Search for eV-scale sterile neutrinos with IceCube DeepCore

Dissertation
zur Erlangung des akademischen Grades
doctor rerum naturalium
(Dr. rer. nat.)

im Fach: Physik
Spezialisierung: Experimentalphysik

eingereicht an der
Mathematisch-Naturwissenschaftlichen Fakultät
der Humboldt-Universität zu Berlin

von
Alexander Trettin M. Sc

Präsidentin der Humboldt-Universität zu Berlin
Prof. Dr. Julia von Blumenthal

Dekanin der Mathematisch-Naturwissenschaftlichen Fakultät
Prof. Dr. Caren Tischendorf

Gutachter:innen:	Prof. Dr. Marek Kowalski Prof. Dr. Walter Winter Prof. Dr. Kathrin Valerius
Datum der mündlichen Prüfung:	28.07.2023

Copyright Notice

This work is licensed under CC BY-SA 4.0.

To view a copy of this license, visit:
https://creativecommons.org/licenses/by-sa/4.0/

Colophon

This document was typeset with the help of KOMA-Script and LaTeX using the open-source kaobook template class.

The source code of this thesis is available at:
https://github.com/atrettin/PhD-Thesis

Abstract

Neutrino oscillations are the only phenomenon beyond the Standard Model that has been confirmed experimentally to a very high statistical significance. This work presents a measurement of atmospheric neutrino oscillations using eight years of data taken by the IceCube DeepCore detector between 2011 and 2019. The event selection has been improved over that used in previous DeepCore measurements with a particular emphasis on its robustness with respect to systematic uncertainties in the detector properties. The oscillation parameters are estimated via a maximum likelihood fit to binned data in the observed energy and zenith angle, where the expectation is derived from weighted simulated events. This work discusses the simulation and data selection process, as well as the statistical methods employed to give an accurate expectation value under variable detector properties and other systematic uncertainties. The measurement is first performed first under the standard three-flavor oscillation model, where the atmospheric mass splitting and mixing angle are estimated to be $\Delta m^2_{32} = 2.42^{+0.77}_{-0.75} \times 10^{-3}$ eV^2 and $\sin^2\theta_{23} = 0.507^{+0.050}_{-0.053}$, respectively. The three-flavor model is then extended by an additional mass eigenstate corresponding to a sterile neutrino with mass splitting $\Delta m^2_{41} = 1$ eV^2 that can mix with the active ν_μ and ν_τ flavor states. No significant signal of a sterile neutrino is observed and the mixing amplitudes between the sterile and active states are constrained to $|U_{\mu 4}|^2 < 0.0534$ and $|U_{\tau 4}|^2 < 0.0574$ at 90% C.L. These limits are more stringent than the previous DeepCore result by a factor between two and three and the constraint on $|U_{\tau 4}|^2$ is the strongest in the world.

Zusammenfassung

Neutrinooszillationen sind das einzige Phänomen jenseits des Standardmodells, das experimentell mit hoher statistischer Signifikanz bestätigt wurde. Diese Arbeit präsentiert eine Messung der atmosphärischen Neutrinooszillationen unter Verwendung von acht Jahren an Daten, die zwischen 2011 und 2019 vom IceCube DeepCore-Detektor aufgenommen wurden. Die Ereignisauswahl wurde im Vergleich zu früheren DeepCore-Messungen verbessert, wobei ein besonderes Augenmerk auf ihre Robustheit gegenüber systematischen Unsicherheiten in den Detektoreigenschaften gelegt wurde. Die Oszillationsparameter werden über eine Maximum-Likelihood-Fit an gebinnte Daten in der gemessenen Energie und Zenitwinkel geschätzt, wobei die Erwartungswerte aus gewichteten simulierten Ereignissen abgeleitet werdem. Diese Arbeit diskutiert den Simulations- und Datenauswahlprozess sowie die statistischen Methoden, die verwendet werden, um einen genauen Erwartungswert unter variablen Detektoreigenschaften und anderen systematischen Unsicherheiten zu liefern. Die Messung wird zunächst unter Verwendung des Standardmodells der Drei-Flavor-Oszillation durchgeführt, wobei das atmosphärische Massensplitting und der Mischwinkel auf $\Delta m^2_{32} = 2.42^{+0.77}_{-0.75} \times 10^{-3}; \mathrm{eV}^2$ und $\sin^2\theta_{23} = 0.507^{+0.050}_{-0.053}$ geschätzt werden. Das Drei-Flavor-Modell wird dann um einen zusätzlichen Masseneigenzustand erweitert, der einem sterilen Neutrino mit Massensplitting $\Delta m^2_{41} = 1; \mathrm{eV}^2$ entspricht und mit den aktiven ν_μ- und ν_τ-Flavorzuständen mischen kann. Es wird kein signifikantes Signal eines sterilen Neutrinos beobachtet, und die Mischungsamplituden zwischen den sterilen und aktiven Zuständen werden auf $|U_{\mu 4}|^2 < 0.0534$ und $|U_{\tau 4}|^2 < 0.0574$ bei 90% C.L. begrenzt. Diese Grenzwerte sind um den Faktor zwei bis drei strenger als das vorherige DeepCore-Ergebnis, und die Einschränkung von $|U_{\tau 4}|^2$ ist die stärkste der Welt.

Contents

Abstract iii

Zusammenfassung iv

Contents v

1 Introduction 1

2 Neutrinos in the Standard Model 3
- 2.1 Standard Model Particles 3
  - 2.1.1 Spin, Helicity and Chirality 4
  - 2.1.2 Electroweak Symmetry Breaking 5
  - 2.1.3 Charged Fermion Masses 6
  - 2.1.4 Neutrino Masses 7
  - 2.1.5 See-Saw Mechanism 9
  - 2.1.6 Radiative Neutrino Masses 10
- 2.2 Neutrino Properties 10
  - 2.2.1 Quantum Numbers 10
  - 2.2.2 Mass 11
  - 2.2.3 Active Neutrino Flavors 11
- 2.3 Neutrino Interactions 12
  - 2.3.1 Weak interactions after symmetry-breaking 12
  - 2.3.2 Neutrino-Lepton Scattering 13
  - 2.3.3 Neutrino Interactions with Nuclei 14

3 Neutrino Oscillations 19
- 3.1 Neutrino Oscillations in Vacuum 19
  - 3.1.1 Two-neutrino mixing 20
  - 3.1.2 Decoherence 21
- 3.2 Oscillations in matter 22
  - 3.2.1 Effective potentials 22
  - 3.2.2 The MSW Effect 24
  - 3.2.3 Parametric resonance 25
  - 3.2.4 Neutrino Mass Ordering 25
- 3.3 Standard three-flavor oscillations 26
  - 3.3.1 Solar neutrinos 27
  - 3.3.2 Reactor neutrinos 28
  - 3.3.3 Atmospheric Neutrino Oscillations 29

- 3.3.4 Accelerator neutrinos 32
- 3.4 Anomalies in neutrino oscillation measurements 33
  - 3.4.1 Reactor neutrino anomaly 33
  - 3.4.2 Gallium anomaly 35
  - 3.4.3 LSND and MiniBooNE Anomalies 35
  - 3.4.4 Global picture of oscillation anomalies 37
  - 3.4.5 Other constraints on sterile neutrinos 38
- 3.5 Open Questions in Neutrino Oscillation Physics 40

4 Neutrinos in IceCube and DeepCore **41**
- 4.1 The IceCube in-ice Array and DeepCore 41
  - 4.1.1 The Antarctic Ice 41
  - 4.1.2 In-Ice Array 42
  - 4.1.3 IceTop 44
  - 4.1.4 Digital Optical Modules 44
- 4.2 Propagation of particles in ice 45
  - 4.2.1 Cherenkov Effect 45
  - 4.2.2 Muons 46
  - 4.2.3 Electromagnetic Showers 46
  - 4.2.4 Hadronic Showers 47
- 4.3 Particle Signatures in IceCube 48
  - 4.3.1 Neutrinos 48
  - 4.3.2 Atmospheric muons 49

5 Simulation and Data Processing **51**
- 5.1 Event Simulation 51
  - 5.1.1 Neutrino Interactions 51
  - 5.1.2 Atmospheric muons 54
  - 5.1.3 Photon Propagation 56
  - 5.1.4 Simulation of Detector Response 56
- 5.2 Data Processing 58
  - 5.2.1 Trigger 58
  - 5.2.2 Online Filter 58
  - 5.2.3 Offline Filter 59
  - 5.2.4 Event Reconstruction 62
  - 5.2.5 Signature Classification 67
  - 5.2.6 Final Sample Selection and Binning 69
  - 5.2.7 Seasonal Stability 72
- 5.3 Implementation of systematic uncertainties 73
  - 5.3.1 Variation of Detector Properties 73
  - 5.3.2 Variation of the Atmospheric Neutrino Flux 76

6 Three-flavor oscillation measurement **79**
- 6.1 Statistical Analysis 79
  - 6.1.1 Definition of test statistic 79
  - 6.1.2 Modeling of Detector Response 80
  - 6.1.3 Selection of Free Parameters 81
- 6.2 Analysis Checks 83
  - 6.2.1 Robustness of the minimization 83
  - 6.2.2 Ensemble tests 84

6.3 Results 86
6.3.1 Measured Nuisance Parameter Values 86
6.3.2 Oscillation parameters 87
6.3.3 Post-fit Data/MC agreement 88
6.3.4 Likelihood Coverage Test 88

7 Search for sterile neutrino mixing 91
7.1 Atmospheric oscillations in the presence of an eV-scale sterile neutrino 91
7.1.1 The 3+1 model 91
7.1.2 Neutrino production height effects 92
7.1.3 Oscillation signal for large mass splittings 92
7.1.4 Oscillation signal for small mass splittings 93
7.2 Oscillation Probability Calculation with nuSQuIDS 93
7.2.1 Node placement 94
7.2.2 Production height averaging 94
7.2.3 Low-pass filtering 95
7.3 Nuisance oscillation parameters 97
7.4 Statistical Analysis 98
7.4.1 Signal in Analysis Binning 98
7.4.2 Definition of test statistic 99
7.4.3 Modeling of detector response via likelihood-free inference . 99
7.4.4 Selection of Free Parameters 104
7.4.5 Likelihood Optimization 105
7.4.6 Analysis Checks 107
7.5 Results 108
7.5.1 Best Fit Parameters 108
7.5.2 Likelihood Scan and Contour 110

8 Summary and Outlook 113
8.1 Summary 113
8.1.1 Three-Flavor Oscillation Measurement 113
8.1.2 Sterile Neutrino Search 114
8.2 Outlook 114
8.2.1 Reconstruction Improvements 115
8.2.2 Ice model 115
8.2.3 Treatment of systematic uncertainties 116
8.2.4 IceCube Upgrade 116

9 Conclusion 119

APPENDIX 121

A Mathematical derivations 123
A.1 Detector systematics via Likelihood-free Inference 123
A.1.1 Linear correction for KNN bias 123

B Additional Figures 125
B.1 Event Selection 125
B.1.1 Seasonal Stability of Variables (KS Tests) 125
B.1.2 PID Variables 125

B.2 Three-flavor analysis 128
  B.2.1 Seasonal Stability of Nuisance Parameter Fits 128

C Additional tables **131**
C.1 Three-flavor oscillation measurement 131

Bibliography **133**

List of Figures

2.1 Neutral-current lepton interaction vertices. 12
2.2 Charged-current lepton interaction vertices. 12
2.3 Feynman diagram for neutrino-electron scattering. 13
2.4 Spin configuration for particle-antiparticle interactions in the center-of-mass frame. Thin arrows (→) indicate momentum, thick arrows (⇒) show angular momentum. 14
2.5 Deep inelastic scattering of a muon neutrino in the quark-parton model. 16
2.6 Inclusive cross sections for neutrinos and antineutrinos. Figure taken from [22]. 17

