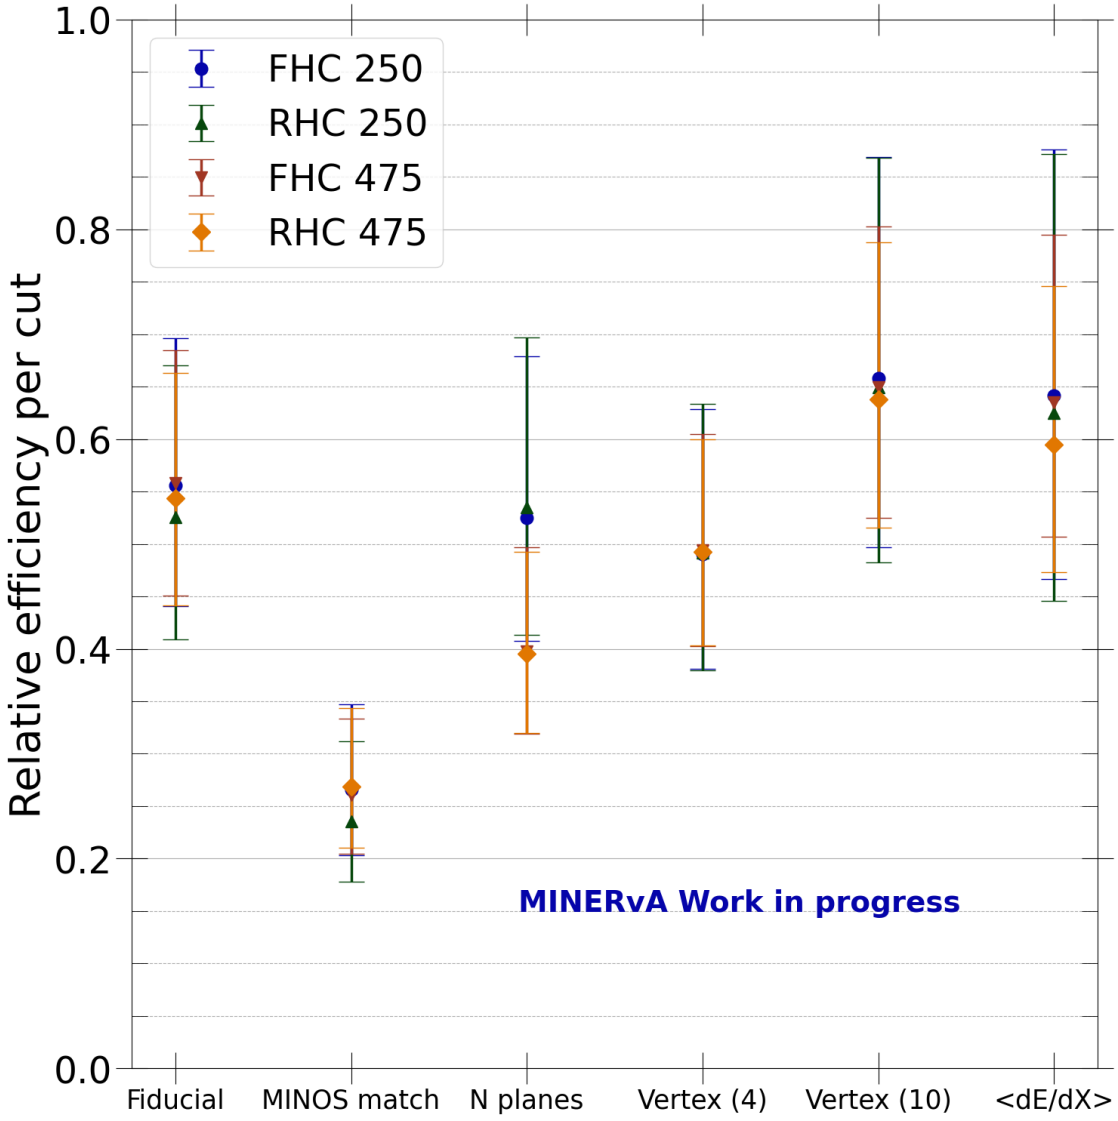


Figure 7.27: Relative efficiencies for the characteristic HNL samples.

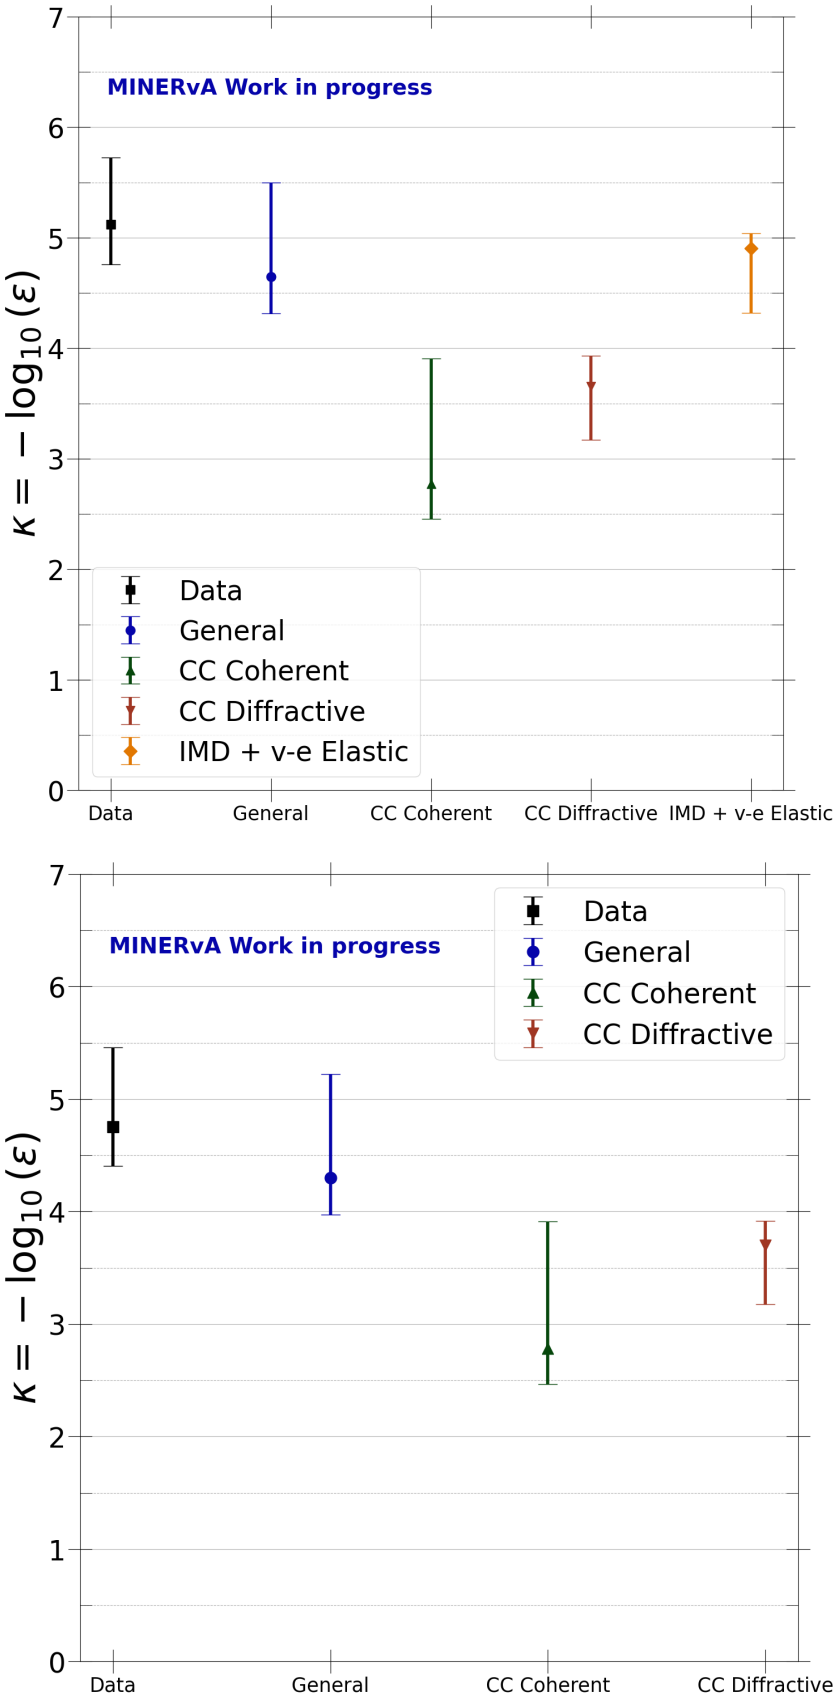


Figure 7.28: Discrimination capability for the SM backgrounds, FHC (top) and RHC (bottom).

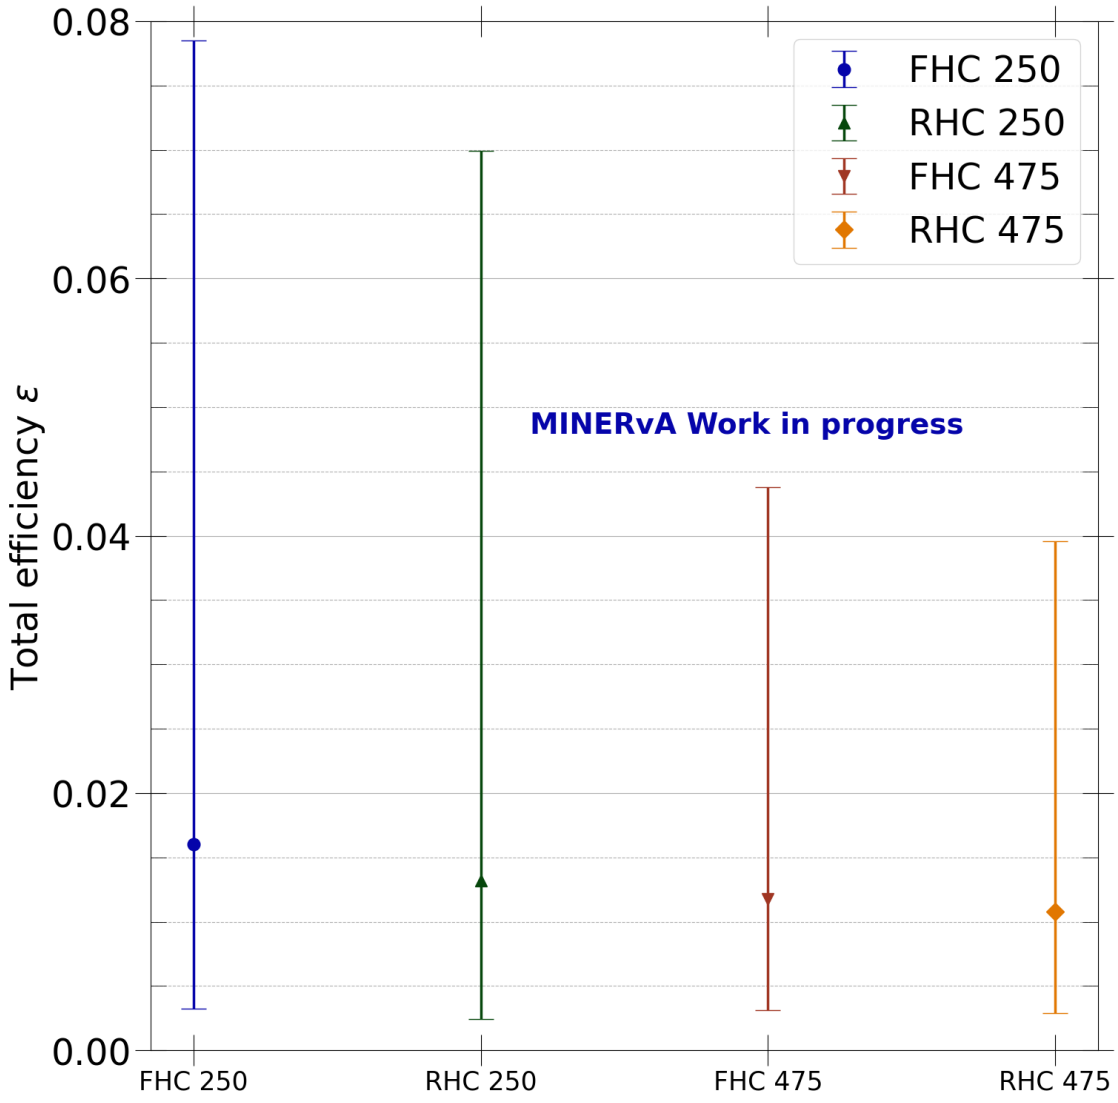


Figure 7.29: Total efficiency for the characteristic HNL samples.

8 Background estimation and systematic uncertainties

Contents

8.1 Backgrounds to HNL decay 155
8.1.1 Coherent pion production in tracker 155
8.2 Simulation of Standard Model background sources . . 157
8.2.1 Mainline GENIE simulation 157
8.2.2 MINERνA tunes 159
8.3 Systematic evaluation 160
8.3.1 Many-universe method 160
Error matrix calculation 169
8.3.2 Systematic uncertainties 169
8.4 Background tuning 171
8.4.1 Sideband fitting 171
8.4.2 Tuning the background 177
8.4.3 Fit results 180
Unconstrained $N = 5$ fit 182
Reduced fit: $N = 1$ 189
Reduced fit: $N = 2$ 190
8.4.4 Conclusions 193

The $N_4 \to \mu^{\mp}\pi^{\pm}$ one-track signal signature can be mimicked by various Standard Model processes which happen to produce a collimated muon-pion pair. These interactions, although rarely occurring in the region of phase space which survives the selection process, are frequent enough to mimic, and potentially mask, the HNL decay that is being selected. For this reason, suitable signal-poor control regions must be identified to constrain the Monte Carlo prediction for each of the background interaction types. This Chapter will detail the techniques used to accomplish this task.

8.1 Backgrounds to HNL decay

8.1.1 Coherent pion production in tracker

The "irreducible" component of the background is the CC coherent pion production process $\nu_\alpha(\overline{\nu}_\alpha) + A \to l_\alpha^{\mp} + \pi^{\pm} + A$. The kinematics of this interaction are dominated

by forward-going pions, since almost none of the four-momentum transfer from the lepton ends up going to the nucleus, which remains in its ground state. In other words, the neutrino interacts coherently with the entire nucleus. This means that the outgoing pion is typically forward-going [226]. Typically, the coherent cross section is written down as codified by Adler, under the hypothesis of a partially conserved axial current (PCAC) [280]. The PCAC relates the forward scattering amplitude at zero four-momentum transfer, $Q^2 = 0$, to the divergence of the axial current which is estimated from the elastic pion-nucleus cross section [281]. The neutrino interacts with hadronic strongly charged fluctuations in the presence of the nuclear field, which all "act like" a pion [282] and scatter by exchanging a pomeron, leading to a prediction that the cross section is peaked at low values of the four-momentum transfer to the nucleus [283].

The precise way in which one must extrapolate away from $Q^2 = 0$ is not known. The most popular model is the Rein-Sehgal model [283], but other PCAC-compatible approaches have been derived, e.g. [284–286]. Another class of approaches, called the *microscopic models* of coherent pion production, attempt to describe the coherent process as the coherent sum of resonant pion production off of one nucleon, for all the nucleons in a nucleus; see [280] for an overview. The CC coherent pion production process on various nuclear targets and beam energies has been the subject of multiple measurements by MINERνA (see for example [272, 273, 287]), but it is still currently unknown exactly how the coherent cross section behaves with neutrino energy. The MINERνA baseline simulation uses the Rein-Sehgal model for the coherent cross section, with a correction applied for lepton mass effects [264]; further tuning will have to be specific to the region of phase space probed in this analysis. Finally, there are also backgrounds not from neutrino interactions in the tracker at all; cosmic rays and through-going rock muons. These are removed at the reconstruction stage by running specialised algorithms that identify particles of non-detector origin, and remove them from a physics event in a readout gate (by excising the clusters associated with the energy deposits from those signatures), without discarding the physics event.

How might one separate the coherent background from true HNL decays? In the previous Chapter, a selection for two MIPs was performed with the aim of separating out as much of the background as possible. As seen in Figs. 7.23 and 7.24, there is a noticeable "second peak" at high $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$. This causes complications, as it is difficult to disentangle the coherent mimic from the true HNL signal sample, especially given the mass of the HNL (and therefore its precise kinematics) is unknown.

8.2 Simulation of Standard Model background sources

8.2.1 Mainline `GENIE` simulation

As discussed, the event generator used to simulate neutrino interactions is `GENIE 2.12.6`. The MINERνA "nominal" flux energy spectrum encoded in `dk2nu` (see Chatper 3) is used as an input to the generator, along with a detailed description of the detector geometry, to generate events in the full detector volume according to the appropriate mixture of nuclei in each subdetector. Details of the "baseline" cross-section model for the simulated interactions can be found in [242], only the essential components of the simulation will be repeated here.

Quasielastic (QEL) scattering is simulated using the formalism of Llewellyn Smith [288], with vector form factors modelled using the BBBA05 model [289], and the axial form factor using a dipole form with mass $M_A = 0.99$ GeV/c^2 [214]. Resonant (RES) scattering is simulated using the Rein-Sehgal model [263] with an axial mass $M_A^{\mathrm{RES}} = 1.12$ GeV/c^2, and deep inelastic scattering (DIS) is simulated using the Bodek-Yang prescription [290]. Finally, scattering off of correlated pairs of nucleons in the initial state ("2p2h"/"MEC" scattering) is modelled using the Valencia model [291–293]. The inelastic processes (RES, DIS) become more dominant as neutrino energy increases.

A schematic overview of the primary contributing scattering processes can be seen in Fig. 8.1; in "transition" regions where there is an overlap between processes,

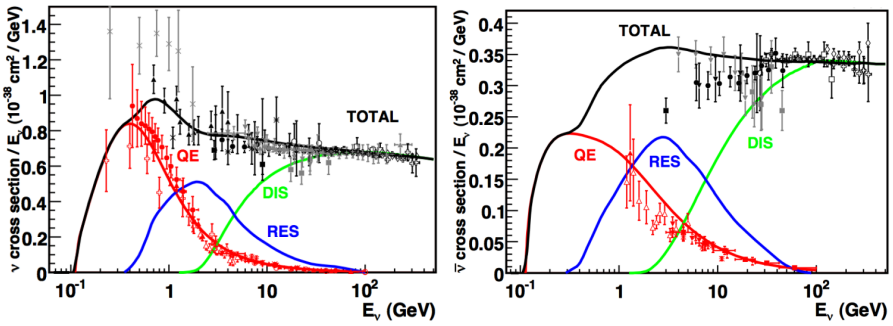


Figure 8.1: Cross-sections per nucleon for neutrino (right) and antineutrino (left) CC interactions, for an isoscalar target, divided by neutrino energy. Data from [262]; colour version from [294].

GENIE adopts the approach of [295, 296] to approximate the total cross-section as

$$\sigma^{\text{tot}} = \sigma^{\text{QEL}} + \sigma^{\text{RES}} + \sigma^{\text{DIS}} + \dots$$

The GENIE nuclear model is the Smith-Moniz relativistic Fermi gas [297] which includes Fermi motion (at a peak momentum of about 220 MeV/c) and Pauli blocking, with an added Bodek-Ritchie high-Q^2 tail [298], to incorporate short-range correlations between nucleons. A nuclear modification factor accounts for structure functions of a nucleon in a nuclear environment, such as shadowing or anti-shadowing. A significant complication for neutrino interactions is encountered in so-called "final-state interactions" (FSI), which describe the collective effects of the propagation of the products of the "intrinsic" neutrino interaction on a nucleon through the nuclear remnant. Effectively, a detector may only observe *asymptotic* final states, that is particles that have exited the nucleus, however good its resolution; but it has no way to know if, say, a neutrino interaction which led to one muon, one positive pion, one proton and one recoil remnant being detected truly corresponded to those particles being produced from the interaction. A true interaction that produces that topology could be resonant $\nu_\mu + \text{p} \to \mu^- + \Delta^{++} \to \mu^- + \text{p} + \pi^+$; however, with FSI, reactions such as $\nu_\mu + \text{n} \to \mu^- + \Delta^+ \to \mu^- + \text{p} + \pi^0$ (intrinsic interaction), $\pi^0 + \text{A} \to \pi^+ + \text{A}'$ (FSI) become possible. This kind of FSI is known as a "charge exchange" interaction, where the pion exchanges its charge with one of the nucleons in the nuclear remnant. It is typical of the kind of "fates" the

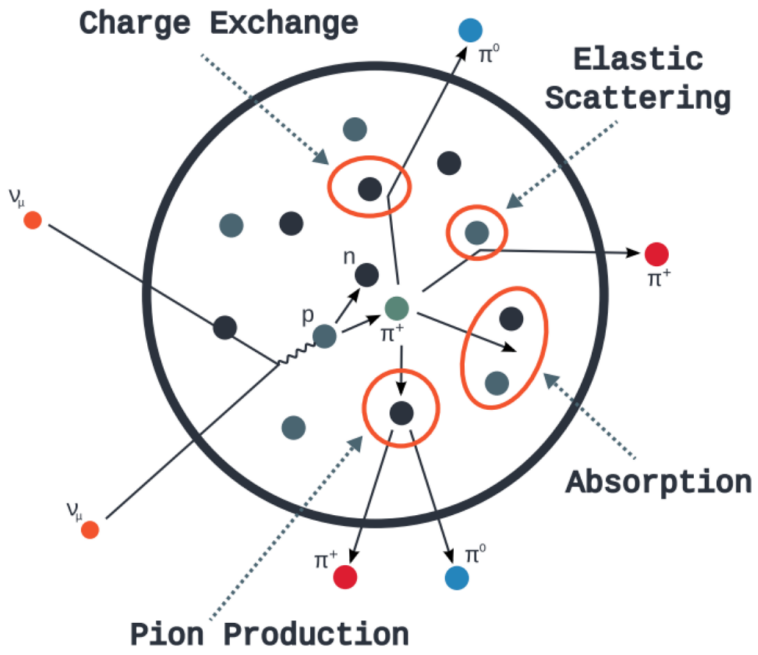


Figure 8.2: Final-state interactions (FSI) for products of "intrinsic" neutrino scattering off a nucleon in a nuclear environment. Note the modification of the interaction topology in the cases of charge exchange or pion production/absorption. Figure from Tomasz Golan.

products of the intrinsic neutrino-nucleus interaction might experience; see Fig. 8.2. If the neutron from the charge exchange does not have sufficient energy to escape the nuclear remnant, then there is no way to distinguish whether an event with detected $\mu^- + \mathrm{p} + \pi^+$ corresponds to the true $\nu_\mu + \mathrm{p} \to \mu^- + \Delta^{++} \to \mu^- + \mathrm{p} + \pi^+$ interaction. Figure 8.2 shows an overview of possible "fates" that can occur to final-state particles from intrinsic neutrino interactions (see Chapter 6 for an explanation of final-state interactions).

8.2.2 MINERνA tunes

The baseline `GENIE` simulation is modified to better describe MINERνA data. There are a series of "MINERνA tunes", generally made to optimise the data/MC agreement in the region of phase space that each analysis is probing. Because this analysis is exploring a particular region of the pion phase space, the minimal "MnvTunev1" modification is chosen, which includes the following modifications:

1. The QEL cross section is modified as a function of energy and three-momentum transfer, based on the random phase approximation, as identified in [219, 299].

2. The 2p2h cross section is increased for specific regions of energy and three-momentum transfer to match MINERνA Low-Energy data [219].

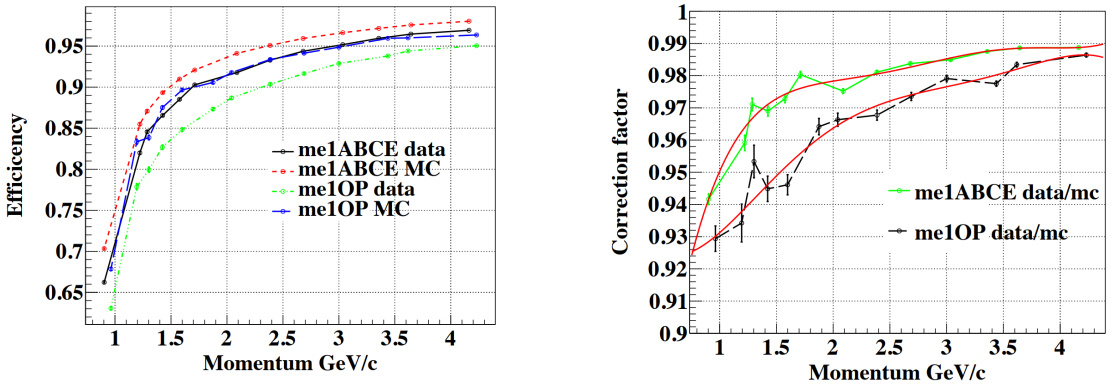
3. The non-resonant pion production is decreased by 43%, along with the corresponding uncertainty [300].

Apart from these modifications, two further alterations are made to the MC. First, an overestimation in the tracking efficiency of MINOS-matched muons in MINERνAwas noticed in the MC, compared to data, attributed to the MINOS reconstruction of pile-up events. Fig. 8.3 (a) shows this overestimation; it is a function of muon momentum, and lower efficiency at low p_μ. indicates the importance of multiple scattering of muons inside MINERνAin lowering the tracking efficiency. The effect of pile-up is also evidenced by the lowering of the efficiency in higher-intensity runs (ME1O,P) compared to the lower-intensity ones (ME1A,B,C,E). A requisite correction factor, (b), is estimated from the "low-intensity" and "high-intensity" run periods of MINERνA- see Fig. 6.1, and used to estimate the correction as required per run period applied. In addition, for all the SM backgrounds, a further constraint from the flux prediction, as discussed in Chapter 3, is applied; this significantly reduces the associated flux uncertainties in background events. It should be noted that, due to the nature of the selection cuts in this thesis, the region of phase space selected (high E_π, low $\theta_{\mu\pi}$) corresponds to a region poorly constrained by prior data, as it is not a region usually studied due to the very small likelihood that events from these processes will be in that region of phase space (cf. Fig. 7.6). This point will be revisited in Section 8.4.2.

8.3 Systematic evaluation

8.3.1 Many-universe method

For this analysis, the MINERvA Analysis Toolkit (MAT) [209] is used to propagate systematic uncertainties through the signal and the background. It is a tool to "wrap



(a) (b)

Figure 8.3: (a): Data/MC tracking efficiencies for the "low-intensity" (black and red) and "high-intensity" (green and blue) run periods in MINOS, as a function of muon momentum. (b): Correction curves for MINOS tracking efficiency, "low-intensity" (green points) and "high-intensity" (black points).

up" systematic uncertainties of the *same physics kind* into *error bands*, exploiting the "many-universe" method to populate histograms with consistent choices of MC parameter variations, or "systematic universes", and construct from many random throws of correlated parameters an uncertainty, for histograms of some suitably binned observable. For an example of how error bands look, the reader can refer to Section 8.3.2, especially Figs. 8.10, 8.11, 8.12; the various coloured histograms each correspond to one error band. Universes from the simultaneous variation of the same physics parameters are collected into the same error band b, and a single systematic uncertainty $\sigma_{i;b}$ is derived in each bin i (to borrow notation from the one-dimensional case; this generalises straightforwardly for two or three dimensions of variables). Henceforth, the terms "error bands" and "uncertainty sources" will be used interchangeably. For each i, for a collection of error bands $\{b\}$, the total systematic uncertainty σ_i^{syst} is obtained as the summation-in-quadrature of the individual systematics, as

$$\sigma_i^{\text{syst}} = \left(\sigma_{i;1}^2 + \cdots + \sigma_{i;N_b}\right)^2 . \tag{8.1}$$

The full uncertainty is then given as the summation in quadrature of the statistical and systematic uncertainties,

$$\sigma_i = \sqrt{(\sigma_i^{\text{stat}})^2 + (\sigma_i^{\text{syst}})^2}. \tag{8.2}$$

Mechanically, this is automatically handled by the use of the `MnvH*D` class, which extends the `TH*D` class of ROOT [249] to include collections of systematic variations $\{\kappa_b\}_{b\in\{b\}}$, for a collection of error bands $\{b\}$. The primary interface for obtaining errors from the `MnvH*D` is the `GetCVHistoWithError()` method, which returns a single `TH*D` histogram whose bins are populated with the counts from the "central value" (CV) universe, corresponding to the nominal MC, and whose errors in each bin have been calculated according to Eqs. 8.1, 8.2. For the rest of this thesis, the three variables of interest are defined to be the following:

- Muon track angle with respect to nominal beam direction, $\theta_\mu := \cos^{-1}\left(\hat{\boldsymbol{p}}_\mu \cdot \hat{\boldsymbol{z}}\right)$,
- Muon transverse momentum with respect to BEAM z axis, $p_{\text{T}}^{(\mu)} := \left|\boldsymbol{p}_\mu \times \hat{\boldsymbol{z}}\right| = \left|\boldsymbol{p}_\mu\right| \sin\theta_\mu$,
- Muon longitudinal momentum with respect to BEAM z axis, $p_z^{(\mu)} := \left|\boldsymbol{p}_\mu \cdot \hat{\boldsymbol{z}}\right| = \left|\boldsymbol{p}_\mu\right| \cos\theta_\mu$.

Distributions for these variables in FHC and RHC, for the various samples, can be seen in Figs. 8.4 - 8.9. The binning for these variables was chosen such that running a sideband fit over these variables would be practicable. In order to have a well-behaved Pearson's χ^2 for the fit, it is important to have enough entries in each bin to make the assumptions of the χ^2 test valid (the "grouping is not too fine") [301]; simultaneously, the binning should be fine enough that the shape of the MC is retained. Therefore, the finest possible binning was chosen in each distribution that satisfies the following empitical rule: not more than 1 in 5 bins contains fewer than 5 entries [302]. The tuning will be explained in more detail in Section 8.4.2. Table 8.1 shows the analysis binning used.

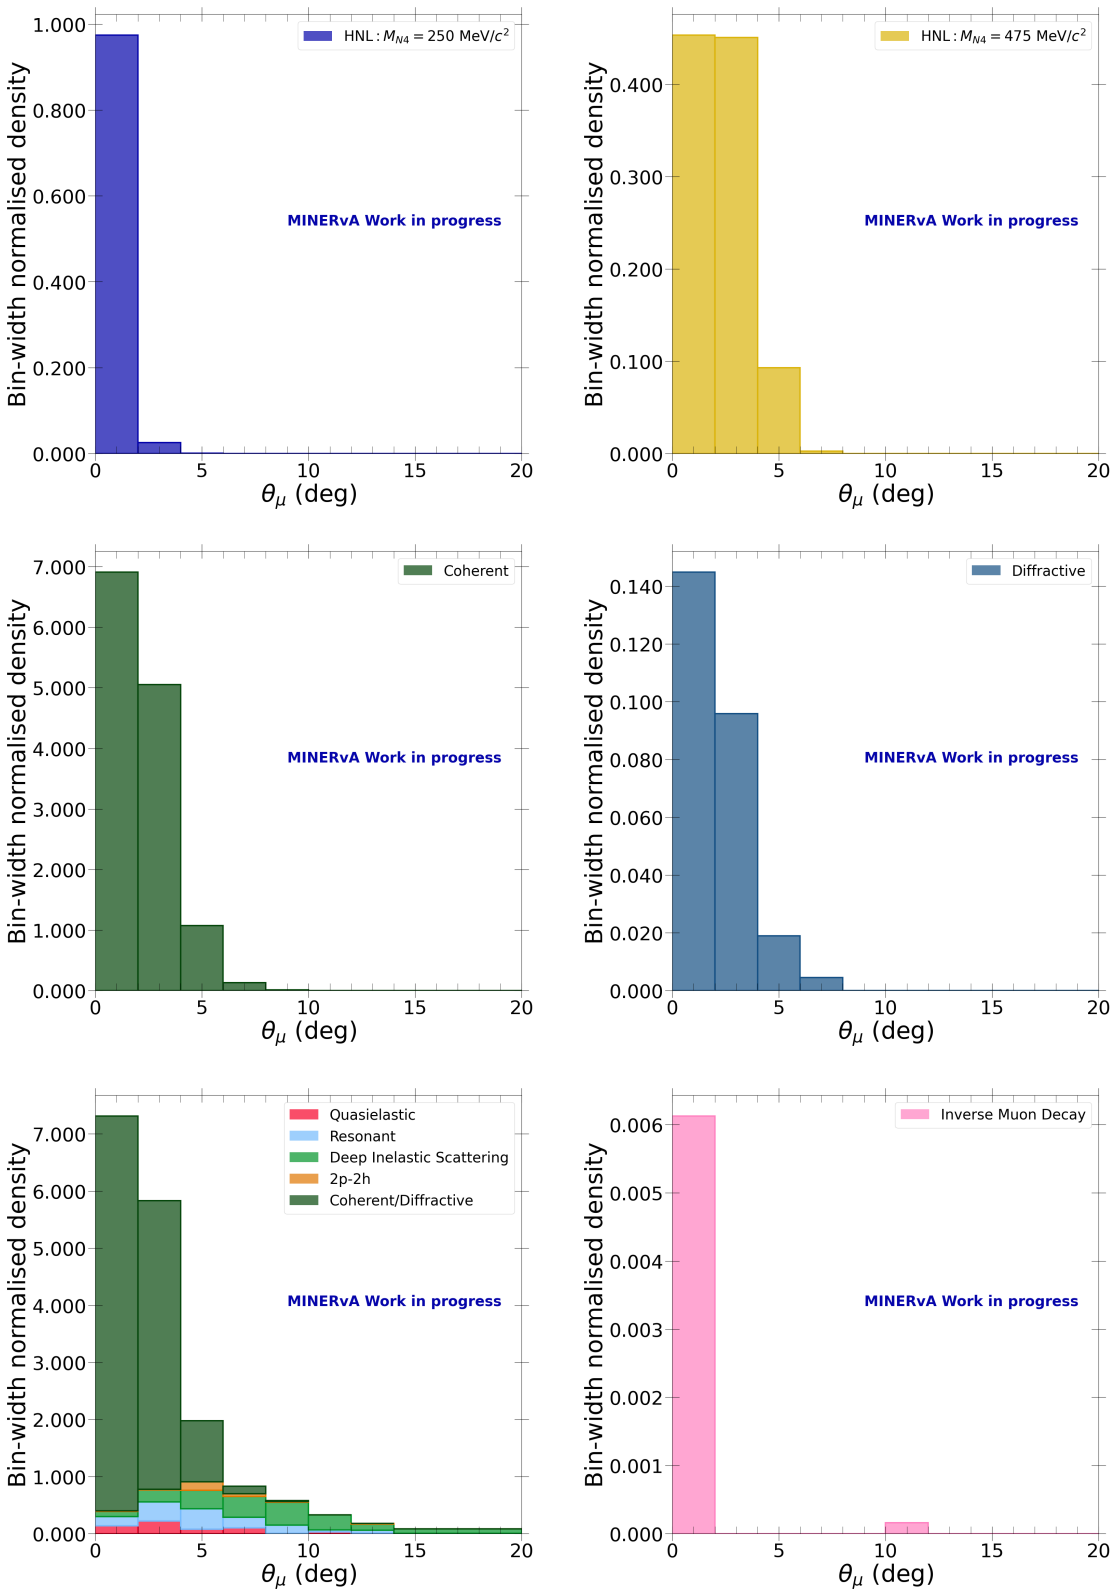


Figure 8.4: Distributions of θ_μ, FHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General, IMD.

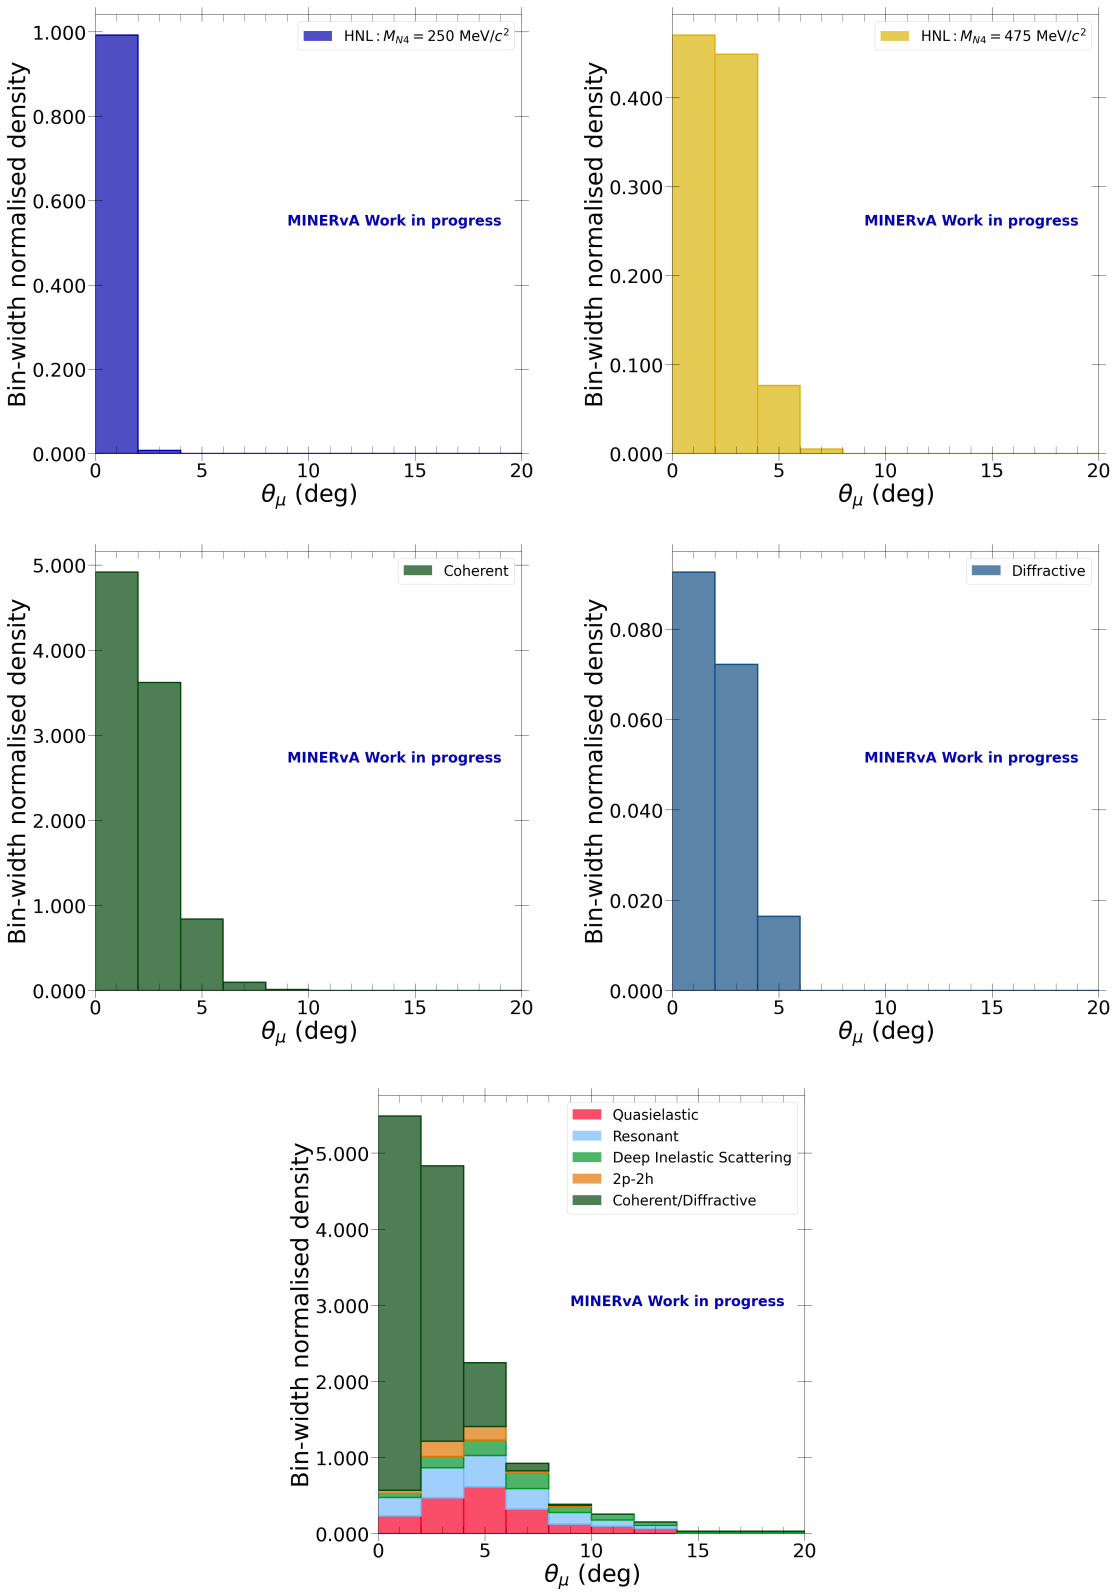


Figure 8.5: Distributions of θ_μ, RHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General.