3.1 Transition probabilities of an initial electron neutrino in vacuum, calculated at global best-fit[24] oscillation parameters. The distance is normalized such that the slow oscillations due to the smaller mass splitting Δm^2_{21} have a period of one. 21
3.2 Coherence lengths for neutrinos produced in pion decays at different mass splittings. The black solid lines indicate the baseline and energy ranges of some existing and proposed accelerator experiments. Figure taken from [25]. 22
3.3 Feynman diagrams of the coherent forward scattering processes for neutrinos travelling through Earth. 23
3.4 Parametric resonance for a simplified mantle-core-mantle propagation with two flavors. The assumed mass-splitting is $\Delta m^2 = 2.5\times10^{-3}\,\mathrm{eV}^2$, the mixing angle is $\theta = \arcsin(0.1)$, and the energy is $3.4\,\mathrm{GeV}$. The matter potential used for the calculation. It is an approximation of the density profile of Earth in the 12-layer PREM model, shown in gray. 26
3.5 Energy ranges and baselines for a selection of neutrino oscillation experiments, operational or planned. The range of each experiment is shown either as a line for experiments with a fixed oscillation distance such as T2K or DUNE, or as a box for experiments with both a variable baseline and energy range such as IceCube. Figure taken from[32]. 28
3.6 Neurino flux at the South Pole, averaged over all directions. Figures taken from [40]. 30
3.7 Illustration of the geometry of atmospheric oscillation measurements. The detector is shown as the gray box near the bottom and θ_z indicates the observed zenith angle of a neutrino. 30
3.8 Muon-neutrino survival probability calculated at NuFit 4.0[24] global best fit parameters. 31
3.9 Neutrino Main Injector (NuMI) facility producing neutrinos for the MINOS experiment. Figure taken from[48]. 33
3.10 Exclusion contours for the sterile mass splitting and mixing angle from the STEREO experiment. The effective mixing angle θ_{ee} is equivalent to θ_{13} in Equation 3.39. The red solid line shows the (FC corrected) exclusion contour, the dashed line shows the expected sensitivity. Blue lines show the sensitivity and exclusion limits using a different statistical method. Gray solid lines show the preferred values of the RAA. Figure taken from [50]. 34

3.11 Deficit in $\bar{\nu}_e$ flux from commercial nuclear reactors referred to as Reactor Antineutrino Anomaly. Figure taken from [49]. 34
3.12 Experimental setup of the BEST experiment. The diameters of the inner and outer volumes are $133.5\,\mathrm{cm}$ and $218\,\mathrm{cm}$, respectively. Figure taken from [54]. 35
3.13 Contours delimiting the 2σ allowed regions of the Gallium anomaly and the reactor anomaly in mass splitting and the effective mixing angle $\sin^2(2\theta_{ee}) = 4(1-|U_{e4}|^2)|U_{e4}|^2$. The reactor contours are shown for different flux predictions and the Gallium anomaly contours are shown under different neutrino cross-section models. Figure taken from [53]. 36
3.14 MiniBooNE neutrino mode E_ν^{QE} for ν_e CCQE data (points with error bars) and background (histograms). The dashed line shows the best fit of a two-flavor oscillation hypothesis. Figure taken from[58]. 36
3.15 MiniBooNE allowed regions for a combined neutrino and antineutrino dataset within a two-neutrino oscillation model. The shaded areas show the 90% and 99% C.L. LSND $\bar{\nu}_\mu \to \bar{\nu}_e$ allowed regions. Figure taken from [58]. 37
3.16 Contours delimiting the 99.73% C.L. (3σ) regions in mass splitting and the effective mixing angle $\sin^2(2\theta_{\nu e}) \equiv 4\left|U_{\mu 4}\right|^2|U_{e4}|^2$ from the appearance ($\nu_\mu \to \nu_e$) channel and the disappearance ($\nu_e \to \nu_e$) channel. Figure taken from [65]. 38
3.17 Test statistic for the magnitudes of unitary condition violations calculated in [67]. The shown quantities correspond to those in Equations 3.42 to 3.45. 39

4.1 An overview of the IceCube detector 42
4.2 Schematic view of the IceCube detector as seen from the top (upper panel) and the side(lower panel). The DeepCore fiducial volume is indicated by the hexagon in the upper panel and the green shaded area in the bottom panel. The side-band on the lower panel shows the scattering and absorption coefficients as a function of depth. 43
4.3 Schematic of a DOM, taken from [74]. 44
4.4 Schematic of the cable assembly of a DOM. Figure taken from [74]. . . . 44
4.5 An electrically charged particle emitting light while travelling below (upper panel) and above (lower panel) the speed of light in a medium. Image taken from [81]. 45
4.6 Angular profile of the Cherenkov emission of an electromagnetic cascade (e^-) and a hadronic cascade (π^+) using the parametrization from [86]. . . 47
4.7 An idealized cascade event (left) and starting track event (right) seen from the side. Each DOM that has received light is highlighted with a colored bubble, where the size is proportional the total charge seen by the DOM and the color indicates the time of the hits relative to the time at which the neutrino interaction happened. 48

5.1 Simulated MC livetime as a function of energy, calculated using the HKKM[40] model flux with NuFit 2.2[88] oscillation parameters. 52
5.2 GENIE interaction weights as a function of the pull of the axial mass term M_A^{CCRES}, for five ν_e CC events produced via resonance interactions. Each dot represents a discrete point for which the event's cross section is computed in GENIE. The line represents the quadratic fit made used to interpolate the weight value over the continuous range allowed for the systematic parameter. 53

5.3 Fractional difference in event rates between (top)M_A^{RES} (bottom) dis_csms at 1σ and at nominal value for both PID bins. 54
5.4 Impact on the final histograms when the muon normalization is increased by 50%. The largest impact is seen above the horizon in the mixed PID channel with a change in bin count of 5%. 55
5.5 Impact on the final histograms when the muon spectral index is increased by 1σ. 55
5.6 Scattering and absorption lengths as a function of depth in the South Pole Ice (SPICE) model that is used to produce the simulation for this work. . 56
5.7 The green (yellow) regions show the 68% (90%) spread in the SPE charge templates for a given charge. Superimposed are the average SPE charge templates for the variety of hardware configurations shown in the black dotted, dashed, and solid lines. The TA0003 distribution, shown in red, originates from laboratory measurements. Figure taken from [78]. 57
5.8 Efficiency of the IceCube and DeepCore triggers as a function of the primary neutrino energy. Figure taken from [75]. 58
5.9 Example of an event that would be rejected by the online filter algorithm. DOMs that have observed light are highlighted in color depending on time from red (early hits) to blue (late hits). DOMs that have not observed any light are shown as black dots. Figure taken from [75]. 59
5.10 Distribution of one of the variables used in the L3 offline filter, the time between the last hit and the first hit after noise cleaning (left) and the z-position of the first HLC hit (right). Histograms show the distributions in simulated data separated by event type, data points with error bars show the distribution of real data. The bottom panel shows the ratio between data and simulation. Events falling on the "signal" side of the histogram are passed to the next filter level. 60
5.11 Distribution scores for the noise (left) and muon (right) BDT. The distributions of the muon classifier are shown for events where the score of the noise BDT is greater than 0.7. Histograms show the distributions in simulated data separated by event type, data points with error bars show the distribution of real data. The bottom panel shows the ratio between data and simulation. Events falling on the "signal" side of the histogram are passed to the next filter level. 61
5.12 Distributions for two of the L5 corridor cut variables. Histograms show the distributions in simulated data separated by event type, data points with error bars show the distribution of real data. The bottom panel shows the ratio between data and simulation. Events falling on the "signal" side of the histogram (or, equivalently, opposite to the "cut" side of the histogram) are passed to the next filter level. 62
5.13 Example of a $\nu_{\mu,\mathrm{CC}}$event reconstructed with SANTA with hits on several strings. Strings 84, 83 and 37 are spaced $\sim 80\,\mathrm{m}$ apart from each other and form a highly obtuse triangle. 64
5.14 Detailed geometry of a light cone created by a track. $\vec{q}$ is the position of the anchor point and $\vec{r}$ is the position of the optical module. $\vec{p}(t_{\mathrm{em}})$ and $\vec{p}(t_{\mathrm{geom}})$ are the positions of the muon at the time the photon is emitted and when it is geometrically expected to arrive, respectively. 65
5.15 Median error on the reconstructed zenith angle at the final level of the sample selection as a function of the true simulated neutrino energy. . . 65
5.16 Median fractional error on the reconstructed energy at the final level of the sample selection as a function of neutrino energy. 67

5.17 Distribution and data/MC comparison for the two most important input variables into the classification BDT. 68
5.18 PID score distribution for simulated neutrino events at the final level of the event selection, weighted according to the HKKM flux model [40] and neutrino oscillations with NuFit 4.0[24] global fit parameters, normalized to unity. The BDT score ranges from the most cascade-like at 0 to the most track-like event signature at 1. 69
5.19 Distribution of the "z-travel" variable calculated for the uppermost 15 layers of IceCube DOMs. Only events with at least 4 hits in the uppermost 15 layers of DOMs are included in the histogram. 70
5.20 Distribution of the SANTA goodness-of-fit variable and the reconstructed zenith angle at L5 of the event selection process. 70
5.21 Expected event counts in 7.5 years of live time assuming no sterile mixing and NuFit 4.0 [24] global best fit parameters at Normal Ordering. 71
5.22 Muon template before (top) and after (bottom) the application of KDE smoothing. The shown values are the average of 20 KDE evaluations on different bootstrap samples. 72
5.23 Two parameter model used to parametrize the optical efficiency in this analysis (left) and the positions in this two-dimensional space where older hole ice models are located (right). The relative optical efficiency curves are normalized to have the same area, which can lead to acceptance values greater than 1. 75
5.24 Perturbation of the scattering and absorption coefficients with respect to the nominal ice model applied in additional MC sets. 75
5.25 Bin-wise p-value of the fitted slopes as a function of the step-function ice model variation. 76
5.26 Fully correlated regions of uncertainties in the hadronic interaction model. Figure taken from [120]. 77
5.27 Relative uncertainty assigned to each region of hadron phase space in percent. Figure taken from [120]. 77

6.1 Example of a linear regression in one bin of the analysis projected onto the dimension of the DOM efficiency. Data points with translucent error bars originate from MC sets where one or more parameters besides DOM efficiency are at off-nominal points and are projected along the fitted surface to the nominal point. 80
6.2 Gradient of the relative bin count with respect to DOM efficiency. 81
6.3 Grid of $\Delta\chi^2_{\mathrm{mod}}$ values showing the impact of the axial mass M_A^{CCQE} being pulled by 1σ. 82
6.4 Ranking of $\Delta\chi^2_{\mathrm{mod}}$ values for all nuisance parameters considered for the three-flavor oscillation analysis. 83
6.5 Asimov inject/recover test result for the three-flavor oscillation analysis. 84
6.6 Observed test statistic value of the three-flavor oscillation analysis compared to expected distribution from ensemble. 85
6.7 Contribution of every bin to the over-all test statistic in the three-flavor analysis (left) and their observed distribution compared to the expected distribution from pseudo-data trials (right). 85
6.8 Observed best fit values of the two-flavor fit compared to the distribution from pseudo-data trials. 86

6.9 Angular acceptance curves corresponding to the best fit point of the three flavor and sterile oscillation fits compared to the results of LED flasher calibration studies. 87
6.10 Contours showing the 90% C.L. allowed region for the physics parameters of the three-flavor analysis and of other experiments[128–131]. The sensitivity for this work is calculated at the best fit point and the cross shows the best fit value. Other experiments shown in dotted lines are accelerator results, while solid lines are atmospheric oscillation results. All contours shown assume Normal Ordering. 88
6.11 Observed contour in $\sin^2(\theta_{23})$ (solid) compared to the Asimov expectation (dashed) and the distribution of 100 pseudo-data trials (yellow and green bands) produced at the best fit point of the three-flavor oscillation analysis. The observed contour is fully contained within 68% of the fluctuations of the trials. 89
6.12 Oscillation argument L/E calculated from reconstructed quantities at the best fit point of the three-flavor oscillation analysis. 89
6.13 Fraction of trials below the 90% threshold expected from Wilks' theorem for a range of points in mass splitting (left) and mixing angle (right). . . 90

7.1 Muon neutrino survival probability in the presence of a fourth mass eigenstate with $\Delta m^2_{41} = 1$ eV^2 and $\theta_{24} = 15°$ with a fixed production height of 20 km (left) and with production heights averaged between 10 km and 30 km (right). 92
7.2 ν_μ survival probability at $\Delta m^2_{41} = 1$ eV^2 (left) and at $\Delta m^2_{41} = 0.1$ eV^2 (right). The mixing angle is $\theta_{24} = 15°$ in both cases 93
7.3 Optimized placement of NUSQUIDS nodes (black dots) with extra dense node spacing in the three critical regions. The injected value for Δm^2_{41} is 0.1 eV^2. 94
7.4 Pull introduced into each bin of the analysis histogram when changing the distance over which the production height is averaged from 20 km down to 1 km at sterile mixing angles of $\theta_{24} = \theta_{34} = 15°$. 95
7.5 Muon neutrino survival probability in the presence of a sterile neutrino after application of both height averaging and low-pass filtering. 96
7.6 Results of the inject/recover test (left) and the corresponding mis-modeling values, $\Delta\chi^2_{\mathrm{mod}}$, attributable to low-pass filtering in NUSQUIDS (right). A negative value of $\Delta\chi^2_{\mathrm{mod}}$ indicates that the fit did not converge to the true global optimum. 97
7.7 Muon neutrino survival probability for a directly up-going neutrino as a function of energy in the presence of sterile neutrinos at different values of the sterile CP violating phase. 98
7.8 Signal in the analysis binning produced by different combinations of θ_{24}, θ_{34}, and δ_{24} as a fraction of the Poisson error in each bin. The mass splitting of the sterile state is $\Delta m^2_{41} = 1$ eV^2. 99
7.9 Idealized sketch of the effect of DOM efficiency on the distribution of the energy reconstruction error. The blue line shows the distribution of the error in the nominal MC set, the orange line that of events in an MC set with an increased DOM efficiency. The black line shows the ratio between the two PDFs. 100

7.10 Illustration of the re-weighting process used to model changes in the detector response. The axis labeled Δx stands in for the parameters that characterize each individual event such as true reconstructed energy and zenith angle. The axis labeled $\Delta \alpha$ stands in for the detector parameters that vary between systematic sets, such as the DOM efficiency. The distribution of event parameters, x, are shown for the nominal MC set in blue, the off-nominal sets in yellow. The black line shows the function along which an event at a particular x is re-weighted as a function of α. 101
7.11 Binwise pulls between the nominal set after re-weighting according to eq. 7.11 and the systematic MC sets 0001, 0002, 0003, and 0004 representing DOM efficiency values of 90%, 95%, 105%, and 110%, resepctively. The 1D histogram in each row shows the distribution of the pulls over all bins. . 104
7.12 Prediction of bin counts in one bin of the analysis as a function of the DOM efficiency scale, ϵ_{DOM}. The error band in the left panel corresponds to the error on the nominal MC prediction without errors on the event-wise gradients. 105
7.13 Result of the ensemble test with randomly injected nuisance parameters. This test was run before the priors on barr_i_Pi, barr_z_K, and barr_z_antiK have been inflated as described in Section 5.3.2. Parameters framed in red have been deemed to be negligible for the analysis. The color scale shows the likelihood difference between the free fit and a fit in which the physics parameters (θ_{24}, θ_{34} have been fixed to their true value. If this number is negative, the trial is circled in red. In such cases, the minimizer failed to find the correct global optimum in the free fit. 106
7.14 Observed test statistic compared to the distribution expected from pseudo-data trials for the sterile oscillation analysis. 108
7.15 Scan of the χ^2_{mod} difference with respect to the best fit point with 90% and 99% C.L. contours assuming Wilks' theorem with two degrees of freedom (left) and ensemble test of the coverage on three points along the 90% C.L. line (right). The test points indicated in the left panel correspond to the points at which the ensemble test were produced in the right panel. . . . 110
7.16 Contour of the 90% C.L. of this analysis compared to measurements from the ANTARES[142], Super-Kamiokande[143] and NOνA[144] experiments and a previous high-energy IceCube oscillation study[29] (left), and compared to the previous DeepCore study[109] (right). 111

8.1 Projected sensitivity of the sterile neutrino search when using the table-based reconstruction method. 115
8.2 Expected sensitivity to the atmospheric neutrino oscillation parameters with three years of Upgrade data compared to recent results from DeepCore and other experiments. Figure taken from [148]. 117

B.1 Kolmogorov-Smirnov p-values calculated between each season of data for reconstructed quantities used in the fit (top) and control variables (bottom) 126
B.2 Histograms of the likelihood score from the energy reconstruction (left) and the goodness-of-fit ratio of the zenith reconstruction (right) in simulation. 126
B.3 Histograms of the radius with respect to string 36 of the vertex (left) and the endpoint of the reconstructed track (right) in simulation. 127
B.4 Histograms of the z-coordinate of the vertex (left) and the endpoint of the reconstructed track (right) in simulation. 127
B.5 Histogram of the reconstructed track length in simulation. 127

B.6 P-value of the blind ensemble test to ensure the compatibility of fit results of the three-flavor analysis between seasons. P-values are calculated using histograms of the fit results from an ensemble of pseudo-data where one year of live time is assumed for each trial. Season 2011 is excluded because its live time is smaller. 129

List of Tables

2.1 Fermions in the Standard Model. The electric charge, Q, is the conserved charge of the $U(1)_Q$ symmetry group. 4
2.2 Eigenvalues of the weak isospin I, of its third component I_3 and the hypercharge $Y = 2(Q - I_3)$ for one generation of fermions. Reproduced from [7]. 4

3.1 Best-fit values of the three-flavor oscillation parameters and the oscillation channel that is primarily involved in their measurement. Values taken from [31]. 28
3.2 Definitions of effective mixing angles in the "3+1" oscillation model for oscillation channels relevant to neutrino oscillation anomalies. 37

4.1 Secondary particles and signatures produced by each type of neutrino interaction. 49

5.1 Table of generation volumes used for GENIE neutrino simulation. The cylinder is centered in DeepCore in all cases. 52
5.2 Parameters used in the noise simulation. Typical values taken from [100], actual values are fit for each DOM individually. 57
5.3 Summary of the rates obtained after each level of selection. Neutrinos are weighted to an atmospheric spectrum with oscillations included. 62
5.4 Successively applied cuts on the data sample. The bottom row corresponds to the final rates in the sample after all cuts have been applied. The total rate of the data and simulation does not match, which is expected since there is a large amount of uncertainty in the total normalization. Muon contamination is the muon rate divided by the total event rate. Numbers calculated at the NuFit 4.0 global best-fit point. 71
5.5 Expected event rate with 8 years livetime broken down in event types and PID bins, calculated at NuFit 4.0 global best fit parameters. 72

6.1 Fitted values of all nuisance parameters from the all-season three-flavor fit. The pull of the best fit value is shown for parameters with a defined prior. 87
6.2 Contribution of each category of systematic uncertainties to the total error budget in each physics parameter. 88

7.1 Oscillation parameters of the 3+1 model and their treatment in this analysis. 98
7.2 List of all free parameters in the sterile oscillation analysis with their respective ranges and priors (if applicable). 106
7.3 Fitted values of all nuisance parameters from the all-season sterile oscillation fit. The pull of the best fit value is shown for parameters with a defined prior. 109

C.1 List of systematic parameters considered in the three-flavor oscillation analysis along with their priors and allowed ranges. 132

Introduction 1

The neutrino is a nearly massless and electrically neutral particle whose existence was first conjectured by Pauli in the 1930s to explain the fact that the energy spectrum of radiation from nuclear beta decay was continuous. If the only particles produced by the decay were the nucleus and the beta particle, then the energy of the beta particle would have been fixed by conservation of energy and momentum. It was a bold proposition at the time, because there were no observable traces of this particle and Pauli himself feared that it might be unobservable. His fears proved unwarranted when the first direct experimental observation of neutrinos was made in 1956 by Cowan and Reines[1] by detecting the distinct signature of so-called "inverse" beta-decay reaction

$$\bar{\nu} + p \longrightarrow n + e^{+} \tag{1.1}$$

inside a water tank close to a nuclear reactor. About a decade later, in 1960, the Homestake experiment was able to measure the flux of neutrinos from the Sun. However, the observed rate of electron neutrinos was lower than what was expected from nuclear fusion reactions inside the Sun, leading to the *solar neutrino problem*. The muon neutrino was discovered in 1962 by an experiment at the Brookhaven National Laboratory[2] and the tau neutrino in 2000 by the DONUT experiment at Fermilab[3], completing the current picture of the Standard Model (SM) with three generations of leptons. In this model, neutrinos are described as spin-$1/2$ fermions that only interact via the Weak nuclear force. The Weak force only interacts with left-handed chiral neutrinos and right-handed chiral antineutrinos, and no other neutrino states have so far been observed. This description requires that neutrinos are massless, because the Higgs mechanism that produces the masses of all other particles in the SM requires an interaction involving both right-handed and left-handed chiral fields.

The solution to the solar neutrino problem accepted today is that neutrinos have the ability to oscillate from one flavor to another. In this way, the electron neutrinos that are initially produced by the Sun can turn into a different flavor to which the Homestake experiment was not sensitive. This phenomenon of *neutrino oscillations* was first demonstrated by the Super-Kamiokande experiment for muon neutrinos that are produced in the Earth's atmosphere[4]. In 2002, the SNO experiment provided the first direct evidence that this flavor conversion was also happening to electron neutrinos from the Sun[5]. The existence of neutrinos oscillations has profound implications, because it means that neutrinos cannot be massless.

Neutrino oscillations can be explained by postulating that the flavor eigenstates with which the Weak force interacts are mixtures of different mass eigenstates. The mass eigenstates can be described as wave packets with slightly different frequencies that overlap. These wave packets travel at different speeds due to their mass differences and therefore interfere with one another, leading to the phenomenon of neutrino oscillations. However, the SM provides no explanation of how neutrinos acquire their masses. The Higgs mechanism requires couplings of both left-handed and right-handed chiral fields to the Higgs field, but there are no right-handed neutrinos in the SM. The extreme lightness of the neutrinos compared to other particles suggests that the process that produces them might be different altogether from the process that generates the masses of all other particles. Neutrino oscillations, therefore, are direct evidence of physics beyond the Standard Model (BSM) and motivate the search for new particles and forces that might be involved in the process of neutrino mass genereration.

This work describes a neutrino oscillation measurement using the IceCube Neutrino Observatory, a neutrino detector located at the geographic South Pole. It uses 5160 optical sensors deployed in a volume of one cubic kilometer of the Antarctic glacier to measure faint flashes of Cherenkov light that is produced when neutrinos interact with the ice. It can detect neutrinos in a wide energy range starting from atmospheric neutrinos at a few GeV up to the PeV energy range of cosmic neutrinos. The data analysis presented in this work uses observations from the DeepCore sub-array of IceCube that is specifically optimized for the detection of neutrinos that are produced in the atmosphere of the Earth. These neutrinos consist mostly of muon neutrinos that travel through the Earth and have a chance to oscillate into other neutrino flavors before they are detected at the South Pole. After collecting tens of thousands of neutrinos over the course of several years, the muon neutrino survival probability can be estimated to a high precision. The results allow inferences about the mass differences between different neutrino states and can be probed for signs of BSM physics.

Chapter 2 of this thesis summarizes how neutrinos are described in the Standard Model. A particular emphasis of this chapter is the famous Higgs mechanism by which all particles in the SM acquire their mass and how this mechanism can be expanded to include neutrino masses. In Chapter 3, we will describe how the mass differences between neutrinos lead to neutrino oscillations and how these oscillations manifest in experimental observations. The chapter also describes the experimental anomalies of neutrino oscillation measurements that motivate the search for additional heavy neutrino states of the eV scale. The IceCube Neutrino Observatory is introduced in Chapter 4 with special focus on the DeepCore sub-array. The chapter will also describe in detail how the DeepCore data is filtered to produce a data sample with a high purity of muon neutrinos. The data sample is then used for two different measurements: The first, described in Chapter 6, measures the parameters that characterize the muon neutrino survival probability in the picture of three oscillating neutrino flavors. The second measurement, shown in Chapter 7, is probing the observed oscillation pattern for signs of an additional heavier neutrino mass eigenstate that is associated with an otherwise non-interacting *sterile* neutrino. The result of this measurement places limits on the amount of mixing that is allowed between the hypothetical sterile neutrino state and the three known active neutrino flavors. With these constraints, this work provides another puzzle piece in the search for the origin of neutrino masses.

Neutrinos in the Standard Model

2

2.1 Standard Model Particles 3
2.1.1 Spin, Helicity and Chirality 4
2.1.2 Electroweak Symmetry Breaking 5
2.1.3 Charged Fermion Masses 6
2.1.4 Neutrino Masses 7
2.1.5 See-Saw Mechanism . . . 9
2.1.6 Radiative Neutrino Masses 10
2.2 Neutrino Properties . . . 10
2.2.1 Quantum Numbers . . . 10
2.2.2 Mass 11
2.2.3 Active Neutrino Flavors 11
2.3 Neutrino Interactions . . 12
2.3.1 Weak interactions after symmetry-breaking . . . 12
2.3.2 Neutrino-Lepton Scattering 13
2.3.3 Neutrino Interactions with Nuclei 14

The Standard Model (SM) of particle physics is a relativistic quantum field theory based on the gauge symmetry group $\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, where the subscripts C, L and Y correspond to the conserved quantities *color*, *left-handed chirality* and *weak hypercharge*, respectively. In this model, all matter particles are described as fermions, that is, excitations of Dirac-type fermion fields permeating space-time. The forces acting between fermions are mediated by an exchange of bosons, and all interactions must preserve the over-all symmetry of the theory. Since its completion in the early 1970s, it has been shown that it accurately describes the interactions between elementary particles due to the Strong Force, the Weak Force and the electromagnetic force to an impressive degree of precision. It can also explain how quarks and leptons acquire their masses via the Higgs mechanism, whose byproduct, the Higgs boson, was detected at the LHC in 2012[6]. Despite its success, the Standard Model has some known shortcomings such as its incompatibility with General Relativity and inability to explain cosmological phenomena most commonly interpreted as Dark Matter and Dark Energy. Most relevantly for this work, it predicts that neutrinos should be massless and therefore does not allow for neutrino oscillations. Since neutrino oscillations can be experimentally observed at very high statistical significance[4], it is clear that the SM has to be extended to accomodate neutrino masses. There are several candidate theories for such an extension, but none of them could so far be experimentally confirmed. This chapter describes the properties and interactions of neutrinos in the SM and briefly outlines the Higgs mechanism. It will then show how it can be extended to produce neutrino masses via the introduction of right-handed "sterile" neutrino states.