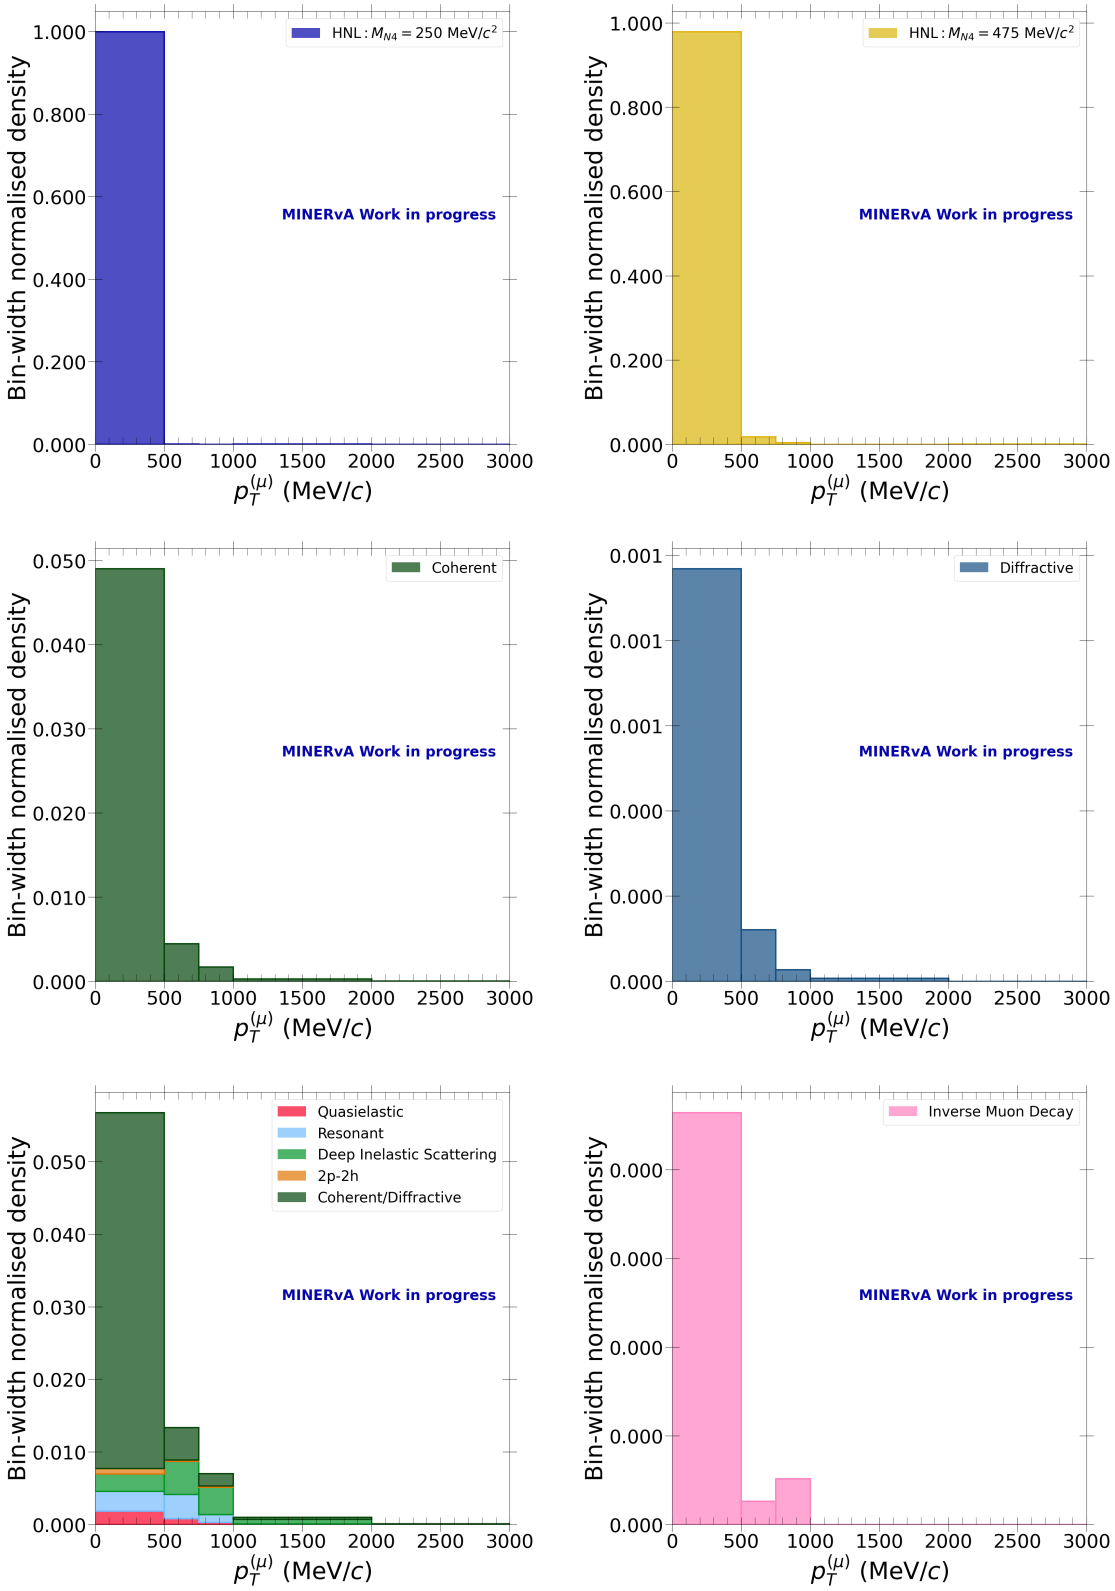


Figure 8.6: Distributions of $p_{\mathrm{T}}^{(\mu)}$, FHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General, IMD.

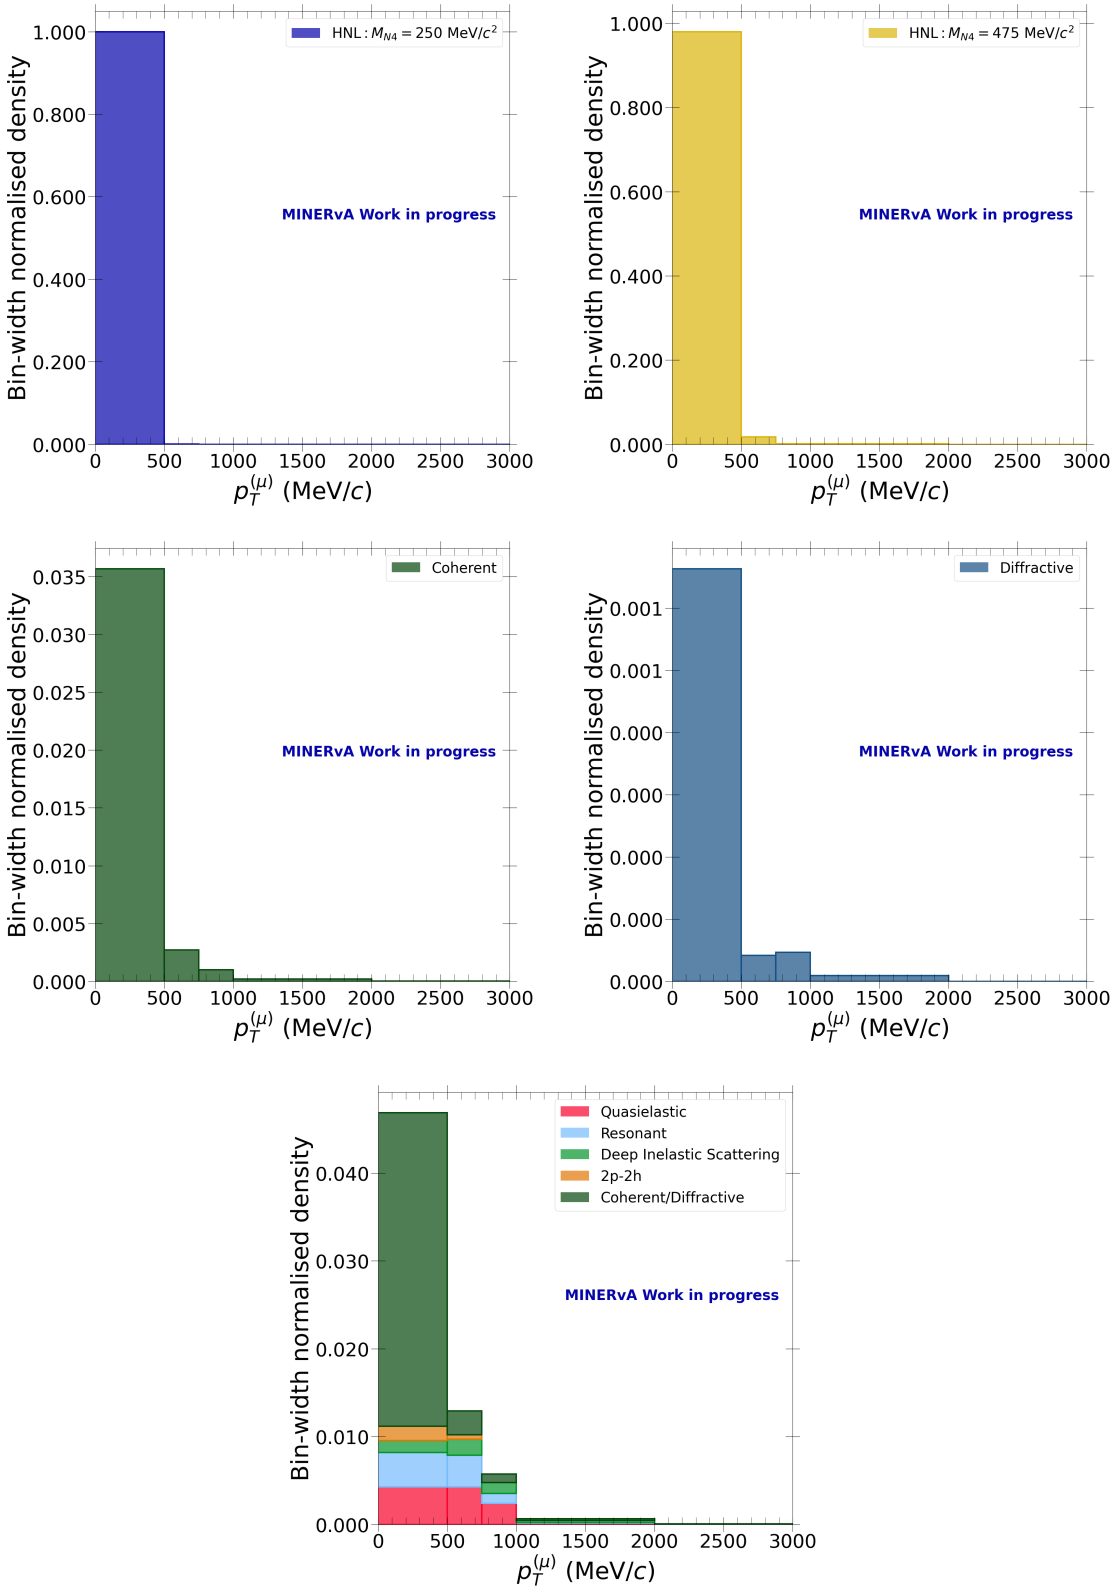


Figure 8.7: Distributions of $p_{\mathrm{T}}^{(\mu)}$, RHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General.

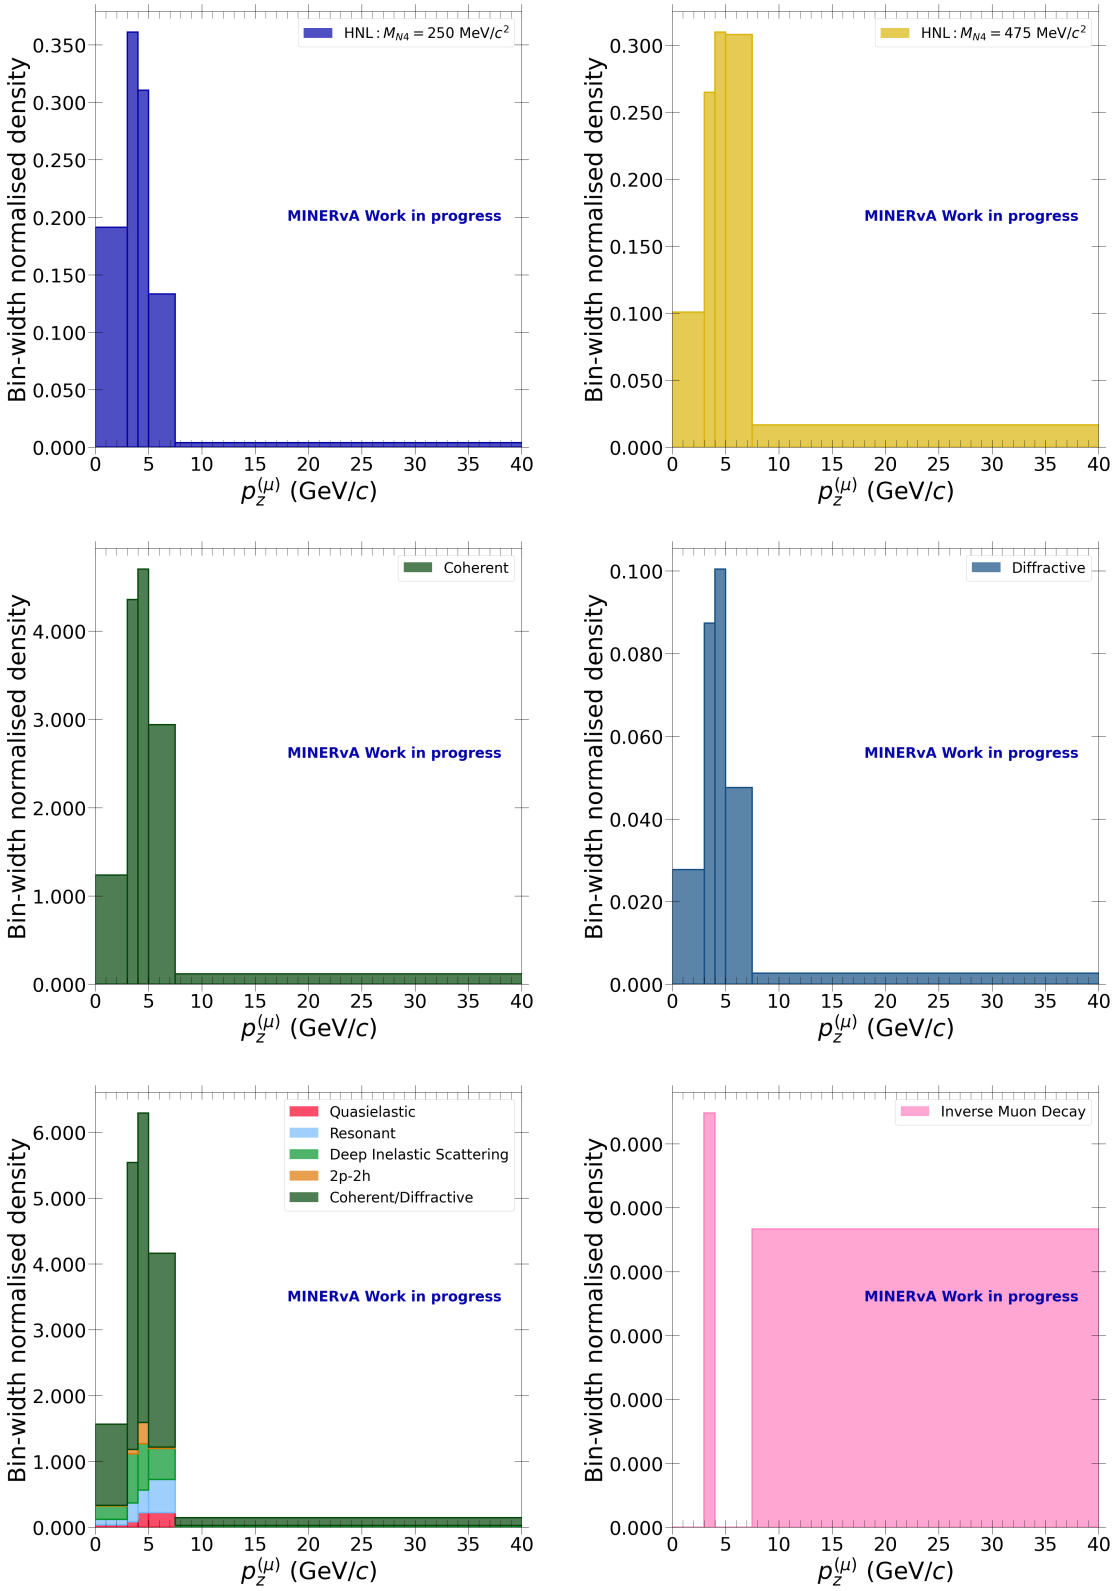


Figure 8.8: Distributions of $p_z^{(\mu)}$, FHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General, IMD.

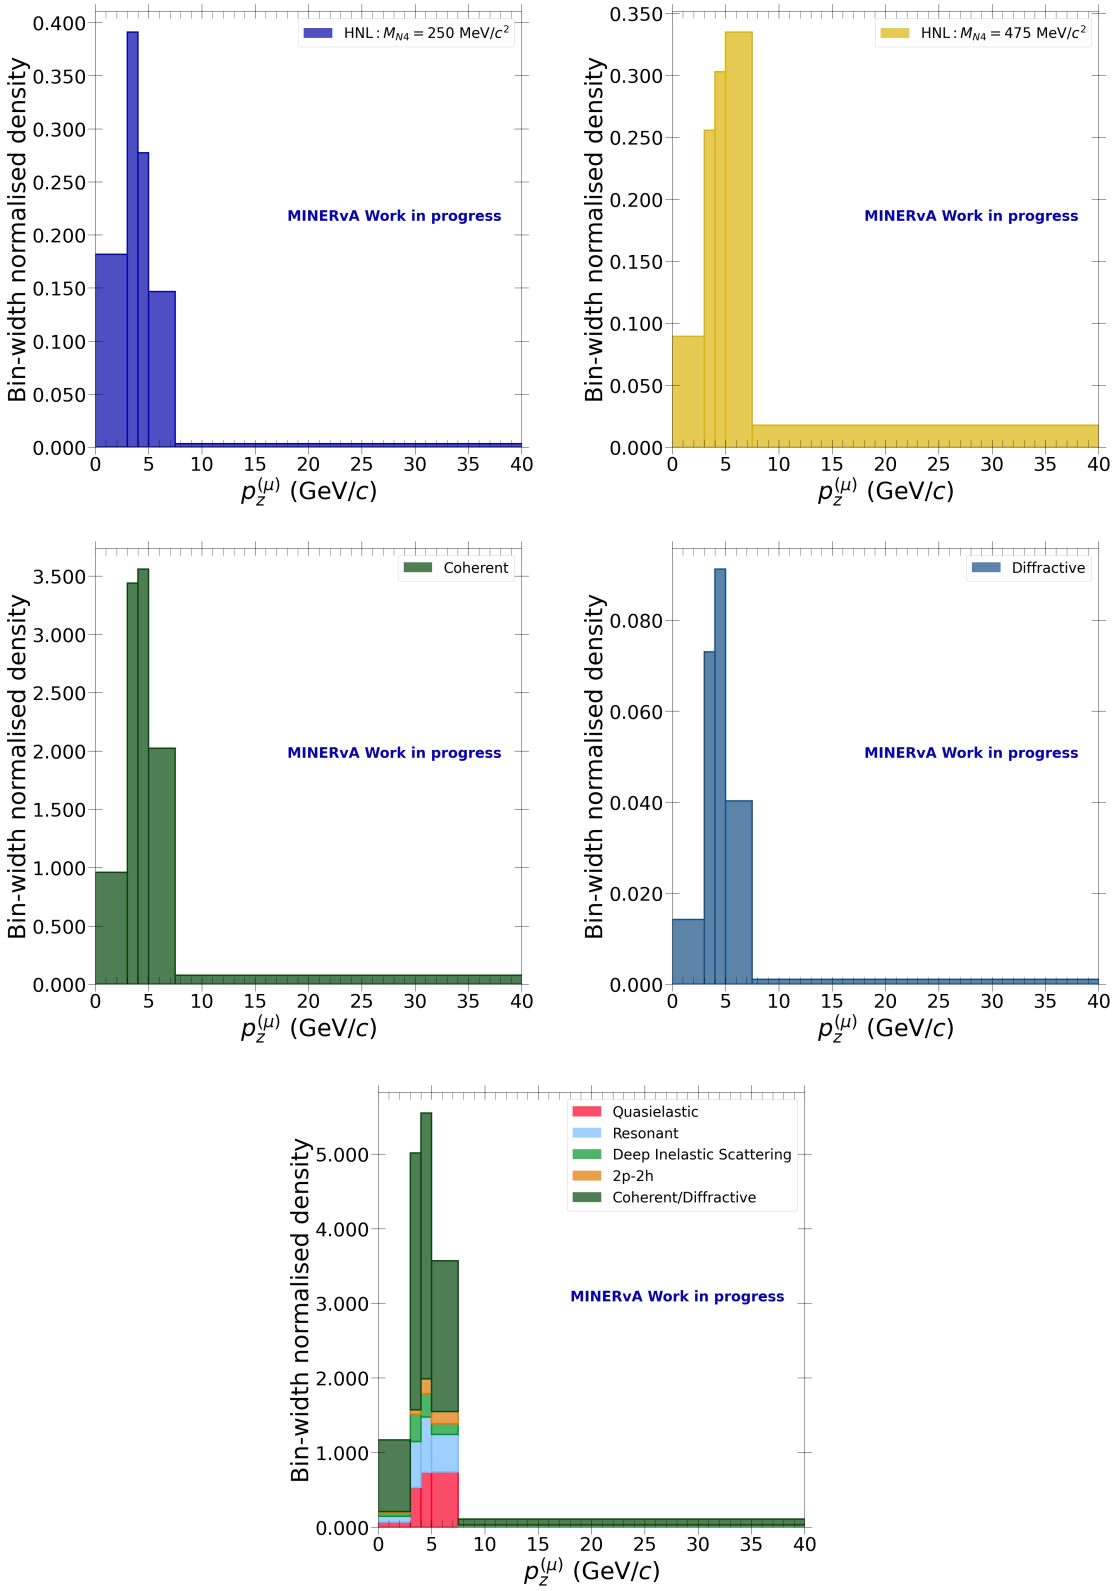


Figure 8.9: Distributions of $p_z^{(\mu)}$, RHC. From top left: HNL ($M_{\mathrm{N4}} = 250$ MeV/c^2), HNL ($M_{\mathrm{N4}} = 475$ MeV/c^2), CC Coherent, CC Diffractive, General.

$\theta_\mu[\circ]$	$[0, 2, 4, \ldots, 14, 20, 25]$
$p_{\mathrm{T}}^{(\mu)}[\mathrm{MeV}/c]$	$[0, 500, 750, 1000, 2000, 3000]$
$p_z^{(\mu)}[\mathrm{GeV}/c]$	$[0, 3, 4, 5, 7.5, 60]$

Table 8.1: Analysis binning choices.

Error matrix calculation

To get the total error σ_i implementing Eq. 8.2, a series of steps is called. This technical subsection may not be of interest to all readers. (In all subsequent notation, the CV universe will be indexed by 0.) Algorithm 7 summarises the main high-level procedure for the calculation.

Data: `MnvH1D` with bins $\{i\}$, error bands $\{b\}$
Result: Error matrix ε_{ij}
$\forall$ error band b obtain an error matrix ε_{ij}^b depending on error band type;
switch *type of* b **do**
  when b not uncorrelated $\Rightarrow$ See Algorithm 8;
  else $\Rightarrow \varepsilon_{ij}^b = \mathrm{diag}\,(\sigma_i^2)$, σ_i the histogram error in each bin i;
end
return $\varepsilon_{ij} = \sum_b \varepsilon_{ij}^b$;

Algorithm 7: High-level overview of error calculation.

8.3.2 Systematic uncertainties

There are a number of systematic uncertainty sources relevant to this analysis. In summary, they may be described as follows:

1. Flux: The flux uncertainty (discussed in Chapter 3 3.4) is evaluated by using 100 universes where the beamline focusing parameters are randomly and simultaneously altered according to their variances using a correlation matrix for all of them.

2. Interaction model: These correspond to parameter uncertainties in the `GENIE` interaction model used for event generation, and as such are relevant only to the background samples. Each parameter is varied by $\pm 1\sigma$ around the central value, creating a corresponding universe.

Data: Universes ℓ binned in $\{i\}$, with counts $\left\{c_i^{(\ell)}\right\}$, and a weight w_ℓ for each universe in the band
Result: Band's error matrix ε_{ij}^b
Calculate $\forall \ell$ the normalised weight $W_\ell := w_\ell / \sum_p w_p$;
if *spread* **then**
  $\forall$ bins i obtain the largest and smallest counts, M_i and m_i;
  switch N *universes* **do**
    when $N = 2 \Rightarrow \epsilon_{ij}^b = (M_i - m_i)(M_j - m_j)$;
    when $2 < N \leq 11 \Rightarrow \epsilon_{ij}^b = (M_i - m_i)(M_j - m_j)/4$;
    else Sort the counts and obtain the "interquartile range", i.e. F_i, f_i such that 1/4 of entries is smaller than f_i and 3/4 of entries is smaller than F_i. $\epsilon_{ij}^b = (F_i - f_i)(F_j - f_j)/1.34896^2$;
  end
end
else
  For each bin i, universe ℓ calculate deviation $\delta_i^{(\ell)} := c_i^{(\ell)} - c_i^{(0)}$;
  $\varepsilon_{ij}^b = \sum_\ell W_\ell \cdot \left(\delta_i^{(\ell)} \delta_j^{(\ell)}\right)$;
end
return ε_{ij}^b;

Algorithm 8: Detailed calculation of error matrix in an error band b of correlated parameters.

3. Detector model: Uncertainties in the detector response function are encapsulated in this band. These may be briefly summarised as:

   - Muon energy and tracking efficiency: The muon energy uncertainty comes from adding in quadrature the uncertainties derived by the energy loss inside MINERνA, as well as the MINOS momentum-by-range and momentum-by-curvature methods. The tracking efficiency uncertainty is got by evaluating the errors on the paramters of the correction function shown in Fig. 8.3(b).
   - Beam direction: Uncertainties in the nominal beam direction are encapsulated here.
   - GEANT modelling of hadrons: For protons, pions, and neutrons, the paramters governing propagation of these particles through matter are varied in the underlying GEANT simulation. For neutrons, the path

length is varied according to neutron kinetic energy, whereas for pions and protons the inelastic cross section is varied according to the nucleus with which they interact, the total path length of the particle, and the particle energy. These variations are determined by comparing model predictions to hadron scattering data.

4. RPA and low-recoil: These parametrise the uncertainty associated with the MINERνA tunes, which modify the nuclear model of GENIE (RPA) and the 2p2h cross section.

The effect of the uncertainties, statistical and systematic, after all the selections may be succinctly seen in Figs. 8.10, 8.11, 8.12. As can be immediately deduced, the dominant uncertainties are statistical, rather than systematic; this is unsurprising, given the extremely stringent selection criteria. It is encouraging that the systematic uncertainties appear to be "well-behaved", in that the detector response is parametrised well enough to make reasonable statements about the total number of counts one sees in the MINERνA dataset; however, it is immediately apparent that a selection this stringent will struggle to accumulate significant enough statistics for a fine-grained search (such as might produce a probability distribution extending to the edges of the kinematic phase space probed by this search). This will become important in the next steps, as large uncertainty in the background component will be a major systematic uncertainty in the search.

8.4 Background tuning

8.4.1 Sideband fitting

To better constrain the background MC, it is necessary to identify a *sideband region* in observable space which is simultaneously signal-poor and background-rich, to use as a control. The rationale behind this procedure is to compare a background-dominated region of phase space with data, and obtain from that data constraints on the background MC. The sideband-constrained MC is then a better description of the

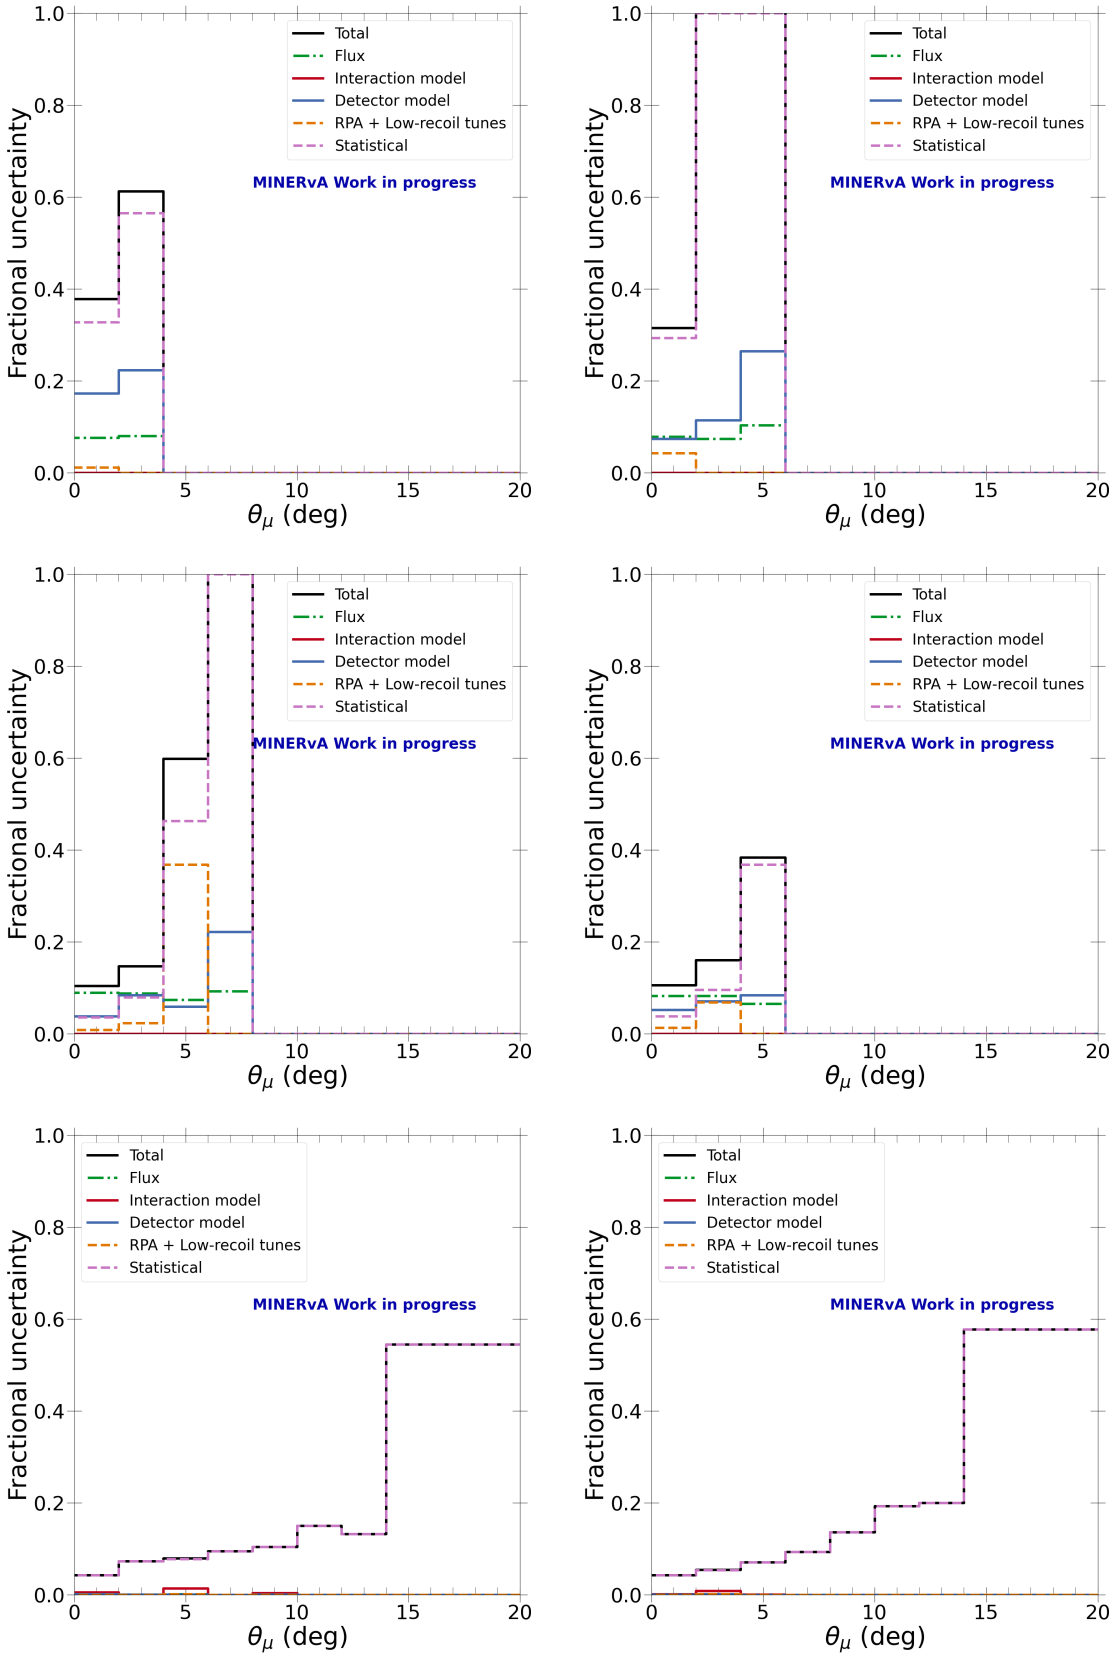


Figure 8.10: Fractional uncertainties, statistical and systematic, in θ_μ. Left: FHC, right: RHC. Top to bottom: $M_{\mathrm{N4}} = 250$ MeV/c^2, $M_{\mathrm{N4}} = 475$ MeV/c^2, SM inclusive MC.

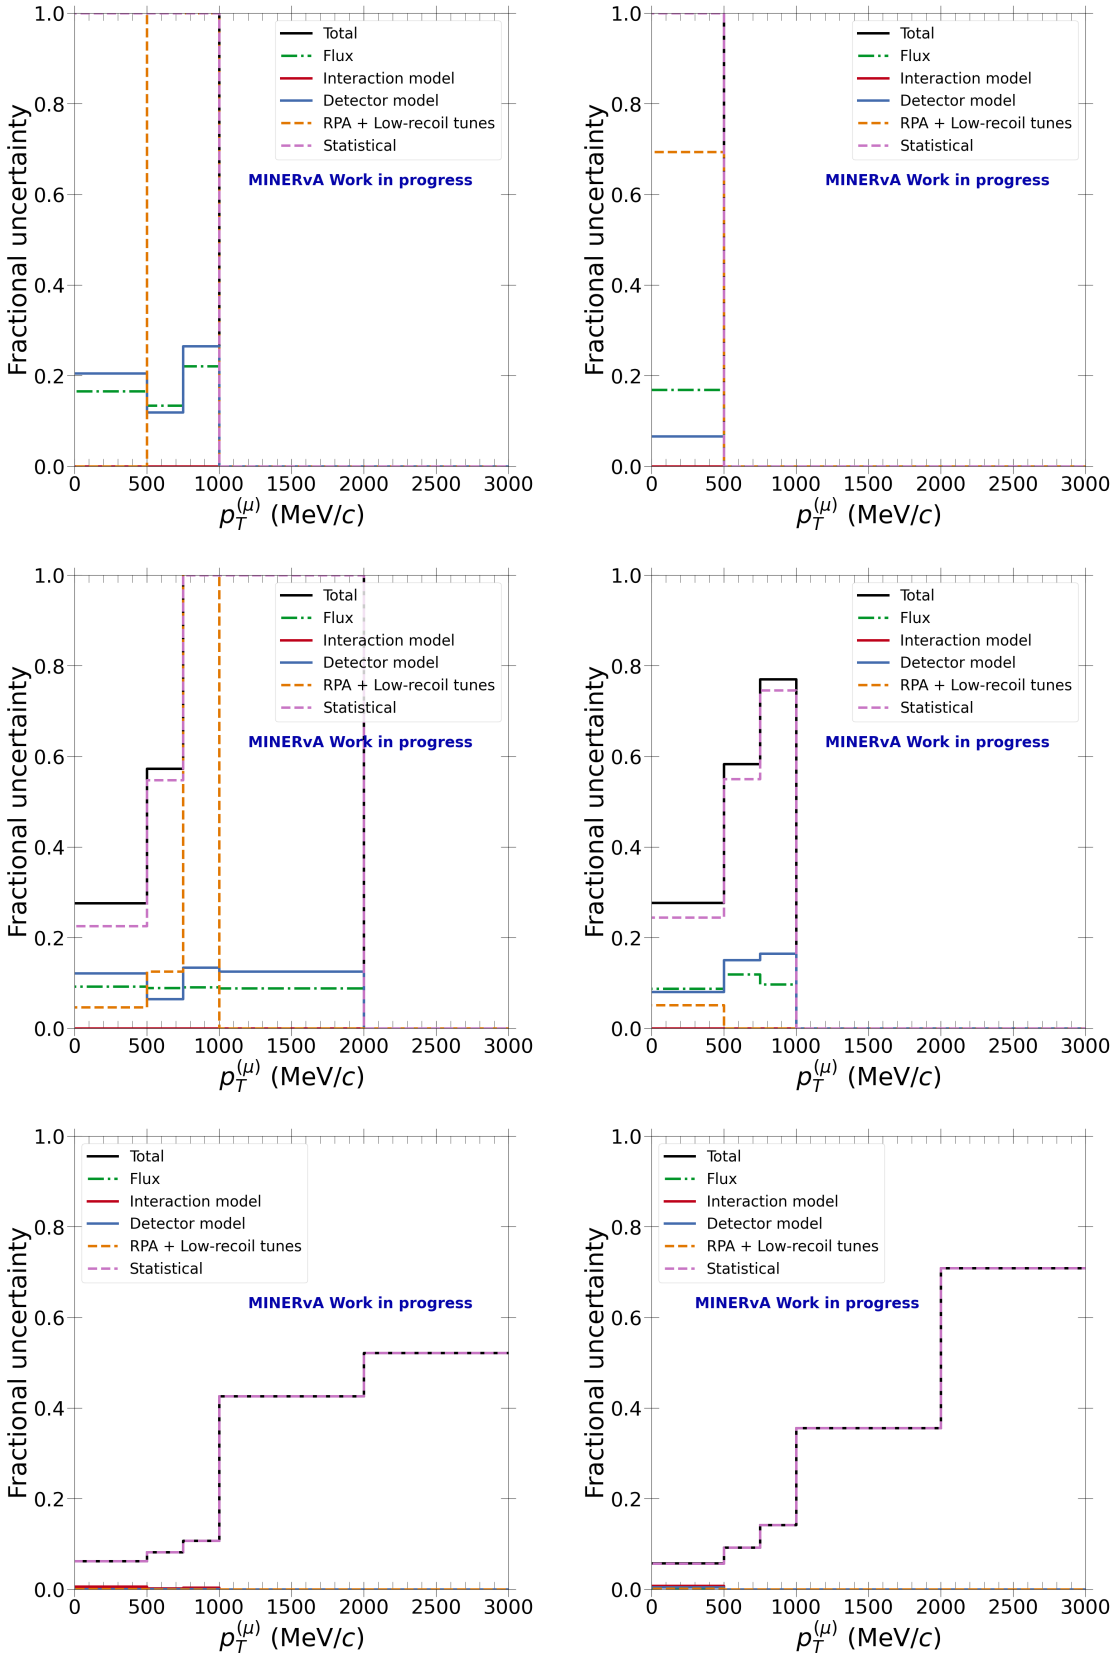


Figure 8.11: Fractional uncertainties, statistical and systematic, in $p_{\mathrm{T}}^{(\mu)}$. Left: FHC, right: RHC. Top to bottom: $M_{\mathrm{N4}} = 250$ MeV/c^2, $M_{\mathrm{N4}} = 475$ MeV/c^2, SM inclusive MC.