The derivations of the SM, Weak interactions and the extended Higgs mechanism presented in this chapter follow the explanations in [7] and [8].

[7]: Giunti et al. (2007), *Fundamentals of Neutrino Physics and Astrophysics*

[8]: Schwartz (2013), *Quantum Field Theory and the Standard Model*

2.1 Standard Model Particles

The elementary particles of the SM are organized into fermions and bosons, where fermions make up the observable matter while bosons are the particles that mediate forces. The number of force-mediating bosons is determined by the generators of the symmetry groups that all interactions must obey, while the strength of each force is determined by a *coupling constant* that has to be estimated experimentally. There are eight massless gluons that correspond to the generators of the $\mathrm{SU}(3)_C$ group and mediate the Strong force. All Strong interactions conserve the so-called *color* charge of the involved particles. The symmetry group $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ is the combined symmetry of the *electroweak* force and produces the gauge boson fields W_1, W_2, W_3 and B. The electroweak symmetry group is broken by interactions of fermions with the Higgs field that mix the W and B fields into massive $W^{\pm}$ and Z^0 bosons,

$$Z = \cos\theta_W W_3 - \sin\theta_W B \tag{2.1}$$

$$W^{\pm} = \frac{1}{\sqrt{2}}(W_1 \mp iW_2)\ , \tag{2.2}$$

where θ_W is the *Weinberg angle*. After symmetry breaking, the Lagrangian is symmetric under the $\mathrm{U}(1)_Q$ group where the electric charge Q is conserved and the massless gauge boson is the photon,

$$\gamma = \sin\theta_W W_3 + \cos\theta_W B\ , \tag{2.3}$$

The fermions of the SM are divided into quarks and leptons. Quarks participate in all strong, weak and electromagnetic interactions and are always found in combinations that form baryons (protons, neutrons) or mesons (kaons, pions). The leptons, on the other hand, do not participate in strong interactions. Charged leptons are massive and participate in both the weak and electromagnetic interactions, while neutral leptons, the neutrinos, are massless and participate only in weak interactions. The fermions can be grouped into three *generations* of quarks and leptons that are only distinguished by their masses, leading to a convenient arrangement of quarks and leptons in a 3×4 scheme, as shown in Table 2.1. For each (massive) fermion, there exists a left-handed and a right-handed component. The left-handed components of each generation form a doublet of the $\mathrm{SU}(2)_L$ group with weak isospin $\frac{1}{2}$, while the right-handed components are singlets. The right-handed and left-handed fields for one generation and their charges are summarized in Table 2.2.

Table 2.1: Fermions in the Standard Model. The electric charge, Q, is the conserved charge of the $\mathrm{U}(1)_Q$ symmetry group.

	generation			
	1	2	3	Q
quarks	u	c	t	+2/3
	d	s	b	−1/3
leptons	ν_e	ν_μ	ν_τ	0
	e	μ	τ	−1

Table 2.2: Eigenvalues of the weak isospin I, of its third component I_3 and the hypercharge $Y = 2(Q - I_3)$ for one generation of fermions. Reproduced from [7].

	I	I_3	Y
$L_L \equiv \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}$	1/2	1/2 −1/2	-1
e_R	0	0	-2
$Q_L \equiv \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	1/2	1/2 −1/2	1/3
u_R	0	0	4/3
d_R	0	0	−2/3

2.1.1 Spin, Helicity and Chirality

The states that describe fundamental particles must, by definition, belong to irreducible representations of the Poincaré group. All possible representations can be classified by so-called Casimir operators, which are operators that are invariant under the group transformations. The first Casimir operator for the Poincaré group is the square of the four-momentum

$$P^2 = P_\mu P^\mu \tag{2.4}$$

with the eigenvalue $p^2 = m^2$ being the rest mass of the particle. The second one is the square of the so-called Pauli-Lubanski four-vector

$$W^2 = W_\mu W^\mu = -m^2 \vec{S}^2\ , \tag{2.5}$$

which has the eigenvalue $w^2 = -m^2 s(s+1)$, where s is the spin of the particle. Thus, all representations of the Poincaré group can be classified by their rest mass and spin. Given m and s for a particle, different states of that particle can be distinguished by their momentum $\vec{p}$ and the zeroth component of W^μ, which is $W^0 = \vec{S}\vec{p}$. A convenient quantum number to define is the *helicity* of a particle, which is the eigenvalue of the operator

$$\hat{h} = \frac{\vec{S}\vec{p}}{s|\vec{p}|} \tag{2.6}$$

that can have values of ±1. The relativistic properties of a particle with known mass and spin are fully described by a measurement of its momentum, $\vec{p}$, and helicity, h.

The states of spin-1/2 fermions, such as neutrinos, are representations of the Poincaré group that are constructed out of two spinors, Ψ_L and Ψ_R, that are called *left-handed* and *right-handed*, respectively. Their propagation is governed by the Dirac equation, which reads in Fourier space

$$\begin{aligned}(E - \vec{\sigma} \cdot \vec{p})\Psi_R &= m\Psi_L \\ (E + \vec{\sigma} \cdot \vec{p})\Psi_L &= m\Psi_R\end{aligned} . \tag{2.7}$$

Whether a particle is left-handed or right-handed is referred to as its *chirality*. If the mass in Equation 2.7 is zero, the two fields describe independent particles that are eigenstates of the helicity operator, which is $\hat{h} = \frac{\vec{\sigma}\vec{p}}{|\vec{p}|}$ in this case. This means that helicity and chirality eigenstates are identical for massless particles, but for massive particles, helicity eigenstates are superpositions of Ψ_L and Ψ_R.

2.1.2 Electroweak Symmetry Breaking

The process of breaking the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ symmetry group deserves special attention for the purposes of this work, because it is the process by which the exchange bosons of the Weak force acquire their mass. If the symmetry was unbroken, as it is the case for the SU(3) group of the Strong force, then the exchange bosons would all remain massless, just like the gluons. To simplify the discussion, the process can be illustrated using only the first generation of SM fermions. The starting point is to introduce the Higgs doublet of complex scalar fields

$$\Phi = \begin{pmatrix} \Phi^+ \\ \Phi^0 \end{pmatrix} , \tag{2.8}$$

where Φ^+ is charged and Φ^0 is neutral[1]. The Lagrangian describing the dynamics of this field,

$$\mathscr{L}_{\text{Higgs}} = (D_\mu \Phi^\dagger)(D^\mu \Phi) - \lambda \left(\Phi^\dagger \Phi - \frac{v^2}{2} \right)^2 , \tag{2.9}$$

with the covariant derivative

$$D_\mu \Phi = \partial_\mu \Phi - igW_\mu^a \tau^a \Phi - \frac{1}{2} ig' B_\mu \Phi \tag{2.10}$$

is invariant under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ symmetry and adds a quartic self-interaction potential with the parameters λ and v, where λ is taken to be positive, such that the potential is bounded from below. The fields W_μ^a in the covariant derivative correspond to the gauge bosons of the $\mathrm{SU}(2)_L$ group whose generators are $\tau^a = \frac{\sigma^a}{2}$, where σ^a are the Pauli matrices. The field B_μ is the boson of the $\mathrm{U}(1)_Y$ group. The factors g

[1] In a more general discussion, the Higgs doublet would be written down without assigning the charges a priori, they would be derived later. See [8] for a more rigorous derivation.

and g' are the $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ coupling constants, respectively, and are related to the Weinberg angle by

$$\tan\theta_W = \frac{g'}{g} \ . \tag{2.11}$$

Because the potential has a minimum that is not at zero, the field Φ acquires a non-zero *vacuum expectation value* (VEV) where $\Phi^\dagger\Phi = \frac{v^2}{2}$. Since the vacuum is electrically neutral, this VEV can only come from the neutral part, Φ^0, of the doublet and can be written as

$$\Phi_{\mathrm{VEV}} = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix} \ . \tag{2.12}$$

This vacuum expectation value is no longer symmetric under the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ group, but is still symmetric under the $\mathrm{U}(1)_Q$ group in which the electric charge is conserved. To see what happens to the Lagrangian, the field Φ can be expressed in the unitary gauge as a variation around the VEV such that

$$\Phi(x) = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v + H(x)\end{pmatrix} \ . \tag{2.13}$$

Plugging this into the Lagrangian in Equation 2.9 and re-writing the W^i_μ and B_μ fields in terms of Z and $W^\pm$ using the relationships given in Equation 2.2 and Equation 2.11 we find

$$\begin{aligned}\mathscr{L}_{\mathrm{Higgs}} = \ & \frac{1}{2}(\partial H)^2 - \lambda v^2 H^2 - \lambda v H^3 - \frac{\lambda}{4}H^4 \\ & + \frac{g^2 v^2}{4} W^\dagger_\mu W^\mu + \frac{g^2 v^2}{8\cos^2\theta_W} Z_\mu Z^\mu \\ & + \text{Higgs vertices} \ ,\end{aligned} \tag{2.14}$$

where Higgs vertices are 3-vertices and 4-vertices between the Higgs field and the W and Z. The notable part is that the W and Z bosons have acquired a kinetic term in the second line of Equation 2.14 with a mass that is proportional to the VEV of the Higgs field, giving massive exchange bosons to the Weak force[2].

[2] The massless photon field is found by expanding the full electroweak Lagrangian in the same way, which we neglect here for the sake of brevity.

2.1.3 Charged Fermion Masses

In Quantum Electrodynamics, a Lorentz-invariant mass term for spin-$\frac{1}{2}$ fermions can be written as a product of left-handed and right-handed Weyl spinors, also known as the Dirac mass

$$\mathscr{L}_{\mathrm{Dirac}} = m(\bar{\Psi}_R\Psi_L - \bar{\Psi}_L\Psi_R) \ . \tag{2.15}$$

However, such a term is not invariant under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ and therefore cannot be added to the SM Lagrangian directly. Fortunately, masses for fermions can be recovered if we add a Yukawa coupling term between the fermions and the Higgs field, such as

$$\mathscr{L}_{\mathrm{Yuk}} = -y\bar{L}_L\Phi e_R + \text{h.c.} \ , \tag{2.16}$$

where L_L denotes the $\mathrm{SU}(2)_L$ doublet listed in Table 2.2 and y is the Yukawa coupling constant. When the VEV is inserted into this term, it produces a mass term $-m_e(\bar{e}_L e_R + \bar{e}_R e_L)$ with $m_e = \frac{y}{\sqrt{2}} v$ for the charged leptons and the down-type quarks d, s, and b. A similar term that is also invariant under $\mathrm{SU}(2)_L$ and generates masses for the up-type quarks is $-y\bar{L}_L\tilde{\Phi}u_R$, where we defined $\tilde{\Phi} \equiv i\sigma_2\Phi$.

2.1.4 Neutrino Masses

The Higgs mechanism described in Section 2.1.3 requires both left-handed and right-handed Weyl spinors to interact with the Higgs field. Since there are no right-handed neutrinos in the SM, it predicts that they should be massless, in contradiction to experimental evidence. However, if we add right-handed neutrino fields into the model, then neutrino masses can be generated in a way that is tantalizingly similar to that of up-type quarks by adding interactions of the form $Y^{\nu}_{ij}\bar{L}^i\tilde{\Phi}\nu^j_R$ to the Lagrangian. Such a right-handed field would be uncharged with respect to all symmetry groups of the SM and therefore would not interact with any other particle, and hence it is called a *sterile neutrino*. Because neutrinos are electrically neutral, another possibility for a mass term that is allowed by the symmetry of the SM is the so-called *Majorana mass*, $m\nu^c_R\nu_R$, in which $\nu^c_R = \nu^T_R\sigma_2$ is the charge conjugate Weyl spinor.