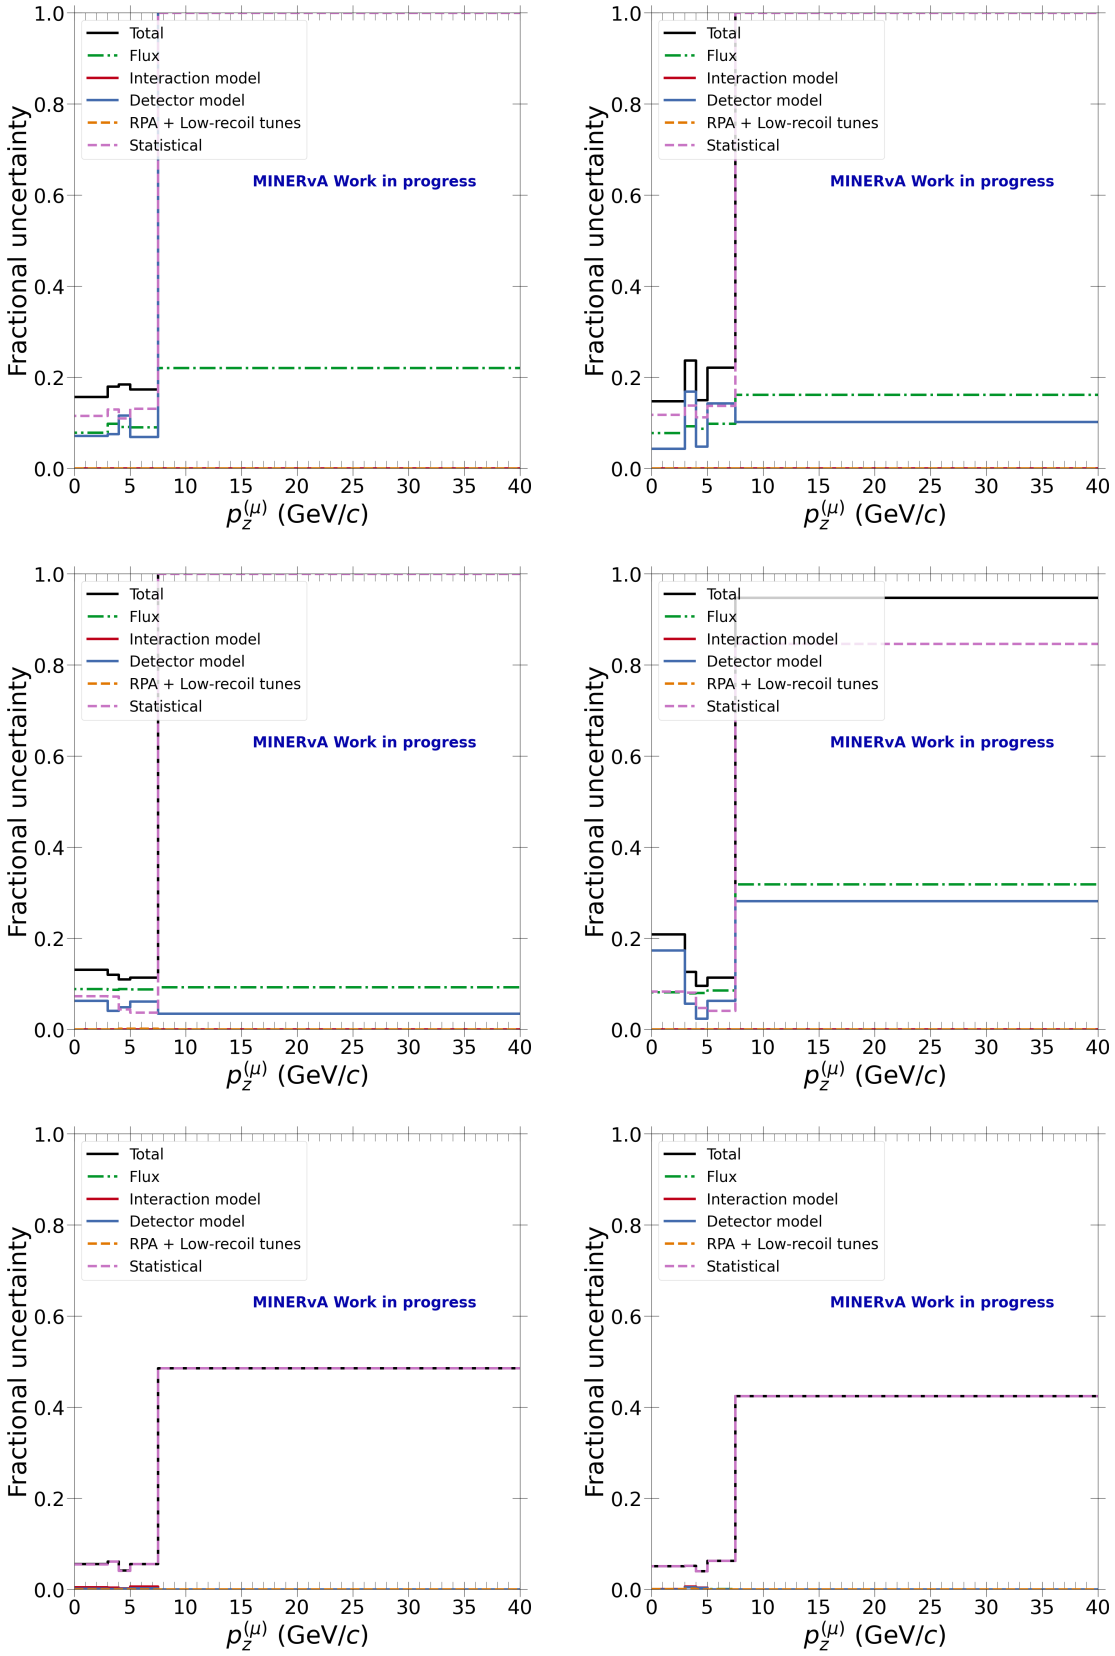


Figure 8.12: Fractional uncertainties, statistical and systematic, in $p_z^{(\mu)}$. Left: FHC, right: RHC. Top to bottom: $M_{\mathrm{N4}} = 250$ MeV/c^2, $M_{\mathrm{N4}} = 475$ MeV/c^2, SM inclusive MC.

MC inside the *signal region*, where the presence of a background must be estimated. Generally, the way to proceed with a sideband fit is to minimise the quantity

$$\chi^2 = \sum_i \frac{\left(N_i^{\mathrm{DATA}} - \sum_b \alpha_b N_i^{\mathrm{MC};b}\right)^2}{\sum_b \alpha_b N_i^{\mathrm{MC};b}}, \tag{8.3}$$

where the index i refers to bins in the variable $\mathcal{V}$, b refers to interaction type in the background, and α_b is the *scale factor* for the interaction type's contribution. For the purposes of this discussion, b runs over the indices "QEL" (quasi-elastic), "RES" (resonant pion production), "DIS" (deep inelastic scattering), "MEC" (2p2h scattering), "COH" (CC coherent pion production and CC diffractive pion production, where the event selection histograms of the coherent and diffractive samples have been added together). The "IMD" sample has an extremely low selection efficiency, and small cross section, meaning a negligible number of entries survive selection; the sample is therefore dropped from the list of model parameters. There are, in other words, $N = 5$ model parameters to constrain.

One sideband, named here "C_1", immediately presents itself upon inspection of the θ_μ distributions in Figs. 8.4, 8.5; the region $\theta_\mu \geq 6^\circ$ is extremely signal-poor, with the highest signal contamination (defined as the integral of the signal distribution that lies in C_1 divided by the integral over the entire phase space) in the RHC $M_{\mathrm{N4}} = 475$ MeV/c^2 sample of about 0.5%. Though this provides a suitable level of statistics for the "QEL", "RES", and "DIS" interaction types, the "MEC", "COH", and "DIF" modes are highly suppressed in this region and will therefore not be able to be constrained in here. Consequently, one would have to trust the `GENIE` calculation of cross-sections for these processes at that specific region of phase space for our signal region, which is not certain to have good agreement with data. The "IMD" type is similarly suppressed, but so few events survive cuts relative to the POT processed (see Table 6.1) that, to first order, the contribution from "IMD" is negligible.

A second signal-poor region with a substantially larger contribution of all the above modes, named "C_2", comes from inspecting the "double-peak" structure of Figs. 7.23, 7.24; one may immediately isolate a region $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} \leq 2.0$ MeV/cm

which is signal-poor (again, the RHC $M_{\mathrm{N4}} = 475$ MeV/c^2 sample is the largest contributor, with 0.4% of the sample in this region), and the added benefit of having a much larger "MEC", "COH", and "DIF" component. In Table 8.2, a summary of the contaminations in either sideband is shown. Apart, however, from

All contaminations in %.				
Mode	C_1^{FHC}	C_2^{FHC}	C_1^{RHC}	C_2^{RHC}
HNL-250	0.00	0.0123	0.0208	0.0453
HNL-475	0.297	0.250	0.502	0.401
QEL	26.1	81.7	32.9	78.9
RES	33.9	71.6	34.9	82.7
DIS	67.9	55.8	52.4	75.1
MEC	29.4	85.3	10.5	85.1
COH	1.15	22.3	1.22	30.8
DIF	1.70	18.7	0.00	21.7
IMD	2.58	59.0	N/A	N/A

Table 8.2: Contaminations of background sources in the two signal-poor regions "C_1" and "C_2". "COH" refers to the dedicated CC Coherent sample, rather than the component of the General sample (which actually refers to CC coherent pion production from carbon, i.e. the "CC Coherent", and from hydrogen, i.e. the "CC Diffractive" modes). "Contamination" of a process P in a sideband S is defined as the integral of the normalised distribution of P in S, divided by the integral of the normalised distribution of P over all phase space. For C_1, these distributions are as in Figs. 8.4, 8.5; for C_2, they are as in Figs. 7.23, 7.24.

the contaminations alone, there is a subtlety with the extrapolation to the signal region which arises in C_2. Effectively, since C_2 corresponds to a part of phase space where the outgoing pion has relatively little kinetic energy (and as such "ranges out" quickly, without causing Algorithm 6 to detect a pion endpoint), it is not immediately apparent how well the MC simulation can translate it to the high-energy signal region $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} \geq 2.7$ MeV/cm. In particular, the "transition region" 2.0 MeV/cm $\leq \langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} \leq 2.7$ MeV/cm is a difficult region to model, as it sits on the edge of the signal region's phase space (see Fig. 7.6, the left-most regions of the bottom row). Therefore, extracting scale factors from C_2 will hinge upon the assumption that the MINERνA simulation models the transition region well enough to preserve the smoothness of the underlying phase space, to justify using

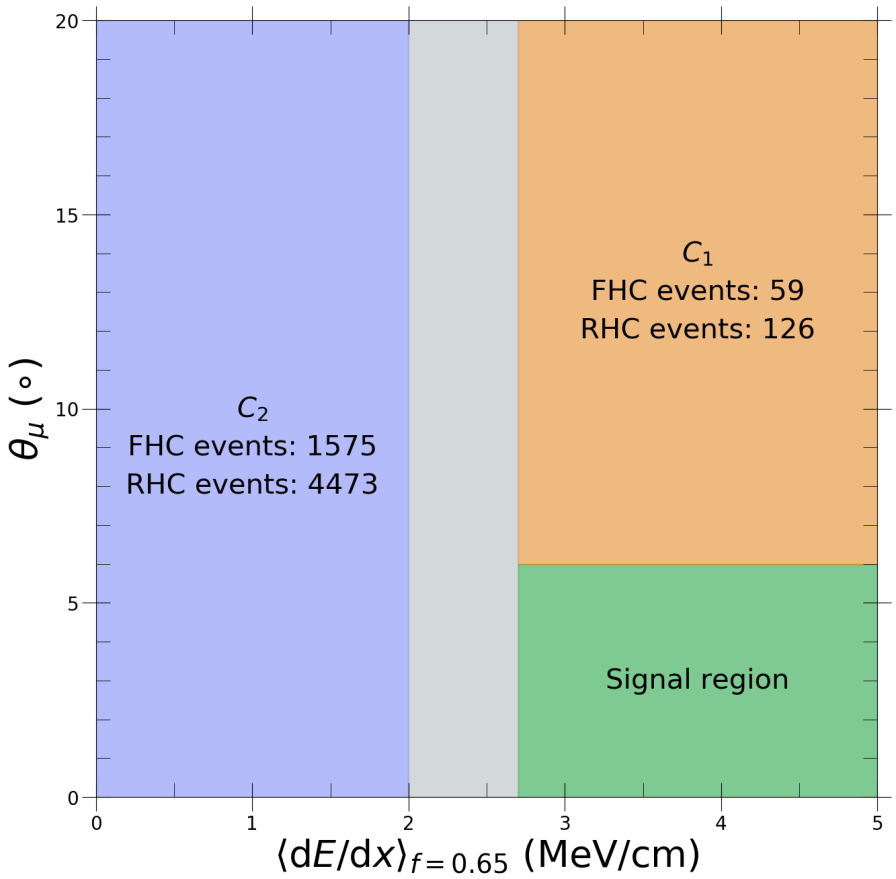


Figure 8.13: Schematic illustration of sidebands C_1 (top right) and C_2 (left band), as well as the signal region (bottom right) as a function of $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65}$ and θ_μ. The number of events in data in each sideband is also shown.

the scale factors here. On the contrary, C_1 utilises muon kinematic information. Schematically, $C_{1,2}$ and the signal region are shown in Figure 8.13.

8.4.2 Tuning the background

The SM background processes under study occupy a hitherto poorly constrained region of phase space, which necessitates a way to find *modifications of the background predictions* around their nominal value, to correct for model dependencies as best as possible.

Pearson's χ^2 test statistic for binned counts provides a useful metric for evaluating the agreement between the MC and the data in the sideband region. In order to have a well-behaved χ^2, the number of counts in each bin must be large enough to make the assumptions in Pearson's chi-squared test valid; specifically, not more

than 1 in 5 bins contains fewer than 5 entries [302]. For each of the sidebands, a simultaneous fit across all three variables was performed minimising Pearson's χ^2 using the `minos` routine [303] implemented in `MINUIT` [304, 305], which calculates 1σ errors on each scale factor.

There are few events in data that survive the sideband cuts (about 100 events per current in C_1, and about 5000 per current in C_2), which makes the ideal procedure of constructing a full three-dimensional distribution in $(\theta_\mu \times p_{\mathrm{T}}^{(\mu)} \times p_z^{(\mu)})$ impracticable. Instead, a singular one-dimensional distribution with all the bins enumerated 1 through 19 is constructed, and a fit run through the entire series of bins. This means that every single data point (and MC entry) is *counted three times* (once for each variable) in the distribution. In the case where each distribution were identical, and one would not expect any new information from including a new variable, this would mean the total χ^2 would be three times as large as one would expect from perfectly uncorrelated measurements. Therefore, the errors extracted would be a factor $\sqrt{3}$ larger than they should be. However, this is not the case, as the different MC contributions have different shapes in each variable; therefore, the total χ^2 is indeed larger than the "ideal" χ^2, but by a factor $N < 3$; so the normalisation factor of the errors is strongly bounded upper by $\sqrt{3}$. In line with the conservative treatment to be employed, the errors as obtained by this one-dimensional fitting will be kept unmodified, with the understanding that they introduce an overcoverage of $< \sqrt{3}$ in the final result. It is, however, an interesting open question to consider how an optimal binning choice might be designed for multi-dimensional fits, given the small number of events and the large number of bins (for this thesis's binning choice, for example, there would be $9 \times 5 \times 5 = 225$ bins in a full 3-D distribution).

In the frequentist picture, the computation of errors depends on the landscape of the $\chi^2_N(\boldsymbol{X})$, where N is the number of free parameters in the fit and $\boldsymbol{X}$ is the vector of data. Specifically, to calculate the error for a confidence level α for one single fit variable ξ, i.e. $N = 1$, one finds the values ξ_-, ξ_+ such that

$$\Delta\chi^2_{N=1}(\xi; \boldsymbol{X}) = \delta, \tag{8.4}$$

where δ is a function of α and the degrees of freedom N as prescribed by Wilks' theorem [306]. For a 1σ error, $\alpha = \text{erf}\left(1/\sqrt{2}\right) \simeq 0.6827$. At $N = 1$, the equation 8.4 yields the well-known answer $\delta = 1$ at $\alpha = 1\sigma$; however, for more degrees of freedom, the error definitions grow. The results for various combinations of α, N are summarised in Table 8.3. Based on this error definition, the `minos`

N \ α	1σ	0.90	2σ
1	1	2.71	4
2	2.30	4.61	6.18
3	3.53	6.25	8.02
4	4.72	7.78	9.72
5	5.89	9.24	11.31

Table 8.3: $\Delta\chi^2$ values for various errors and number of fit parameters.

routine constructs contours of $\Delta\chi^2$ in parameter space and reports asymmetric errors on each parameter as the values (ξ_-, ξ_+) which bound the contour. If the parameter in question has physical limits (such as in our case where the scale factors are ≥ 0), the routine will terminate at the limits in parameter space.

The reason behind the choice of error definition comes from the parameter counting: initially, all of the scale factors start off variable. The fit is run with all the parameters, with the best-fit values $\{\hat{a}_i\}$ saved.

This "inflated" choice of errors is made, to account for parameter correlations in the fit. Ideally, one would be able to apply the $N = 1$ error definition, but that is only correct in the case of no correlations between parameters. In the presence of correlations, the $N = 1$ definition is insufficient to capture the simultaneous variation of all the parameters. An example is shown in Fig. 8.14, where the best-fit point (star) is plotted in the space of two scale factors, as well as the $\Delta\chi^2$. There are various contours: first, a dashed "inner" contour, corresponding to $\Delta\chi^2 = 1$ and the "traditional" 1σ error definition with $N = 1$ free parameter. However, as discussed before, the inner contour is not the appropriate one to use for a multi-parameter fit; what one requires is the contour which encloses enough of the parameter space such that the likelihood of the "true" value of the scale factors

being in the contour, *allowing all scale factors to vary simultaneously*, is $P = 0.6827$. To that effect, the contours for a 1σ definition of errors for $N = 1, \ldots, 5$ are shown; for both modes in the full fit, the outer solid contour at $N = 5(\Delta\chi^2 = 5.89)$ is used. The effects of the contour's non-symmetric shape (reflecting the correlations between parameters) are obvious in the FHC contour, justifying our approach of using `minos` errors (i.e. profile likelihood). Essentially, using the inflated errors (corresponding to a 5-dimensional hypersurface in the space of the scale factors being fitted) yields a region of parameter space such that the likelihood of the "true" value is $P \geq 0.6827$. One could reduce the degree of overcoverage by employing an alternative parametrisation of the parameter space with linear combinations $\{\beta_j\} : \beta_j = \sum_i w_j^{(i)} \alpha_i$, such that the correlation matrix of the parameters β were diagonal. This could, for example, be achieved using a suitable decomposition of the correlation matrix. However, as will become clear in Section 8.4.3, the full five-dimensional fit yields results that motivate a simpler parametrisation at the expense of overcoverage, as described here.

It also becomes apparent how the physical limit $\alpha_i \geq 0 \; \forall \; i$ becomes relevant: the contour intersects the limit at $\alpha_{\text{QEL}} = 0(\alpha_{\text{RES}} = 0)$, and therefore `MINOS` terminates at those points, assigning the appropriate errors.

8.4.3 Fit results

Throughout this Section, there will be three fits performed. At the beginning, the full unconstrained fit with $N = 5$ parameters, where the index b in Eq. 8.3 runs over all five background sources. The reason for this attempt is to estimate the effect of parameter correlations on the fit from the $C_{1,2}$ sidebands, and to test if the two sidebands are yielding a consistent picture of this search's model dependence on the backgrounds. To test the agreement of the SM background simulated shape against the data, a reduced $N = 1$ fit will be run, where b will have only one label: "SM". This implies all the components are added up with the nominal scale factors from the MC, and only their normalisation is varied. Finally, to better isolate the contribution of the "COH" component, a fit at $N = 2$ with the "QEL", "RES",

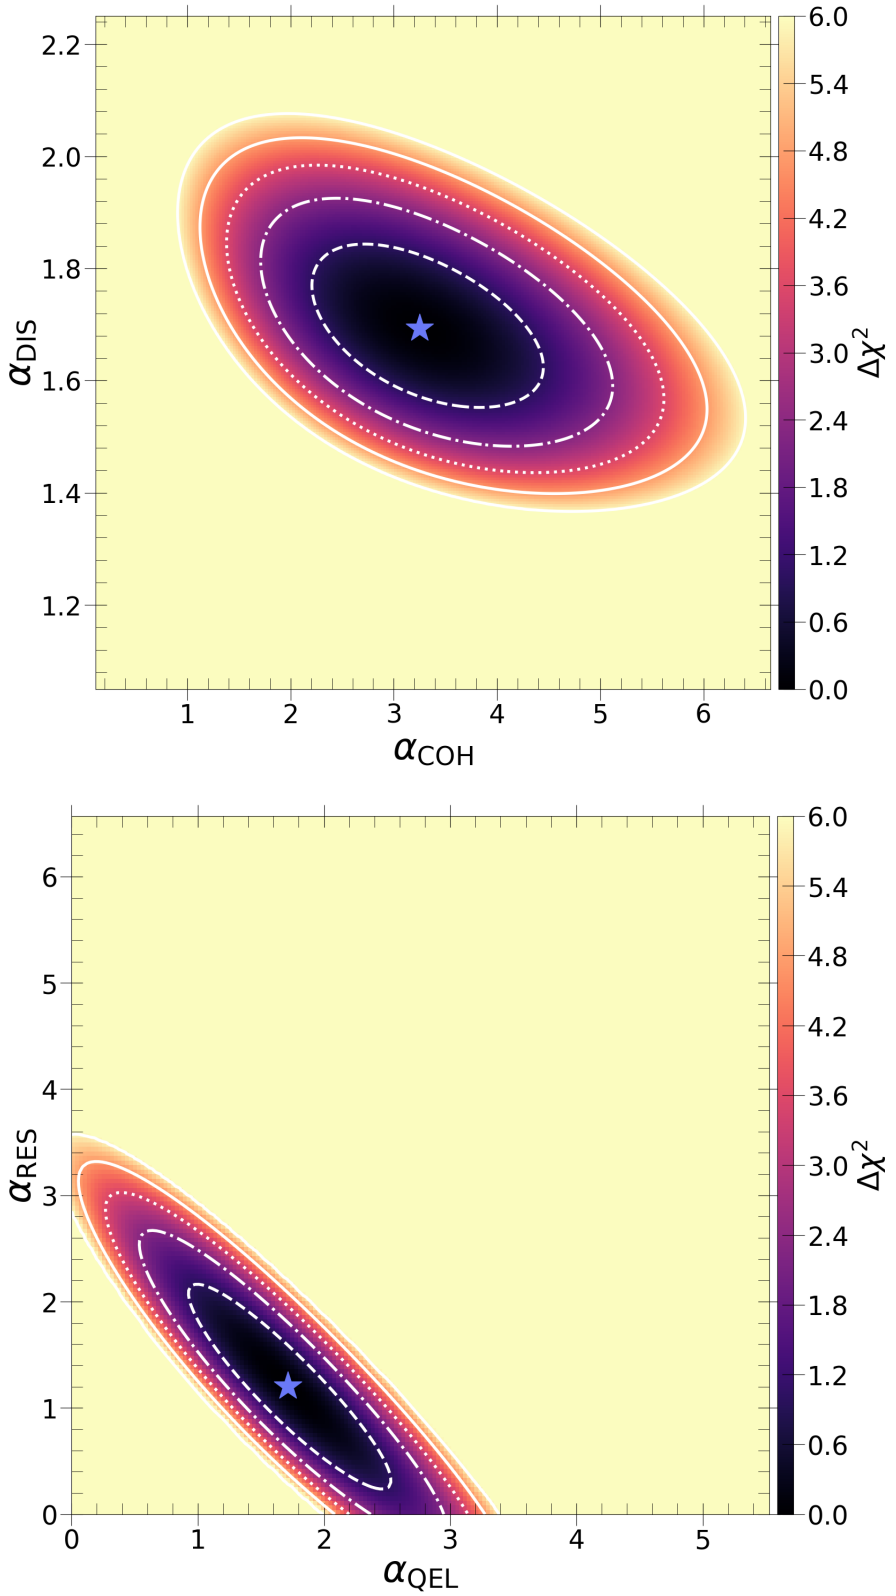


Figure 8.14: $\Delta\chi^2$ landscape plots for the C_1 sideband, FHC (top) and RHC (bottom). The best-fit point is represented as a star. The various contours are, from inside to outside, the 1σ contours for $N = 1, \ldots, 5$. $\Delta\chi^2$ values greater than 6 are truncated. The parameters not on the axes are fixed to their best-fit values.

"DIS", and "MEC" components added up into a single "other" component and the "COH" component left alone. Table 8.4 shows a summary of these definitions.

Fit	$\Delta\chi^2$	COH	QEL	RES	DIS	MEC
$N = 5$	5.89	✓	✓	✓	✓	✓
$N = 2$	2.30	✓	Merged into "other"			
$N = 1$	1.00		Merged into "SM"			

Table 8.4: Definition of free parameters used in the full unconstrained and the reduced fits in this Chapter.

Unconstrained $N = 5$ fit

To estimate the background contribution, a combined fit across the three kinematic variables accessible to this analysis (θ_μ, $p_{\mathrm{T}}^{(\mu)}$, $p_z^{(\mu)}$) was performed, by concatenating the three 1D distributions into a larger 1D distribution (see discussion above) for each of the two sidebands, for FHC and RHC. The fit results are summarised in Table 8.5. The comparisons of the untuned and tuned MC backgrounds in the C_1 and C_2 sidebands can be seen in Figs. 8.15-8.18. The choice of binning in each variable is made to ensure there are adequate statistics in each bin to have a well-behaved χ^2, as discussed above, with sufficient granularity to capture the shape of each distribution. In the plots, each bin has the number of data events superimposed in brackets. The $\chi^2/N_{\mathrm{d.o.f}}$ in each figure refer to the figure of merit for that single distribution only, not the global $\chi^2/N_{\mathrm{d.o.f}}$ obtained by `MINUIT`.

The results from the $N = 5$ fit lead to two main conclusions. Firstly, the untuned Monte Carlo does not do a particularly good job of modelling the inherent background contribution to either sideband. With the exception of the C_1 FHC sideband, where at least the figure of merit $\chi^2/N_{\mathrm{d.o.f}}$ is $\mathcal{O}(1)$, the MC background is significantly underpredicted (and even the C_1 FHC sideband is pulled towards a higher contribution after tuning). The post-fit prediction, on the other hand, approximates the data rather well. For the C_1 sideband, where all cuts including $\langle \mathrm{d}E/\mathrm{d}x \rangle_{f=0.65} \geq 2.7$ MeV/cm have been applied, the RHC modelling is decidedly worse; this region of phase space corresponds to high-energy pions accompanying

Component	Best-fit	Low error ξ_-	High error ξ_+
C_1 FHC; $\Delta\chi^2 = 5.89$.			
$\chi^2/N_{\text{d.o.f}} = 13.49/10 = 1.35$; Untuned $\chi^2/N_{\text{d.o.f}} = 51.33/15 = 3.42$.			
QEL	0.0	0.0	5.67
RES	0.0	0.0	1.67
DIS	1.65	0.60	0.39
MEC	0.0	0.0	13.4
COH	3.38	3.00	3.27
C_1 RHC; $\Delta\chi^2 = 5.89$.			
$\chi^2/N_{\text{d.o.f}} = 4.84/10 = 0.48$; Untuned $\chi^2/N_{\text{d.o.f}} = 1214.18/15 = 80.95$.			
QEL	1.71	1.71	3.82
RES	1.20	1.20	5.37
DIS	5.26	5.26	6.65
MEC	31.2	31.2	40.7
COH	3.94	3.94	11.9
C_2 FHC; $\Delta\chi^2 = 5.89$.			
$\chi^2/N_{\text{d.o.f}} = 5.72/14 = 0.41$; Untuned $\chi^2/N_{\text{d.o.f}} = 2531.86/19 = 133.26$.			
QEL	2.49	1.92	1.90
RES	1.86	1.42	1.38
DIS	1.65	1.10	1.25
MEC	0.97	0.97	3.35
COH	0.90	0.90	1.26
C_2 RHC; $\Delta\chi^2 = 5.89$.			
$\chi^2/N_{\text{d.o.f}} = 4.31/14 = 0.31$; Untuned $\chi^2/N_{\text{d.o.f}} = 13974.82/19 = 735.52$.			
QEL	2.94	1.41	2.03
RES	2.32	2.32	1.70
DIS	1.76	1.76	3.08
MEC	1.09	1.09	2.01
COH	0.0	0.0	1.71

Table 8.5: Scale factors for each background component as reported by the `minos` routine from `MINUIT`. Errors are asymmetric, see main text. The minimum $\chi^2/N_{\text{d.o.f}}$ and error definition based on Table 8.3 are reported. Values of background at 0.0 result from the `minos` routine terminating at a physical limit. The error definition $\Delta\chi^2 = 5.89$ (cf. Table 8.3) is made as appropriate for five simultaneously varying scale factors.

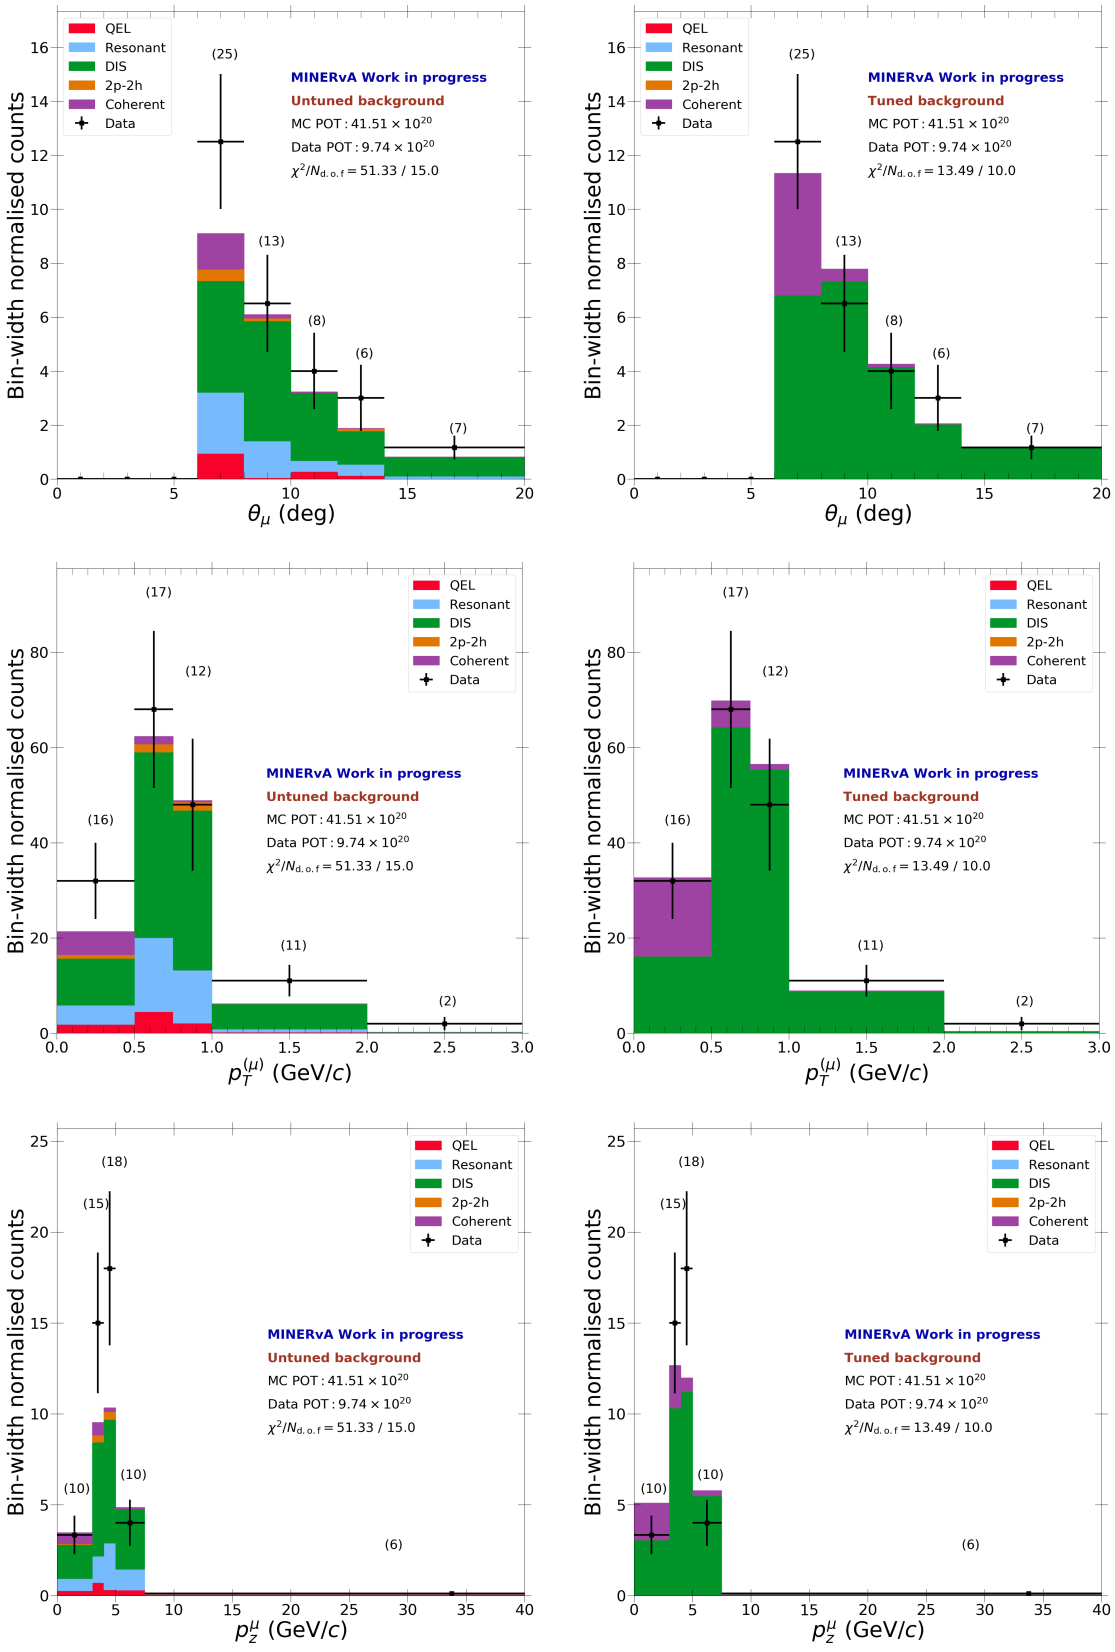


Figure 8.15: Comparison of untuned (left) and tuned (right) MC in the FHC C_1 sideband. Number of data events in each bin in brackets.

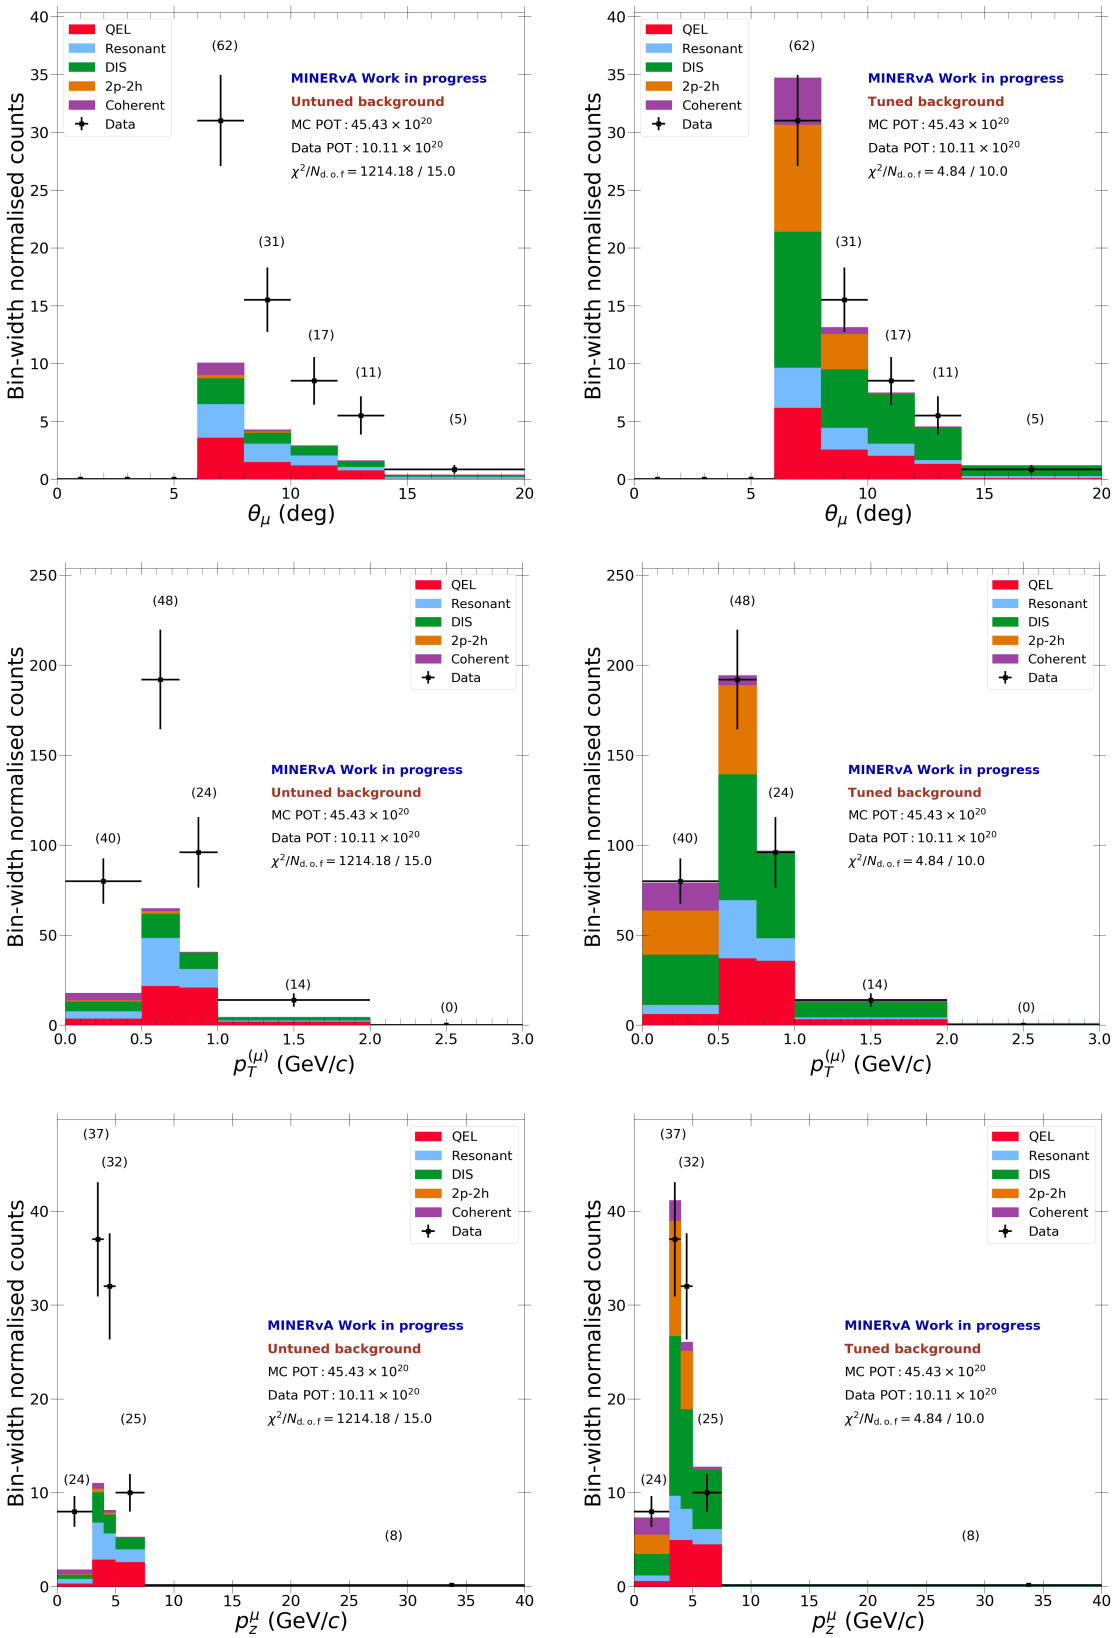


Figure 8.16: Comparison of untuned (left) and tuned (right) MC in the RHC C_1 sideband. Number of data events in each bin in brackets.

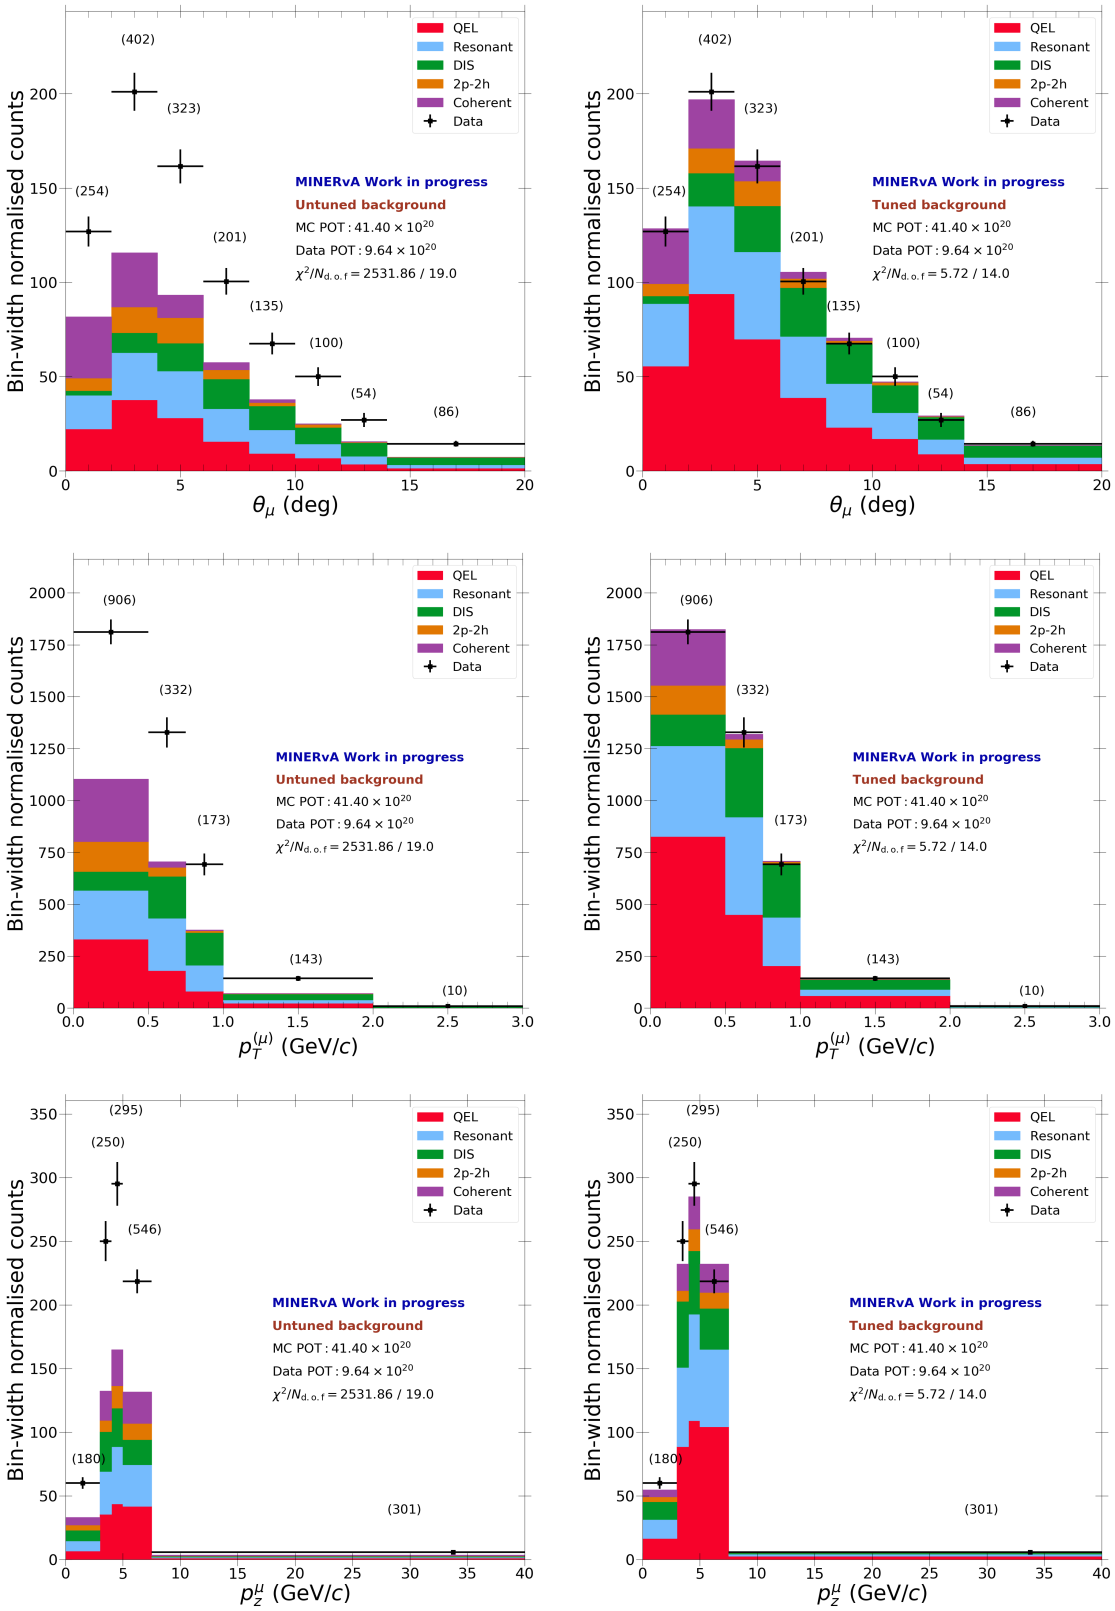


Figure 8.17: Comparison of untuned (left) and tuned (right) MC in the FHC C_2 sideband. Number of data events in each bin in brackets.

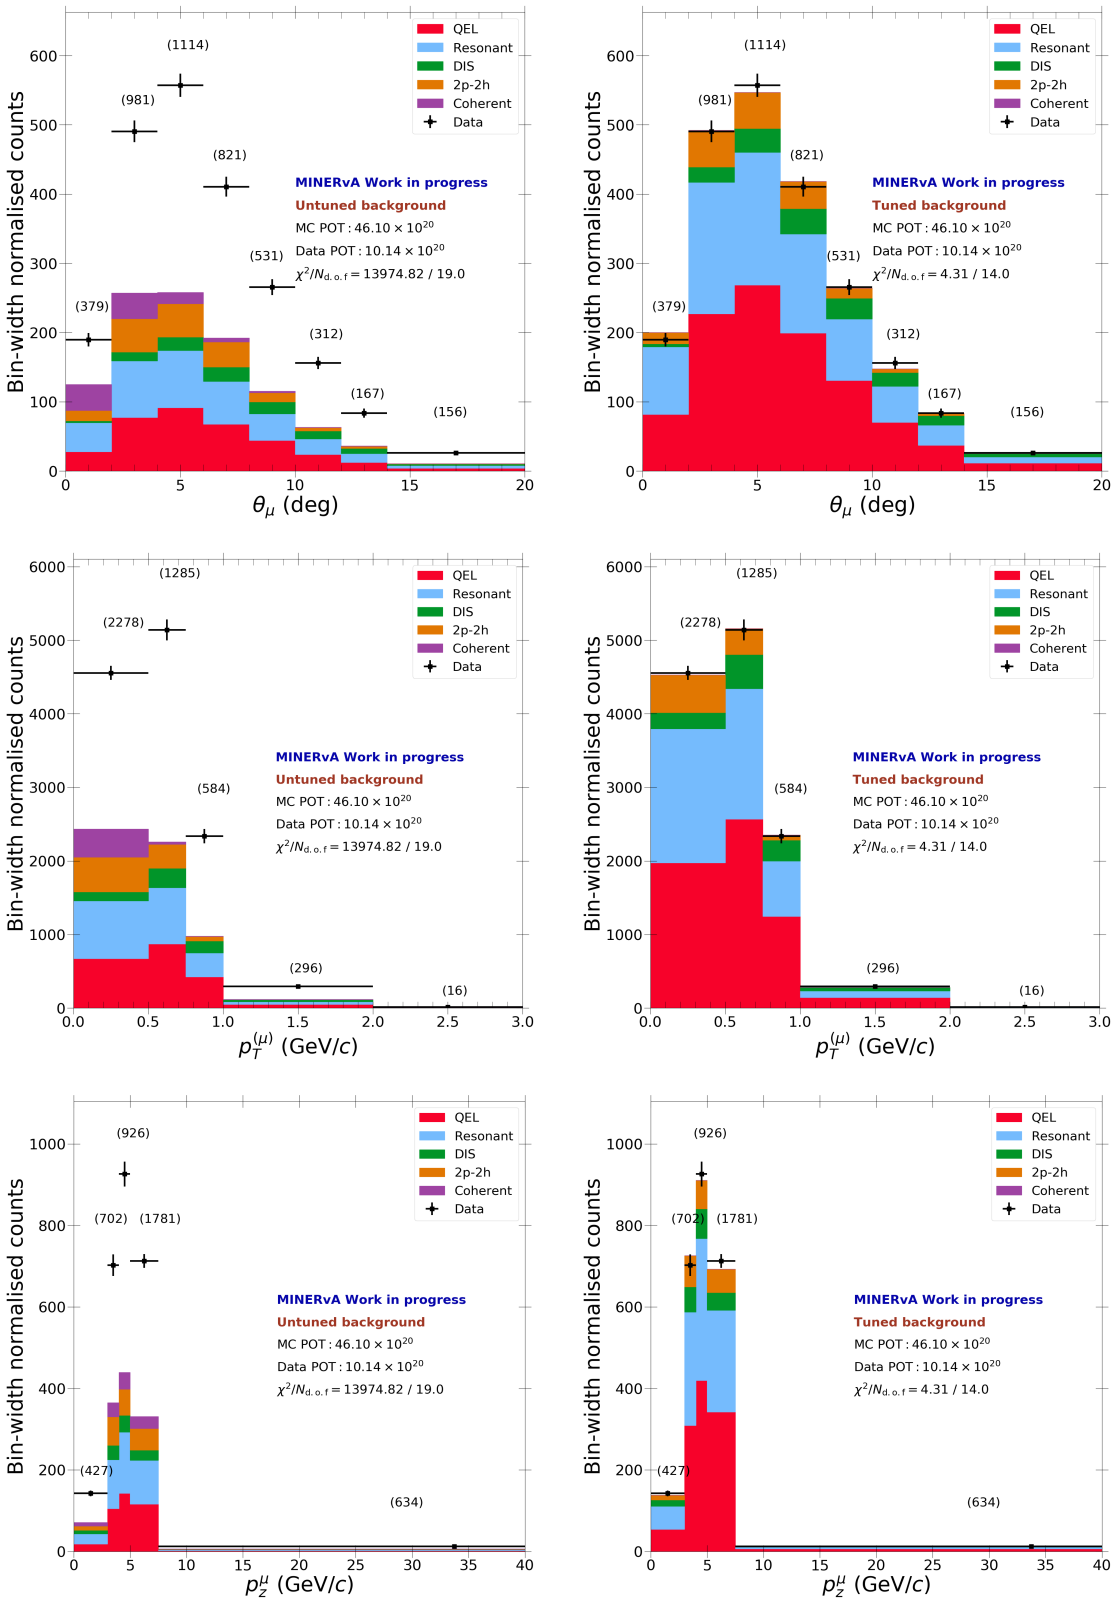


Figure 8.18: Comparison of untuned (left) and tuned (right) MC in the RHC C_2 sideband. Number of data events in each bin in brackets.

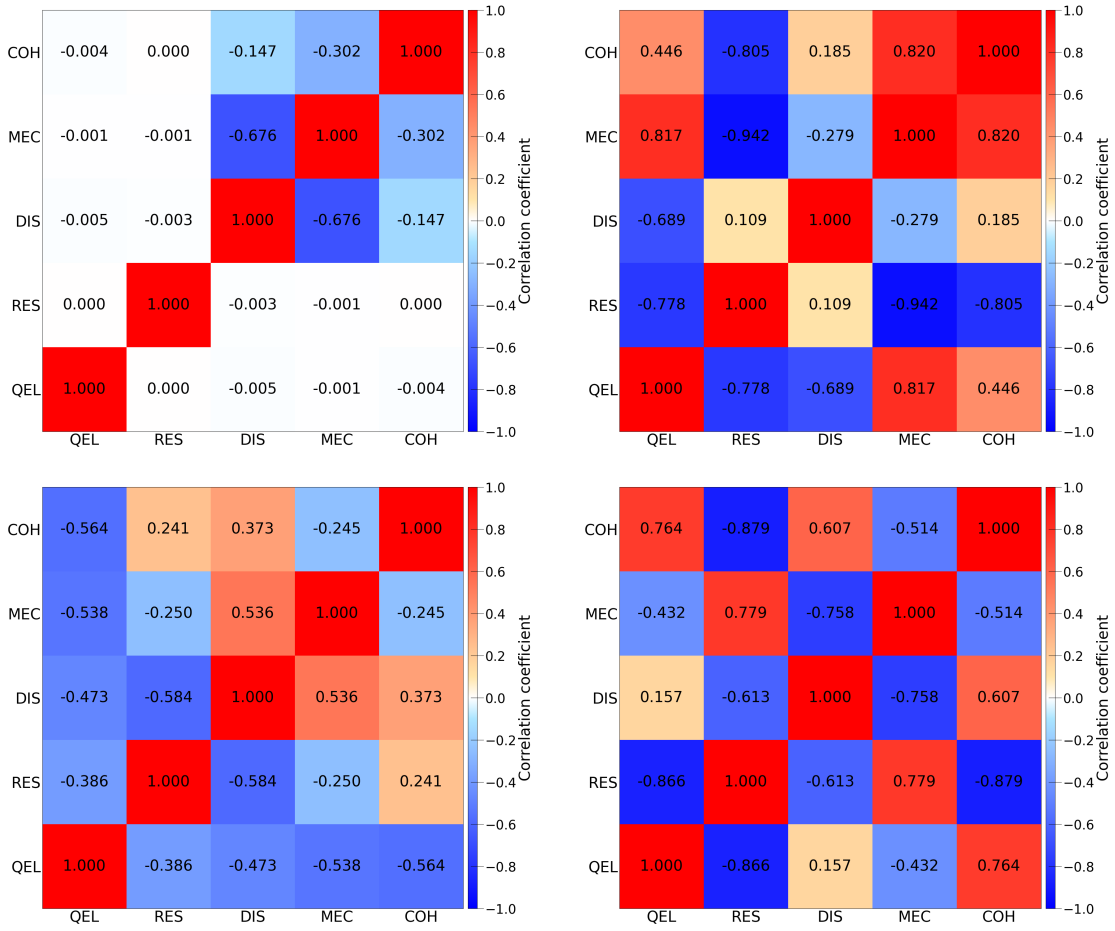


Figure 8.19: Correlation matrices for the SM background components, for the C_1 (top) and C_2 (bottom) sidebands, with FHC (left) and RHC (right).

the muons and therefore registering in the detector as a double-MIP pair. This suggests that there needs to be an enhancement in the forward-going high-energy pions across both FHC and RHC modes, though it is unclear from this fit which processes need to be modified. Secondly, the parameters are correlated, which complicates the fit; the correlation matrices are shown in Figure 8.19. Due to the relatively small number of non-zero bins and the weak shape discrimination of the various background components, especially in the C_1 sideband, the fit is very underconstrained. This is evident in the particularly large errors in the scale factors. Things are yet worse for the scale factor α_{COH}, since the contamination in C_1 is very small and the range of the CC Coherent contribution is restricted to fewer bins still. The improvement of the χ^2 is huge; in all but the FHC C_1 sideband, the untuned

prediction is not even close to being $\mathcal{O}(1)$; after the fit, the minimum chi-squared is even smaller than the degrees of freedom. By inspecting Figs. 8.15-8.18, one may see the untuned MC significantly underpredicts the data for the selection; effectively, all interaction modes must be scaled up, with the curious exception of the C_2 RHC sideband where the best-fit point causes α_{COH} to intersect the physical limit $\alpha_{\text{COH}} = 0$. It is, naturally, extremely unlikely that there should be no contribution at all in the low-$\langle \mathrm{d}E/\mathrm{d}x \rangle$ from coherently produced pions; the attached uncertainty $\delta\alpha_{\text{COH}} = 1.71$ is another clue that the best-fit point, taken as is, is not the best predictor for the true value of α_{COH}. In other words, there is not a consistent picture between FHC and RHC, between C_1 and C_2. This implies that a reparametrisation of the parameter space with the aim of using the smaller $N = 1$ errors is not justified, and the full fit would require the inflated, overconservative $N = 5$ errors.

For these reasons, then, a reduced fit was explored. The aim of the reduced fit (which was attempted using both at $N = 1$ using a single, "SM" backgroundm and at $N = 2$ isolating the "COH" component from the "non-COH" component) is to investigate if it is possible to reduce the overcoverage introduced in the full fit.

Reduced fit: $N = 1$

First, to test the hypothesis of a normalisation only causing the discrepancy, a single-parameter fit ($N = 1$) with all the MC contributions weighted the same way (see Tab. 8.4, "SM" component) was performed. The results of these fits with the improvement in the χ^2 are shown in Table 8.6. Of the four combinations, each one is substantially improved even with a single-parameter fit, and the errors of each component are now far smaller. Given that there is now only one freely varying parameter, the appropriate definition is $\Delta\chi^2 = 1$; it is not the $\Delta\chi^2$ that causes the dramatic decrease of the error, though. To illustrate this point, in Figure 8.20 the $\Delta\chi^2$ for the parameters $N = 1, \ldots, 5$ is shown, with the intersection points defining the `minos` errors for each N; the errors are the same order of magnitude as the $N = 1$ case. This could be expected from the fact that the $\Delta\chi^2$ curve is almost parabolic.

Component	Best-fit	Low error ξ_-	High error ξ_+
Tuned $\chi^2/N_{\text{d.o.f}} = 20.70/14 = 1.48$; Untuned $\chi^2/N_{\text{d.o.f}} = 36.67/15 = 2.44$			
C_1 FHC	1.268	0.079	0.085
Tuned $\chi^2/N_{\text{d.o.f}} = 25.17/14 = 1.80$; Untuned $\chi^2/N_{\text{d.o.f}} = 867.27/15 = 57.82$			
C_1 RHC	3.12	0.13	0.14
Tuned $\chi^2/N_{\text{d.o.f}} = 24.84/18 = 1.38$; Untuned $\chi^2/N_{\text{d.o.f}} = 1969.24/19 = 103.64$			
C_2 FHC	1.763	0.023	0.024
Tuned $\chi^2/N_{\text{d.o.f}} = 179.58/18 = 9.98$; Untuned $\chi^2/N_{\text{d.o.f}} = 10869.31/19 = 572.07$			
C_2 RHC	2.120	0.016	0.016

Table 8.6: Single-parameter fit results. Errors are reported using `minos` under the appropriate error definition $\Delta\chi^2 = 1$.

All of the sidebands have $\chi^2/N_{\text{d.o.f}}$, min $\simeq 1$; the notable exception of C_2 RHC is still $\mathcal{O}(1)$. This suggests that the shape of the background contribution, at least, does an adequate job of describing the shape of the data. This characterisation, though, is impeded by the very low statistics in this region, leading to large errors on both the data and the MC; in other words, there is a low degree of confidence in the shape of the data and the MC in this region. Figures 8.21-8.23 showcase the shape of the distribution with the error on the MC also shown; the errors on data and MC are notable. Qualitatively, though, it becomes apparent that the MC distribution broadly follows the shape of the data in all but the RHC C_2 sideband, where there is a definite overprediction of very beam-collimated, low-$p_{\text{T}}^{(\mu)}$ events compared to data. This behaviour of the C_2 sideband cautions one to not use the $N = 1$ errors obtained by this fit; since the fit result from the $N = 5$ case pointed to the coherent contribution having anomalous behaviour, the next step is to excise it from the other backgrounds and allow it to vary simultaneously from the other four modes.

Reduced fit: $N = 2$

Finally, to isolate the CC Coherent irreducible HNL background, a two-parameter ($N = 2$) fit was performed, once again summing up all the other components and varying just that sum, as well as the CC Coherent / Diffractive combination (see Tab. 8.4, "COH" and "other" components). The result of that fit is shown in

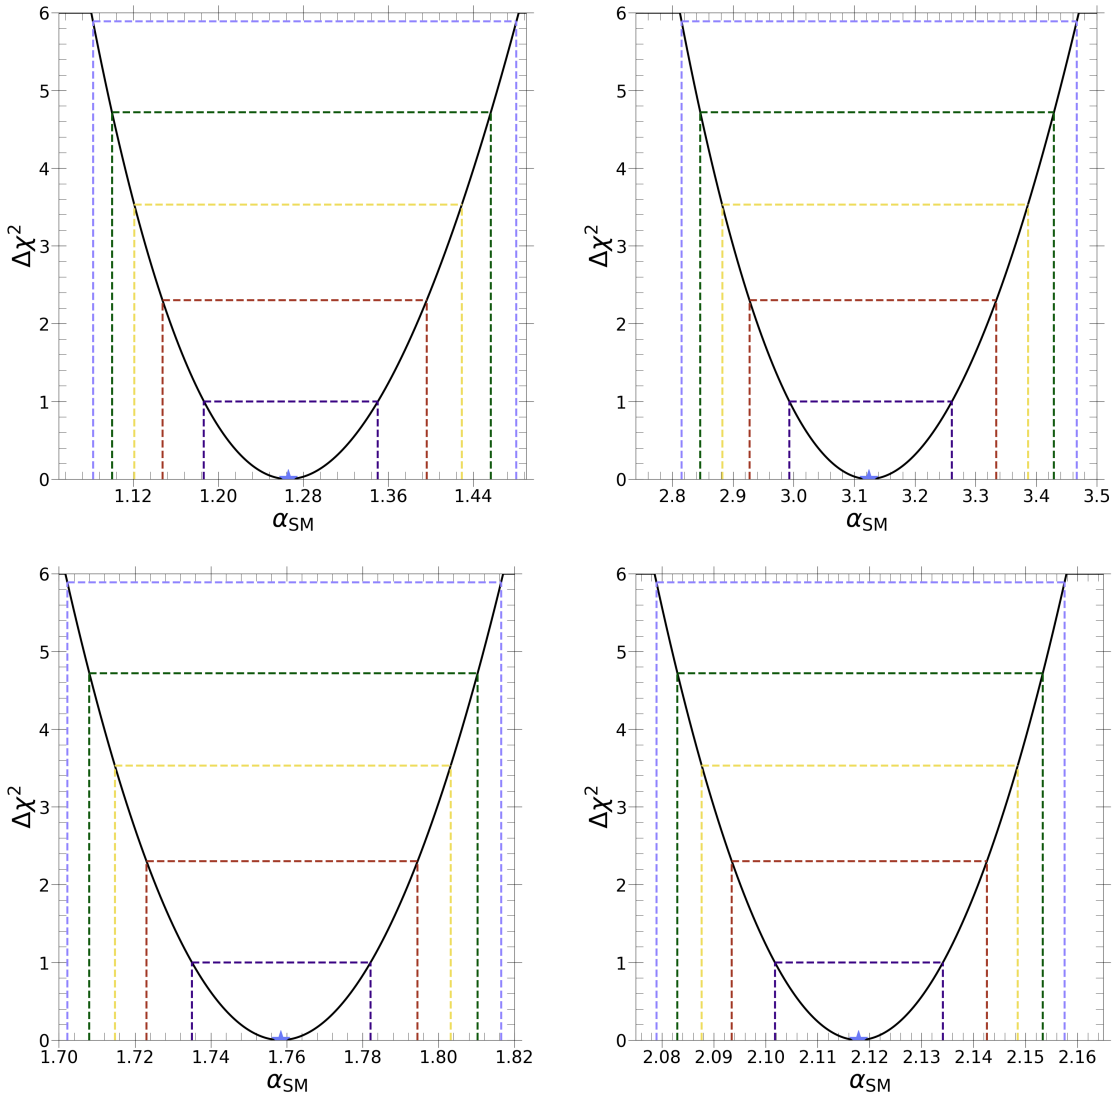


Figure 8.20: Errors resulting from the $\Delta\chi^2$ curve, for the single-parameter fit. C_1 sideband (top) and C_2 sideband (bottom), FHC (left) and RHC (right). Dashed lines indicate the $\Delta\chi^2$ for $N = 1, \ldots, 5$ (horizontal) and the coordinates of the resulting scale factors (vertical). For a one-parameter fit the inner-most dashed lines are appropriate.

Table 8.7. The error profiles for the best-fit points are shown in Figure 8.24, and the effect on the distributions in Figures 8.25-8.27. From the $N = 2$ fit, it is clear that there is still anomalous behaviour in the coherent contribution in the C_2 RHC sideband. Given that this contribution was the most problematic, a compromise between the large degree of overcoverage demonstrated in the $N = 5$ fit and the overoptimistic $N = 1$ fit, whose normalisation-only fit yielded poor agreement with data in C_2 RHC, is to allow the simultanous variation of the coherent component

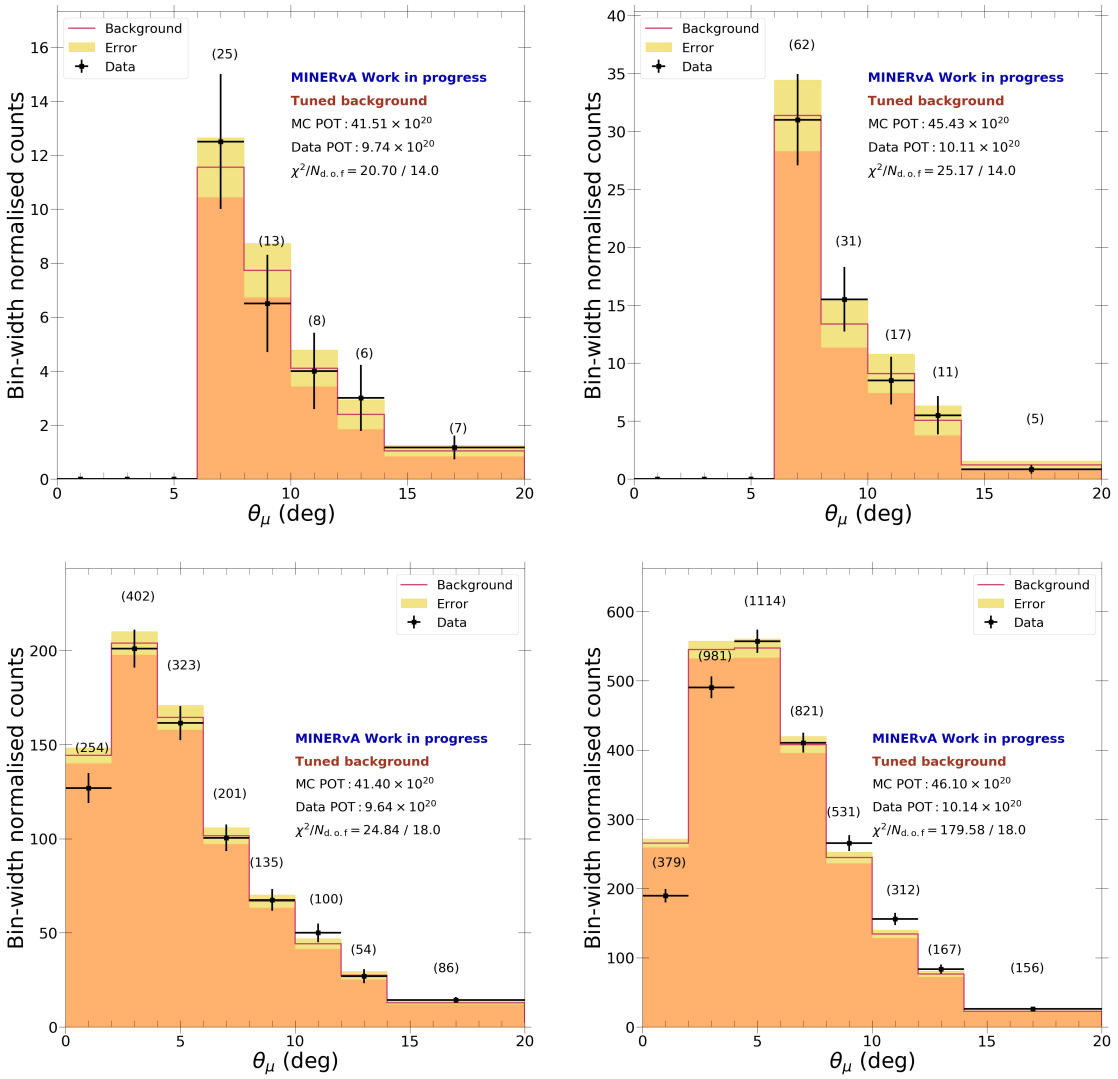


Figure 8.21: Single-parameter fit result on the MC sideband distributions for θ_μ. The differently-coloured bands on the top of each bin are the range of the MC uncertainty bin-by-bin. C_1 (top) and C_2 (bottom) sidebands are shown, FHC (left) and RHC (right). Number of data events in each bin in brackets.

with the other components. The large variation of α_{COH} between C_1 and C_2 in the RHC, which is an order of magnitude larger than for the α_{Other} component, is certainly puzzling; therefore, one has to accept the errors got from setting $N = 2$. These errors are likely to be in the right range for each individual process that enters the "Other" component; given the difficulty of modelling these processes in that region of phase space and the relative sparsity of data in the C_1 sideband, for the purposes of this thesis these errors will be taken as-is.

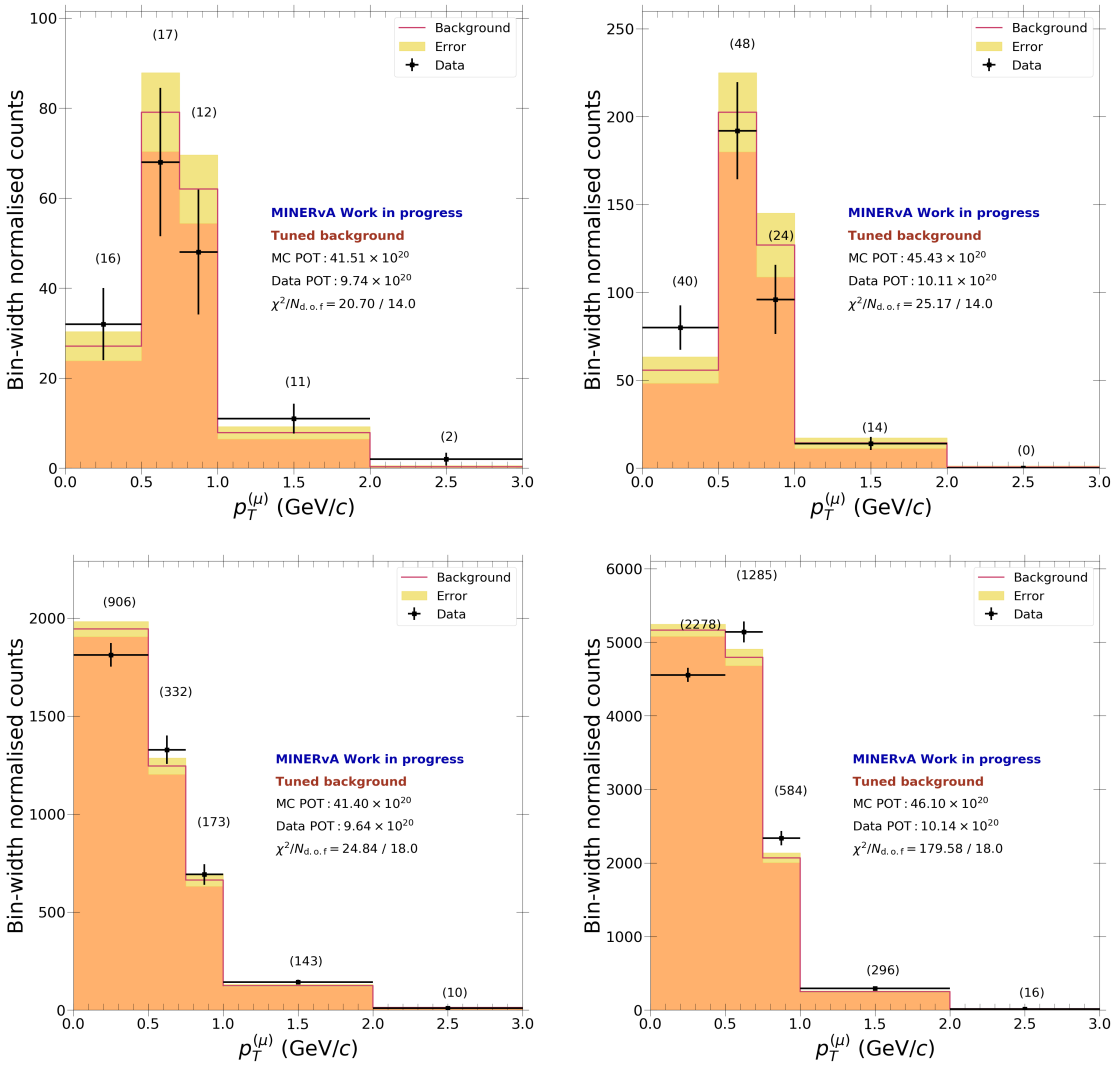


Figure 8.22: Single-parameter fit result on the MC sideband distributions for $p_{\mathrm{T}}^{(\mu)}$. The differently-coloured bands on the top of each bin are the range of the MC uncertainty bin-by-bin. C_1 (top) and C_2 (bottom) sidebands are shown, FHC (left) and RHC (right). Number of data events in each bin in brackets.

8.4.4 Conclusions

In the end, the interesting result from RHC is the great improvement in $\chi^2/N_{\mathrm{d.o.f}}$ with the best-fit point $\alpha_{\mathrm{COH}} = 0$, with the minimum chi-squared being reconciled with $\mathcal{O}(1)$. The best-fit is consistent with the full fit α_{COH}, with a far smaller error at 1σ. However, in the C_1 RHC sideband, the best-fit in both the two-parameter and full fits indicate there is significant *under*-, not over-, prediction of the CC Coherent contribution. The error is again far smaller in the two-parameter fit than

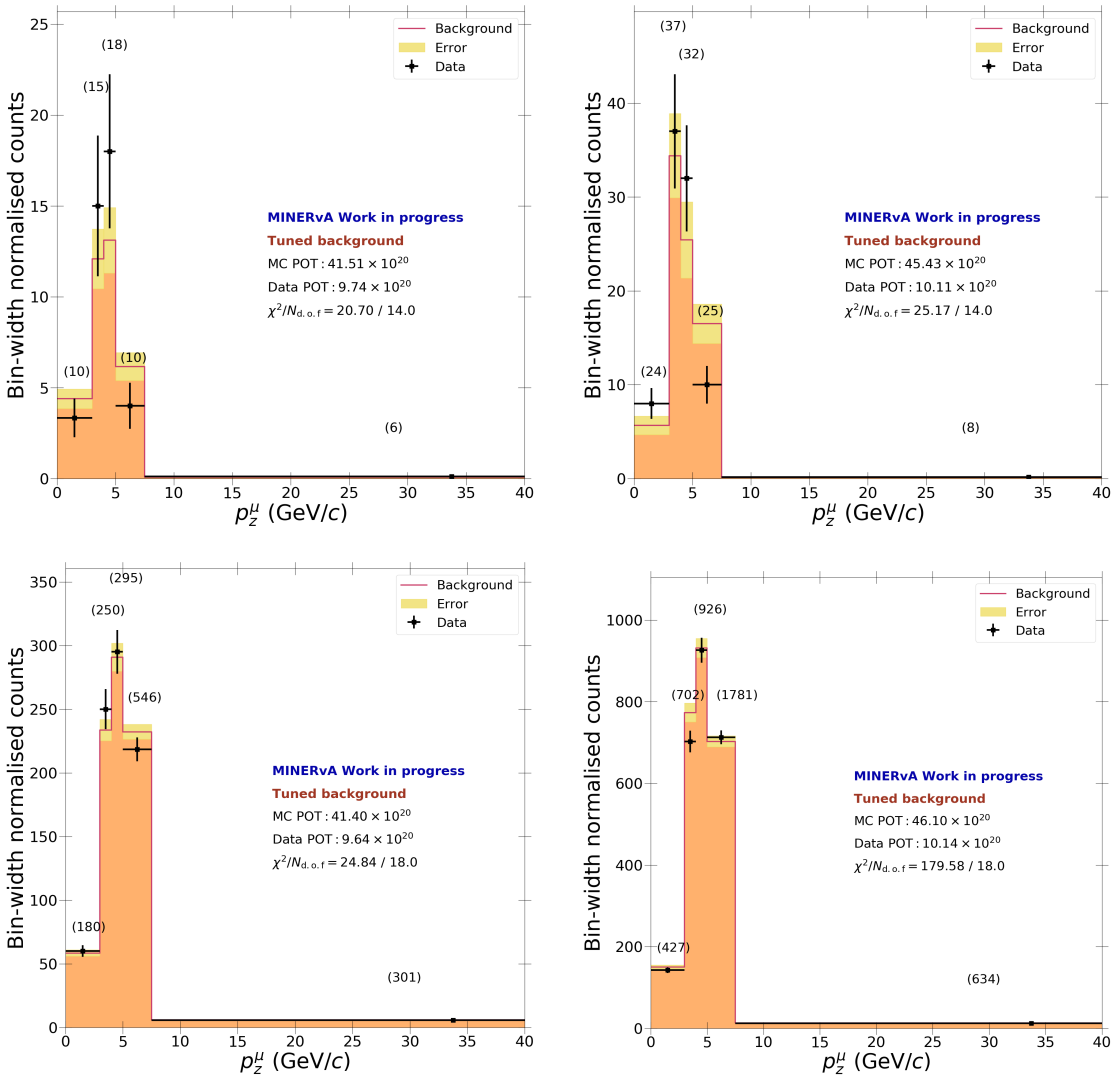


Figure 8.23: Single-parameter fit result on the MC sideband distributions for $p_z^{(\mu)}$. The differently-coloured bands on the top of each bin are the range of the MC uncertainty bin-by-bin. C_1 (top) and C_2 (bottom) sidebands are shown, FHC (left) and RHC (right). Number of data events in each bin in brackets.

in the full fit, and even the $\Delta\chi^2 = 5.89$ contour in the two-parameter fit would yield $\alpha_{\mathrm{COH}} \gtrsim 5$. In contrast, the FHC case is far more regular; the α_{COH} contribution is about $2-3$ in the C_1 sideband, and consistent with 1 in C_2. These conclusions are the same in the two-parameter fit as well as the full fit.

For the purposes of this thesis, the average value of α_{COH} between the C_1 and C_2 sidebands will be assumed, i.e. $\alpha_{\mathrm{COH}} = \frac{1}{2} \cdot \left(\alpha_{\mathrm{COH}}^{(C_1)} + \alpha_{\mathrm{COH}}^{(C_2)}\right)$, reflecting the transition from low- to high-energy pions. This is a very simplistic assumption, which can

Component	Best-fit	Low error ξ_-	High error ξ_+
C_1 FHC. Tuned $\chi^2/N_{\text{d.o.f}} = 18.99/13 = 1.46$; Untuned $\chi^2/N_{\text{d.o.f}} = 39.49/15 = 2.63$.			
COH	2.72	1.67	1.99
Other	1.18	0.15	0.16
C_1 RHC. Tuned $\chi^2/N_{\text{d.o.f}} = 9.01/13 = 0.69$; Untuned $\chi^2/N_{\text{d.o.f}} = 933.99/15 = 62.27$.			
COH	11.04	3.31	3.75
Other	2.69	0.24	0.25
C_2 FHC. Tuned $\chi^2/N_{\text{d.o.f}} = 6.37/17 = 0.37$; Untuned $\chi^2/N_{\text{d.o.f}} = 2085.06/19 = 109.74$.			
COH	1.02	0.26	0.26
Other	1.897	0.061	0.063
C_2 RHC. Tuned $\chi^2/N_{\text{d.o.f}} = 20.62/17 = 1.21$; Untuned $\chi^2/N_{\text{d.o.f}} = 11508.68/19 = 100.45$.			
COH	0.000	0.000	0.073
Other	2.291	0.026	0.027

Table 8.7: Double-parameter fit results. Errors are reported using `minos` under the appropriate error definition $\Delta\chi^2 = 2.30$, for two varying parameters. The best-fit point for RHC C_2 remains at $\alpha_{\text{COH}} = 0$ with significantly smaller error than the full fit.

definitely be improved on and significantly reduce the large uncertainty from the C_1 sideband. In summary, the following values for the two-parameter fit are carried forward into the next Chapter: The C_2 RHC is clearly still a bad fit, which means

Current	α_{COH}	α_{Other}	$C_1\chi^2/N_{\text{d.o.f}}$	$C_2\chi^2/N_{\text{d.o.f}}$
FHC	1.87 ± 0.99	1.54 ± 0.09	19.92/13	49.04/17
RHC	5.52 ± 1.88	2.49 ± 0.13	19.00/13	1081.51/17

Table 8.8: Final values of scaling factors used in Chapter 9.

the values obtained in this part should be taken with some caution.

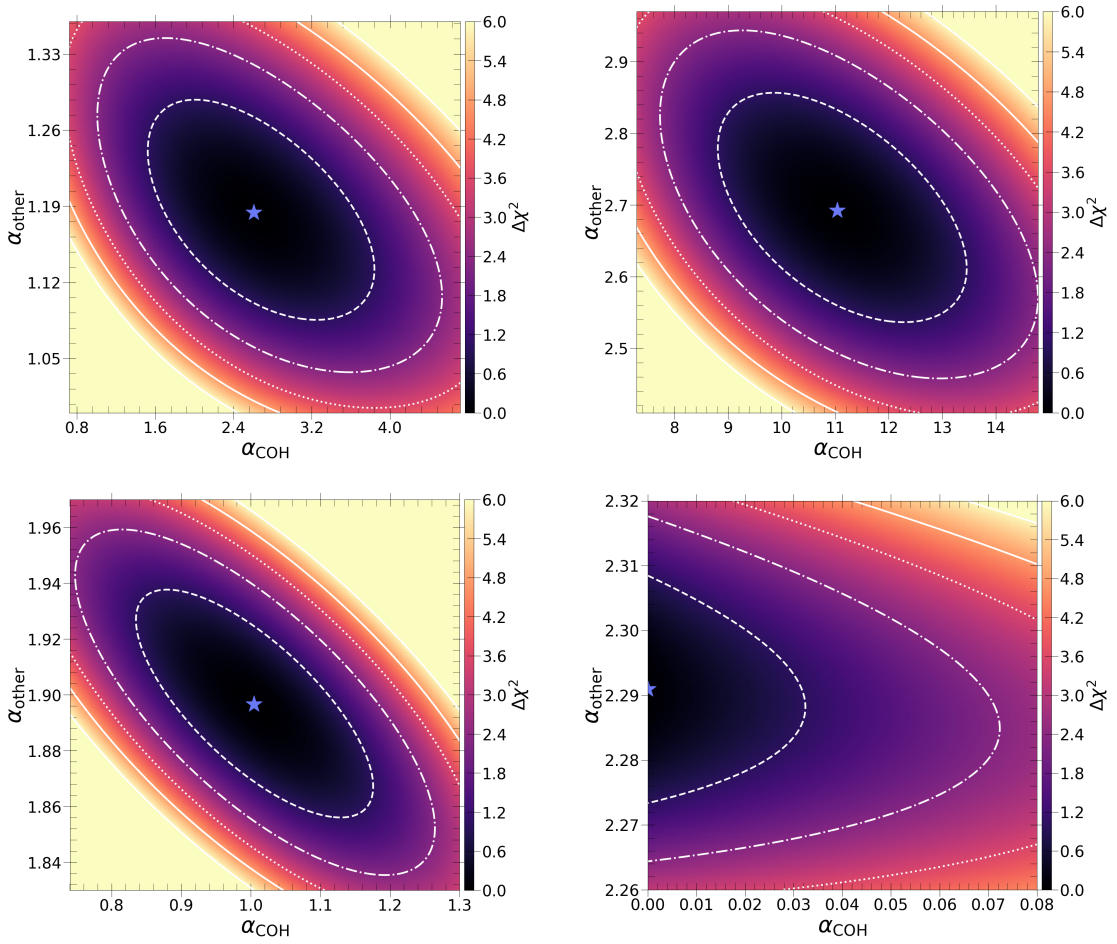


Figure 8.24: Errors resulting from the $\Delta\chi^2$ curve, for the double-parameter fit. C_1 sideband (top) and C_2 sideband (bottom), FHC (left) and RHC (right). The best-fit point is represented as a star. The various contours are, from inside to outside, the 1σ contours for $N = 1, \ldots, 5$. $\Delta\chi^2$ values greater than 6 are truncated. For a two-parameter fit the second-to-inner-most dashed lines are appropriate.