Dirac Mass

Dirac masses are produced by adding Yukawa couplings

$$\mathscr{L}_{\mathrm{Yuk}} = -Y^e_{ij}\bar{L}^i\Phi e^j_R - Y^{\nu}_{ij}\bar{L}^i\tilde{\Phi}\nu^j_R + \mathrm{h.c.} \tag{2.17}$$

to the Lagrangian, where the indices i and j run over the generations e, μ, and τ and the matrices $Y_{e,\nu}$ contain the complex Yukawa coupling constants that are free parameters of the model. After symmetry breaking, the mass terms become

$$\begin{aligned}\mathscr{L}_{\mathrm{mass}} &= -\frac{v}{\sqrt{2}}\left[Y^e_{ij}\bar{e}^i_L e^j_R + Y^{\nu}_{ij}\bar{\nu}^i_L\nu^j_R\right] + \mathrm{h.c.} \\ &= -\frac{v}{\sqrt{2}}\left[\bar{e}_L Y_e e_R + \bar{\nu}_L Y_\nu \nu_R\right] + \mathrm{h.c.}\ .\end{aligned} \tag{2.18}$$

To find the physical fields with definite masses, the $Y_{e,\nu}$ matrices are diagonalized with two unitary matrices

$$Y_e = U_e M_e K^{\dagger}_e \qquad\qquad Y_\nu = U_\nu M_\nu K^{\dagger}_\nu\ , \tag{2.19}$$

such that $M_{e,\nu}$ are diagonal and contain the physical masses of leptons and neutrinos. By applying the transformation

$$\begin{aligned} e_L &\longrightarrow U_e e_L, \quad e_R \longrightarrow K_e e_R \\ \nu_L &\longrightarrow U_\nu \nu_L, \quad \nu_R \longrightarrow K_\nu \nu_R\ , \end{aligned} \tag{2.20}$$

the Lagrangian can be written in the *mass basis*

$$\mathscr{L}_{\text{mass}} = \frac{v}{\sqrt{2}}\left[\bar{e}_L M_e e_R - \bar{\nu}_L M_\nu \nu_R\right] \ . \tag{2.21}$$

The application of these transformations to the lepton fields will also affect the interaction part of the electroweak Lagrangian for leptons, but not all of them are of any physical consequence. Since there is no coupling to right-handed neutrinos, the transformation of ν_R has no effect at all, and we can simply ignore K_ν. The transformation of the right-handed charged leptons enters the coupling term to the B gauge boson, but it exactly cancels since $K^\dagger K = \mathbb{1}$. The transformations $e_L \to U_e e_L$ and $\nu_L \to U_\nu \nu_L$ on their own enter the coupling to the combined B and W_3 bosons (which become the Z boson after symmetry breaking), but they also cancel because neutral-current interactions do not mix neutrinos and charged leptons[3]. The only physically relevant effect is on the charged-current interactions for which the Lagrangian in the mass basis becomes

[3] This cancellation is also known as the GIM mechanism[9].

$$\mathscr{L}_{\text{CC, mass basis}} = \frac{g}{\sqrt{2}}\left[W^+_\mu \bar{\nu}^i_L \gamma^\mu (U)^{ij} e^j_L + W^-_\mu \bar{e}^i_L \gamma^\mu (U^\dagger)^{ij} \nu^j_L\right] , \tag{2.22}$$

where $U \equiv U^\dagger_\nu U_e$ is the PMNS matrix[10], which is a complex unitary matrix analogous to the CKM matrix[11] in the quark sector. The properties of the PMNS matrix will be discussed in detail in Chapter 3. As a consequence of the mixing between generations in Equation 2.22, charged-current weak interactions that produce a neutrino always generate a superposition of mass eigenstates, leading to neutrino oscillations. If we define neutrino *flavor eigenstates* as

[10]: Maki et al. (1962), *Remarks on the Unified Model of Elementary Particles*

[11]: Kobayashi et al. (1973), *CP-Violation in the Renormalizable Theory of Weak Interaction*

$$\begin{pmatrix} \nu^e_L \\ \nu^\mu_L \\ \nu^\tau_L \end{pmatrix} = U^\dagger \begin{pmatrix} \nu^1_L \\ \nu^2_L \\ \nu^3_L \end{pmatrix} , \tag{2.23}$$

then the Lagrangian in Equation 2.22 can be written in the *flavor basis* as in Equation 2.37. The flavor states of the charged leptons are identical to their mass eigenstates and therefore do not mix among each other. It is worth noting that, in this picture of pure Dirac neutrino masses, there is no mixing between the left-handed and the right-handed states, and therefore there is no observable oscillation effect that could be measured between active and sterile neutrinos.

Dirac-Majorana Masses

There exists another term that can generate fermion masses that does not require an interaction between left-handed and right-handed spinors that was discovered by E. Majorana[12]. If Ψ_L is a left-handed chiral Weyl spinor, then the term

[12]: Majorana (2008), *Teoria simmetrica dell'elettrone e del positrone*

$$m\Psi^T_L \mathscr{C}^\dagger \Psi_L = m\Psi^c_L \Psi_L \tag{2.24}$$

is a Lorentz-invariant mass term, in which $\mathscr{C}$ is the charge conjugation operator. The reason why this term cannot be used to give masses to all fermions in the SM is that it is not invariant under any of the SM symmetries. However, if we propose that right-handed neutrinos exist and that they are singlets under all SM symmetries, then a Majorana mass term can be added for them without breaking the global SM symmetry. The most general Lagrangian including all Yukawa couplings and Majorana mass terms (excluding left-handed neutrino Majorana masses) of the lepton sector is

$$\mathscr{L}_{\text{mass}} = -Y^e_{ij}\bar{L}^i\Phi e^j_R - Y^\nu_{ij}\bar{L}^i\tilde{\Phi}\nu^j_R - \frac{1}{2}M_{kl}(\nu^k_R)^c\nu^l_R + \text{h.c.}\,, \tag{2.25}$$

where the indices i and j run over the generations e, μ and τ and the matrix Y^e_{ij} contains the Yukawa coupling constants, while the matrix M_{kl} contains the Majorana masses with the indices k and l running over the N_s sterile states s_1 to s_{N_s}. Although it seems natural to associate exactly one sterile right-handed state with each active flavor, the model in principle allows for any number N_s of sterile neutrinos. By placing all the left-handed and charge-conjugated right-handed fields in column vectors

$$\boldsymbol{N}_L = \begin{pmatrix}\boldsymbol{\nu}_L \\ \boldsymbol{\nu}^c_R\end{pmatrix} \quad \boldsymbol{\nu}^c_R = \begin{pmatrix}\nu^c_{s_1R} \\ \vdots \\ \nu^c_{s_{N_s}R}\end{pmatrix}, \tag{2.26}$$

the Dirac and Majorana mass terms that are produced after symmetry breaking can be written in a unified notation as

$$\mathscr{L}^{\text{D+M}}_{\text{mass}} = \frac{1}{2}\boldsymbol{N}^T_L\mathscr{C}^\dagger M^{\text{D+M}}\boldsymbol{N}_L + h.c. \tag{2.27}$$

with the symmetric mass matrix

$$M^{\text{D+M}} \equiv \begin{pmatrix}0 & (M^D)^T \\ M^D & M^R\end{pmatrix}. \tag{2.28}$$

The mass eigenstates are linear combinations of states that linearize the mass matrix in Equation 2.28. In contrast to the pure Dirac mass, this linear combination allows for mixing between active and sterile states, and therefore oscillation effects such as those that are tested in this work are possible. Another consequence of this mixing is that the GIM mechanism no longer works, because the part of the mixing matrix that applies to the active flavors is no longer unitary. It is therefore possible for transitions between active and sterile eigenstates to occur in neutral-current interactions, which can lead to the production of Heavy Neutral Leptons (HNL)[13].

[13]: Abdullahi et al. (2022), *The Present and Future Status of Heavy Neutral Leptons*

2.1.5 See-Saw Mechanism

One of the puzzles of the SM is the question of why the neutrino masses are much smaller than those of the other fermions. If neutrino masses are purely Dirac, then the masses are proportional to the VEV of the Higgs field and their Yukawa couplings as can be seen in Equation 2.22. For neutrino masses to be small, the couplings would have to be fine tuned to be very small. If, on the other hand, neutrino masses are the result of combined Dirac-Majorana masses, then the mass of the active neutrino flavors is suppressed if the mass of the right-handed neutrinos is large. This can be understood very simply by looking at the mass term for only one generation where the mass matrix in

Equation 2.28 is a 2×2 matrix with a Dirac mass m_D in the off-diagonal entries and a Majorana mass, M, in the lower right corner. The physical masses, that is, the eigenvalues of the mass matrix, are

$$m_{1,2} = \sqrt{m_D^2 + \frac{1}{4}M^2} \pm \frac{1}{2}M \; . \tag{2.29}$$

In the limit where $M \gg m_D$, the two solutions are approximately $m_{\text{heavy}} \approx M$ and $m_{\text{light}} \approx \frac{m_D^2}{M}$. This is, in essence, the *see-saw mechanism*: As M increases, the masses of the active flavors decrease. The picture becomes more complicated when all three flavors are involved, but if the Majorana masses are much larger than the Dirac masses, the mass matrix can be approximately diagonalized by blocks for a similar effect. It is also possible to build a model with a small but non-zero left-handed Majorana mass, M^L, that replaces the zero in the upper left corner of the mass matrix in Equation 2.28, which leads to the *type-II see-saw mechanism*[14]. Such models may be attractive because they could explain the parity violation of weak interactions by spontaneous symmetry breaking of the $L \leftrightarrow R$ symmetric $SU(2)_R \times SU(2)_L \times U(1)$ gauge group. However, they also require new interactions beyond the Standard Model involving a new Higgs triplet to be added to the Lagrangian in Equation 2.25.

[14]: Mohapatra et al. (1980), *Neutrino Mass and Spontaneous Parity Nonconservation*

2.1.6 Radiative Neutrino Masses

An alternative method to naturally generate small neutrino masses is to add a contribution to the Lagrangian that generates a mass term as a loop-level correction. Such models can explain the smallness of neutrino masses without the need for additional sterile neutrinos, but they require the imposition of a new symmetry that suppresses the Yukawa coupling from Equation 2.17 at tree level. A consequence of some radiative neutrino mass models are Non-Standard Interactions (NSI) that violate either flavor conservation or lepton universality. The existence of such interactions would change the effective potential (see Section 3.2.1) felt by neutrinos passing through matter and could therefore be probed with atmospheric neutrino oscillation measurements. The discussion of radiative neutrino mass models is beyond the scope of this thesis, a comprehensive summary can be found in [15].

[15]: Cai et al. (2017), *From the Trees to the Forest: A Review of Radiative Neutrino Mass Models*

2.2 Neutrino Properties

2.2.1 Quantum Numbers

Neutrinos in the SM are described as having a spin of 1/2 and exclusively left-handed chirality. Since they are massless, their helicity is also fixed to $h = -1$. Antineutrinos are conversely right-handed with a helicity of $h = 1$. As evidenced by neutrino oscillations, they do have a very small mass, but so far only upper limits could be measured experimentally.

2.2.2 Mass

Due to their high abundance in the universe, neutrinos play an important role in the evolution of large structures in our cosmos in a way that is dependent on their masses. Cosmological observations and fits assuming the ΛCDM model[16] provide the most stringent limit on the sum of neutrino masses at

[16]: Carroll (2001), *The Cosmological constant*

$$\sum m_i < 0.12\,\mathrm{eV} \tag{2.30}$$

at 95% C.L.[17][18]. While these limits are very constraining, they strongly depend on the underlying cosmological assumptions.

[17]: Planck Collaboration et al. (2020), *Planck 2018 results - VI. Cosmological parameters*

[18]: Alam et al. (2021), *Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory*

The most sensitive direct neutrino mass measurement to date comes from the KArlsruhe TRItium Neutrino (KATRIN) experiment[19], which measures the energy spectrum of the β-decay of molecular tritium

[19]: Aker et al. (2022), *Direct neutrino-mass measurement with sub-electronvolt sensitivity*

$$T_2 \longrightarrow {}^3\mathrm{HeT}^+ + e^- + \bar{\nu}_e \ . \tag{2.31}$$

A non-zero mass m_ν of the neutrino reduces the maximum possible energy that the outgoing electron can carry, and therefore shifts the endpoint of the energy spectrum. As described in Section 2.1.4, the flavor eigenstate ν_e is not a mass eigenstate, but is instead a superposition of several mass eigenstates whose observable mass can be approximated as $m_\nu^2 = \sum_i |U_{ei}|^2 m_i^2$, where U is the PMNS matrix (see Section 3.1). The 90% C.L. upper bound on the neutrino mass from the latest KATRIN result[19] is

$$m_\nu < 0.8\,\mathrm{eV} \ . \tag{2.32}$$

This limit has the benefit of being independent of assumptions about whether neutrino masses are Dirac or Majorana and independent of cosmological models.