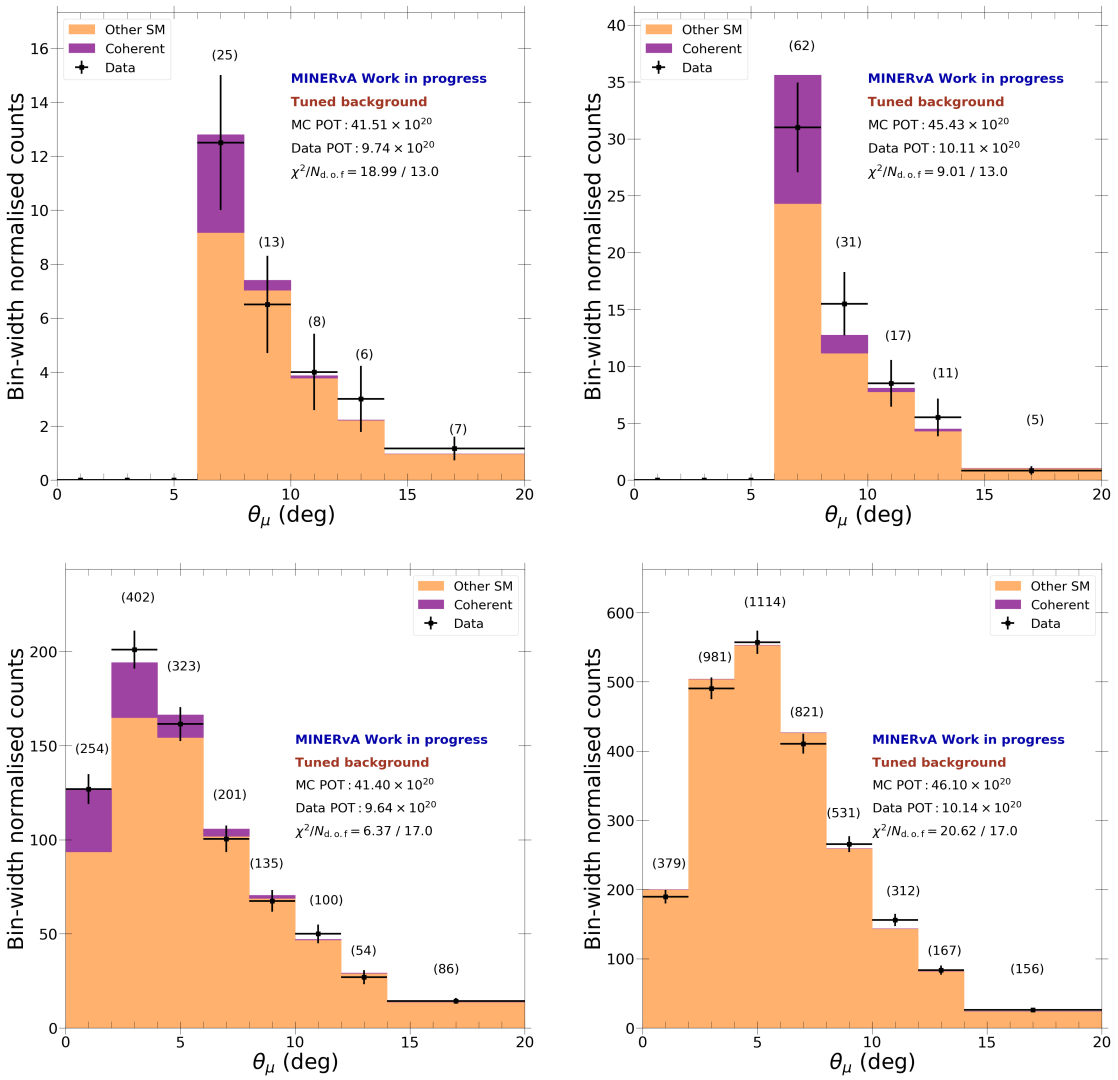


Figure 8.25: Double-parameter fit result on the MC sideband distributions for θ_μ. C_1 (top) and C_2 (bottom) sidebands are shown, FHC (left) and RHC (right). Number of data events in each bin in brackets.

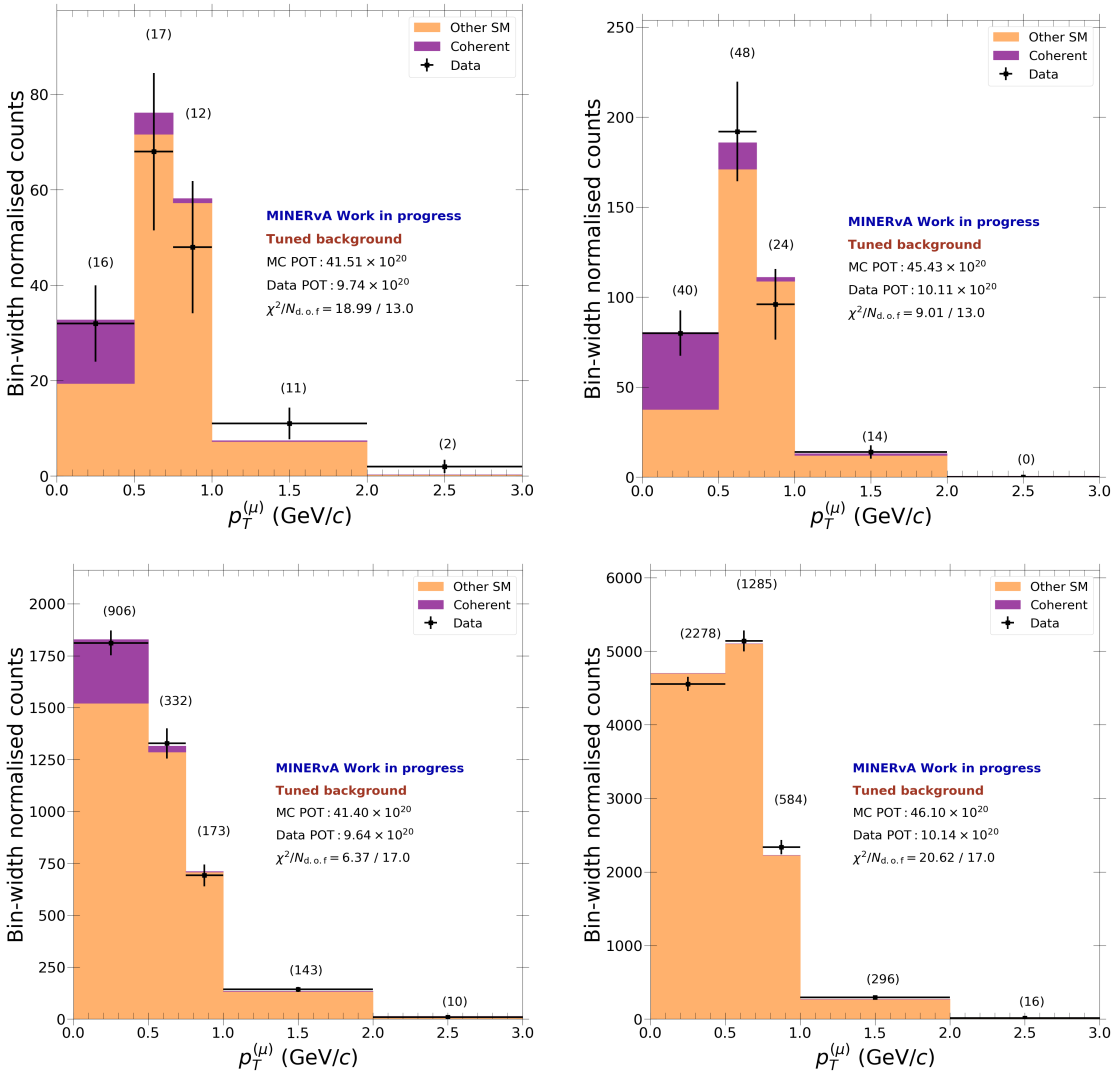


Figure 8.26: Double-parameter fit result on the MC sideband distributions for $p_{\mathrm{T}}^{(\mu)}$. C_1 (top) and C_2 (bottom) sidebands are shown, FHC (left) and RHC (right). Number of data events in each bin in brackets.

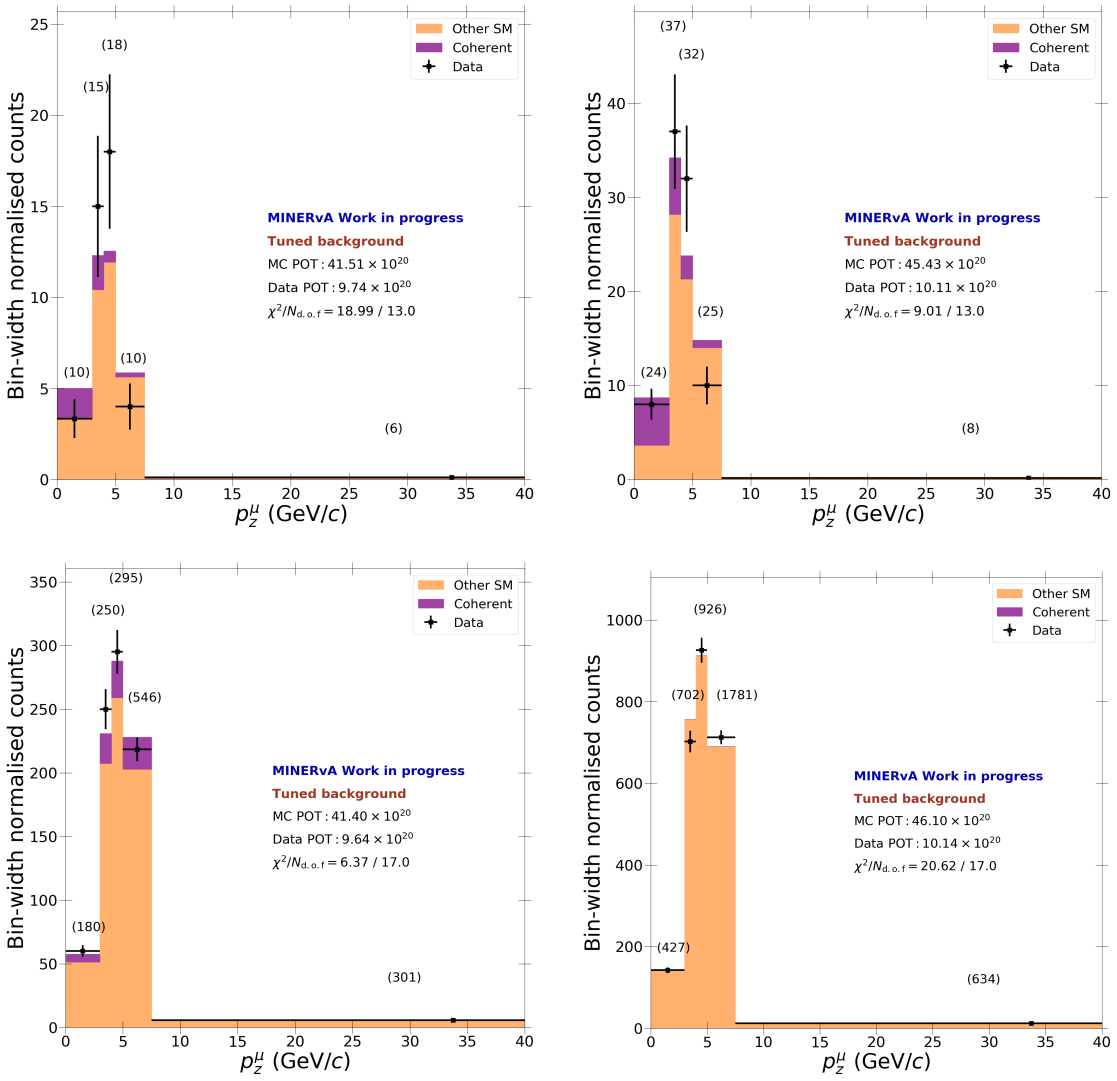


Figure 8.27: Double-parameter fit result on the MC sideband distributions for $p_z^{(\mu)}$. C_1 (top) and C_2 (bottom) sidebands are shown, FHC (left) and RHC (right). Number of data events in each bin in brackets.

9 Inference of signal and limit extraction

Contents

9.1 Introduction 200
9.2 Overview of inference procedure 201
9.3 Feldman-Cousins limit extraction 204
9.4 Fake data study for HNL 207
9.4.1 Fake data study 211
9.5 Setting exclusion limits on the HNL hypothesis 215

9.1 Introduction

The landscape for the HNL parameter space is set by stringent prior measurements (see Fig. 2.3), which motivates a small reference value $\alpha_0 = 8 \leftrightarrow |U_{\alpha 4}|^2 = 1-^{-8}$ in Equation 6.2. Based on the exposure of the analysed ME samples ($9.33 \cdot 10^{20}$ and $10.52 \cdot 10^{20}$ POT for FHC and RHC), at this reference value one would expect a small number of $N_4 \to \mu^{\mp}\pi^{\pm}$ to happen at MINERνA; see Fig. 9.1. Given this observation, it is therefore expected that the signal will be small ($\mathcal{O}(10)$ events), with respect to the irreducible background, which is expected under the previous analysis cuts to have $\mathcal{O}(500)$ events. Therefore, one must focus here on *limit setting* as the strategy for inferring results. For completeness, a strategy for reference cases corresponding to larger expected signal rates will be discussed in Section 9.4, to illustrate the significance in the variation of the scale factors extracted at the end of Chapter 8. This variation is quite large, and therefore enters as a significant systematic uncertainty.

The strategy for the analysis is to estimate the number of signal events in a mass-independent way using frequentist statistics; then by comparing with the number of expected signal events and the signal efficiency for each mass, to establish the mixing parameter using Eq. 6.2. Essentially, this Chapter will seek to establish the limits on new-physics signals first, before making any mass hypothesis, and then

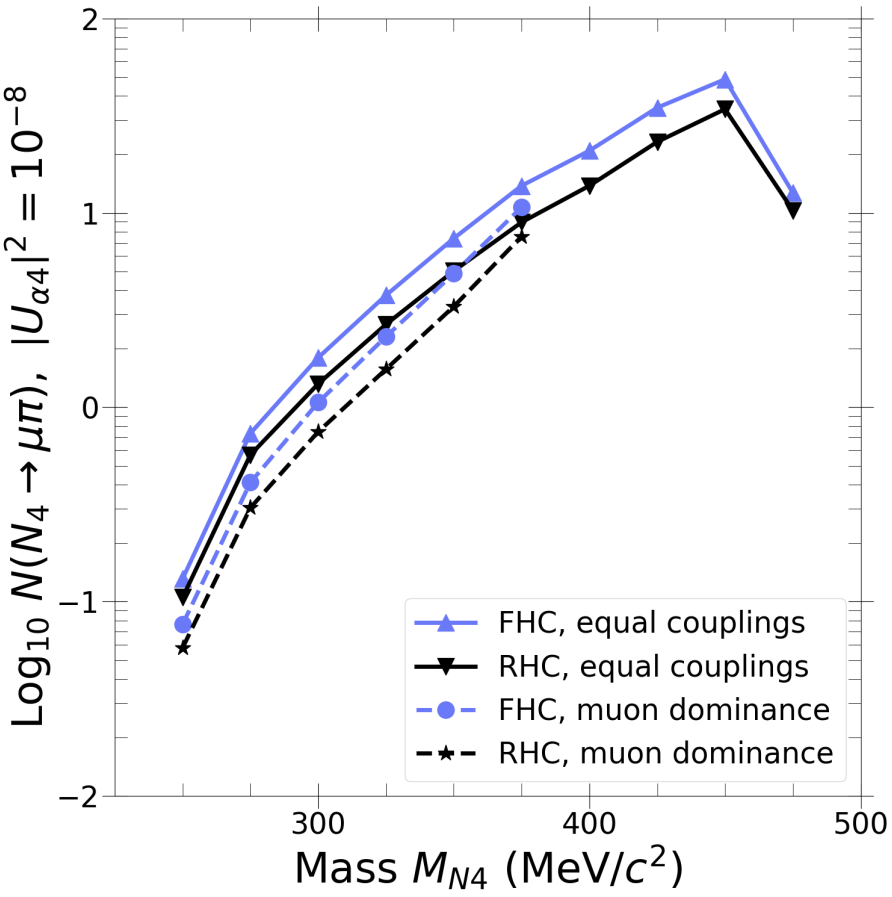


Figure 9.1: Expected number of $N_4 \to \mu\pi$ decays in MINERνA, for mixing element $|U_{\alpha 4}|^2 = 10^{-8}$. The reduction in expected decay numbers at $M_{\mathrm{N4}} = 475$ MeV/c^2 is due to the Lorentz boost of the parents completely dominating the neutrino kinematics close to kinematic threshold.

translate those limits into a constraint on the HNL parameter space, once for each mass simulated.

9.2 Overview of inference procedure

After characterising the background, the next step is to test the data in the fully selected sample against a suitably constructed hypothesis. It is customary to present a hypothesis as a choice between a *signal plus background*, or HNL, hypothesis (H_1) and a *background-only*, or null, hypothesis (H_0). These can be summarised as:

1. **HNL hypothesis**: The data n_{obs} can be explained by adding up the signal (s) and the background (b) distribution.

2. **Null hypothesis**: The data can be explained by just the background distribution, without any HNL signal.

There are various strategies one may pursue, in general, to test H_1 with data. Firstly, one must decide beforehand (based on the expected number of signal events with respect to the background) whether to pursue a "bump-hunting" or a "limit-setting" strategy. An important difference in the treatment of these two frameworks illustrates this need. In the case of a large excess, and the inference of its significance ("bump-hunting"), the framework is to construct a likelihood using the binned analysis variables; this is to test if the anomalous events correspond to a model hypothesis, that is, one seeks to confirm that H_1 is true. In the case of an insignificant number of signal events, however ("limit-setting"), one merely wishes to count the total number of events and juxtapose it with the expected number of background events, in order to infer an upper value on the signal rate in the detector. Essentially, one seeks to calculate the maximal value of one's model parameter, such that one cannot reject H_0. Limit-setting frameworks may also be distinguished based on the statistical framework that each utilises:

1. A *frequentist* approach, using the method of Feldman and Cousins [307], in which the likelihood of the signal+background hypothesis, given counts for data and expected background, is compared against the background-only hypothesis;

2. A *modified frequentist* approach, as defined by Read [308, 309] (also known as the CL_s technique), in which the ratio of likelihoods $\mathcal{L}_s := \mathcal{L}_{s+b}/\mathcal{L}_b$ is constructed as a function of a test statistic Q. The value of Q given data is then used to determine how "signal-like" (as opposed to how "background-like") the data are;

3. A *Bayesian* approach, where a posterior probability given the data is constructed from suitably chosen priors on the nuisance parameters and scaling factors controlling event rates, multiplied by the appropriate likelihood

function [159]. It is a useful way to combine likelihoods from multiple analyses, though Bayesian analyses have to consider a number of subtleties that do not arise in the frequentist context [310].

In this thesis, the Feldman-Cousins frequentist approach to set limits is utilised, which "automatically" constructs two-sided confidence intervals at arbitrary confidence level α. Taking a Poisson process with expectation value $\lambda = \mu + b$ and observed counts n_{obs},

$$\mathcal{P}\left(\lambda, n_{\text{obs}}\right) := \frac{\lambda^{-n_{\text{obs}}} e^{-\lambda}}{n_{\text{obs}}!} \ , \tag{9.1}$$

one can make a statement about the range of μ given b and n_{obs}:

$$\mathcal{P}\left(\mu_1 \leq \mu \leq \mu_2\right) \geq \alpha. \tag{9.2}$$

For the sake of brevity, in this thesis the focus will be on finding the upper end μ_2 of the Feldman-Cousins interval. This choice is motivated by the fact that a future measurement would still likely correspond to the limit-setting scenario, which is to say there would be a finite upper limit $\mu_2 > 0$, and a lower limit $\mu_1 = 0$. Hence, the equation 9.2 is recast into a one-sided interval, at the expense of overcoverage:

$$\mathcal{P}\left(\mu \leq \mu_2\right) \geq \mathcal{P}\left(\mu_1 \leq \mu \leq \mu_2\right) \geq \alpha. \tag{9.3}$$

From the parameter-space-*independent* upper limit μ_2 on the pure signal process, one can construct the parameter-space-*dependent* corresponding upper limit $\left(|U_{\alpha 4}|^2\right)_{\text{max}}^{(M_{\text{N4}},\text{coupling})}$, using the selection efficiencies calculated in Chapter 7 and the expected number of HNL decays.

The straightforward relation 6.2 can be used to extract the number of expected decays $n_{\text{exp}}^{(\rho)}$ in MINERνA from the calculated number of expected decays $n_{\text{exp}}^{(0)}$ at the reference value $|U_{\alpha 4}|^2 = 10^{-8}$, as

$$n_{\text{exp}}^{(\rho)} = \left(10^{\rho}\right)^2 n_{\text{exp}}^{(0)}, \quad \rho = -8 + \log_{10}\left(\frac{|U_{\alpha 4}|^2}{10^{-8}}\right). \tag{9.4}$$

Notice the square in $\left(10^{\rho}\right)^2$, which occurs from the scaling of both the HNL production and decay rates as 10^{ρ}. From the known exposure at MINERνA,

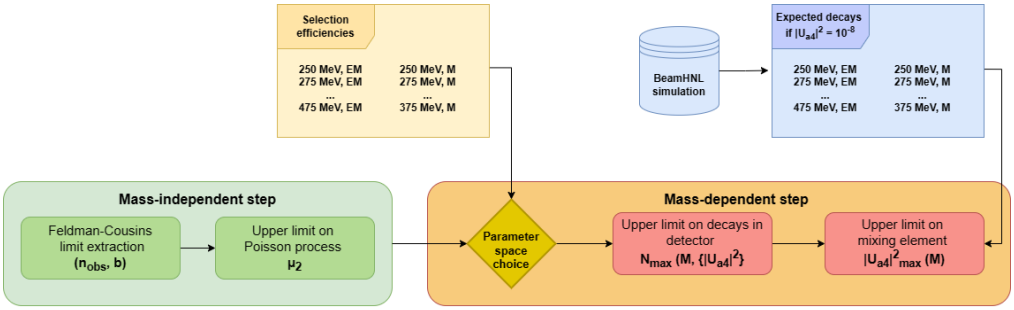


Figure 9.2: High-level overview of limit extraction procedure.

the number of expected decays in the detector with the real exposure can be obtained and with the corresponding efficiency transformed to the actual signal rate μ_0. Lastly, one can use the output of `BeamHNL` to estimate the POT necessary to produce a known number of signal events (cf. Figure 5.15), at $|U_{\alpha 4}|^2 = 10^{-8}$. With that information, one compares μ_0 to the number of expected decays if $|U_{\alpha 4}|^2 = 10^{-8}$ to extract the final deliverable $\rho_{\max}$,

$$\rho_{\max} = -8 + \frac{1}{2}\log_{10}\left(\frac{\mu_0}{n_{\exp}^{(0)}}\right). \tag{9.5}$$

This is to say, if the allowed signal rate μ_0 is 100 times greater than what one would have expected from the simulation $n_{\exp}^{(0)}$, $\rho_{\max} = -7$ and the quoted upper limit would be $|U_{\alpha 4}|^2_{\max} = 10^{-7}$; if μ_0 is 100 times smaller, $\rho_{\max} = -9$ and $|U_{\alpha 4}|^2_{\max} = 10^{-9}$. At a high level, the logic flow is encapsulated in Figure 9.2.

9.3 Feldman-Cousins limit extraction

The Feldman-Cousins two-sided interval is based on an *ordering principle* that is based on the ratio of two likelihoods. Essentially, one works in a parameter space where the number of background events B and the number of signal (not background) events S can vary, under the constraint that the experiment one has has observed N events. *A priori*, there is no relation between N, B, and S. For

each point in the parameter space (B, S), one orders the likelihood of various N:

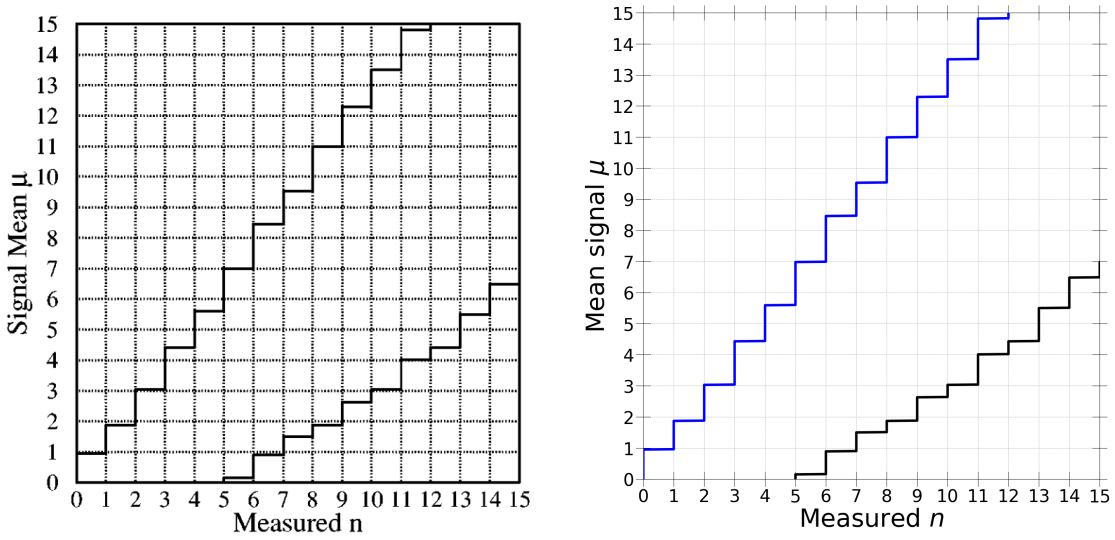
$$\begin{aligned} \mathcal{L}(x, y; \kappa) &:= \mathcal{P}(x + y, \kappa), \\ \mathcal{L}_0(N) &= \mathcal{L}(\widehat{S}, B; N), \\ \mathcal{L}_1(N) &= \mathcal{L}(S, B; N), \\ \mathcal{R}(N) &= \frac{\mathcal{L}_1(N)}{\mathcal{L}_0(N)}, \\ \widehat{S} &= \max(0, N - B). \end{aligned} \tag{9.6}$$

($\widehat{S}$ is the maximum-likelihood estimator for Poisson under the constraint $S \geq 0$, corresponding to presence of a true signal). One then constructs the reordered cumulative mass function

$$\mathcal{C}(r) = \sum_{i=0}^{r} [\mathcal{P}(S + B; N_i)] \tag{9.7}$$

such that $\mathcal{C}(r) \geq \alpha$. Then, for an experiment that knows the number N of events it did observe and the number B of background events it expects, a number of signal events S is included in the Feldman-Cousins interval at confidence level α if the term for N, $\mathcal{P}(S + B; N)$ is included in the sum of Eq. 9.7. Also, the fact that the map $\mathcal{R}$ is not monotonic ($i_1 < i_2$ does not mean $\mathcal{R}(i_1) < \mathcal{R}(i_2)$) automatically makes two-sided intervals.

To validate the implementation of the Feldman-Cousins ordering principle, first the reference case of $B = 3$ was plotted for various N, S at confidence level $\alpha = 0.9$ and compared with [307]. The result is shown in Figure 9.3. One notices the two curves enclosing an allowed space between them, excluding all on either side; this is the automatic two-sided interval construction. For each choice of $S(\mu)$, there exists a maximum (black rightmost curve) number of events $N^{(S)}_{\text{max}}$ and a minimum (blue leftmost curve) $N^{(S)}_{\text{min}}$ (this is zero for $\mu \leq 0.95$). One can interpret this equivalently by freezing a choice of N (i.e. an experiment has already observed N events in its selected signal); in this picture, the upper limit μ_2 is the blue uppermost curve. In the case where the background is expected to be large, one may make the assumption that the lower limit μ_1 will be zero (cf. Fig. 9.3), which means the Feldman-Cousins interval will translate into an exclusion of all rates above μ_2. Furthermore, an



(a)Original plot from [307] (b)Validation plot

Figure 9.3: Validation of the Feldman-Cousins ordering implementation. The reference case of $B = 3, \alpha = 0.9$ is plotted for varying N (horizontal) and S (here denoted as μ for compatibility with the original notation, vertical). Note the two-sided interval, with the upper (blue) and lower (black) curve delimiting the allowed range of values such that Equation 9.2 is satisfied.

extension to the $\alpha = 0.9$ case was performed to include the standard $1, 2, 3$-sigma confidence levels to show the widening of the confidence intervals as α gets larger; this is intuitive in a frequentist sense. The result is shown in Figure 9.4. Further, the plot of the upper end μ_2 from [307] was regenerated and compared with the results from this thesis's implementation. A pathology connected to Eq. 9.7 is first corrected for; due to the discreteness of $\mathcal{C}$, a distinctive "sawtooth" pattern emerges if the ranking scheme is used "as-is"; see Fig. 9.5 (b). This sawtooth is due to the degree of overcoverage not being constant; as the background varies, the uppermost S such that a particular choice of N is the maximum allowed choice does not monotonically decrease, as one would expect if the overcoverage was constant. Intuitively, for a particular choice of N, an increasing background puts *stronger* constraints on the inferred signal S; the curves of given N should thus be monotonically *decreasing*. Feldman and Cousins noted this in their paper, and treated the pathology, which is visible in the horizontal edges in Fig. 9.5 (a). In Figure 9.6, the validation plot has had the pathology treated by enforcing monotonicity of $\mathcal{C}$, as in [307].

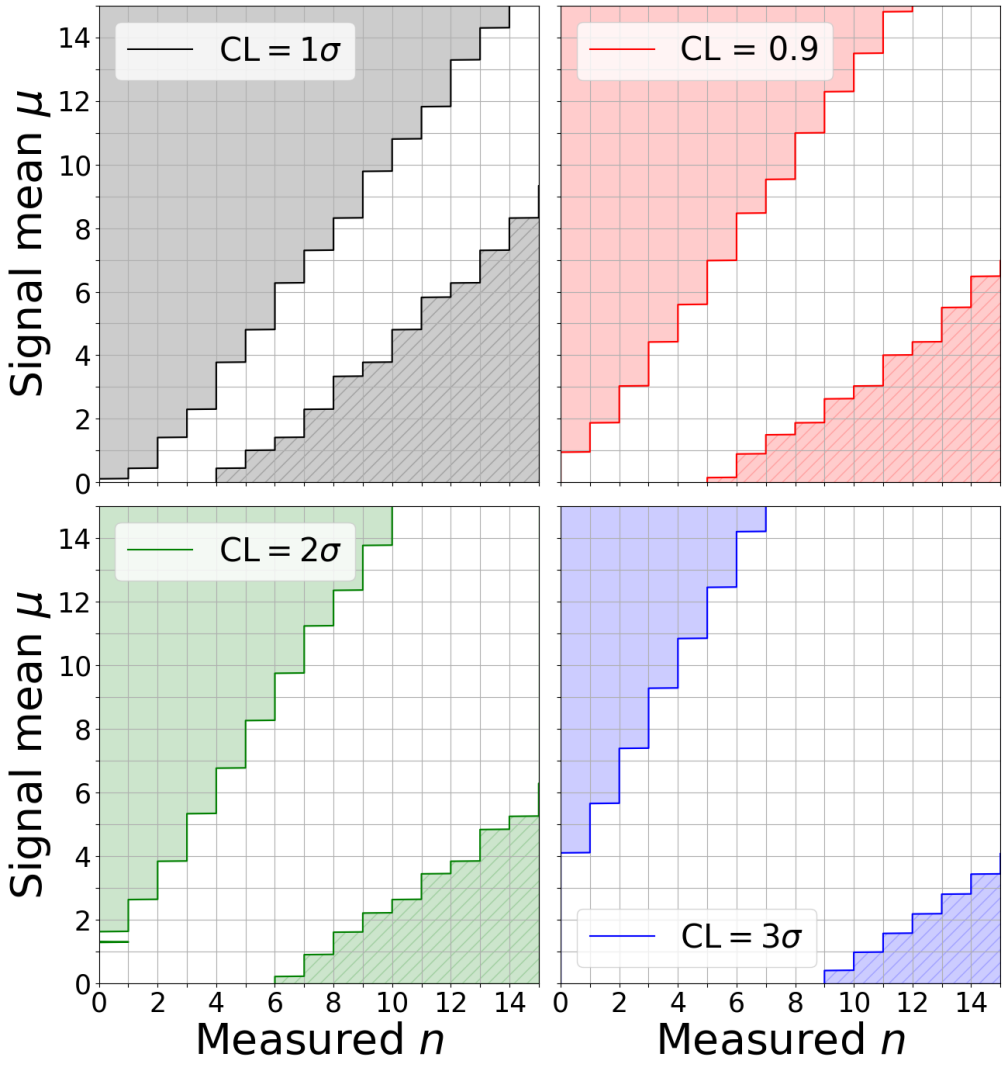
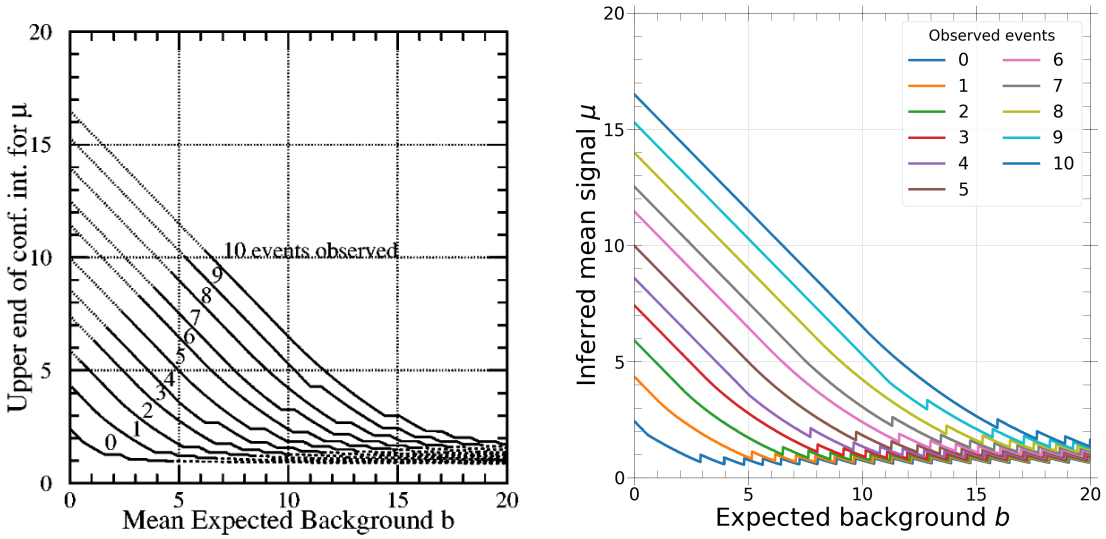


Figure 9.4: Feldman-Cousins confidence intervals at various choices of α.

9.4 Fake data study for HNL

Neglecting for the moment the constraints on the HNL parameter space from Fig. 2.3, one may wonder what the appropriate strategy for searching for an HNL-like excess might be for a large signal. For example, if one were expecting a large number of events from signal compared to the background (such as might happen if an HNL with mixing $|U_{\alpha 4}|^2 = 10^{-5}$, say), then this would be the appropriate method; A fake data study using the average of the scale factors from the two-parameter study in the $C_{1,2}$ sidebands (Table 8.7) was performed to investigate three benchmark excess scenarios:



(a) (b)

Figure 9.5: (a): The original plot from [307] where the values of μ_2 from various numbers of N are shown. Horizontal edges are visible in the plot, where the monotonicity of $\mathcal{C}$ has been enforced. (b): A Poissonian pathology is evident in the direct computation of $\mu_2 = \mu_2(B, N)$; this is due to varying degrees of overcoverage from Eq. 9.7 leading to non-smooth changes in the maximum number of events compatible with a given combination of $B, \mu = S$.

- *Small signal*: 100 events;
- *Medium signal*: 300 events;
- *Large signal*: 500 events.

Each of these scenarios was considered in both FHC and RHC. The fake data is obtained by adding together a background and a signal component in each bin. The background component is drawn from a Gaussian distribution with mean and variance the bin content and error, respectively, and rounded to the nearest integer. The bin content and error are drawn from the COH and Other components of the MC scaled by the scale factors derived in Chapter 8 (see Table 8.8). The signal component is drawn from a randomly selected HNL sample, set such that the desired number of events is added in each distribution.

Because one does not *a priori* know if the data is compatible with any of the HNL signals, or which HNL mass the data is most compatible with, a combined log-likelihood ratio (LLR) $\mathcal{R}$ is calculated as follows: First, the *global likelihood*

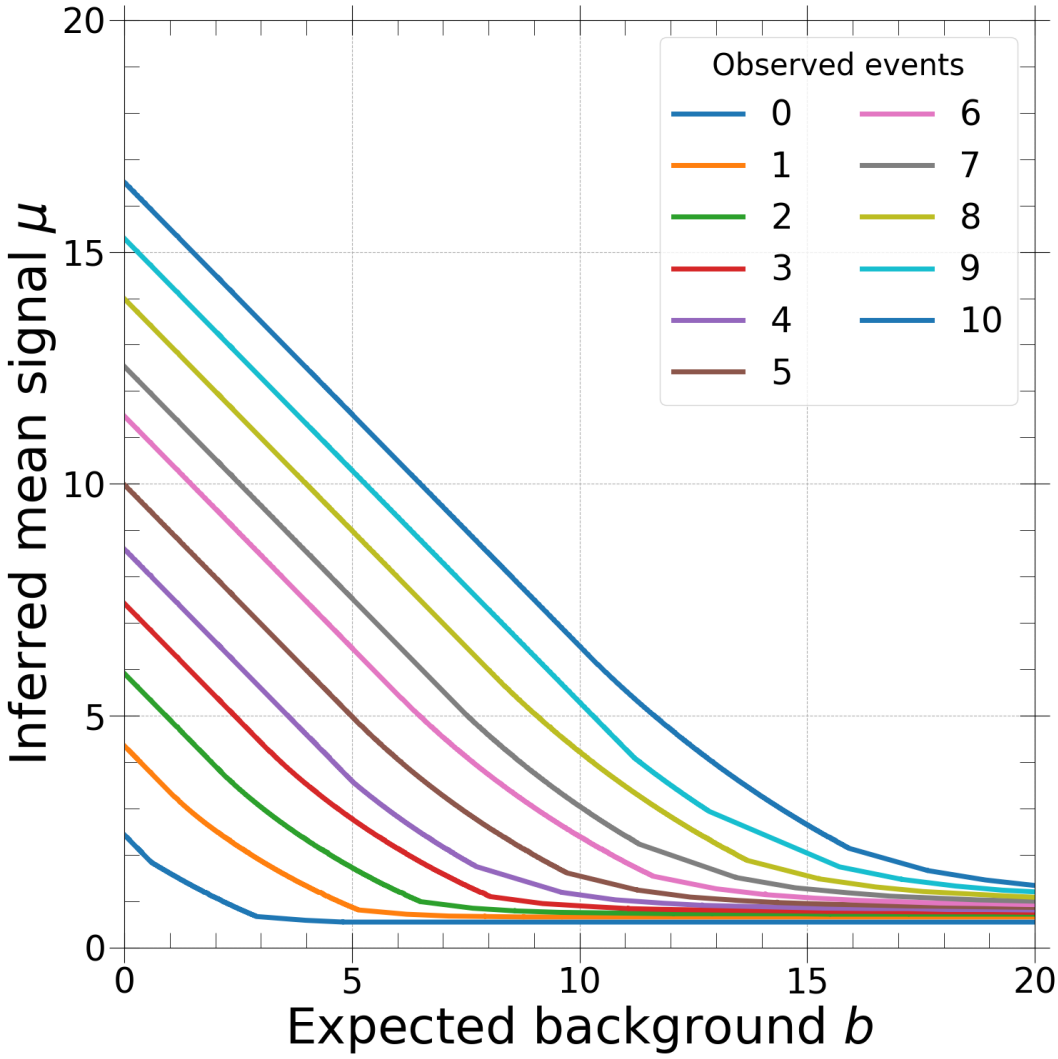


Figure 9.6: The same plot as Fig. 9.5 (b), but with the pathology treated by enforcing the monotonicity of $\mathcal{C}$.

of a mass hypothesis labelled by υ, with signal contribution $(s_i^\upsilon \pm \delta s_i^\upsilon)$ over all available distributions with bins i is defined as

$$\begin{aligned} \mathcal{L}_\upsilon\left(\{s_i, \delta s_i\}\right) &= \prod_i \left[\mathcal{L}_\upsilon\left(s_i, \delta s_i\right)\right] \\ &= \prod_i \left[\frac{1}{\sigma_i\sqrt{2\pi}} \exp\left(-\frac{\left[n_i - (s_i + b_i)\right]^2}{2\sigma_i^2}\right)\right], \\ \sigma_i &= \sqrt{(\delta s_i)^2 + (\delta b_i)^2}, \end{aligned} \tag{9.8}$$

where $\delta s_i^{(\upsilon)}$ is the uncertainty in the signal distribution's bin content s_i^υ for a mass hypothesis υ, and δb_i is the corresponding uncertainty in the *total* background distribution. This uncertainty in the background bears commenting on, as there

are two sources of error contributing to it; the statistical+systematic uncertainty of each background process as obtained from the multi-universe method, and the systematic uncertainty in the scale factors of each background process. Because there are large uncertainties in the scale factors $\{\alpha\}$ (particularly in α_{COH}), the significance what naïvely appears to be a significant excess is strongly diluted.