2.2.3 Active Neutrino Flavors

The number of active neutrino flavors (that is, flavors that interact via the Weak force) can be constrained by measurements of the decay width of the Z^0 boson. A Z^0 can decay into hadrons, charged leptons and neutrinos, with respective decay widths Γ_h, Γ_ℓ and Γ_ν. The total decay width scales with the number of active flavors, because each flavor provides an additional decay channel that increases the decay probability. Since neutrinos are invisible to the detector, the invisible part of the total decay width, Γ_{inv}, can be attributed to decays to neutrinos and the number of active flavors is

$$N_\nu = \frac{\Gamma_{\mathrm{inv}}}{\Gamma_\ell} \left(\frac{\Gamma_\ell}{\gamma_\nu} \right)_{\mathrm{SM}} \ , \tag{2.33}$$

where Γ_ℓ / γ_ν is calculated from the SM. This measurement has been done at the LEP using $e^+ + e^-$ collisions at the Z^0 resonance energy with the result[20]

[20]: ALEPH Collaboration et al. (2006), *Precision electroweak measurements on the Z resonance*

$$N_\nu = 2.9840 \pm 0.0082 \ , \tag{2.34}$$

which is compatible with the assumption of three active neutrino flavors. The result only constrains the number of neutrino states that are weakly interacting and that are lighter than the Z^0 mass, and therefore does not preclude the existence of additional sterile neutrinos.

2.3 Neutrino Interactions

Left-handed neutrinos and right-handed antineutrinos can interact with quarks and leptons via the exchange of Z^0 (neutral-current) and $W^\pm$ (charged-current) bosons. In practice, neutrinos are observed to either scatter off of electrons in reactions such as

$$\nu_e + e^- \longrightarrow \nu_e + e^- \ , \tag{2.35}$$

or to interact with nucleons. While the calculation of electron scattering is straight-forward from the electroweak Lagrangian, the scattering off nuclei is rather complicated and requires different approximations depending on the energy scale. This chapter first describes the scattering processes that are most important for the purpose of neutrino detection at energies of >1 GeV, where the total cross-section is dominated by interactions with nuclei. It then briefly summarizes the process of coherent forward scattering that influences neutrino oscillations during propagation through bulk material.

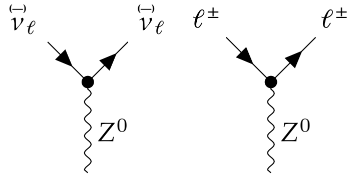

(a) neutrinos **(b)** charged leptons

Figure 2.1: Neutral-current lepton interaction vertices.

2.3.1 Weak interactions after symmetry-breaking

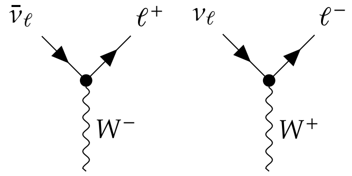

(a) W^- vertex **(b)** W^+ vertex

Figure 2.2: Charged-current lepton interaction vertices.

Neutrino interactions with matter are described by Weak force interactions after electroweak symmetry-breaking described in Section 2.1.2. The Lagrangian for these interactions can be written as the sum of the neutral-current (NC) and charged-current (CC) interactions. The NC part describes the exchange of neutral Z^0 bosons, which couples to all quarks and leptons except for right-handed neutrinos (if they exist). For leptons, the NC Lagrangian reads

$$\begin{aligned}\mathscr{L}_{\mathrm{NC,L}} = -\frac{g}{2c_W^2} \sum_{\ell=e,\mu,\tau} (&\bar{\nu}_{\ell,L}\gamma^\mu \nu_{\ell,L} + (2s_W^2 - 1)\bar{e}_{\ell,L}\gamma^\mu e_{\ell,L} \\ &+ 2s_w^2 \bar{e}_{\ell,R}\gamma^\mu e_{\ell,R})Z_\mu^0 \ ,\end{aligned} \tag{2.36}$$

where ν denotes a neutrino field, e a lepton field and the subscripts L and R denote left-handed and right-handed fields, respectively. The coefficient s_W (c_W) is the sine (cosine) of the Weinberg angle and g is the coupling constant that determines the overall strength of the electroweak force. This Lagrangian leads to the trilinear couplings shown in Figure 2.1. The couplings to quarks have the same form as those to the charged leptons up to a difference in coupling strength[4]. Neutral-current interactions conserve both the electric charge and lepton number, such that a neutral-current interaction of a neutrino will always produce a neutrino of the same flavor.

The charged-current (CC) part of the Weak Lagrangian in the flavor basis is

$$\mathscr{L}_{\mathrm{CC}} = -\frac{g}{\sqrt{2}} \sum_{\ell=e,\mu,\tau} \bar{\nu}_{\ell,L}\gamma^\mu e_{\ell,L} W_\mu^+ + \bar{e}_{\ell,L}\gamma^\mu \nu_{\ell,L} W_\mu^- + \text{h.c.} \ . \tag{2.37}$$

[4] Quark mixing has no effect on neutral current interactions due to the GIM mechanism.

In contrast to neutral current interactions, the charged current interactions couple exclusively to left-handed fields[5]. The associated lepton interaction vertices are shown in Figure 2.2.

The weak CC interactions with quarks are affected by quark mixing as a result of their mass generation via the Higgs mechanism. After electroweak symmetry breaking, the mass eigenstates and the flavor eigenstates of quarks are not identical but are instead mixed with a unitary matrix, V, that is also called the Cabbibo-Kobayashi-Maskawa (CKM) matrix. In the basis of mass eigenstates, the Lagrangian for weak CC interactions with quarks is

$$\mathscr{L}_{\mathrm{CC,Q}} = \frac{g}{\sqrt{2}} \sum_{\alpha=1}^{3} \sum_{\beta=1}^{3} \bar{u}_{\alpha,L} \gamma^{\mu} V_{\alpha\beta} d_{\beta,L} W_{\mu}^{+} + \bar{d}_{\alpha,L} \gamma^{\mu} V_{\beta\alpha}^{*} u_{\beta,L} W_{\mu}^{-} + \mathrm{h.c.}\ , \tag{2.38}$$

where the indices α and β run over the generations.

2.3.2 Neutrino-Lepton Scattering

The simplest process to consider is that of a neutrino scattering off of a single lepton, such as

$$\nu_{\mu} + e^{-} \longrightarrow \mu^{-} + \nu_{e} \tag{2.39}$$

via the exchange of a W^{+} boson. The Feynam diagram for this process is shown in Figure 2.3. To calculate the kinematics of this process, it is useful to define variables that are invariant under Lorentz transformations

$$\begin{aligned} s &= (p_{\nu} + p_{e})^{2} \\ Q^{2} &= (p_{\nu} - k_{\nu})^{2} \\ y &= \frac{p_{e} \cdot q}{p_{e} \cdot p_{\nu}}\ , \end{aligned} \tag{2.40}$$

where s is the center of mass energy, Q^{2} is the 4-momentum transfer, and y is the inelasticity. The inelasticity in the laboratory frame is the fraction of the energy carried by the outgoing lepton.

For a two-body collision between a neutrino with a negligibly small mass and a stationary target electron, the differential cross-section is

$$\frac{\mathrm{d}\sigma}{\mathrm{d}y} = \frac{m_{e} E_{\nu}}{8\pi} \frac{|\mathscr{M}|^{2}}{(s - m_{e}^{2})^{2}}\ , \tag{2.41}$$

in which $\mathscr{M}$ is the matrix element of the interaction. The matrix element can be calculated at tree level from the Feynman diagram in Figure 2.3 and the charged-current Lagrangian from Equation 2.37. In the case that $Q^{2} \ll m_{W}$ and that the energy is well above the electron production limit, the matrix element is

$$\mathscr{M}_{CC} = -\frac{G_{F}}{\sqrt{2}} \left\{ [\bar{e} \gamma^{\mu} (1 - \gamma^{5}) \nu_{e}][\bar{\nu}_{e} \gamma_{\mu} (1 - \gamma^{5}) e] \right\}\ . \tag{2.42}$$

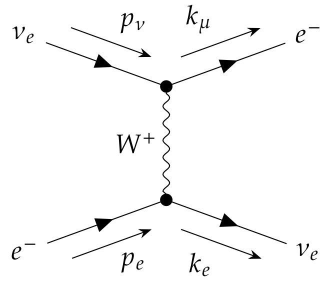

Figure 2.3: Feynman diagram for neutrino-electron scattering.

[5] The left-handed fields in the flavor basis are superpositions of mass eigenstates that may contain a (charge-conjugated) right-handed Majorana component as described in Section 2.1.4

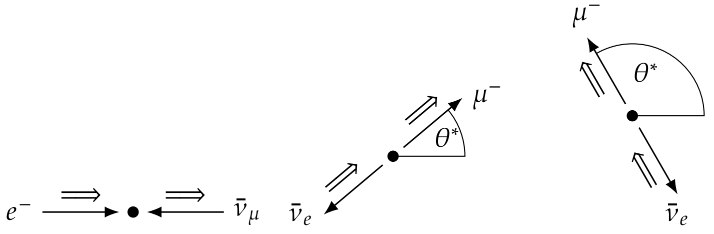

(a) Initial state **(b)** Final state with small angle **(c)** Final state with large angle

Figure 2.4: Spin configuration for particle-antiparticle interactions in the center-of-mass frame. Thin arrows (→) indicate momentum, thick arrows (⇒) show angular momentum.

Here, the constant G_F is the fermi constant[21]

$$G_F = \frac{g^2}{2\sqrt{2}M_W^2} = (1.166\,378\,8 \pm 0.000\,000\,6) \times 10^{-5}\,\mathrm{GeV}^{-2}\,. \tag{2.43}$$

The integrated cross-section for this process is

$$\sigma \simeq \frac{G_F^2 s}{\pi}\,. \tag{2.44}$$

Particle-antiparticle scattering

In a reaction between a particle and an antiparticle, such as

$$\bar{\nu}_\mu + e^- \longrightarrow \mu^- + \bar{\nu}_e\,, \tag{2.45}$$

the different spin states introduce a preferred direction to the scattering process due to the conservation of angular momentum. This can be seen easily when the reaction is illustrated in the center-of-mass frame as in Figure 2.4. The scattering process is kinematically suppressed by a factor of $(1 + \cos\theta^*)^2/4$, even though the matrix element for the reaction is the same as for particle-particle scattering. The scattering angle in the center-of-mass frame, θ^*, is related to the inelasticity by $y = (1 - \cos\theta^*)/2$, and thus the cross-section picks up an additional factor of $(1 - y)^2$. After integrating over y, the total cross-section is reduced by a factor of 3. In practice, this spin suppression of large scattering angles means that antineutrinos produce on average more highly energetic secondary leptons, while the over-all cross-section is smaller than that of neutrinos.

2.3.3 Neutrino Interactions with Nuclei

At energies of ≥1 GeV, the total cross-section of neutrinos is dominated by interactions with nuclei, while scattering off electrons can be effectively neglected. There are three processes that each have different characteristic energy ranges. The descriptions of these processes and their cross-sections largely follow those in [22].

[22]: Formaggio et al. (2012), *From eV to EeV: Neutrino Cross Sections Across Energy Scales*

Charged-current Quasi-elastic Scattering

At energies below 1 GeV, neutrinos do not resolve the inner structure of a nucleon and the scattering process can be described as an interaction with the nucleon as a whole as

$$\begin{aligned} \nu_\ell + n &\longrightarrow p + \ell^- \ , \\ \bar{\nu}_\ell + p &\longrightarrow n + \ell^+ \ . \end{aligned} \tag{2.46}$$

The differential cross-section for this process as a function of the neutrino energy E_ν is

$$\frac{\mathrm{d}\sigma}{\mathrm{d}Q^2} = \frac{G_F^2 M^2 |V_{ud}|^2}{8\pi E_\nu^2} \left[A \pm \frac{s-u}{M^2} B + \frac{(s-u)^2}{M^4} C \right] , \tag{2.47}$$

in which V is the CKM matrix, G_F is the Fermi constant, Q^2 is the squared four-momentum transfer and M is the nucleon mass. The ± sign is positive for neutrinos and negative for antineutrinos. The variables s and u are Mandelstam variables that are functions of the momentum transfer and the factors A, B and C are functions of the form factor of the nucleon. In practice, these factors depend largely only on the vector (F_1 and F_2) and axial-vector (F_A) form factor, the latter of which is

$$F_A(Q^2) = \frac{g_A}{\left(1 + \frac{Q^2}{(M_A^{CCQE})^2} \right)^2} , \tag{2.48}$$

where M_A^{CCQE} is the *axial mass*. Since the vector form factors and the coupling constant g_A are well constrained from electron scattering and nuclear beta decay, respectively, the only free parameter that must be constrained directly from measurements of the neutrino cross-sections[23] is the axial mass.

Neutral-current Elastic Scattering

Cross-sections of neutral-current interactions of the form

$$\begin{aligned} \nu p &\longrightarrow \nu p \\ \nu n &\longrightarrow \nu n \end{aligned} \tag{2.49}$$

are described by an equation that has the same form as Equation 2.47, albeit with different coupling constants, and their most important uncertainties can be parametrized with the same axial mass as in Equation 2.48. For this reason, experiments have usually measured the ratio between CC and NC interactions.