An indicative example with a toy model: one generates fake data by overlaying 100 events drawn from the $M_{\mathrm{N4}} = 250$ MeV/c^2 HNL sample on top of the MC background with nominal scale factors, such that the data exactly matches the background in bins where $s_i = 0$. The distribution of such fake data in θ_μ is shown in Figure 9.7. The naïve expectation would be 3.4σ in the first bin; when the spread in the scale factors is taken into account, however, the excess completely disappears; in Figure 9.8 (a), the naïve excess is evident, and in Figure 9.8 (b), the zero line lies within the red band which indicates the uncertainty due to the scale factors. Therefore, for each bin, the proper value δb_i must be taken as the direct sum

$$\delta b_i = \left(\sum_j \left[\alpha_j \cdot \delta b_i^{(j)} \right]^2 \right)^{1/2} \oplus \left(\sum_j \left[\delta\alpha_j \cdot b_i^{(j)} \right] \right), \tag{9.9}$$

where $\delta b_i^{(j)}$ is the uncertainty of the j-th MC component *with nominal scaling factors* and $\delta\alpha_j$ is the ξ_+ for each component from Table 8.7. In this way, each SM source indexed by j has a contribution in the bin i, which is within the interval $\left[(\alpha_j - \delta\alpha_j) \left((b_i^{(j)} - \delta b_i^{(j)}) \right), (\alpha_j + \delta\alpha_j) \left((b_i^{(j)} + \delta b_i^{(j)}) \right) \right]$.

For this scenario to be valid, there must exist at least one signal hypothesis υ for which there is a meaningful excess given the uncertainties in the SM background. The signal hypothesis that looks most like the fake data is therefore first searched for using Alg. 9, and the normalisation c_υ of that signal fit to the fake data, using the nominal MC prediction; the uncertainty δc_υ is then applied to estimate the likelihood in each bin. From these, by Eq. 9.8 the combined LLR is constructed according to Algorithm 9. One can translate the LLR to a confidence level (p-value) using Wilks' theorem [306, 311]:

$$\mathcal{P}\left(H_0 \text{ true} \leq 1 - \alpha\right) = \mathcal{P}\left(\chi^2_{d=1} \geq 2 \cdot \mathcal{R}_\alpha\right), \tag{9.10}$$

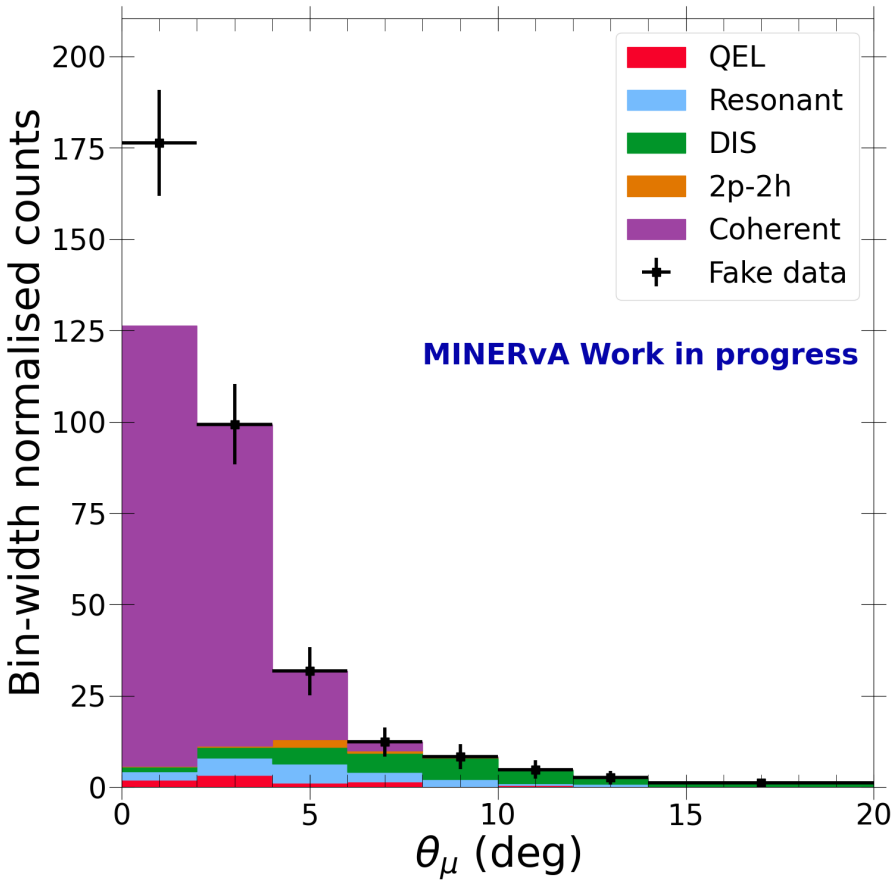


Figure 9.7: Toy fake data model: FHC SM background at nominal scale factors overlaid with 100 events from the $M_{\mathrm{N4}} = 250$ MeV/c^2 HNL sample. The distribution is over the muon angle with respect to beam direction, θ_μ

since there is only one more free parameter under any H_1 than in H_0 (the magnitude of the signal). Using this theorem, one can immediately read off the critical values of the LLR, $\mathcal{R}_\alpha$, that give the following confidence levels from Table 9.1.

9.4.1 Fake data study

For each representative scenario as outlined above, first the fake data set $\{d_i\}$ is generated using Algorithm 10.

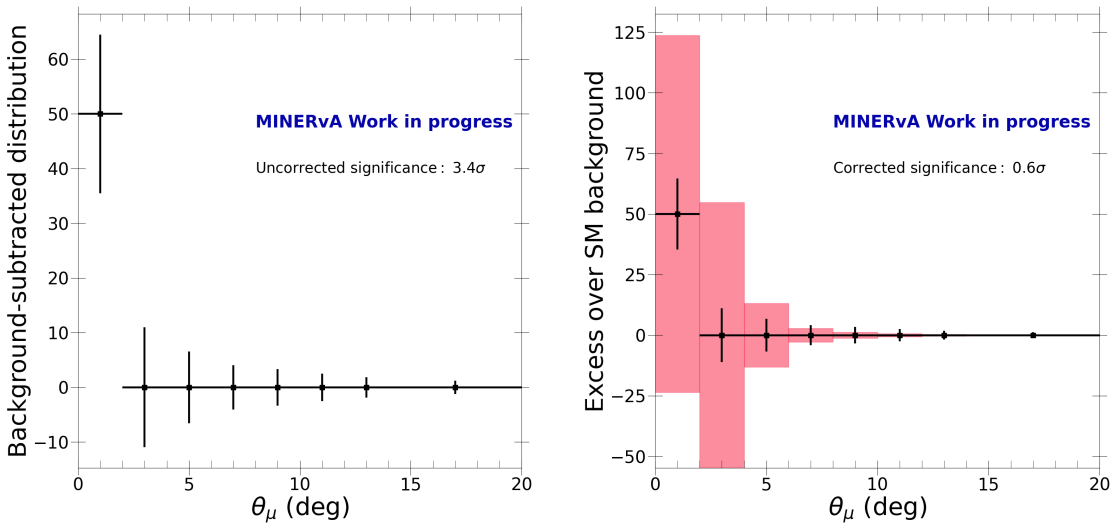
After the fake data have been generated, a fit is performed for each available sample, and the best-fit signal scale factor c_υ is obtained, which is then used to generate the mass-dependent prediction $n_i^{(\upsilon)}$ for each bin. As a sanity check, it was verified that, when a mass hypothesis υ was chosen to be added onto the nominal background such that each bin content was exactly the central value,

Data: Distributions with MC prediction $\left\{b_i^{(a)} \pm \delta b_i^{(a)}\right\} \oplus \{\delta a\}$ and observed data $(n_i \pm \delta n_i)$ in each bin i; Signal distributions $(s_i^{(\upsilon)} \pm \delta s_i^{(\upsilon)})$ for $\upsilon \in \{\upsilon_j\}$ mass points
Result: Set of LLR, $\{\mathcal{R}\}_\upsilon$
Let $\boldsymbol{\mathcal{L}} = (\mathcal{L}_0, \mathcal{L}_{\upsilon_0}, \dots)$ be the vector of likelihoods from null hypothesis ($\mathcal{L}_0$) and signals $\upsilon \in \{\upsilon_j\}$;
For each mass hypothesis $\upsilon \in \{\upsilon_j\}$ perform global fit to data and extract the normalisation $\{c_\upsilon \pm \delta c_\upsilon\}$;
Set $\mathcal{L} = 1 \; \forall \mathcal{L} \in \boldsymbol{\mathcal{L}}$;
foreach *bin* i **do**
  //Calculate null likelihood for bin;
  $\mu_i^{(0)} = n_i - b_i$;
  $\mathcal{L}_0 \mapsto \frac{1}{\delta b_i \sqrt{2\pi}} \exp\left(-\frac{(n_i - b_i))^2}{2\delta b_i^2}\right) \mathcal{L}_0$;
  foreach υ **do**
    // Calculate the mean and error;
    $\mu_i^{(\upsilon)} = n_i - b_i - c_\upsilon \cdot s_i^{(\upsilon)}$;
    $\sigma_i^{(\upsilon)} = \sqrt{(\delta b_i)^2 + \left(c_\upsilon \cdot \delta s_i^{(\upsilon)}\right)}$;
    // Calculate the likelihood of hypothesis;
    $\mathcal{L}_\upsilon \mapsto \frac{1}{\sigma_i^{(\upsilon)}\sqrt{2\pi}} \exp\left(-\frac{\left(n_i - b_i - c_\upsilon \cdot s_i^{(\upsilon)})\right)^2}{2\left(\sigma_i^{(\upsilon)}\right)^2}\right) \mathcal{L}_\upsilon$;
  end
end
// Find the maximal LLR and attached hypothesis;
return υ_{best}, $\mathcal{R}_{\text{best}} = \ln\left(\mathcal{L}_{\text{best}}/\mathcal{L}_0\right)$

Algorithm 9: Excess hunting algorithm.

α	$\mathcal{R}_\alpha$
0.5	0.227468
1σ	0.5
0.9	1.35277
0.95	1.92073
2σ	2
3σ	4.5

Table 9.1: Critical values for LLR at various CLs.



(a) (b)

Figure 9.8: Effect of scale factor uncertainty on discovery potential, using a toy fake data model over the muon angle with respect to beam direction, θ_μ. (a): Naïve excess of 3.4σ with nominal scale factors. (b): The excess disappears when the uncertainty of the background scaling factors (red bands) is considered.

Data: Nominal MC prediction $\left\{b_i^{(a)}\right\}$ with MC errors $\left\{\delta b_i^{(a)}\right\}$, and nominal scale factors $\{a\}$. Number of excess events N
Result: Fake data f_i
// Construct "fake background";
foreach *reconstructed variable bin* i **do**
  Set nominal value $B_i = \sum_j \alpha_j \cdot b_i^{(\alpha_j)}$;
  Set nominal error $\delta B_i = \left(\sum_j \alpha_j^2 \cdot \left(\delta b_i^{(\alpha_j)}\right)^2\right)^{1/2}$;
  Vary nominal value by Gaussian, $\delta e_i \sim \mathcal{N}\left(B_i, 3\delta B_i\right) - B_i$;
  $e_i = B_i + \delta e_i$;
end
// Set signal;
Randomly select one HNL mass from available samples;
For each reconstructed variable X obtain the signal shape, g_X;
while $n_{\text{exc}} < N$ **do**
  For each X, from the bins $\{i\}$ in the appropriate range add one event at random $q_i \sim g_X$;
  $n_{\text{exc}} + +$;
end
return $f_i = e_i + q_i$;

Algorithm 10: Fake data generation algorithm. 3σ errors are thrown, so as to capture 99% of the possible values of the background in each bin.

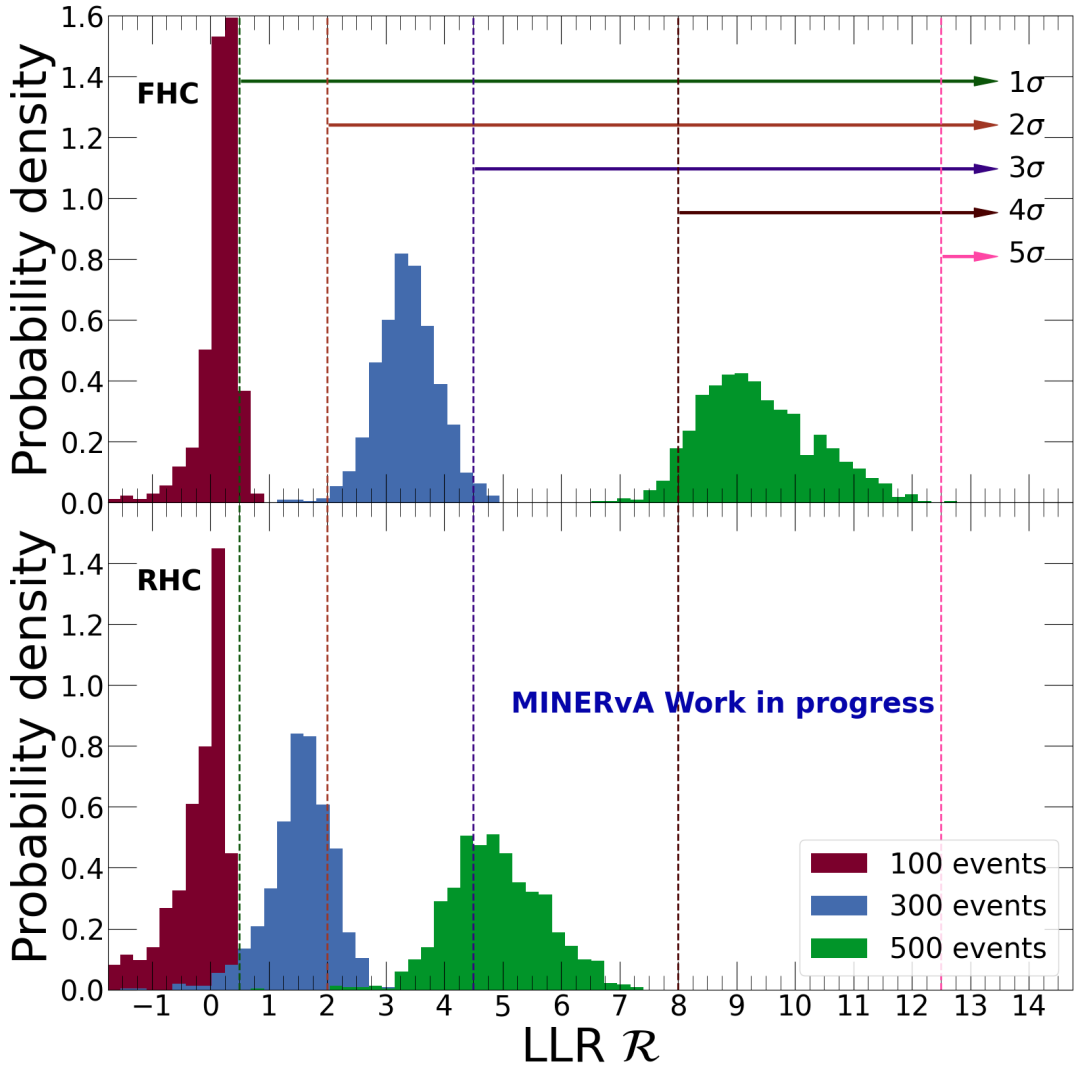


Figure 9.9: Discovery potential of benchmark excess scenarios, FHC (top) and RHC (bottom). Each coloured band, left to right, corresponds to 100, 300, and 500 excess events. Dashed lines indicate critical LLR values for $N = 1, \ldots, 5\sigma$ C. L. The reduced power of the RHC samples, compared to FHC, is caused by the larger errors in α_{COH}.

Algorithm 9 correctly identified hypothesis υ to be the most data-like, and extracted a scale factor consistent with 1. From this fit, the mass-dependent fake data $n_i^{(\upsilon)}$ are fed into Algorithm 9 and the LLR $\mathcal{R}_\upsilon$ is calculated. To estimate the discovery potential associated with each of the three benchmark scenarios, 1000 pseudoexperiments were thrown for each scenario, for FHC and for RHC; the resulting distributions in $\mathcal{R}$ are shown in Fig. 9.9. For ease of reference, the critical points for $N = 1, \ldots, 5\sigma$ discovery thresholds are plotted alongside the distributions. From Figure 9.9, it becomes apparent that, in order to claim a significant HNL-like

signal at 3σ, one would need at least 300 events of excess over the SM background. This is a very large signal indeed; one must turn toward the limit-setting procedure discussed in Section 9.3 as the more suitable approach for small signals.

9.5 Setting exclusion limits on the HNL hypothesis

Returning to the existing constraints on parameter space and Fig. 9.1, it becomes apparent that the appropriate inference strategy is not the "bump-hunting" strategy described in the previous Section, but rather to *setting exclusion limits*. In brief, the method consists of calculating the highest value of $|U_{\alpha 4}|^2$ that allows for the rejection of H_1. The background contamination determines the strength of these limits. A seemingly counter-intuitive feature of the Feldman-Cousins construction must be addressed here: the larger the background, the more restricted the inferred Poisson signal, and the smaller the upper limit on the parameter space for a given mass. This serves to highlight a subtle point about interpreting this scenario: one must be cautious to not interpret this feature thoughtlessly. The small-background case is not, *per se*, applicable for limit-setting, as it would be more appropriate there to employ the procedure used in the discovery scenario outlined above; to put it differently, given a large number of data points over a small background, one should interpret the data as being compatible with the test hypothesis H_1. In the small-signal scenario, one cannot distinguish any excess from a fluctuation in the background; that is, one's data *reject* the test hypothesis, and accept H_0. With this in mind, the figure of merit is the "flat" region (or *floor*) of the Feldman-Cousins construction. Upward fluctuations of the background around the nominal value (towards the right on Figure 9.10) will bring the inferred signal closer to its floor limit value μ_{D} (down on the Figure), whereas downward fluctuations will bring it closer to the zero-background value μ_0. In Figure 9.10, where for two cases of events seen, $N = 0$ and $N = 10$, the "critical point" at which the upper end of the Feldman-Cousins interval becomes flat with respect to changes in the expected background. An analysis will have a fixed number of events in its signal region,

so it will lie on one curve. Higher background moves an experiment further along the right on the plot, which corresponds to a smaller value for μ.

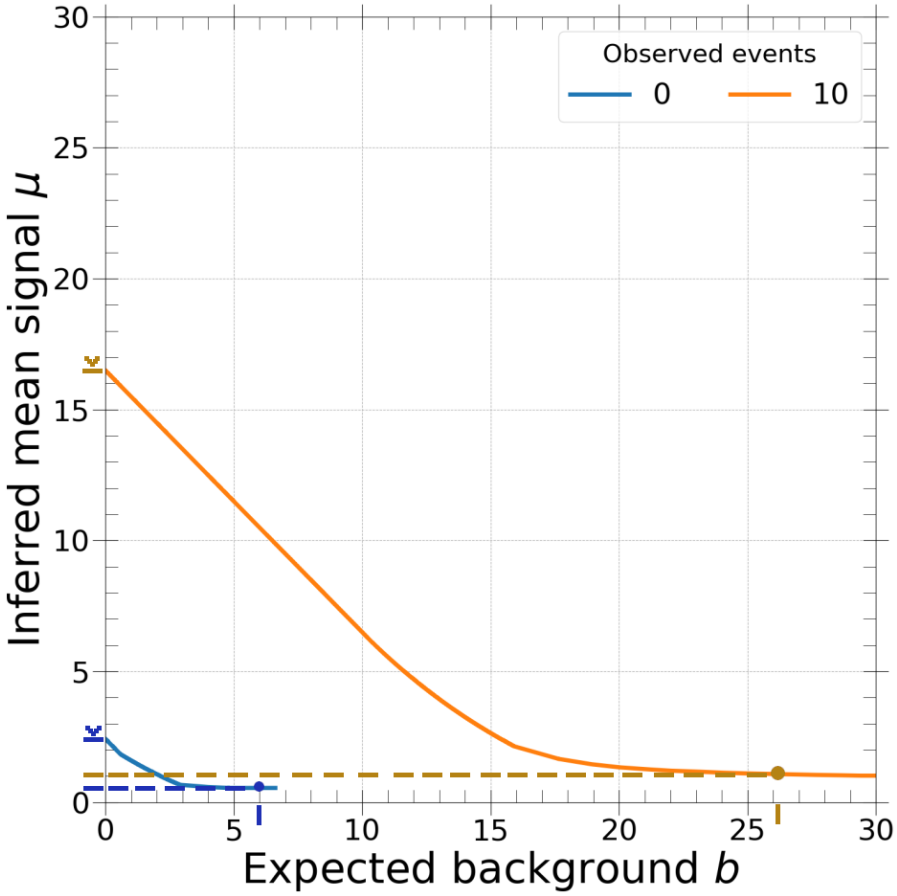


Figure 9.10: Diagram showing critical points for the upper end of the Feldman-Cousins interval, for 0 and 10 events observed. The curves become degenerate (flat) after the points indicated in colour; the coordinates of these points are $(b_{\mathrm{D}}, \mu_{\mathrm{D}})$. Also shown are the intercepts μ_0, indicated on the y-axis by coloured tickmarks with triangular oversets.

Evidently, the driver for limit-setting, in the most conservative case, must be the value $\mu_+(b - \delta b)$, which corresponds to the largest possible $|U_{\alpha 4}|^2$ compatible with the rejection of H_1 (equivalently, the least restriction of the allowed space for HNL to exist). Effectively, a downward fluctuation in the expected background which drives one up on y axis of Fig. 9.10 gives one the largest upper limit μ_+ that could be compatible with rejecting H_1, and therefore the largest (worst) limit on $|U_{\alpha 4}|^2$. This reconciles our intuitive understanding of "good background control $\Leftrightarrow$ more sensitive limits" with the procedure outlined above. To accommodate the large irreducible background associated with the CC Coherent interaction mode,

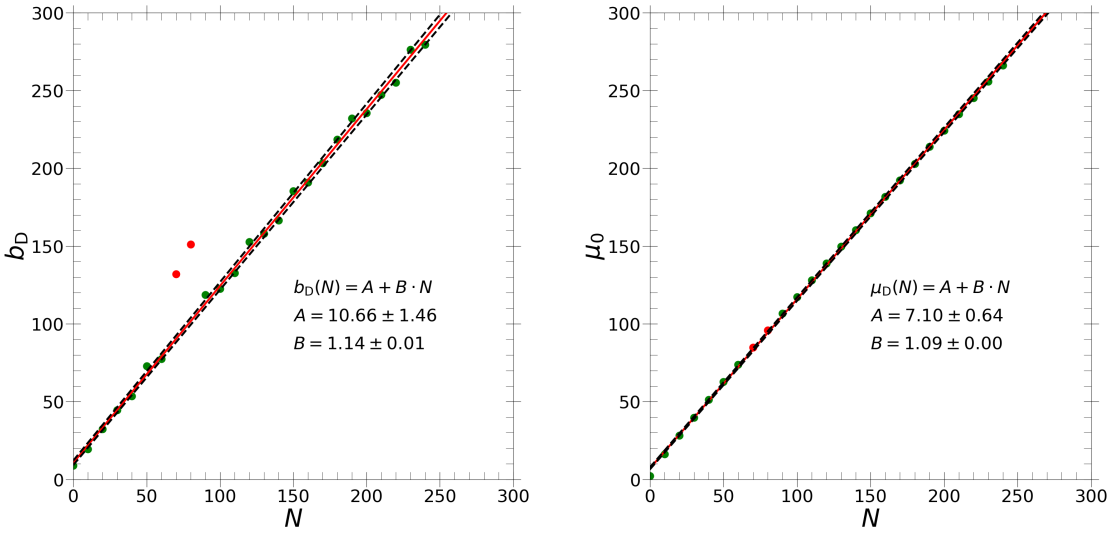
the Feldman-Cousins procedure was first performed over $N \in \{0, 10, \ldots, 240\}$ events, and the value of the critical point $\left(b_{\mathrm{D}}^{N}, \mu_{\mathrm{D}}^{N}\right)$ was obtained. The critical point is the point at which the Feldman-Cousins distribution "flattens out", and achieves a degenerate value μ_{D}. Because of the Poissonian pathology, the extraction is not exact; instead, by plotting the points an empirical $f(x) = A + B \cdot \ln(x)$ relation was determined between the values $(b_{\mathrm{D}}, \mu_{\mathrm{D}})$ (with the presence of two outliers at $N = 70, 80$ where the extraction appears to have failed). The critical background b_{D} at which degeneracy is achieved follows a simple linear relation $b_{\mathrm{D}}(N) = A + B \cdot N$. Similarly, the inferred signal at zero background $\mu_0(N)$ follows a linear distribution. From these three distributions, it is possible to extrapolate the value μ_+ from the background $b \pm \delta b$; the fits are arrayed in Fig. 9.11. An illustrative example for $N = 1000$ observed events, showing the spread of values arising from the fit errors, is shown in Fig. 9.12.

Using the estimate for the background variance in FHC and RHC, one can calculate the relevant parametrisations for the Feldman-Cousins curves; from those, one obtains the upper limit $\{\mu^{(100)}, \delta\mu_{\pm}^{(100)}\}$ for each case assuming 100 excess events. The spread in $\mu^{(100)}$ is not symmetric, as $\delta\mu_{-}^{(100)}$ is bounded by $\mu^{(100)} - \mu^{(\mathrm{D})}$. In Table 9.2, the curve parameters for 100 excess events, background nominal value and variance, and values $\mu^{(100)}, \delta\mu_{\pm}^{(100)}$ are shown. It is also required to know the

Current	Background	$b_{\mathrm{D}}^{(100)}$	$\mu_{\mathrm{D}}^{(100)}$	$\mu_{0}^{(100)}$	$\mu^{(100)}$
FHC	578.6 ± 294.3	781.9 ± 9.1	6.57 ± 1.34	742.5 ± 3.3	$197.6^{+282.4}_{-192.4}$
RHC	1353.9 ± 395.7	1663 ± 16.2	7.63 ± 1.40	1582.0 ± 6.4	$300.3^{+386.9}_{-294.1}$

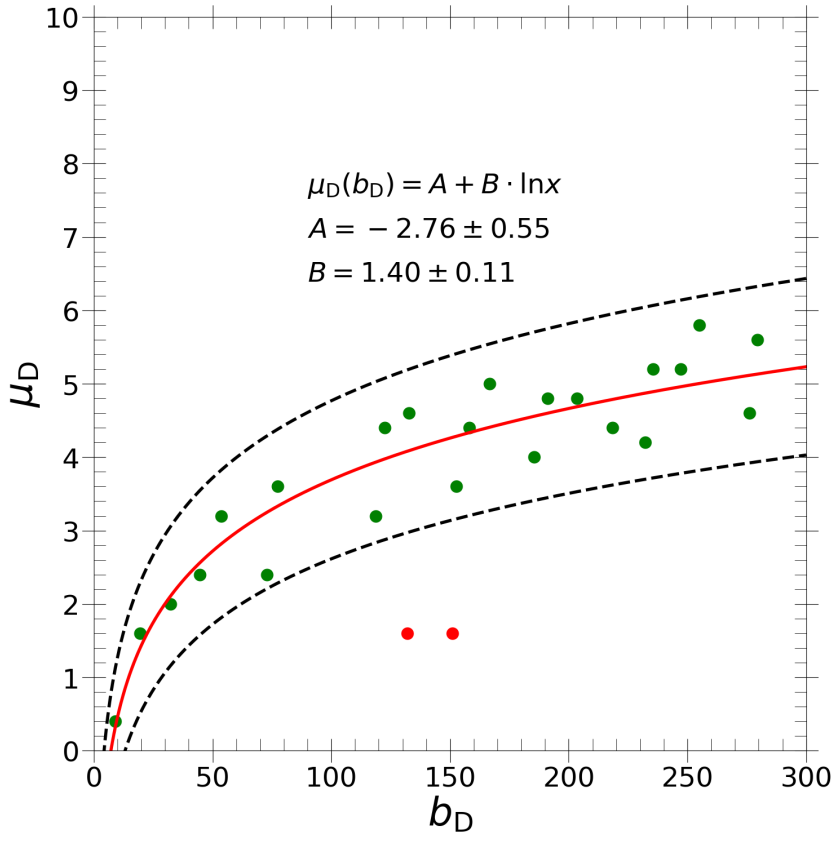
Table 9.2: Parameters for Feldman-Cousins curve extraction.

signal efficiency for each HNL hypothesis (for a definition of the signal efficiency and how the values and their errors were calculated, see Section 7.3.4). Then, one converts the mass-independent limit $\mu^{(100)}$ to a mass-dependent efficiency-corrected limit $N_{\upsilon} = \mu^{(100)}/\epsilon_{\upsilon}$, which one can then compare to the expected number of decays from Fig. 9.1 and apply Equation 6.2 to obtain the limit on the mixing, $|U_{\alpha 4}|^2$. The efficiencies of the various samples are shown in Table 9.3.



(a)

(b)



(c)

Figure 9.11: Fits of Feldman-Cousins procedure at high background values, showing the value b_D of background at floor (a), inferred upper limit on signal at zero background (b), and the value of inferred upper limit on signal at floor (c). Markers correspond to $N = 0, 10, \ldots, 240$. The red markers $N = 70, 80$ are not used in the fit.

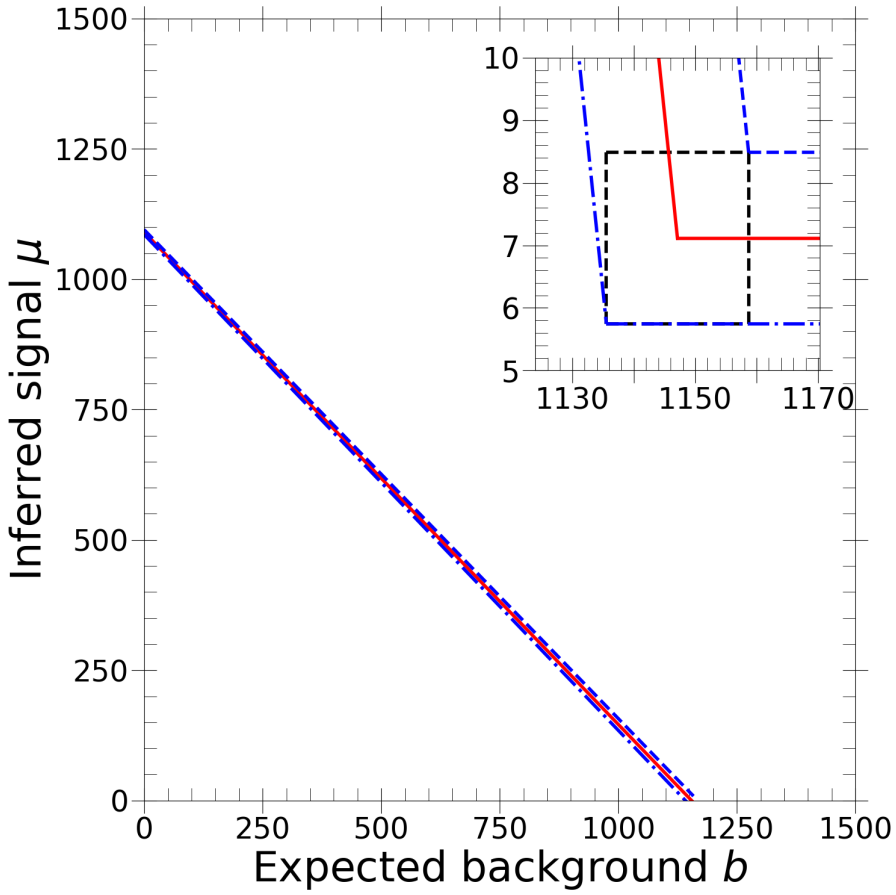


Figure 9.12: Example of $N = 1000$ Feldman-Cousins schema obtained using extrapolation. The nominal extrapolated value is the solid red line, and the dashed blue lines are the errors on that extrapolation. In the inset, the critical point $(b_{\mathrm{D}}, \mu_{\mathrm{D}})$ lies in the centre of the square defined by $(b_{\mathrm{D}} \pm \delta b_{\mathrm{D}}) \times (\mu_{\mathrm{D}} \pm \delta \mu_{\mathrm{D}})$ (dashed black square).

The number of expected events is estimated by tallying up the expected number of POT for the full number of signal events in each sample, and then normalising to the exposure recorded in data. To prevent biasing the number of expected decays by events which, due to their kinematics, have a very large weight (expected POT count), a maximum weight is defined to be that such that the top 1% of the distribution N_{POT} is excluded. The total number of events not discarded is then taken to calculate the normalisation to data exposure. A summary is shown in Table 9.4, and Figure 9.1.

Applying the limit-setting procedure from Section 9.3, one first obtains the efficiency-corrected events and then uses Equation 9.5 to get the ρ for each sample. The uncertainty on ρ, in each band, is conservatively given by assuming the floor

Mass [MeV/c^2]	FHC (equal)	RHC (equal)	FHC (muon)	RHC (muon)
250	$1.60^{+6.22}_{-1.28}$	$1.32^{+5.66}_{-1.07}$	$1.64^{+5.92}_{-1.29}$	$1.37^{+5.73}_{-1.11}$
275	$1.61^{+5.84}_{-1.27}$	$1.25^{+4.88}_{-1.00}$	$1.51^{+5.11}_{-1.17}$	$1.34^{+4.71}_{-1.05}$
300	$1.56^{+5.16}_{-1.20}$	$1.27^{+4.44}_{-0.99}$	$1.54^{+4.90}_{-1.17}$	$1.27^{+4.32}_{-0.99}$
325	$1.52^{+4.95}_{-1.17}$	$1.22^{+4.10}_{-0.94}$	$1.50^{+4.60}_{-1.13}$	$1.35^{+4.26}_{-1.03}$
350	$1.51^{+4.54}_{-1.13}$	$1.19^{+3.56}_{-0.89}$	$1.51^{+4.33}_{-1.12}$	$1.33^{+3.73}_{-0.99}$
375	$1.47^{+4.32}_{-1.11}$	$1.32^{+3.79}_{-0.92}$	$1.60^{+4.38}_{-1.17}$	$1.36^{+3.54}_{-0.98}$
400	$1.50^{+4.39}_{-1.12}$	$1.24^{+3.59}_{-0.92}$	–	–
425	$1.44^{+4.12}_{-1.07}$	$1.25^{+3.49}_{-0.92}$	–	–
450	$1.40^{+3.94}_{-1.04}$	$1.25^{+3.43}_{-0.92}$	–	–
475	$1.18^{+3.19}_{-0.86}$	$1.08^{+2.86}_{-0.79}$	–	–

Table 9.3: Selection efficiencies for the various samples, by mass (rows) and coupling hierarchy (columns). All efficiencies in %.

Mass (MeV/c^2)	FHC (equal)	RHC (equal)	FHC (muon)	RHC (muon)
250	0.131	0.105	0.0763	0.0578
275	0.734	0.569	0.410	0.304
300	1.80	1.32	1.06	0.748
325	3.77	2.67	2.31	1.57
350	7.40	5.03	4.88	3.30
375	13.8	8.95	10.7	7.57
400	20.9	13.8	–	–
425	34.8	23.2	–	–
450	48.7	34.2	–	–
475	12.7	10.3	–	–

Table 9.4: Expected number of $N_4 \to \mu\pi$ decays in MINERνA, for mixing element $|U_{\alpha 4}|^2 = 10^{-8}$. Data exposures are $9.33 \cdot 10^{20}$ POT (FHC), $10.52 \cdot 10^{20}$ POT (RHC).

signal $\mu_{\mathrm{D}}^{(100)} = \mu^{(100)} - \delta\mu_{-}^{(100)}$ then dividing by the upper bound on efficiency (for the "optimistic scenario"), and by assuming the upper signal $\mu^{(100)} + \delta\mu_{+}^{(100)}$ then dividing by the lower bound on efficiency. To account for the normalisation uncertainty on the parent kaon multiplicity (hence, number of expected HNL events) from the `GEANT4` model spread (`QGSP_BERT` vs `FTFP_BERT`), an extra uncertainty is compounded with the "conservative" scenario. One takes the multiplicities $n(K^+), n(K^-)$ from Table 5.1 and calculates their fractional difference $\mathcal{F}$ between FTFP_BERT and QGSP_BERT, $\mathcal{F}_{\pm} = [n_F(K^{\pm}) - n_Q(K^{\pm})]/n_F(K^{\pm})$. The larger difference between

$\mathcal{F}_+$ and $\mathcal{F}_-$ in either mode becomes the representative model spread. These spreads are 32.7% and 34.4% respectively. Finally, a reminder that all these limits are *Dirac* limits: a Majorana HNL would produce equal rates of neutrinos decaying to $\mu^-\pi^+$ and $\mu^+\pi^-$. Hence, the number of expected decays doubles, and the limits are better than the Dirac limits by a factor of $\sqrt{2}$.

In Figure 9.13, the limits obtained at 90% C. L. with the Feldman-Cousins procedure are displayed as coloured bands. The cases of equal couplings and muon dominance are considered separately, in keeping with the different numbers of expected neutrino decays. Since in the equal-coupling case the number of decays is greater, stricter limits are placed in that regime; the muon dominance scenario also cannot place limits above the mass threshold $m_\mathrm{K} - m_\mu \simeq 388\ \mathrm{MeV}/c^2$.