Resonant Scattering

At energies of a few GeV, neutrino scattering can produce excited states of the nucleus in reactions such as

$$\begin{aligned} \nu_\ell + N &\longrightarrow \ell^- + X^* \ , \\ \bar{\nu}_\ell + N &\longrightarrow \ell^+ + X^* \ , \end{aligned} \tag{2.50}$$

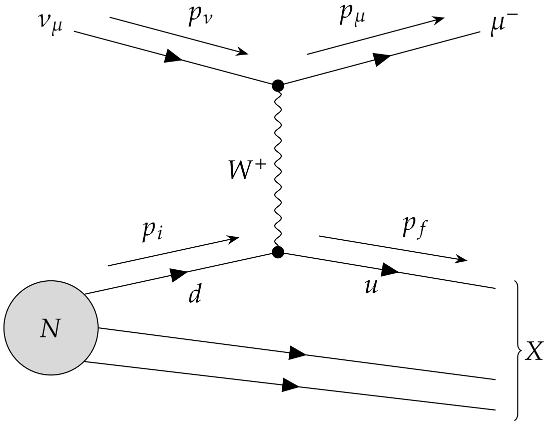

Figure 2.5: Deep inelastic scattering of a muon neutrino in the quark-parton model.

where X^* is the excited state. The differential cross-section for these processes can be approximated with an expression like Equation 2.47, with different form factors and an independent axial mass parameter, M_A^{CCRES}. The large number of possible resonant states and the necessary nuclear corrections make the correct calculation of the cross-section in this energy regime particularly challenging.

Deep Inelastic Scattering

At energies >10 GeV, the scattering process begins to resolve the inner structure of nuclei and neutrinos can scatter off of single quarks via the exchange of a $W^\pm$ or Z^0 boson, breaking up the struck nucleon into a shower of hadrons. For muon neutrinos, the possible DIS processes are

$$\begin{aligned} \nu_\mu + N \longrightarrow \mu^- X \quad & \bar{\nu}_\mu N \longrightarrow \mu^+ + X \\ \nu_\mu + N \longrightarrow \nu_\mu X \quad & \bar{\nu}_\mu N \longrightarrow \bar{\nu}_\mu + X \ . \end{aligned} \tag{2.51}$$

The Feynman diagram of the charged-current process is shown in Figure 2.5. The struck quark initiates a hadronization process that can produce multiple mesons in the output shower X. The kinematics of this reaction can be calculated using the same Lorentz-invariant quantities that are defined for neutrino-lepton scattering in Equation 2.40 and one additional variable,

$$x \equiv \frac{Q^2}{p_N \cdot q} = \frac{Q^2}{(s - m_N^2)y} \tag{2.52}$$

that is referred to as Bjorken-x. In the lab frame, the inelasticity is the fraction of the energy of the initial neutrino that is converted into hadrons

$$y_{\text{lab}} = 1 - \frac{E_\mu}{E_\nu} = E_{\text{hadrons}}/E_\nu \ . \tag{2.53}$$

The conditions for deep inelastic scattering to occur are

$$\begin{aligned} Q^2 &\gg m_N^2 \\ p_N \cdot q &\gg m_N^2 \\ s &\gg m_N^2 \ . \end{aligned} \tag{2.54}$$

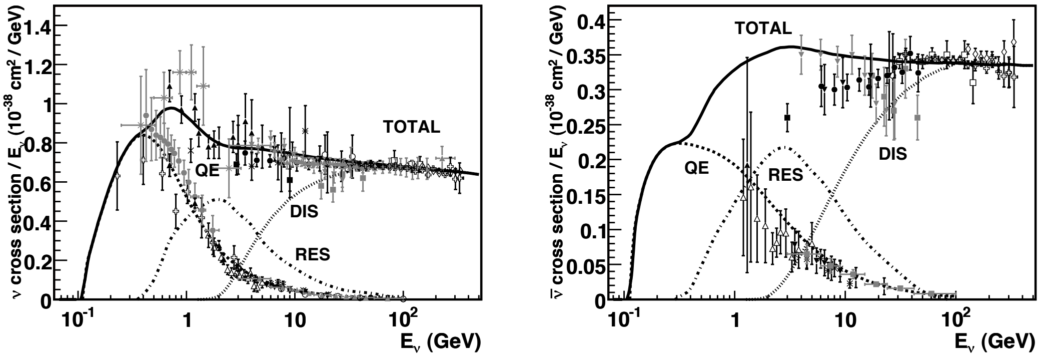

(a) neutrinos

(b) antineutrinos

Figure 2.6: Inclusive cross sections for neutrinos and antineutrinos. Figure taken from [22].

Under these conditions, the differential cross-sections for DIS for neutrinos and antineutrinos scattering off of a nucleon are

$$\begin{aligned}\frac{\mathrm{d}^2\sigma^{\nu N}_{\mathrm{CC}}}{\mathrm{d}x\mathrm{d}y} &= 2x\sigma^0_{\mathrm{CC}}\left[\sum_{q=d,s} f^N_q(x) + (1-y)^2 \sum_{\bar{q}=\bar{u},\bar{c}} f^N_{\bar{q}}(x)\right] \\ \frac{\mathrm{d}^2\sigma^{\bar{\nu} N}_{\mathrm{CC}}}{\mathrm{d}x\mathrm{d}y} &= 2x\sigma^0_{\mathrm{CC}}\left[\sum_{\bar{q}=\bar{d},\bar{s}} f^N_{\bar{q}}(x) + (1-y)^2 \sum_{q=u,c} f^N_q(x)\right] ,\end{aligned} \tag{2.55}$$

with

$$\sigma^0_{\mathrm{CC}} = \frac{G_F^2}{2\pi} s \left(1 + \frac{Q^2}{m_W^2}\right)^{-2} . \tag{2.56}$$

The functions $f^N_{\overset{(-)}{q}}(x)$ are the parton distribution functions for each type of quark inside a nucleus. The factor of $(1-y)^2$ for scattering processes that mix particles and antiparticles comes from conservation of angular momentum, as outlined earlier for neutrino-lepton scattering. This kinematic suppression reduces the inclusive cross-section for antineutrinos by a factor of approximately two. The cross-section for neutral-current interactions has the same form as Equation 2.55, only that the mass m_W is replaced by m_Z and the sum runs over left-handed and right-handed quark states that each have different coupling constants.

Neutrino Oscillations | 3

3.1 Neutrino Oscillations in Vacuum 19
3.1.1 Two-neutrino mixing . . 20
3.1.2 Decoherence 21
3.2 Oscillations in matter . . 22
3.2.1 Effective potentials 22
3.2.2 The MSW Effect 24
3.2.3 Parametric resonance . . 25
3.2.4 Neutrino Mass Ordering 25
3.3 Standard three-flavor oscillations 26
3.3.1 Solar neutrinos 27
3.3.2 Reactor neutrinos 28
3.3.3 Atmospheric Neutrino Oscillations 29
3.3.4 Accelerator neutrinos . . 32
3.4 Anomalies in neutrino oscillation measurements 33
3.4.1 Reactor neutrino anomaly 33
3.4.2 Gallium anomaly 35
3.4.3 LSND and MiniBooNE Anomalies 35
3.4.4 Global picture of oscillation anomalies 37
3.4.5 Other constraints on sterile neutrinos 38
3.5 Open Questions in Neutrino Oscillation Physics 40

As outlined in Section 2.1.4, flavor eigenstates of neutrinos are not identical to their mass eigenstates. The mass and flavor eigenstates mix among each other via the PMNS matrix U that, in the case of pure Dirac masses, is a 3×3 unitary matrix. Thus, the neutrino flavor eigenstates can be described as a superposition of mass eigenstates as

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k} |\nu_k\rangle \ , \tag{3.1}$$

where $|\nu_k\rangle$ is the kth mass eigenstate and $|\nu_\alpha\rangle$ is the neutrino state of the flavor α. As a matter of experimental fact, the mass eigenstates that make up the three known flavor eigenstates have different masses. This has profound consequences for the propagation of neutrinos, because the wavepackets of the different mass eigenstates do not travel at exactly the same speed. The lighter states are faster than the heavier ones, which causes the waves to interfere constructively or destructively as long as the wave packets still overlap. This chapter describes how this interference of eigenstates leads to the phenomenon of neutrino oscillations and how these oscillations can be described in vacuum and in matter. The derivations are presented under the simplifying assumption that the mass eigenstates are ideal plane waves. The fact that they are actually wave packets with an uncertain energy and a finite extent leads to decoherence for very large propagation distances or mass differences. While this is irrelevant for the standard three-flavor oscillation result of this work, it does put an upper limit on the mass splitting for which the sterile oscillation result is valid. This will be described briefly in Section 3.1.2.

3.1 Neutrino Oscillations in Vacuum

The simplest case to describe is that of neutrino oscillations in vacuum. The propagation of these states is governed by the Schrödinger equation with the Hamiltonian $\mathscr{H}$

$$i\frac{\mathrm{d}}{\mathrm{d}t} |\nu_k(t)\rangle = \mathscr{H} |\nu_k(t)\rangle \ . \tag{3.2}$$

This equation has the plane wave solution

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle \ , \tag{3.3}$$

where E_k is the energy of the state $|\nu_k\rangle$ and t is the propagation time. Substituting the flavor eigenstate from Equation 3.1 into Equation 3.3 using the relation $U^\dagger U = \mathbb{1}$, the propagation becomes

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_k U^*_{\alpha k} e^{-iE_k t} U^*_{\beta k} \right) |\nu_\beta\rangle \ . \tag{3.4}$$

This leads directly to the expression for the probability of measuring one flavor after a given time

$$P_{\nu_\alpha \to \nu_\beta} = \left|\left\langle \nu_\beta \middle| \nu_\alpha(t) \right\rangle\right|^2 = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} e^{-i(E_k - E_j)t} \; . \tag{3.5}$$

If we assume that all mass eigenstates have the same momentum[1] and that they are highly relativistic, we can approximately express the energy in terms of the mass of each state

$$E_k = \sqrt{\vec{p}^{\,2} + m_k^2} \simeq E + \frac{m_k^2}{2E} \tag{3.6}$$

and write the transition probability in terms of the energy and the differences of squared masses between the mass eigenstates

$$\Delta m^2_{kj} \equiv m_k^2 - m_j^2 \tag{3.7}$$

and the distance traveled, L, as

$$P_{\nu_\alpha \to \nu_\beta}(L, E) = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \exp\left(-i \frac{\Delta m^2_{kj} L}{2E}\right) . \tag{3.8}$$

This state evolution can also be expressed with the effective Hamiltonian

$$\mathscr{H}_{\text{eff}} = \frac{1}{2E} U \text{diag}(0, \Delta m^2_{21}, \Delta m^2_{31}) U^\dagger \tag{3.9}$$

and one can define the characteristic distance for oscillations between mass eigenstates k and j at which the oscillation amplitude is maximal as

$$L^{\text{osc}}_{kj} = \frac{4\pi E}{\Delta m^2_{kj}} \; . \tag{3.10}$$

Antineutrino flavor eigenstates are superpositions of the corresponding antineutrino mass eigenstates, and are related by the complex-conjugated PMNS matrix such that

$$|\bar{\nu}_\alpha\rangle = \sum_k U_{\alpha k} |\bar{\nu}_k\rangle \; . \tag{3.11}$$

This leads to the same expression for their oscillation probability as in Equation 3.8, except that all elements of U are complex-conjugated.

3.1.1 Two-neutrino mixing

Neutrino oscillation experiments are typically limited to a certain oscillation length and energy range that they can probe. Because the two mass splittings between the three known neutrino flavors are two orders of magnitude apart ($\mathcal{O}(10^{-5}\,\text{eV}^2)$ for Δm^2_{21} vs. $\mathcal{O}(10^{-3}\,\text{eV}^2)$ for Δm^2_{32})[24], each experiment is in practice much more sensitive to one of them than to the other as will be discussed in Section 3.3. This can also be seen in the oscillation plot in Figure 3.1, which shows the transition probabilities of an initial electron neutrino into the other flavors in vacuum as a function of oscillation distance. It is therefore a good first approximation to calculate neutrino oscillation probabilities is to only

[1] The assumption of equal momentum is not exactly realistic, but more detailed derivations can show that deviations from it do not cause an observable effect. See also Chapter 8.1.2 in [7].