The *width* of the band is controlled primarily by the uncertainty in the efficiencies, as inherited from the total uncertainty of the signal sample; it is also contributed to heavily by the uncertainty in $\delta\mu_\pm^{(100)}$, which corresponds to poor constraints on the irreducible CC Coherent background. As for the *position* of the band, it is convenient to think that it is controlled primarily by the large background contribution in the signal region. However, the main driver for sensitivity, with a fixed number of expected decays, is either the Poisson mass-independent upper limit μ_2, or the mass-dependent signal efficiency ϵ. From Fig. 9.11, it is apparent that even with small background μ_2 can only be brought down to $\mathcal{O}(10)$; that is to say, a factor perhaps $50-100$ times better than the current control. This would shift these limit curves down by $7-10$ times, i.e. $\rho \mapsto \rho + \varpi, \varpi \in [0.85, 1.0]$. Simultaneously, an increase in efficiency by $5-10$ times would shift the limits as $\rho \mapsto \rho + \varpi, \varpi \in [0.35, 1.0]$ (as indeed is shown by the width of the band, which is about that wide). This is to say, the controlling factors in these limit curves are both the background contamination and the efficiency, at the same time. Consequently, it is crucial for searches to isolate mostly background-free signal regions, as well as guarantee a selection procedure which preserves as much of the efficiency as possible. The high presence of ill-constrained background in this analysis serves to obfuscate the limit-setting procedure, which means the limits are ultimately

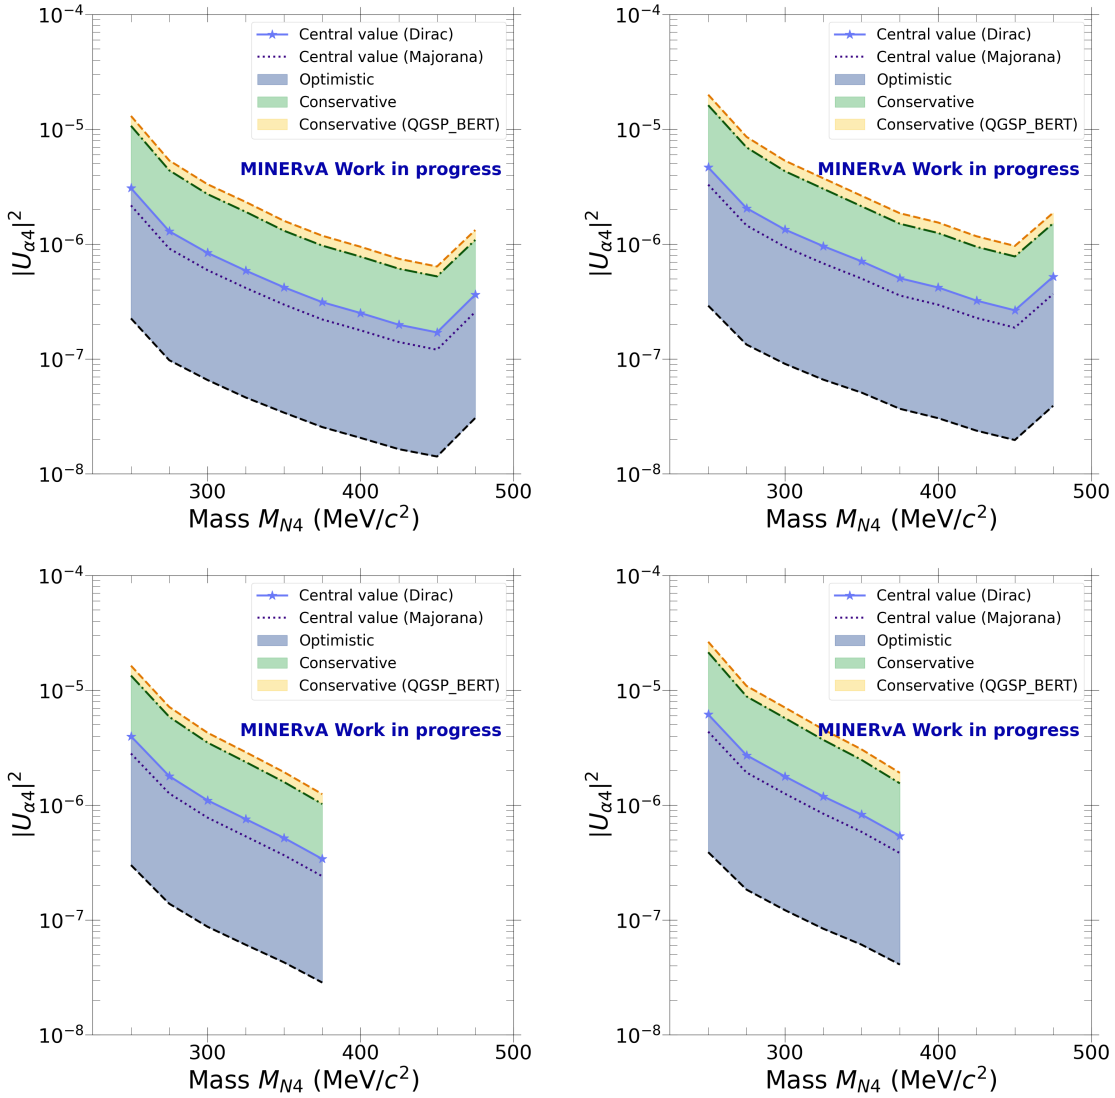


Figure 9.13: Feldman-Cousins 90% limits on HNL in the small-signal scenario. FHC (left) and RHC (right), equal coupling (top) and muon dominance (bottom) scenarios presented. "Optimistic" and "conservative" scenarios indicated by coloured bands, with "conservative" scenario highlighted before and after applying the error from the `GEANT4` model spread. The nominal value for Majorana neutrinos, shifted down by a factor $\rho \mapsto \rho + 0.15$, is also shown. See text for details.

not competitive: however, a subsequent search with a suitable definition of a (mostly) background-free region stands to produce competitive limits (such as using a nanosecond-scale timing trigger, as in MicroBooNE [256]).

10 Conclusion and future directions

The limits obtained in the previous chapter, though not competitive with those produced by the T2K experiment or modern liquid argon detectors, are representative of MINERνA's capability to discriminate HNL from an interesting and poorly constrained Standard Model process in a topology where the irreducible background is poorly constrained. In particular, MINERνA possesses a strong capability to differentiate tracks with two MIPs, highlighting its good $\mathrm{d}E/\mathrm{d}x$ measuring capabilities. Additionally, the presence of MINOS downstream and of the ECAL and HCAL do offer a useful handle to separate muons from pions. Nonetheless, a central consideration is the enduring presence of the SM background in the two-MIP, one-track region which, for this topology, is the most-suited event signature for an HNL search. Unlike the lower-energy beams of T2K and MicroBooNE, the NuMI beam delivers very high-energy neutrino parents which cause the one-track region to be significantly preferred for HNL production.

With this topology, a first instructive series of limits can be obtained. Additionally, a robust discovery test for HNL detection can be performed on MINERνA data, when that is unblinded. However, the MINERνA data can be exploited further to enhance the predictive power of the experiment. Specifically, other decay channels with significantly smaller background may be employed to produce a combined constraint on the HNL signal. Out of the available channels, an immediate first candidate would be to target the $N_4 \to \nu\mu^+\mu^-$ signature, similar to [161]; there is only one irreducible background. This is the *neutrino trident production* $\nu + \mathrm{A} \to \nu + \mu^+ + \mu^- + \mathrm{A}$, which is extremely rare; about $20-30$ dimuon trident events *over the entire running period of MINERνA* are expected to occur [312]. There are a few more events in the other dilepton channels, but still well within $\mathcal{O}(100)$. This is decidedly smaller than the $\mathcal{O}(10^4 - 10^5)$ CC Coherent events expected for the muon-pion channel.

But commensurately the dilepton channels come with their own problems. First and foremost, there will be far fewer expected HNL decays to this channel in the detector (cf. Fig. 5.2), which provides a caveat to this "background-free" search. Moreover, the reconstruction of such signals in MINERνA will be challenging; for example, as discussed in Chapter 8, the MINOS reconstruction rejects dimuon events, which means a custom matching of a dimuon track to all possible (diverging) MINOS-ND tracks will need to be developed (as in [161]). Custom reconstruction will also need to be developed for the other dilepton signatures.

Despite that, a combined analysis of multiple channels and their combination using a suitable modified-frequentist or Bayesian framework (as in T2K or MicroBooNE [159, 313]) could yield very strong limits in the range below the kaon mass. Especially important is the observation that sensitivity is driven from the most background-free regions; these can exploited either by enhancing the selection in MINERνA to make use of two-track $N_4 \to \mu\pi$ events, such that reconstruction of more kinematic quantities is possible, or by considering alternate channels such as the one just described and performing a multi-channel analysis.

With that in mind, there is considerable effort to study HNL as specific realisations of long-lived particles at various current- and future-generation experiments. The Short Baseline Program (SBN) at Fermilab, for example, is performing searches for HNL at the Booster Neutrino Beam [181], and MicroBooNE is continuing searches for HNL [313]. Further on the liquid argon front, there is interest to search for HNL in the current iteration of ProtoDUNE [314], and to evaluate the sensitivity of DUNE's multi-purpose near detector to generic long-lived particles [315, 316] and specific model realisations, HNL among them [317]. This interest is coupled with ongoing work on LArTPC technology in ProtoDUNE [318], and to understand the performance of the eventual high-pressure gas time-projection chamber at DUNE [319, 320], whose lower tracking threshold will be crucial to probe the low-energy component of HNL signatures. This motivates the continued development of analysis techniques to exploit data at neutrino detectors.

Adding further, it is made clear that, in the case of rare searches, background control is crucial for statistical sensitivity. In regions of kinematic phase space where there are few measurements, the constraint on rare SM processes is, as seen previously, inadequate for rare searches. To further complicate matters, current treatments of nuclear effects and final-state interactions of SM neutrino interaction products do not yet precisely model existing data; a promising technique is via Transverse Kinematic Imbalance (TKI) measurements which are sensitive to final-state interactions [318, 321, 322]. Specifically for the case of the coherent background, there is too little data in the very forward region with high pion energy to constrain the irreducible background for an HNL search. Next-generation detectors with high position and angular resolutions, with sufficient constraints on the beam direction and flux, might be able to provide more measurements in this region of kinematic phase space. Better constraints of the coherent kinematics would be valuable input for future dark-sector searches.

In this thesis, a full analysis of a HNL signal at MINERνA, with support from a completely generalisable module to simulate HNL and a consistent search and discovery strategy, has been performed. Such an analysis is typical of what might be becoming a standardised pipeline for BSM discovery physics in the neutrino sector.

Appendices

A The BeamHNL GENIE simulation package

BeamHNL is an experiment-agnostic, configurable simulation package designed to allow for a generic description of HNL from an accelerator beamline. The detailed physics of the simulation is described in Chapter 5: in particular, a high-level overview may be seen at a glance in Figure 5.16. In this appendix, some details about the configuration file and GENIE executables are presented.

BeamHNL is not enabled by default. It may be switched on by adding a line

```
-enable-heavy-neutral-lepton \
```

to one's configure script. A suite of validation tools is also available, to test the flux generation, decay, and vertex assignment operations of the package. It may be built by adding the line `-enable-hnl-validation \` .

A.1 The configuration file

There is a single configuration file, config/CommonHNL.xml, which contains all the paramteres parsed at runtime by each executable. It is organised as a set of tags, which are as follows:

- ParameterSpace: Controls HNL mass, couplings $\left\{|U_{\alpha 4}|^2\right\}_{\alpha=e,\mu,\tau}$, and a Dirac versus Majorana switch.

- InterestingChannels: Of the ten kinematically allowed channels in the region below the kaon mass, the user may well not be interested to simulate all. Each channel may be *allowed* or *inhibited*; though the HNL lifetime is set by constructing the total decay with of all kinematically allowed channels, inhibited or not, the actual event generation will only produce events from the pool of allowed channels by evaluating the conditional probability of decay to a specific channel. In this way, single- or few-channel production campaigns may be run without compromising the POT calculation from Fig. 5.15.

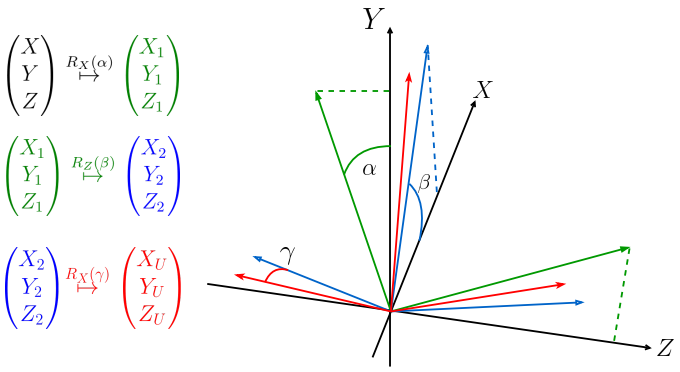


Figure A.1: Euler rotations, in the extrinsic $X - Z - X$ convention.

- `CoordinateXForm`: Determines the translations and Euler angles between the NEAR (flux-file coordinate system), BEAM (same origin as NEAR but rotated), and USER (geometry-file coordinate system) coordinate frames. The Euler angles follow the extrinsic $X - Z - X$ convention, which is to say that if three angles α, β, γ are specified, then the overall rotation matrix $R(\alpha, \beta, \gamma) = R_X(\alpha)R_Z(\beta)R_X(\gamma)$ with the X, Y, Z coordinate system fixed. It is [243]

$$R(\alpha, \beta, \gamma) = \begin{pmatrix} c_\beta & -c_\alpha s_\beta & s_\alpha s_\beta \\ s_\beta c_\gamma & c_\alpha c_\beta c_\gamma - s_\alpha s_\gamma & -s_\alpha c_\beta c_\gamma - c_\alpha s_\gamma \\ s_\beta s_\gamma & s_\alpha c_\gamma + c_\alpha c_\beta s_\gamma & c_\alpha c_\gamma - s_\alpha c_\beta s_\gamma \end{pmatrix}. \tag{A.1}$$

  Also see Fig. A.1.

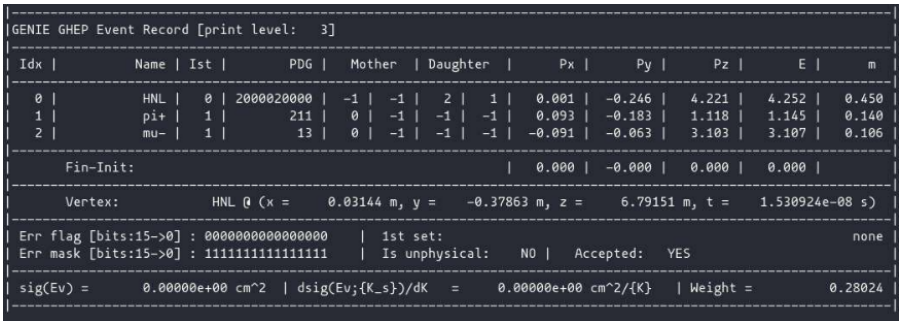
- `FluxCalc`: The main entry here is the number of expected POT per each parent that produces a neutrino; in other words, the inverse multiplicity of each particle species.

- `ParticleGun`: Parameters relating to the origin point and momentum of the `gevgen_pghnl` particle gun application.

A.2 Running the code

Archetypally, the invocation of a `BeamHNL` run is

```
gevgen_hnl -n <nEvts> -f <path/to/flux/file/directory>     \
-g <path/to/geom/file.root>.
```



```
GENIE GHEP Event Record [print level:   3]
| Idx |          Name | Ist |        PDG |   Mother  | Daughter  |      Px |      Py |      Pz |       E |      m  |
|   0 |           HNL |   0 | 2000020000 |  -1 |  -1 |   2 |   1 |   0.001 |  -0.246 |   4.221 |   4.252 |   0.450 |
|   1 |           pi+ |   1 |        211 |   0 |  -1 |  -1 |  -1 |   0.093 |  -0.183 |   1.118 |   1.145 |   0.140 |
|   2 |           mu- |   1 |         13 |   0 |  -1 |  -1 |  -1 |  -0.091 |  -0.063 |   3.103 |   3.107 |   0.106 |
      Fin-Init:                                          |   0.000 |  -0.000 |   0.000 |   0.000 |
      Vertex:           HNL @ (x =     0.03144 m, y =    -0.37863 m, z =     6.79151 m, t =    1.530924e-08 s)
Err flag [bits:15->0] : 0000000000000000    |  1st set:                                                    none
Err mask [bits:15->0] : 1111111111111111    |  Is unphysical:    NO |   Accepted:   YES
sig(Ev) =       0.00000e+00 cm^2  | dsig(Ev;{K_s})/dK   =     0.00000e+00 cm^2/{K}   | Weight =          0.28024
```

Figure A.2: `GENIE EventRecord` for an example $N_4 \to \pi^+\mu^-$ decay in MINERνA.

There is also a particle gun utility, `gevgen_pghnl`, which takes an origin point, directional cosines for momentum, and HNL energy supplied in the config file. It only requires a ROOT geometry file.

Finally, a validation app `gevald_hnl` was developed to check the performance of each of the three monadic operations of event generation. These are implemented as individual tests, or *modes*, in `gevald_hnl`:

```
gevald_hnl -n <nEvts> -M <mode> -f <path/to/flux/file/directory>     \
-g <path/to/geom/file.root>.
```

`-M 1, 2, 3` diagnose the flux generation, decay, and vertex assignment routines respecively by producing validation ROOT files containing trees and histograms of the quantities needed for the full event generation application.

The `GENIE EventRecord` output can be straightforwardly passed to an experimental simulation to then propagate downsteam, as detailed in Chapter 6. An example decay $N_4 \to \pi^+ + \mu^-$ in MINERνA is shown in Figure A.2. The incoming HNL has four-momentum $p_N^\alpha = (4.252, 0.001, -0.246, 4.221)$ GeV, and decays to a 1.145 GeV pion and a 3.107 GeV muon. The decay happens in local MINERνA coordinates $\boldsymbol{r}_\mathrm{D} = (0.03, -0.38, 6.79)$ m, and the delay of the HNL to arrive with respect to a SM neutrino is $\delta t = 15.3$ ns. Finally, the expected number of POT for this one signal event, assuming $|U_{\alpha 4}|^2 = 10^{-8}$, was $0.28 \cdot 10^{20}$.

B The dk2nu tuple structure

A more complete record of this documentation can be found at the LBNF Redmine page.

B.1 The main tuple `dk2nutree`

All positions are given in cm, all momenta in GeV/c, and all references to coordinates use the BEAM coordinate system where z is parallel to the beam axis.
The members of `nuray` are vectors, corresponding to the SM neutrino kinematics *under the constraint that it traverses the point specified in the corresponding metadata entry.* Effectively, this is a constraint on the angular region of emission from each neutrino, which corresponds to a probability that this constraint can be satisfied; this is stored as the parent-rest-frame angular acceptance ω, as a weight in `nuray`.

Branch	Content
decay	Information about the neutrino production
decay.nimpwt	"Importance weight": statistical multiplier from `g4numi`
decay.vx	Neutrino production vertex x
decay.vy	Neutrino production vertex y
decay.vz	Neutrino production vertex z
decay.pdpx	Parent p_x at neutrino production
decay.pdpy	Parent p_y at neutrino production
decay.pdpz	Parent p_z at neutrino production
decay.ptype	PDG code of parent
decay.necm	Neutrino energy in parent rest frame
decay.ppdxdz	Parent p_x/p_z at parent production
decay.ppdydz	Parent p_y/p_z at parent production
decay.pppz	Parent p_z at parent production
decay.ppenergy	Parent E at parent production
decay.norig	Code for neutrino origin (target/baffle, muon decay, other)
decay.ndecay	Decay code of neutrino-producing decay
decay.ntype	PDG code of neutrino
job	Simulation job number
potnum	Index of primary proton that led to neutrino production
ppvx	Neutrino parent production vertex x
ppvy	Neutrino parent production vertex y
ppvz	Neutrino parent production vertex z

nuray	Vector of neutrino kinematics assuming neutrino forced to location
nuray.px	Vector of neutrino p_x per metadata location
nuray.py	Vector of neutrino p_y per metadata location
nuray.pz	Vector of neutrino p_z per metadata location
nuray.E	Vector of neutrino E per metadata location
nuray.wgt	Vector of rest-frame geometric acceptances ω
ancestor	Vector of all ancestor particles. Proton first, neutrino last.
ancestor.pdg	PDG code of ancestor
ancestor.startx	Ancestor trajectory start, x
ancestor.starty	Ancestor trajectory start, y
ancestor.startz	Ancestor trajectory start, z
ancestor.startt	Ancestor trajectory start, t, ns
ancestor.startpx	Ancestor starting p_x
ancestor.startpy	Ancestor starting p_y
ancestor.startpz	Ancestor starting p_z
ancestor.startpx	Ancestor final p_x
ancestor.startpy	Ancestor final p_y
ancestor.startpz	Ancestor final p_z
ancestor.pprodpx	Parent particle x-momentum prior to secondaries production
ancestor.pprodpy	Parent particle y-momentum prior to secondaries production
ancestor.pprodpz	Parent particle z-momentum prior to secondaries production
ancestor.nucleus	PDG code of nucleus on which interaction happened
ancestor.proc	String describing process that produced this ancestor
ancestor.ivol	String describing volume where this ancestor was created
ancestor.imat	String describing material where this ancestor was created

B.2 The metadata tree dkmetaTree

Branch	Content
job	Simulation job number
pots	Number of POT simulated
beam0X	Primary proton origin, x
beam0Y	Primary proton origin, y
beam0Z	Primary proton origin, z
beamhwidth	Horizontal beam spot size (RMS)
beamvwidth	Vertical beam spot size (RMS)
beamdxdz	Initial proton p_x/p_z
beamdydz	Initial proton p_y/p_z
physics	GEANT4 version + physics list
physcuts	GEANT4 tracking cuts
tgtcfg	Target configuration; me(bar)*abc*z in MINERνA
horncfg	Horn configuration; *jkl*i in MINERνA

dkvolcfg	Decay pipe configuration; Helium in MINERνA
location	Vectors containing reference locations for quick flux calculations
location.x	Reference x
location.y	Reference y
location.z	Reference z
location.name	Reference name

abc refers to the distance of the target from the position $z = 0$, which is the maximum insertion into the first focusing horn such that no electrical arcing between target and horn can occur [189].

jkl refers to the horn current in kA.

The various locations in the nuray vectors are:

- Random: No constraint applied.
- MINERνA
- MINOS Near 1,2
- NOνA Near 1,2 + Near Detecor on Surface (NDOS)
- MiniBooNE
- MicroBooNE
- SciBooNE
- MINOS Far 1,2
- NOνA Far

References

[1] Nobel Prize Outreach AB, ed. *The Nobel Prize in Physics 2015*. URL: https://www.nobelprize.org/prizes/physics/2015/summary/ (visited on May 26, 2024).

[2] Srednicki, M. *Quantum Field Theory*. 1st ed. Cambridge University Press, 2007.

[3] Glashow, S. L. "The renormalizability of vector meson interactions". *Nuclear Physics* 10 (Feb. 1959), pp. 107–117.

[4] Salam, A. and Ward, J. C. "Weak and electromagnetic interactions". *Il Nuovo Cimento* 11.4 (Feb. 1959), pp. 568–577.

[5] Englert, F. and Brout, R. "Broken Symmetry and the Mass of Gauge Vector Mesons". *Physical Review Letters* 13.9 (Aug. 1964), pp. 321–323.

[6] Higgs, P. W. "Broken Symmetries and the Masses of Gauge Bosons". *Physical Review Letters* 13.16 (Oct. 1964), pp. 508–509.

[7] Weinberg, S. "A Model of Leptons". *Physical Review Letters* 19.21 (Nov. 1967), pp. 1264–1266.

[8] Halzen, F. and Martin, A. D. *Quarks and Leptons: An Introductory Course in Modern Particle Physics*. 2nd ed. John Wiley & Sons, Inc., 1984.

[9] Wu, C. *et al.* "Experimental Test of Parity Conservation in Beta Decay". *Physical Review* 105.4 (Feb. 1957), pp. 1413–1415.

[10] Lee, T. D. and Yang, C. N. "Question of Parity Conservation in Weak Interactions". *Physical Review* 104.1 (Oct. 1956), pp. 254–258.

[11] Gell-Mann, M. and Pais, A. "Behavior of Neutral Particles under Charge Conservation". *Physical Review* 97.5 (Mar. 1955), pp. 1387–1389.

[12] Lee, T. D., Oehme, R., and Yang, C. N. "Remarks on Possible Noninvariance under Time Reversal and Charge Conjugation". *Physical Review* 106.2 (Apr. 1957), pp. 340–345.

[13] Christenson, J. H. *et al.* "Evidence for the 2π Decay of the K_2^0 Meson". *Physical Review Letters* 13.4 (July 1964), pp. 138–140.

[14] Cabibbo, N. "Unitary Symmetry and Leptonic Decays". *Physical Review Letters* 10.12 (June 1963), pp. 531–533.

[15] Glashow, S. L., Iliopoulos, J., and Maiani, L. "Weak Interactions with Lepton-Hadron Symmetry". *Physical Review* ***D*** 2 (Oct. 1970), pp. 1285–1292.

[16] Kobayashi, M. and Maskawa, T. "CP-Violation in the Renormalizable Theory of Weak Interaction". *Progress of Theoretical Physics* 49.2 (Feb. 1973), pp. 652–657.

[17] Charles, J. *et al* (CKMFitter Group). "CP violation and the CKM matrix: assessing the impact of the asymmetric B factories". *The European Physcial Journal* ***C*** 41.1 (May 2005), pp. 1–131.

[18] Feynman, R. P. and Gell-Mann, M. "Theory of the Fermi Interaction". *Physical Review* 109.1 (Jan. 1958), pp. 193–198.

[19] Goldhaber, M., Grodzins, L., and Sunyar, A. W. "Helicity of Neutrinos". *Physical Review* 109.3 (Feb. 1958), pp. 1015–1017.

[20] Landau, L. "On the conservation laws for weak interactions". *Nuclear Physics* 3.1 (1957), pp. 127–131.

[21] Salam, A. "On parity conservation and neutrino mass". *Il Nuovo Cimento* 5.1 (Jan. 1957), pp. 299–301.

[22] Schechter, J. and Valle, J. W. F. "Neutrinoless double-β decay in $SU(2) \times U(1)$ theories". *Physical Review* ***D*** 25.11 (June 1982), pp. 2951–2954.

[23] Päs, H. and Rodejohann, W. "Neutrinoless Double Beta Decay". *New Journal of Physics* 17.11 (July 2015), p. 115010.

[24] Bilenky, S. M. "Neutrino. History of a unique particle". *The European Physical Journal* ***H*** 38.3 (Apr. 2013), pp. 345–404.

[25] Bahcall, J. N. and Bethe, H. A. "Solution of the solar-neutrino problem". *Physical Review Letters* 65.18 (Oct. 1990), pp. 2233–2235.

[26] Hirata, K. S. *et al.* "Experimental study of the atmospheric neutrino flux". *Physics Letters* ***B*** 205.2-3 (Apr. 1988), pp. 416–420.

[27] Casper, D. *et al.* "Measurement of Atmospheric Neutrino Composition with the IMB-3 Detector". *Physical Review Letters* 66.20 (May 1991), pp. 2561–2564.

[28] Kajita, T. and Kearns, E. and Shiozawa, M. "Establishing atmospheric neutrino oscillations with Super-Kamiokande". *Nuclear Physics* ***B*** 908 (Apr. 2016), pp. 14–29.

[29] Pontecorvo, B. "Mesonium and anti-mesonium". *Soviet Journal of Experimental and Theoretical Physics* 6.2 (Feb. 1958), pp. 429–431.

[30] Pontecorvo, B. "Inverse beta processes and nonconservation of lepton charge". *Soviet Journal of Experimental and Theoretical Physics* 7.1 (July 1958), pp. 172–173.

[31] Davis, R. "Attempt to Detect the Antineutrinos from a Nuclear Reactor by the $Cl^{37}(\overline{\nu}, e^-)A^{37}$ Reaction". *Physical Review* 97.3 (Feb. 1955), pp. 766–769.

[32] Danby, G. *et al.* "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". *Physical Review Letters* 9.1 (July 1962), pp. 36–44.

[33] Pontecorvo, B. "Neutrino Experiments and the Problem of Conservation of Leptonic Charge". *Soviet Journal of Experimental and Theoretical Physics* 26.5 (May 1968), pp. 984–988.

[34] Fukuda, Y. *et al* (Super-Kamiokande Collaboration). "Evidence for oscillation of atmospheric neutrinos". *Physical Review Letters* 81.8 (Aug. 1998), pp. 1562–1567.

[35] Ahmad, Q. R. *et al* (SNO Collaboration). "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory". *Physical Review Letters* 89.1 (June 2002), p. 011301.

[36] Eguchi, K. *et al* (KamLAND Collaboration). "First Results from KamLAND: Evidence for Reactor Neutrino Disappearance". *Physical Review Letters* 90.2 (Jan. 2003), p. 021802.

[37] Smirnov, A. Yu. "Solar neutrinos: Oscillations or No-oscillations?" Sept. 2016. arXiv: `1609.02386 [hep-ph]`.

[38] Aharmim, B. *et al* (SNO Collaboration). "Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory". *Physical Review* ***C*** 88.2 (Aug. 2013), p. 025501.

[39] Fukuda, S. *et al* (Super-Kamiokande Collaboration). "Tau Neutrinos Favored over Sterile Neutrinos in Atmospheric Muon Neutrino Oscillations". *Physical Review Letters* 85.19 (Nov. 2000), pp. 3999–4003.

[40] Hampel, W. *et al* (GALLEX Collaboration). "GALLEX solar neutrino observations: results for GALLEX IV". *Physics Letters* ***B*** 447.1 (Feb. 1999), pp. 127–133.

[41] Altmann, M. *et al* (GNO Collaboration). "Complete results for five years of GNO solar neutrino observations". *Physics Letters* ***B*** 616.3 (June 2005), pp. 174–190.

[42] Fukuda, Y. *et al* (Kamiokande Collaboration). "Solar Neutrino Data Covering Solar Cycle 22". *Physical Review Letters* 77.9 (Aug. 1996), pp. 1683–1686.

[43] Ahn, M. H. *et al* (K2K Collaboration). "Measurement of neutrino oscillation by the K2K experiment". *Physical Review* ***D*** 74.7 (Oct. 2006), p. 072003.

[44] Adamson, P. *et al* (MINOS Collaboration). "Combined Analysis of ν_μ disappearance and $\nu_\mu \to \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos". *Physical Review Letters* 112.19 (May 2014), p. 191801.

[45] Agafonova, N. *et al* (OPERA Collaboration). "Final Results of the OPERA Experiment on ν_τ appearance in the CNGS beam". *Physical Review Letters* 120.21 (May 2018), p. 211801.

[46] Aguilar, A. *et al* (LSND Collaboration). "Evidence for Neutrino Oscillations from the Observation of Electron Anti-neutrinos in a Muon Anti-Neutrino Beam". *Physical Review Letters* 77.15 (Oct. 1996), pp. 3082–3085.

[47] Armbruster, B. *et al* (KARMEN Collaboration). "Upper limits for neutrino oscillations muon-antineutrino to electron-antineutrino from muon decay at rest". *Physical Review* ***D*** 65.11 (June 2002), p. 112001.

[48] Aguilar-Arevalo, A. A. *et al* (MiniBooNE Collaboration). "Updated MiniBooNE neutrino oscillation results with increased data and new background studies". *Physical Review* ***D*** 103.5 (Mar. 2021), p. 052002.

[49] Apollonio, M. *et al.* "Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station". *The European Physcial Journal* ***C*** 27.3 (Apr. 2003), pp. 331–374.

[50] Boehm, F. *et al.* "Final results from the Palo Verde neutrino oscillation experiment". *Physical Review* ***D*** 64.11 (Nov. 2001), p. 112001.

[51] Adamson, P. *et al* (MINOS+ Collaboration). "Precision constraints for three-flavor neutrino oscillations from the full MINOS+ and MINOS data set". *Physical Review Letters* 125.13 (Sept. 2020), p. 131802.

[52] Abe, K. *et al* (T2K Collaboration). "Improved constraints on neutrino mixing from the T2K experiment with 3.13×10^{20} protons on target". *Physical Review* ***D*** 103.11 (June 2021), p. 112008.

[53] Acero, M. A. *et al* (NOvA Collaboration). "First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA". *Physical Review Letters* 123.15 (Oct. 2019), p. 151803.

[54] Abratenko, P. *et al* (MicroBooNE Collaboration). "Search for an Excess of Electron Neutrino Interactions in MicroBooNE Using Multiple Final State Topologies". *Physical Review Letters* 128.24 (June 2021), p. 241801.

[55] Adey, D. *et al* (Daya Bay Collaboration). "Improved measurement of the reactor antineutrino flux at Daya Bay". *Physical Review* ***D*** 100.5 (Sept. 2019), p. 052004.

[56] Bak, G. *et al* (RENO Collaboration). "Measurement of Reactor Antineutrino Oscillation Amplitude and Frequency at RENO". *Physical Review Letters* 121.20 (Nov. 2018), p. 201801.

[57] de Kerret, H. *et al* (Double Chooz Collaboration). "Double Chooz θ_{13} measurement via total neutron capture detection". *Nature Physics* 16.5 (Apr. 2020), pp. 558–564.

[58] Machado, P. A. N., Palamara, O., and Schmitz, D. W. "The Short-Baseline Neutrino Program at Fermilab". *Annual Reviews* 69.1 (Oct. 2019), pp. 363–387.

[59] Abratenko, P. *et al* (ICARUS Collaboration). "ICARUS at the Fermilab Short-Baseline Neutrino Program – Initial Operation". Jan. 2023. arXiv: 2301.08634 [hep-ex].

[60] Abed Abud, A. *et al* (DUNE Collaboration). "Snowmass Neutrino Frontier: DUNE Physics Summary". Mar. 2022. arXiv: 2203.06100 [hep-ex].

[61] Abe, K. *et al* (Hyper-Kamiokande Working Group). "A Long Baseline Neutrino Oscillation Experiment Using J-PARC Neutrino Beam and Hyper-Kamiokande". Dec. 2014. arXiv: 1412.4673 [ins-det].

[62] An, Fengpeng *et al.* "Neutrino physics with JUNO". *Journal of Physics* ***G*** 43.3 (Feb. 2016), p. 030401.

[63] Workman, R. L. *et al* (Particle Data Group). "Review of Particle Physics Section 14: Neutrino Masses, Mixing, and Oscillations". *Progress of Theoretical and Experimental Physics* 2022 (2022), p. 083C01.

[64] Giganti, C., Lavignac, S., and Zito, M. "Neutrino oscillations: The rise of the PMNS paradigm". *Progress in Particle and Nuclear Physics* 98 (Jan. 2018), pp. 1–54.

[65] de Salas, P. F. *et al.* "2020 global reassessment of the neutrino oscillation picture". *Journal of High Energy Physics* 2021.2 (Feb. 2021), p. 71.

[66] Esteban, I. *et al.* "The fate of hints: updated global analysis of three-flavor neutrino oscillations". *Journal of High Energy Physics* 2020.9 (Sept. 2020), p. 178.

[67] Esteban, I. *et al. Nu-Fit v5.2: Three-neutrino fit based on data available in November 2022*. Nov. 2022. URL: http://www.nu-fit.org/~q=node/256 (visited on May 31, 2023).

[68] Qian, X. and Vogel, P. "Neutrino mass hierarchy". *Progress in Particle and Nuclear Physics* 83 (July 2015), pp. 1–30.

[69] Formaggio, J. A., de Gouvêa, A. L. C., and Robertson, R. G. H. "Direct measurements of neutrino mass". *Physics Reports* 914 (June 2021), pp. 1–54.

[70] Aker, M. *et al* (KATRIN Collaboration). "Direct neutrino-mass measurement with sub-electronvolt sensitivity". *Nature Physics* 18.2 (2022), pp. 160–166.

[71] Schwingenheuer, B. "Status and prospects of searches for neutrinoless double beta decay". *Annalen der Physik* 525.4 (Jan. 2013), pp. 269–280.

[72] Lesgourgues, J. and Pastor, S. "Neutrino Mass from Cosmology". *Advances in High Energy Physics* 2012 (Dec. 2012), p. 608515.

[73] Bond, J. R., Efstathiou, G., and Silk, J. "Massive neutrinos and the large-scale structure of the universe". *Physical Review Letters* 45.24 (Dec. 1980), pp. 1980–1984.

[74] Jimenez, R. *et al.* "Neutrino masses and mass hierarchy: evidence for the normal hierarchy". *Journal of Cosmology and Astroparticle Physics* 9 (Sept. 2022), p. 006.

[75] Workman, R. L. *et al* (Particle Data Group). "Review of Particle Physics Section 26: Neutrinos In Cosmology". *Progress of Theoretical and Experimental Physics* 2022 (2022), p. 083C01.

[76] Palanque-Delabrouille, N. *et al.* "Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman-α and Planck full-survey data". *Journal of Cosmology and Astroparticle Physics* 2020.4 (Apr. 2020), p. 038.

[77] Choudhury, S. R. and Hannestad, S. "Updated results on neutrino mass and mass hierarchy from cosmology with Planck 2018 likelihoods". *Journal of Cosmology and Astroparticle Physics* 2020.7 (July 2020), p. 037.

[78] Sekiguchi, T. and Takahashi, T. "Cosmological bound on neutrino masses in the light of H_0 tension". *Physical Review* ***D*** 103.8 (Apr. 2021), p. 083516.

[79] Bandyopadhyay, A. *et al.* "Physics at a future Neutrino Factory and super-beam facility". *Reports on Progress in Physics* 72.10 (Sept. 2009), p. 106201.

[80] Abada, A. *et al.* "Low energy effects of neutrino masses". *Journal of High Energy Physics* 2007.12 (Dec. 2007), p. 061.

[81] Hettmansperger, H., Lindner, M., and Rodejohann, W. "Phenomenological consequences of sub-leading terms in see-saw formulas". *Journal of High Energy Physics* 2011.4 (Apr. 2011), p. 123.

[82] Drewes, Marco. "The Phenomenology of Right Handed Neutrinos". *International Journal of Modern Physics* ***E*** 22.8 (Aug. 2013), p. 1330019.

[83] Abazajian, K. N. *et al.* "Light Sterile Neutrinos: A White Paper". Apr. 2012. arXiv: 1204.5379 [hep-ph].

[84] Mohapatra, R. N. and Smirnov, A. Y. "Neutrino Mass and New Physics". *Annual Review of Nuclear and Particle Science* 56.1 (Nov. 2006), pp. 569–628.

[85] Weinberg, S. "Baryon- and Lepton-Nonconserving Processes". *Physical Review Letters* 43.21 (Nov. 1979), pp. 1566–1570.

[86] Broncano, A., Gavela, M. B., and Jenkins, E. "Neutrino physics in the seesaw model". *Nuclear Physics* ***B*** 672.1 (Nov. 2003), pp. 163–198.

[87] King, S. F. "Neutrino mass models". *Reports on Progress in Physics* 67.2 (Dec. 2003), p. 107.

[88] Miranda, O. G. and Valle, J. W. F. "Neutrino oscillations and the seesaw origin of neutrino mass". *Nuclear Physics* ***B*** 908 (Jan. 2016), pp. 436–455.

[89] Antusch, Stefan and Biggio, Carla and Fernández-Martínez, Enrique and Gavela, M. Belen and López-Pavón, Jacobo. "Unitarity of the Leptonic Mixing Matrix". *Journal of High Energy Physics* 2006.10 (Oct. 2006), p. 084.

[90] Antusch, S., Baumann, J. P., and Fernández-Martínez, E. "Non-standard neutrino interactions with matter from physics beyond the Standard Model". *Nuclear Physics* ***B*** 810.1 (Mar. 2009), pp. 369–388.

[91] Malinský, M., Ohlsson, T., and Zhang, H. "Nonunitarity effects in a realistic low-scale seesaw model". *Physical Review* ***D*** 79.7 (Apr. 2009), p. 073009.

[92] Antusch, S. and Fischer, O. "Non-unitarity of the leptonic mixing matrix: Present bounds and future sensitivities". *Journal of High Energy Physics* 2014.10 (July 2014), p. 094.

[93] Basso, L., Fischer, O., and van der Bij, J. J. "Precision tests of unitarity in leptonic mixing". *Europhysics Letters* 105.1 (Jan. 2014), p. 11001.

[94] Soumya, C. "Probing nonunitary neutrino mixing via long-baseline neutrino oscillation experiments based at J-PARC". *Physical Review* ***D*** 105.1 (Jan. 2022), p. 015012.

[95] Chatterjee, S. S. *et al.* "Nonunitarity of the lepton mixing matrix at the European Spallation Source". *Physical Review* ***D*** 106.7 (Oct. 2022), p. 075016.

[96] Miranda, O. G. and Nunokawa, H. "Non standard neutrino interactions: current status and future prospects". *New Journal of Physics* 17.9 (Sept. 2015), p. 095002.

[97] Bhupal Dev, P. S. *et al.* "Neutrino non-standard interactions: A status report". *SciPost Physics Proceedings* 2 (Dec. 2019), p. 001.

[98] Drewes, M. and Garbrecht, B. "Combining Experimental and Cosmological Constraints on Heavy Neutrinos". *Nuclear Physics* ***B*** 921 (Feb. 2015), pp. 250–315.

[99] Boyarsky, A. *et al.* "Sterile Neutrino Dark Matter". *Progress in Particle and Nuclear Physics* 104 (Jan. 2019), pp. 1–45.

[100] Zhang, H. and Zhou, S. "The minimal seesaw model at the TeV scale". *Physics Letters* ***B*** 685.4 (Mar. 2010), pp. 297–301.

[101] Mohapatra, R. N. and Valle, J. W. F. "Neutrino mass and baryon-number nonconservation in superstring models". *Physical Review* ***D*** 34.5 (Sept. 1986), pp. 1642–1645.

[102] Barr, S. M. "New Type of Seesaw Mechanism for Neutrino Masses". *Physical Review Letters* 92.10 (Mar. 2004), p. 101601.

[103] Malinský, M., Romão, J. C., and Valle, J. W. F. "Supersymmetric $SO(10)$ Seesaw Mechanism with low $B - L$ Scale". *Physical Review Letters* 95.16 (Oct. 2005), p. 161801.

[104] Deppisch, F. F. *et al.* "Reconciling the 2 TeV excesses at the LHC in a linear seesaw left-right model". *Physical Review* ***D*** 93.1 (Jan. 2016), p. 013011.

[105] Kersten, J. and Smirnov, A. Y. "Right-handed neutrinos at CERN LHC and the mechanism of neutrino mass generation". *Physical Review* ***D*** 76.7 (Oct. 2007), p. 073005.

[106] Ibarra, A., Molinaro, E., and Petcov, S. T. "Low energy signatures of the TeV scale seesaw mechanism". *Physical Review* ***D*** 84.1 (July 2011), p. 013005.

[107] Asaka, T. and Shaposhnikov, M. "The νMSM, dark matter and baryon asymmetry of the universe". *Physics Letters* ***B*** 620.7 (July 2005), pp. 17–26.

[108] Shaposhnikov, M. "The νMSM, dark matter and neutrino masses". *Journal of Physics: Conference Series* 39.1 (May 2006), p. 176.

[109] Gorbunov, D. and Shaposhnikov, M. "How to find neutral leptons of the νMSM?" *Journal of High Energy Physics* 2007.10 (Oct. 2007), p. 015.

[110] Shaposhnikov, M. "The νMSM, leptonic asymmetries, and properties of singlet fermions". *Journal of High Energy Physics* 2008.8 (Aug. 2008), p. 008.

[111] Araki, T. and Li, Y. F. "Q_6 flavor symmetry model for the extension of the minimal standard model by three right-handed sterile neutrinos". *Physical Review* ***D*** 85.6 (Mar. 2012), p. 065016.

[112] Drewes, M. *et al.* "Testing the low scale seesaw and leptogenesis". *Journal of High Energy Physics* 2017.8 (Aug. 2017), p. 18.

[113] Tremaine, S. and Gunn, J. E. "Dynamical Role of Light Neutral Leptons in Cosmology". *Physical Review Letters* 42.6 (Feb. 1979), pp. 407–410.

[114] Markovič, K. *et al.* "Constraining warm dark matter with cosmic shear power spectra". *Journal of Cosmology and Astroparticle Physics* 2011.1 (Jan. 2011), p. 022.

[115] Canetti, L. *et al.* "Dark matter, baryogenesis and neutrino oscillations from right-handed neutrinos". *Physical Review* ***D*** 87.9 (May 2013), p. 093006.

[116] Adikhari, R. *et al.* "A White Paper on keV sterile neutrino Dark Matter". *Journal of Cosmology and Astroparticle Physics* 2017.1 (Jan. 2017), p. 025.

[117] Fukugita, M. and Yanagida, T. "Baryogenesis without grand unification". *Physics Letters* ***B*** 174.1 (June 1986), pp. 45–47.

[118] Dine, M. and Kusenko, A. "Origin of the matter-antimatter asymmetry". *Reviews of Modern Physics* 76.1 (Dec. 2003), pp. 1–30.

[119] Sakharov, A. D. "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe". *Soviet Physics Uspeckhi* 34.5 (May 1991), p. 392.

[120] Shuve, B. and Yavin, I. "Baryogenesis through neutrino oscillations: A unified perspective". *Physical Review* ***D*** 89.7 (Apr. 2014), p. 075014.

[121] Moffat, K., Pascoli, S., and Weiland, C. "Equivalence between massless neutrinos and lepton number conservation in fermionic singlet extensions of the Standard Model". Dec. 2017. arXiv: `1712.07611 [hep-ph]`.

[122] Anamiati, G., Fonseca, R. M., and Hirsch, M. "Quasi-Dirac neutrino oscillations". *Physical Review* ***D*** 97.9 (May 2018), p. 095008.

[123] Hernández, P., Jones-Pérez, J., and Suarez-Navarro, O. "Majorana vs pseudo-Dirac neutrinos at the ILC". *The European Physcial Journal* ***C*** 79.3 (Mar. 2019), p. 220.

[124] Bolton, P. D., Deppisch, F. F., and Dev, P. S. B. "Neutrinoless double beta decay versus other probes of heavy sterile neutrinos". *Journal of High Energy Physics* 2020.3 (Mar. 2020), p. 170.

[125] Bondarenko, K. *et al.* "An allowed window for heavy neutral leptons below the kaon mass". *Journal of High Energy Physics* 2021.7 (Jan. 2021), p. 193.

[126] Tastet, J. L., Ruchayskiy, O., and Timiryasov, I. "Reinterpreting the ATLAS bounds on heavy neutral leptons in a realistic neutrino oscillation model". *Journal of High Energy Physics* 2021.12 (July 2021), p. 182.

[127] Abada, A. *et al.* "Collider searches for heavy neutral leptons: beyond simplified scenarios". *The European Physcial Journal* ***C*** 82.11 (Nov. 2022), p. 1030.

[128] Bondarenko, K. *et al.* "Phenomenology of GeV-scale heavy neutral leptons". *Journal of High Energy Physics* 2018.11 (Nov. 2018), p. 32.

[129] Coloma, P. *et al.* "GeV-scale neutrinos: interactions with mesons and DUNE sensitivity". *The European Physcial Journal* ***C*** 81.1 (Jan. 2021), p. 78.

[130] Coloma, P. *et al.* "Double-Cascade Events from New Physics in Icecube". *Physical Review Letters* 119.20 (Nov. 2017), p. 201804.

[131] Magill, G. *et al.* "Dipole portal to heavy neutral leptons". *Physical Review* ***D*** 98.11 (Dec. 2018), p. 115015.

[132] Plestid, R. "Luminous solar neutrinos. I. Dipole portals". *Physical Review* ***D*** 104.7 (Oct. 2021), p. 075027.

[133] Gustafson, R. A., Plestid, R., and Shoemaker, I. M. "Neutrino portals, terrestrial upscattering, and atmospheric neutrinos". *Physical Review* ***D*** 106.9 (Nov. 2022), p. 095037.

[134] Aguilar-Arevalo, A. A. *et al* (MiniBooNE Collaboration). "Significant Excess of Electronlike Events in the MiniBooNE Short-Baseline Neutrino Experiment". *Physical Review Letters* 121.22 (Nov. 2018), p. 221801.

[135] Gninenko, S. N. "Resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments". *Physical Review* ***D*** 83.1 (Jan. 2011), p. 015015.

[136] Bertuzzo, E. *et al.* "Dark Neutrino Portal to Explain MiniBooNE Excess". *Physical Review Letters* 121.24 (Dec. 2018), p. 241801.

[137] Kamp, N. W. *et al.* "Dipole-coupled heavy-neutral-lepton explanations of the MiniBooNE excess including constraints from MINERvA data". *Physical Review* ***D*** 107.5 (Mar. 2023), p. 055009.

[138] Abratenko, P. *et al* (MicroBooNE Collaboration). "Search for an anomalous excess of charged-current quasielastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction". *Physical Review* ***D*** 105.11 (June 2022), p. 112003.

[139] Abratenko, P. *et al* (MicroBooNE Collaboration). "Search for an anomalous excess of charged-current quasielastic ν_e interactions without pions in the final state with the MicroBooNE experiment". *Physical Review* ***D*** 105.11 (June 2022), p. 112004.

[140] Abratenko, P. *et al* (MicroBooNE Collaboration). "Search for an anomalous excess of inclusive charged-current quasielastic ν_e interactions in the MicroBooNE experiment using Wire-Cell reconstruction". *Physical Review* ***D*** 105.11 (June 2022), p. 112005.

[141] Argüelles, C. A. *et al.* "MicroBooNE and the ν_e Interpretation of the MiniBooNE Low-Energy Excess". *Physical Review Letters* 128.24 (June 2022), p. 241802.

[142] Shrock, R. "General theory of weak processes involving neutrinos. I. Leptonic pseudoscalar-meson decays, with associated test for, and bounds on, neutrino masses and lepton mixing". *Physical Review* ***D*** 24.5 (Sept. 1981), pp. 1232–1274.

[143] Ballett, P., Pascoli, S., and Ross-Lonergan, M. "MeV-scale sterile neutrino decays at the Fermilab Short-Baseline Neutrino Program". *Journal of High Energy Physics* 2017.4 (Apr. 2017), p. 102.

[144] Ballett, P., Boschi, T., and Pascoli, S. "Heavy Neutral Leptons from low-scale seesaws at the DUNE Near Detector". *Journal of High Energy Physics* 2020.3 (Mar. 2020), p. 111.

[145] Coloma, P. *et al.* "New constraints on Heavy Neutral Leptons from Super-Kamiokande data". *The European Physcial Journal* ***C*** 80.3 (Mar. 2020), p. 235.

[146] Plestid, R. "Luminous solar neutrinos. II. Mass-mixing portals". *Physical Review* ***D*** 104.7 (Oct. 2021), p. 075028.

[147] Pilaftsis, A. "Radiatively induced neutrino masses and large Higgs-neutrino couplings in the Standard Model with Majorana fields". *Zeitschrift für Physik* ***C*** 55.2 (June 1992), pp. 275–282.

[148] Datta, A., Guchait, M., and Pilaftsis, A. "Probing lepton number violation via Majorana neutrinos at hadron supercolliders". *Physical Review* ***D*** 50.5 (Sept. 1994), pp. 3195–3203.

[149] Das, A., Konar, P., and Majhi, S. "Production of heavy neutrino in next-to-leading order QCD at the LHC and beyond". *Journal of High Energy Physics* 2016.6 (June 2016), p. 19.

[150] Degrande, C. *et al.* "Fully automated precision predictions for heavy neutrino production mechanisms at hadron colliders". *Physical Review* ***D*** 94.5 (Sept. 2016), p. 053002.

[151] Ruiz, R., Spannowsky, M., and Waite, P. "Heavy neutrinos from gluon fusion". *Physical Review* ***D*** 96.5 (Sept. 2017), p. 055042.

[152] Kopp, J. *et al.* "Sterile neutrino oscillations: the global picture". *Journal of High Energy Physics* 2013.5 (May 2013), p. 50.

[153] Sabti, N., Magalich, A., and Filimonova, A. "An extended analysis of Heavy Neutral Leptons during Big Bang Nucleosynthesis". *Journal of Cosmology and Astroparticle Physics* 2020.11 (Nov. 2020), p. 056.

[154] Boyarsky, A. *et al.* "Improved big bang nucleosynthesis constraints on heavy neutral leptons". *Physical Review* ***D*** 104.2 (July 2021), p. 023517.

[155] de Gouvêa, A. *et al.* "Characterizing heavy neutral fermions via their decays". *Physical Review* ***D*** 105.1 (Jan. 2022), p. 105019.

[156] Abdullahi, A. M. *et al.* "The present and future status of heavy neutral leptons". *Journal of Physics* ***G*** 50.2 (Jan. 2023), p. 020501.

[157] Bellini, G *et al* (Borexino Collaboration). "New limits on heavy sterile neutrino mixing in ^{8}B decay obtained with the Borexino detector". *Physical Review* ***D*** 88.7 (Oct. 2013), p. 072010.

[158] Aad, G. *et al* (ATLAS Collaboration). "Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector". *Journal of High Energy Physics* 2019.10 (Oct. 2019), p. 265.

[159] Abe, K. *et al* (T2K Collaboration). "Search for heavy neutrinos with the T2K near detector ND280". *Physical Review* ***D*** 100.5 (Sept. 2019), p. 052006.

[160] Abratenko, P. *et al* (MicroBooNE Collaboration). "Search for heavy neutral leptons decaying into muon-pion pairs in the MicroBooNE detector". *Physical Review* ***D*** 101.5 (Nov. 2020), p. 052001.

[161] Acciarri, R. *et al* (ArgoNeuT Collaboration). "New Constraints on Tau-Coupled Heavy Neutral Leptons with Masses $m_N = 280 - -970$ MeV". *Physical Review Letters* 127 (Dec. 2021), p. 121801.

[162] Aaij, R. *et al* (LHCb Collaboration). "Search for heavy neutral leptons in $W^+ \to \mu^+\mu^\pm$ jet decays". *The European Physcial Journal* ***C*** 81.3 (Mar. 2021), p. 248.

[163] Kelly, K. J. and Machado, P. A. N. "The MicroBooNE Experiment, the NuMI Absorber, and Heavy Neutral Leptons". *Physical Review* ***D*** 104.5 (Sept. 2021), p. 055015.

[164] Tumasyan, A. *et al* (CMS Collaboration). "Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV". *Journal of High Energy Physics* 2022.7 (July 2022), p. 81.

[165] Bernardi, G. *et al.* "Search for neutrino decay". *Physics Letters* ***B*** 166.4 (Jan. 1986), pp. 479–483.

[166] Bernardi, G. *et al.* "Further limits on heavy neutrino couplings". *Physics Letters* ***B*** 203.3 (Mar. 1988), pp. 332–334.

[167] Artamonov, A. V. *et al* (E949 Collaboation). "Search for heavy neutrinos in $K^+ \to \mu^+ \nu_H$ decays". *Physical Review* ***D*** 91.5 (Mar. 2015), p. 052001.

[168] Aguilar-Arevalo, A. *et al* (PIENU Collaboration). "Improved search for heavy neutrinos in the decay $\pi \to e\ \nu$". *Physical Review* ***D*** 97.7 (Apr. 2018), p. 072012.

[169] Aguilar-Arevalo, A. *et al.* "Search for heavy neutrinos in $\pi \to \mu\nu$ decay". *Physics Letters* ***B*** 798 (Nov. 2019), p. 134980.

[170] Bryman, D. A. and Shrock, R. "Constraints on sterile neutrinos in the MeV to GeV mass range". *Physical Review* ***D*** 100.7 (Oct. 2019), p. 073011.

[171] Fernandez-Martinez, E., Hernandez-Garcia, J., and Lopez-Pavon, J. "Global constraints on heavy neutrino mixing". *Journal of High Energy Physics* 2016.8 (Aug. 2016), p. 33.

[172] Abada, A. and Teixeira, A. M. "Heavy Neutral Leptons and High-Intensity Observables". *Frontiers in Physics* 6 (Dec. 2018), p. 142.

[173] Tommasini, D. *et al.* "Non-decoupling of heavy neutrinos and lepton flavour violation". *Nuclear Physics* ***B*** 444.3 (June 1995), pp. 451–467.

[174] Alonso, R. *et al.* "Muon conversion to electron in nuclei in type-I seesaw models". *Journal of High Energy Physics* 2013.1 (Jan. 2013), p. 118.

[175] Cline, J., Puel, M., and Toma, T. "A little theory of everything, with heavy neutral leptons". *Journal of High Energy Physics* 2020.5 (May 2020), p. 39.

[176] Abe, K. *et al* (Hyper-Kamiokande Proto-Collaboration). "Hyper-Kamiokande Design Report". Nov. 2018. arXiv: `1805.04163 [ins-det]`.

[177] Kling, F. and Trojanowski, S. "Heavy neutral leptons at FASER". *Physical Review* ***D*** 97.9 (May 2018), p. 095016.

[178] Ahdida, C. *et al* (SHiP Collaboration). "Sensitivity of the SHiP experiment to Heavy Neutral Leptons". *Journal of High Energy Physics* 2019.4 (Apr. 2019), p. 77.

[179] Gorbunov, D. *et al.* "Heavy Neutral Leptons from kaon decays in the SHiP experiment". *Physics Letters* ***B*** 810 (Apr. 2020), p. 135817.

[180] Wang, Z. S. and Wang, K. "Physics with far detectors at future lepton colliders". *Physical Review* ***D*** 101.7 (Apr. 2020), p. 075046.

[181] Abi, B. *et al* (DUNE Collaboration). "Prospects for beyond the Standard Model physics searches at the Deep Underground Neutrino Experiment". *The European Physcial Journal* ***C*** 81.4 (Apr. 2021), p. 322.

[182] Carbajal, S. and Gago, A. M. "Indirect search of Heavy Neutral Leptons using the DUNE Near Detector". Feb. 2022. arXiv: `2202.09217 [hep-ph]`.

[183] Cerci, S. *et al.* "FACET: a new long-lived particle detector in the very forward region of the CMS experiment". *Journal of High Energy Physics* 2022.6 (June 2022), p. 110.

[184] Feng, J. L. *et al.* "The Forward Physics Facility at the High-Luminosity LHC". *Journal of Physics* ***G*** 50.3 (Jan. 2023), p. 030501.

[185] van der Meer, S. *A directive device for charged particles and its use in an enhanced neutrino beam.* CERN 61-7. Feb. 1961.

[186] Dore, U., Loverre, P., and Ludovici, L. "History of accelerator neutrino beams". *The European Physical Journal* ***H*** 44.4-5 (Oct. 2019), pp. 271–305.

[187] Kopp, S. E. "Accelerator neutrino beams". *Physics Reports* 439.3 (Feb. 2007), pp. 101–159.

[188] Adamson, P. *et al.* "The NuMI neutrino beam". *Nuclear Instruments and Methods* ***A*** 806 (July 2016), pp. 279–306.

[189] Aliaga, L. "Neutrino Flux Prediction for the NuMI Beamline". PhD thesis. Williamsburg, Virginia: College of William and Mary, Jan. 2016.

[190] Whitehead, L. H. "Neutrino oscillations with MINOS and MINOS+". *Nuclear Physics* ***B*** 908 (July 2016), pp. 130–150.

[191] Adamson, P. *et al* (NOvA Collaboration). "First measurement of muon-neutrino disappearance in NOvA". *Physical Review* ***D*** 93.5 (Mar. 2016), p. 051104.

[192] Lu, X.-G. *et al* (MINERνA Collaboration). "Exploring neutrino-nucleus interactions in the GeV regime using MINERvA". *The European Physical Journal Special Topics* 230.24 (Dec. 2021), pp. 4243–4257.

[193] Guénette, R. "The ArgoNeuT experiment". Oct. 2011. arXiv: 1110.0443 [physics/ins-det].

[194] Abratenko, P. *et al* (MicroBooNE Collaboration). "Measurement of the flux-averaged inclusive charged-current electron neutrino and antineutrino cross section on argon using the NuMI beam and the MicroBooNE detector". *Physical Review* ***D*** 104.5 (Sept. 2021), p. 052002.

[195] Budal, K. "Measurement of Charge Emission from Targets as a Means of Burst Intensity and Beam Intensity Monitoring". *IEEE Transactions on Nuclear Science* 14.3 (June 1967), pp. 1132–1137.

[196] Ruterbories, D. *LE and ME Flux Plots.* MINERνA DocDB 27462-v4. July 2020.

[197] Aliaga, L. *et al* (MINERνA Collaboration). "Neutrino flux predictions for the NuMI beam". *Physical Review* ***D*** 94.9 (Nov. 2016), p. 092005.

[198] Hatcher, R. *Proposal for a Unified "Flux" N-tuple Format.* Tech. rep. 2012.

[199] Kordosky, M. *MINERvA detector position, the flux driver, and changes for eroica.* MINERνA DocDB 10414-v2. Oct. 2014.

[200] Messerly, B. *ME Focusing Unertainties.* MINERνA DocDB 12955-v2. Feb. 2017.

[201] Goldstein, H., Poole, C., and Safko, J. "Classical Mechanics". 3rd ed. Addison Wesley, 2001. Chap. 4.4, pp. 150–154.

[202] Buss, O. *et al.* "Transport-theoretical description of nuclear reactions". *Physics Reports* 512.1 (Mar. 2012), pp. 1–124.

[203] Paley, J. M. *et al* (MIPP Collaboration). "Measurement of charged pion production yields off the NuMI target". *Physical Review* ***D*** 90.3 (Aug. 2014), p. 032001.

[204] Alt, C. *et al* (NA Collaboration). "Inclusive production of charged pions in $p+\mathrm{C}$ collisions at 158 GeV/c beam momentum". *The European Physcial Journal* ***C*** 49.4 (Mar. 2007), pp. 897–917.

[205] Barton, D. S. *et al.* "Experimental studt of the A dependence of inclusive hadron fragmentation". *Physical Review* ***D*** 27.11 (June 1983), pp. 2580–2599.

[206] Feynman, R. P. "Very High-Energy Collisions of Hadrons". *Physical Review Letters* 23.24 (Dec. 1969), pp. 1415–1417.

[207] Ferrari, A. *et al. FLUKA: A multi-particle transport code (program version 2005).* CERN Yellow Reports: Monographs. Geneva: CERN, 2005.

[208] Böhlen, T. T. *et al.* "The FLUKA Code: Developments and Challenges for High Energy and Medical Applications". *Nuclear Data Sheets* 120 (June 2014), pp. 211–214.

[209] Messerly, B., Fine, R., and Olivier, A. "An Error Analysis Toolkit for Binned Counting Experiments". *EPJ Web of Conferences* 251 (Aug. 2021), p. 03046.

[210] Messerly, B. *ME Focusing Uncertainties*. MINERνA DocDB 16092-v2. Nov. 2017.

[211] Bashyal, A. *et al* (MINERνA Collaboration). "Use of neutrino scattering events with low hadronic recoil to inform neutrino flux and detector energy scale". *Journal of Instrumentation* 16.8 (Aug. 2021), P08068.

[212] Valencia, E *et al* (MINERνA Collaboration). "Constraint of the MINERνA medium energy neutrino flux using neutrino-electron elastic scattering". *Physical Review* ***D*** 100.9 (Nov. 2019), p. 092001.

[213] Zazueta, L. *et al* (MINERvA Collaboration). "Improved constraint on the MINERνA medium energy neutrino flux using $\overline{\nu}e^- \to \overline{\nu}e^-$ data". *Physical Review* ***D*** 107.1 (Jan. 2023), p. 012001.

[214] Ruterbories, D. *et al* (MINERνA Collaboration). "Constraining the NuMI neutrino flux using inverse muon decay reactions in MINERvA". *Physical Review* ***D*** 104.9 (Nov. 2021), p. 092010.

[215] Golan, T., Juszczak, C., and Sobczyk, J. T. "Effects of final-state interactions in neutrino-nucleus interactions". *Physical Review* ***C*** 86.1 (July 2012), p. 015505.

[216] Mosel, U. "Neutrino Interactions with Nucleons and Nuclei: Importance for Long-Baseline Experiments". *Annual Review of Nuclear and Particle Science* 66.1 (Oct. 2016), pp. 171–195.

[217] Coplowe, D. *et al* (MINERνA Collaboration). "Probing nuclear effects with neutrino-induced charged-current neutral pion production". *Physical Review* ***D*** 102.7 (Oct. 2020), p. 072007.

[218] Sajjad Athar, M., Fatima, A., and Singh, S. K. "Neutrinos and their interactions with matter". *Progress in Particle and Nuclear Physics* 129 (Mar. 2023), p. 104019.

[219] Rodrigues, P. A. *et al* (MINERνA Collaboration). "Identification of nuclear effects in neutrino-carbon interactions at low three-momentum transfer". *Physical Review Letters* 116 (Feb. 2016). [Addendum: *Physical Review Letters* 121 (Nov. 2018), p. 209902], p. 071802.

[220] Stowell, P. *et al* (MINERνA Collaboration). "Tuning the GENIE pion production model with MINERvA data". *Physical Review* ***D*** 100.7 (Oct. 2019), p. 072005.

[221] Ruterbories, D. *et al* (MINERνA Collaboration). "Measurement of inclusive charged-current ν_μ cross sections as a function of muon kinematics at $<E_\nu> \sim 6$ GeV on hydrocarbon". *Physical Review* ***D*** 104.9 (Nov. 2021), p. 092007.

[222] Tena-Vidal, Júlia *et al* (GENIE Collaboration). "Neutrino-nucleus CC 0π cross-section tuning in GENIE v3". *Physical Review* ***D*** 106.11 (Dec. 2022), p. 112001.

[223] Li, Weijun *et al* (GENIE Collaboration). "First combined tuning on transverse kinematic imbalance data with and without pion production constraints". Apr. 2024. arXiv: 2404.08510 [hep-ex].

[224] Aliaga, L. *et al.* "Design, calibration, and performance of the MINERvA detector". *Nuclear Instruments and Methods* ***A*** 743 (Apr. 2014), pp. 130–159.

[225] Messerly, B. "Single Charged Pion Production by Muon Neutrinos in the MINERvA Detector Using the NuMI Beam". PhD thesis. Pittsburgh, Pennsylvania: University of Pittsburgh, Dec. 2019.

[226] Ramírez Delgado, M. A. "ν_μ-Induced CC Coherent π^+ Production off Carbon, Hydrocarbon, Iron, and Lead using the MINERνA Detector". PhD thesis. León, Guanajato: Universidad de Guanajato, Jan. 2020.

[227] MINERνA Collaboration. *The MINERνA Technical Design Report v. 1.3*. MINERνA public DocDB 700-v28. Dec. 2006.

[228] Perdue, G. N. *et al.* "The MINERνA data acquisition system and infrastructure". *Nuclear Instruments and Methods* ***A*** 694 (2012), pp. 179–192.

[229] Kleykamp, J. *Minerva 101: Calibrations*. MINERνA DocDB 29171-v1. June 2020.

[230] Auchincloss, P. S. *et al.* "A study of the energy dependence of the mean, truncated mean, and most probable energy deposition of high-energy muons in sampling calorimeters". *Nuclear Instruments and Methods* ***A*** 343.2-3 (1994), pp. 463–469.

[231] Marshall, C. "Measurement of charged kaon production by neutrinos at MINERvA". PhD thesis. Rochester, New York: University of Rochester, Jan. 2016.

[232] Paulus, T. J. "Timing Electronics and Fast Timing Methods with Scintillation Detectors". *IEEE Transactions on Nuclear Science* 32.3 (June 1985), pp. 1242–1249.

[233] Perdue, G. N. *Reco Overview*. MINERνA DocDB 11243. July 2015.

[234] Eberly, B. "Characterization of Final State Interaction Strength in Plastic Scintillator by Muon-Neutrino Charged Current Charged Pion Production". PhD thesis. Pittsburgh, Pennsylvania: University of Pittsburgh, May 2014.

[235] Workman, R. L. *et al* (Particle Data Group). "Review of Particle Physics Section 34: Passage of Particles Through Matter". *Progress of Theoretical and Experimental Physics* 2022 (2022), p. 083C01.

[236] "Guide for Physics Lists 11.1 (doc Rev7.1)". `https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/html/index.html`. Accessed 25-09-2023.

[237] Apostolakis, J. e. a. "Geometry and physics of the Geant4 toolkit for high and medium energy applications". *Radiation Physics and Chemistry* 78.10 (Oct. 2009), pp. 859–873.

[238] Tagg, N. *et al* (MINERνA Collaboration). "Arachne - A web-based event viewer for MINERνA". *Nuclear Instruments and Methods* ***A*** 676 (June 2012), pp. 44–49.

[239] Birks, J. B. "Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations". *Proceedings of the Physical Society* ***A*** 64.10 (Oct. 1951), pp. 874–877.

[240] Aliaga, L. *et al* (MINERνA Collaboration). "MINERvA neutrino detector response measured with test beam data". *Nuclear Instruments and Methods* ***A*** 789 (July 2015), pp. 28–42.

[241] Andreopoulos, C. *et al.* "The GENIE neutrino Monte Carlo generator". *Nuclear Instruments and Methods* ***A*** 614.1 (Feb. 2010), pp. 87–104.

[242] Andreopoulos, C. *et al.* "The GENIE Neutrino Monte Carlo Generator: Physics and User Manual". Oct. 2015. arXiv: `1510.05494 [hep-ex]`.

[243] Plows, K.-J. and Lu, X.-G. "Modeling heavy neutral leptons in accelerator beamlines". *Physical Review* ***D*** 107.5 (Mar. 2023), p. 055003.

[244] Alloul *et al*, A. "FeynRules 2.0 — A complete toolbox for tree-level phenomenology". *Computer Physics Communications* 185.8 (Aug. 2014), pp. 2250–2300.

[245] "FeynRules model database". `https://feynrules.irmp.ucl.ac.be/`. Accessed 29-08-2022.

[246] Mertig, R., Böhm, M., and Denner, A. "Feyn Calc - Computer-algebraic calculation of Feynman amplitudes". *Computer Physics Communications* 64.3 (June 1991), pp. 345–359.

[247] Shtabovenko, V., Mertig, R., and Orellana, F. "New developments in FeynCalc 9.0". *Computer Physics Communications* 207 (Oct. 2016), pp. 432–444.

[248] Shtabovenko, V., Mertig, R., and Orellana, F. "FeynCalc 9.3: New features and improvements". *Computer Physics Communications* 256 (Nov. 2020), p. 107478.

[249] Brun, R. and Rademakers, F. "ROOT - An object oriented data analysis framework". *Nuclear Instruments and Methods* ***A*** 389.1-2 (Apr. 1997), pp. 81–86.

[250] James, F. E. *Monte Carlo phase space*. CERN Academic Training Lecture. CERN, Geneva, 1 May 1968. Geneva: CERN, 1968.

[251] Rubbia, A. Cambridge University Press, 2022. Chap. 5, pp. 144–183.

[252] Levy, J.-M. "Rates and differential distributions in heavy neutral leptons production and decays". May 2018. arXiv: `1805.06419 [hep-ph]`.

[253] Formaggio, J. A. *et al.* "Helicity effects in neutral heavy lepton decays". *Physical Review* ***D*** 57.11 (June 1998), pp. 7037–7040.

[254] de Gouvêa, A. *et al.* "Three-body decays of heavy Dirac and Majorana fermions". *Physical Review* ***D*** 104.1 (July 2021), p. 015038.

[255] Adamson, P. *et al* (MINOS Collaboration, NIST, and USNO). "Precision measurement of the speed of propagation of neutrinos using the MINOS detectors". *Physical Review* ***D*** 92.5 (Sept. 2015), p. 052005.

[256] Abratenko, P. *et al* (MicroBooNE Collaboration). "First demonstration of $\mathcal{O}$(1 ns) timing resolution in the MicroBooNE liquid argon time projection chamber". *Physical Review* ***D*** 108.5 (Sept. 2023), p. 052010.

[257] Agostinelli, S. *et al.* "GEANT4 - a simulation toolkit". *Nuclear Instruments and Methods* ***A*** 506.3 (July 2003), pp. 250–303.

[258] Allison, J. *et al.* "Geant4 Developments and Applications". *IEEE Transactions on Nuclear Science* 53.1 (Feb. 2006), pp. 270–278.

[259] Allison, J. *et al.* "Recent developments in Geant4". *Nuclear Instruments and Methods* ***A*** 835 (Nov. 2016), pp. 186–225.

[260] Yonehara, K. "Megawatt upgrade of NuMI target system". *Proceedings of Science* NuFact2021 (Jan. 2022), p. 107.

[261] Illingworth, R. A. "A data handling system for modern and future Fermilab experiments". *Journal of Physics: Conference Series* 513.3 (June 2014), p. 032045.

[262] Formaggio, J. A. and Zeller, G. P. "From eV to EeV: Neutrino cross sections across energy scales". *Reviews of Modern Physics* 84.3 (May 2012), pp. 1307–1341.

[263] Rein, D. and Sehgal, L. M. "Neutrino-excitation of baryon resonances and single pion production". *Annals of Physics* 133.1 (Apr. 1981), pp. 79–153.

[264] Rein, D. and Sehgal, L. M. "PCAC and the deficit of forward muons in π^+ production by neutrinos". *Physics Letters* ***B*** 657.4-5 (Dec. 2007), pp. 207–209.

[265] Wolcott, J. *et al* (MINERνA Collaboration). "Evidence for Neutral-Current Diffractive π^0 Production from Hydrogen in Neutrino Interactions on Hydrocarbon". *Physical Review Letters* 117.11 (Sept. 2016), p. 111801.

[266] Merenyi, R. *et al.* "Determination of pion intranuclear rescattering rates in $\nu_\mu - \mathrm{Ne}$ versus $\nu_\mu - \mathrm{D}$ interactions for the atmospheric ν flux". *Physical Review* ***D*** 45.3 (Feb. 1992), pp. 743–751.

[267] Alvarez-Ruso, L. *et al* (GENIE Collaboration). "Recent highlights from GENIE v3". *The European Physical Journal Special Topics* 230.24 (Dec. 2021), pp. 4449–4467.

[268] Hayato, Y. "Neut". *Nuclear Physics* ***B*** *- Proceedings Supplements* 112.1-3, pp. 171–176.

[269] Cugnon, J. *et al.* "Processes involving few degrees of freedom in the frame of Intranuclear Cascade approaches". *The European Physical Journal Plus* 131.169 (May 2016), p. 169.

[270] Gallagher, H. R. "Neutrino event generators: A review". *AIP Conference Proceedings* 1189.1 (Nov. 2009), pp. 35–42.

[271] Mosel, U. "Neutrino event generators: foundation, status and future". *Journal of Physics* ***G*** 46.11 (Sept. 2019), p. 113001.

[272] Mislivec, A. *et al* (MINERνA Collaboration). "Measurement of total and differential cross sections of neutrino and antineutrino coherent $\pi^{\pm}$ production on carbon". *Physical Review* ***D*** 97.3 (Feb. 2018), p. 032014.

[273] Ramírez, M. A. *et al* (MINERνA Collaboration). "Neutrino-Induced Coherent π^{+} Production in C, CH, Fe, and Pb at $<E_{\nu}> \sim 6$ GeV". *Physical Review Letters* 131.5 (Aug. 2023), p. 051801.

[274] Rein, D. "Diffractive pion production in neutrino reactions". *Nuclear Physics* ***B*** 278.1 (Dec. 1986), pp. 61–77.

[275] Groom, D. E., Mokhov, N. V., and Striganov, S. I. "Muon Stopping Power and Range Tables 10 MeV – 100 TeV". *Atomic Data and Nuclear Data Tables* 78.2 (July 2001), pp. 183–356.

[276] Barkas, W. H., Birnbaum, W., and Smith, F. M. "Mass-Ratio Method Applied to the Measurement of *L*-Meson Masses and the Energy Balance in Pion Decay". *Physical Review* 101.2 (Jan. 1956), p. 778.

[277] Maronna, R. A., Yohai, V. J., and R. Douglas, M. *Robust statistics: theory and methods.* 1st ed. Wiley series in probability and statistics. Chichester: John Wiley & Sons, 2006.

[278] Lu, Xianguo. "Exploring the performance limits of the ALICE Time Projection Chamber and Transition Radiation Detector for measuring identified hadron production at the LHC". PhD thesis. Heidelberg, Germany: Ruprecht-Karls-Universität Heidelberg, Oct. 2013.

[279] Rolandi, L., Riegler, W., and Blum, W. "Particle Identification by Measurement of Ionization". *Particle Detection with Drift Chambers.* Berlin, Germany: Springer Berlin Heidelberg, 2008.

[280] Athar, M. S., Fatima, A., and Singh, S. K. "Neutrinos and their interactions with matter". *Progress in Particle and Nuclear Physics* 129 (Mar. 2023), p. 104019.

[281] Abe, K. *et al* (T2K Collaboration). "Measurements of the ν_{μ} and $\overline{\nu}_{\mu}$-induced coherent charged pion production cross sections on ^{12}C by the T2K experiment". *Physical Review* ***D*** 108.9 (Nov. 2023), p. 092009.

[282] Kopeliovich, B. Z. "PCAC and shadowing of low energy neutrinos". *Nuclear Physics* ***B*** - *Proceedings Supplements* 139 (Feb. 2005), pp. 219–225.

[283] Rein, D. and Sehgal, L. M. "Coherent π^{0} production in neutrino reactions". *Nuclear Physics* ***B*** 223.1 (Aug. 1983), pp. 29–44.

[284] Berger, C. and Sehgal, L. M. "Partially conserved axial vector current and coherent pion production by low energy neutrinos". *Physical Review* ***D*** 79.5 (Mar. 2009), p. 053003.

[285] Paschos, E. A. and Schalla, Dario. "Coherent pion production by neutrinos". *Physical Review* ***D*** 80.3 (Aug. 2009), p. 033005.

[286] Higuera, A. and Paschos, E. A. "Angular dependence of coherent pion production by neutrinos". *The European Physical Journal Plus* 129.43 (Mar. 2014), p. 43.

[287] Higuera, A. *et al* (MINERνA Collaboration). "Measurement of Coherent Production of $\pi^{\pm}$ in Neutrino and Antineutrino Beams on Carbon from E_{ν} of 1.5 to 20 GeV". *Physical Review Letters* 113.26 (Dec. 2014), p. 261802.

[288] Llewellyn Smith,C. H. "Neutrino reactions at accelerator energies". *Physics Reports* 3.5 (June 1972), pp. 261–379.

[289] Bradford, R. *et al.* "A New Parametrization of the Nucleon Elastic Form Factors". *Nuclear Physics* ***B*** *- Proceedings Supplements* 159 (Sept. 2006), pp. 127–132.

[290] Bodek, A., Park, I., and Yang, U. k. "Improved low Q^2 moel for neutrino and electron nucleon cross sections in few GeV region". *Nuclear Physics* ***B*** *- Proceedings Supplements* 139 (Feb. 2005), pp. 113–118.

[291] Nieves, J., Ruiz Simo, I., and Vicente Vacas, M. J. "Inclusive charged-current neutrino-nucleus reactions". *Physical Review* ***C*** 83.4 (Apr. 2011), p. 045501.

[292] Gran, R. *et al.* "Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV". *Physical Review* ***D*** 88.11 (Dec. 2013), p. 113007.

[293] Schwehr, J., Cherdack, D., and Gran, R. "GENIE implementation of IFIC Valencia model for QE-like 2p2h neutrino-nucleus cross section". Jan. 2016. arXiv: 1601.02038 [hep-ph].

[294] Hewett, J. L. *et al.* "Fundamental Physics at the Intensity Frontier". May 2012. arXiv: 1205.2671 [hep-ex].

[295] Gallagher, H. "The NEUGEN neutrino event generator". *Nuclear Physics* ***B*** *- Proceedings Supplements* 1-3 (Nov. 2002), pp. 188–194.

[296] Kuzmin, K. S., Lyubushkin, V. V., and Naumov, V. A. "How to Sum Contributions into the Total Charged-Current Neutrino-Nucleon Cross Section". Nov. 2005. arXiv: 0511308 [hep-ph].

[297] Smith, R. A. and Moniz, E. J. "Neutrino reactions on nuclear targets". *Nuclear Physics* ***B*** 43 (June 1972), pp. 605–622.

[298] Bodek, A. and Ritchie, J. L. "Further studies of Fermi-motion effects in lepton scattering from nuclear targets". *Physical Review* ***D*** 24.5 (Sept. 1981), pp. 1400–1402.

[299] Gran, R. "Model uncertainties for Valencia RPA effect for MINERvA". May 2017. arXiv: 1705.02932 [hep-ex].

[300] Rodrigues, P., Wilkinson, C., and McFarland, K. "Constraining the GENIE model of neutrino-induced single pion production using reanalyzed bubble chamber data". *The European Physcial Journal* ***C*** 76 (Aug. 2016), p. 474.

[301] Pearson, K. "X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling". *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 50.302 (1900), pp. 157–175.

[302] Cochran, W. G. "Some methods for strengthening the common χ^2 tests". *Biometrics* 10.4 (Dec. 1954), pp. 417–451.

[303] James, F. *Statistical Methods in Experimental Physics*. 2nd. WORLD SCIENTIFIC, 2006.

[304] James, F. and Roos, M. "Minuit - a system for function minimization and analysis of the parameter errors and correlations". *Computer Physics Communications* 10.6 (Dec. 1975), pp. 343–367.

[305] James, F. "MINUIT: Function Minimization and Error Analysis Reference Manual". https://cds.cern.ch/record/2296388. CERN Program Library Long Writeups. 1998.

[306] Wilks, S. S. "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses". *The Annals of Mathematical Statistics* 9.1 (Mar. 1938), pp. 60–62.

[307] Feldman, G. J. and Cousins, R. D. "Unified approach to the classical statistical analysis of small signals". *Physical Review* ***D*** 57.7 (Apr. 1998), pp. 3873–3889.

[308] Read, A. L. "Modified frequentist analysis of search results (the CL_s method)". CERN-OPEN-2000-205. 2000.

[309] Read, A. L. "Presentation of search results: the CL_s technique". *Journal of Physics* ***G*** 28.10 (Oct. 2002), pp. 2693–2704.

[310] Heinrich, J. "The Bayesian approach to setting limits: what to avoid". *Statistical Problems in Particle Physics, Astrophysics and Cosmology*. Ed. by. L. Lyons and. M. K. Ünel. PHYSTAT 05. Imperial College Press, 2006.

[311] Wu, Q. and Vos, P. "Inference and Prediction". *Handbook of Statistics*. Ed. by. V. N. Gudivada and. C. R. Rao. Vol. 38. 2018. Chap. 6, pp. 111–172.

[312] Ballett, P. *et al.* "Neutrino trident scattering at near detectors". *Journal of High Energy Physics* 2019 (Jan. 2019), p. 119.

[313] Abratenko, P. *et al* (MicroBooNE Collaboration). "Search for Heavy Neutral Leotons in Electron-Positron and Neutral-Pion Final States with the MicroBooNE Detector". *Physical Review Letters* 132.4 (Jan. 2024), p. 041801.

[314] Coloma, P. *et al.* "New Physics searches using ProtoDUNE and the CERN SPS accelerator". Jan. 2023. arXiv: 2304.06765 [hep-ph].

[315] Batell, B., Huang, W., and Kelly, K. J. "Keeping it simple: simplified frameworks for long-lived particles at neutrino facilities". *Journal of High Energy Physics* 2023 (Aug. 2023), p. 92.

[316] Coloma, P., Martín-Albo, J., and Urrea, S. "Discovering long-lived particles at DUNE". *Physical Review* ***D*** 109.3 (Feb. 2024), p. 035013.

[317] Berryman, J. M. *et al.* "Searches for decays of new particles in the DUNE Multi-Purpose near Detector". *Journal of High Energy Physics* 2020 (Feb. 2020), p. 174.

[318] Yang, K. "Measurement of the pion charge exchange differential cross section on Argon with the PROTODUNE-SP detector". PhD thesis. Oxford, United Kingdom: University of Oxford, Oct. 2023.

[319] Hamacher-Baumann, P., Lu, X., and Martín-Albo, J. "Neutrino-hydrogen interactions with a high-pressure time projection chamber". *Physical Review* ***D*** 102.3 (Aug. 2020), p. 033005.

[320] Battisti, Federico and Ivanov, Marian and Lu, Xianguo. "A Kalman Filter for track reconstruction in very large time projection chambers". Apr. 2024. arXiv: 2404.08614 [hep-ex].

[321] Lu, X.-G. *et al* (MINERνA Collaboration). "Measurement of Final-State Correlations in Neutrino Muon-Proton Mesonless Production on Hydrocarbon at $<E_\nu> = 3$ GeV". *Physical Review Letters* 121.2 (July 2018), p. 022504.

[322] Cai, T. *et al* (MINERνA Collaboration). "Nucleon binding energy and transverse momentum imbalance in neutrino-nucleus interactions". *Physical Review* ***D*** 101.9 (May 2020), p. 092001.