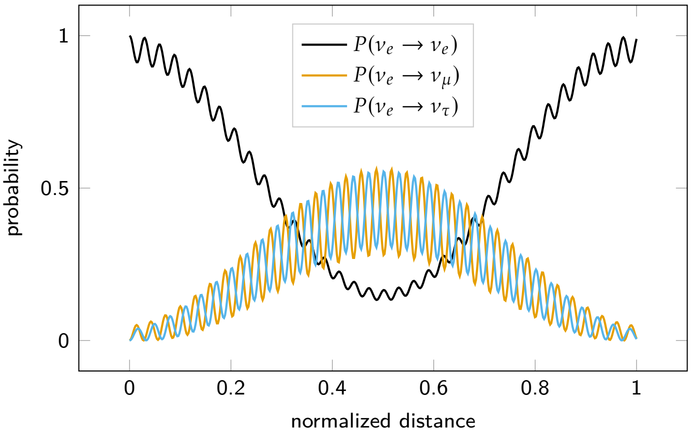

Figure 3.1: Transition probabilities of an initial electron neutrino in vacuum, calculated at global best-fit[24] oscillation parameters. The distance is normalized such that the slow oscillations due to the smaller mass splitting Δm^2_{21} have a period of one.

consider two flavor and two mass eigenstates with a mass splitting of $\Delta m^2 \equiv m_2^2 - m_1^2$ that mix via the rotation matrix

$$U = \begin{pmatrix} \cos\vartheta & \sin\vartheta \\ -\sin\vartheta & \cos\vartheta \end{pmatrix} , \tag{3.12}$$

where the angle ϑ is the *mixing angle* between the two mass eigenstates. The *transition probability* from flavor α to flavor β with $\alpha \neq \beta$ can be quickly derived from the general Equation 3.8 to be

$$P_{\nu_\alpha \to \nu_\beta} = \sin^2 2\vartheta \sin^2\left(\frac{\Delta m^2 L}{4E}\right), \quad \alpha \neq \beta . \tag{3.13}$$

Conversely, the probability that an initial flavor α is still being measured as α after a given propagation distance, also referred to as the *survival probability*, is

$$P_{\nu_\alpha \to \nu_\alpha} = 1 - \sin^2 2\vartheta \sin^2\left(\frac{\Delta m^2 L}{4E}\right) . \tag{3.14}$$

3.1.2 Decoherence

The previous sections treated the neutrino mass eigenstates as if they were plane waves with a fixed energy. This description is an approximation that holds very well for atmospheric neutrino oscillations in the three-flavor picture, but can break down if the mass splitting between states is large. The size of neutrino wave packets, σ_x^ν, is determined by the coherence size of the production process. The effective wave packet size for neutrinos that are produced in pion decays within a gas is calculated in detail in [25]. The competing processes of relativistic wave packet dispersion and scattering in gas leads to an overall energy dependence

$$\sigma_x^\nu \propto E^{-1} . \tag{3.15}$$

[25]: Jones (2015), *Dynamical pion collapse and the coherence of conventional neutrino beams*

Due to the finite size of neutrino wave packets, the vacuum flavor transition probability from Equation 3.8 acquires a dampening term and becomes

$$P_{\nu_\alpha \to \nu_\beta}(L, E) = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \exp\left(-i 2\pi \frac{L}{L^{\text{osc}}_{j,k}} - \left(\frac{L}{L^{\text{coh}}_{j,k}}\right)^2\right) . \tag{3.16}$$

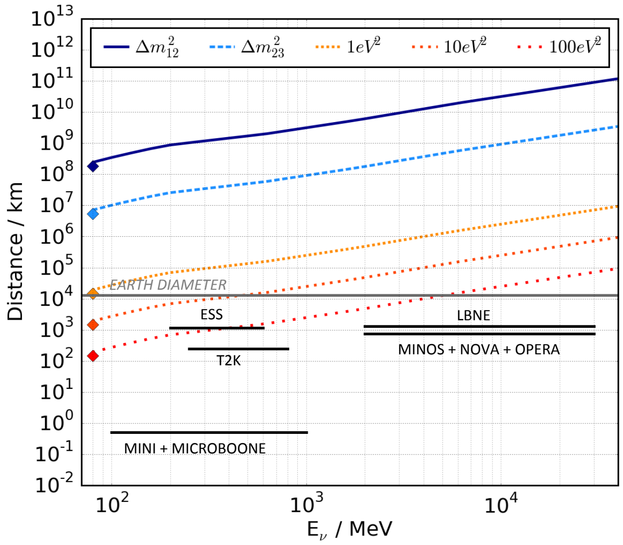

Figure 3.2: Coherence lengths for neutrinos produced in pion decays at different mass splittings. The black solid lines indicate the baseline and energy ranges of some existing and proposed accelerator experiments. Figure taken from [25].

The oscillation length $L^{\mathrm{osc}}_{j,k}$ is the same as in Equation 3.10 and the decoherence length is

$$L^{\mathrm{coh}}_{j,k} = \frac{4\sqrt{2}E^2}{\left|\Delta m^2_{kj}\right|}\sigma^\nu_x \,. \tag{3.17}$$

Taking the expression in Equation 3.17 together with the E^{-1} dependence of the packet size leads to a scaling of the coherence length with E as shown in Figure 3.2. For baselines of the size of one Earth diameter and energies of tens of GeV that are relevant for atmospheric neutrino oscillations, decoherence becomes relevant at mass splitting values approaching $100\,\mathrm{eV}^2$.

3.2 Oscillations in matter

The oscillation probabilities derived in the previous sections for the vacuum are altered significantly when neutrinos pass through large amounts of matter. This is particularly relevant for neutrinos originating in the atmosphere of the Earth that pass through its dense core to be detected at the South Pole. The effect can be described as a continuous potential that is added to the Hamiltonian in the flavor basis, and which corresponds to the isoscalar (nulcei) targets and electrons inside the Earth. The detailed derivation of this matter potential can be found in [26]. Only a brief outline is presented here to illustrate how the matter potential leads to an enhancement of the effective mixing between flavors in the simplified picture of two flavors and a constant matter density.

[26]: Linder (2005), *Neutrino matter potentials induced by Earth*

3.2.1 Effective potentials

At energies relevant to this work[2], neutrinos are mostly affected by coherent forward scattering processes, that is, scattering with a very

[2] Incoherent scattering becomes only important at energies above $\mathcal{O}(10^5\,\mathrm{GeV})$.

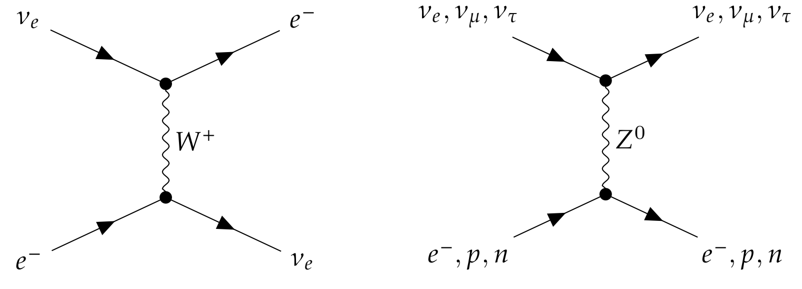

(a) Charged-current **(b)** Neutral-current

Figure 3.3: Feynman diagrams of the coherent forward scattering processes for neutrinos travelling through Earth.

small momentum transfer. Feynman diagrams of the relevant processes are shown in Figure 3.3. All neutrino flavors can interact with electrons and nuclei via the neutral-current (NC) interaction, while only electron-neutrinos can also undergo charged-current (CC) scattering off of electrons. When the exchanged momentum is small, then the propagator in each Feynman diagram can be contracted into a 4-point Fermi interaction with the coupling constant G_F. The effective CC Hamiltonian for the diagram in Figure 3.3a (after a Fierz transformation) becomes

$$\mathcal{H}_{\mathrm{eff}}^{CC}(x) = \frac{G_F}{\sqrt{2}}[\bar{\nu}_e(x)\gamma^\rho(1-\gamma^5)\nu_e(x)][\bar{e}(x)\gamma_\rho(1-\gamma^5)e(x)] \ . \tag{3.18}$$

This Hamiltonian is then averaged over the momenta and helicities of a constant density of electrons with the result

$$\overline{\mathcal{H}_{\mathrm{eff}}^{CC}}(x) = V_{CC}\bar{\nu}_{e,L}(x)\gamma^0\nu_{e,L}(x) \ , \tag{3.19}$$

where V_{CC} is the charged-current potential

$$V_{CC} = \sqrt{2}G_F N_e \tag{3.20}$$

with the electron number density N_e. The derivation of the NC potential follows in similar steps with the addition of a different coupling constant for electrons, protons, and electrons. For electrically neutral media, the potentials for protons and electrons exactly cancel and the only remaining NC potential comes from the neutron density, N_n, and is given by

$$V_{NC} = -\frac{1}{2}\sqrt{2}G_F N_n \ . \tag{3.21}$$

With the addition of these potentials, the effective Hamiltonian governing the evolution of flavor states for neutrinos propagating in matter is

$$\mathcal{H}_{\mathrm{eff}}(E,x) = \mathcal{H}_0(E) + \mathcal{H}_1(E,x) \ , \tag{3.22}$$

with

$$\mathcal{H}_0(E) = \frac{1}{2E}U\begin{pmatrix} 0 & & \\ & \Delta m_{21}^2 & \\ & & \Delta m_{31}^2 \end{pmatrix}U^\dagger \ , \tag{3.23a}$$

$$\mathcal{H}_1(E,x) = \frac{\sqrt{2}}{2}G_F\begin{pmatrix} 2N_e(x)-N_n(x) & & \\ & -N_n(x) & \\ & & -N_n(x) \end{pmatrix} . \tag{3.23b}$$

The NC potential in Equation 3.23b can be neglected in the case of three-flavor oscillations, because a diagonal contribution to the Hamiltonian merely adds an unobservable phase shift that affects all flavors equally. However, if sterile Majorana mass eigenstates are added, their corresponding matter potentials are zero and they only add their respective mass splitting to Equation 3.23a. In that case, the state evolution has to be described using both CC and NC potentials.

3.2.2 The MSW Effect

An interesting consequence of matter potentials is that they can greatly enhance the mixing amplitude between flavors over that of vacuum oscillations, a phenomenon known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect[27][28]. The effect is straight forward to illustrate for the case of two-flavor oscillations, where one flavor feels a potential, V, and the other does not. With the mass-splitting $\Delta m^2 \equiv m_2^2 - m_1^2$ and 2×2 mixing matrix from Equation 3.12, the effective Hamiltonian for this scenario is

[27]: Wolfenstein (1978), *Neutrino oscillations in matter*

[28]: Smirnov (2003), *The MSW effect and solar neutrinos*

$$\begin{aligned}\mathcal{H} &= \frac{1}{4E}\left(U\begin{pmatrix}0 & \\ & \Delta m^2\end{pmatrix}U^\dagger + \begin{pmatrix}A_{\mathrm{CC}} & \\ & 0\end{pmatrix}\right)\\ &= \frac{1}{4E}\begin{pmatrix}-\Delta m^2\cos 2\vartheta + A_{\mathrm{CC}} & \Delta m^2 \sin 2\vartheta \\ \Delta m^2 \sin 2\vartheta & \Delta m^2 \cos 2\vartheta - A_{\mathrm{CC}}\end{pmatrix},\end{aligned} \tag{3.24}$$

where we have subtracted $\frac{1}{4E}(\Delta m^2 + A_{\mathrm{CC}})$ from the diagonal in the second line with $A_{\mathrm{CC}} = 2EV$. This Hamiltonian can be diagonalized by another 2D rotation matrix U_M, with mixing angle ϑ_M, such that

$$U_M^T \mathcal{H} U_M = \mathcal{H}_M = \frac{1}{4E}\begin{pmatrix}-\Delta m_M^2 & \\ & \Delta m_M^2\end{pmatrix} \tag{3.25}$$

is diagonal. The diagonalization is achieved under the condition that

$$\tan 2\vartheta_M = \frac{\tan 2\vartheta}{1 - \frac{A_{\mathrm{CC}}}{\Delta m^2 \cos 2\vartheta}}\,, \tag{3.26}$$

leading to a new effective mass splitting in matter,

$$\Delta m_M^2 = \sqrt{(\Delta m^2 \cos 2\vartheta - A_{\mathrm{CC}})^2 + (\Delta m^2 \sin 2\vartheta)^2}\,. \tag{3.27}$$

As a consequence of the condition in Equation 3.26, the effective mixing angle in matter can become maximal under the condition that

$$V = \frac{\Delta m^2 \cos 2\vartheta}{2E_{\mathrm{res}}} \tag{3.28}$$

with the resonance energy E_{res}. This enhancement of mixing between flavors due to the presence of a matter potential is known as the *MSW effect*, named after Mikheev, Smirnov and Wolfenstein. It is primarily responsible for the fact that there is a sizable disappearance of electron neutrinos from the Sun despite the fact that the relevant vacuum mixing angle is rather small. The potential for isoscalar targets is

$$V = \sqrt{2} G_F N_e \approx Y_e \times \left[\frac{\rho}{\mathrm{g\,cm^{-3}}}\right] \times 7.63 \times 10^{-14}\,\mathrm{eV}\,, \tag{3.29}$$

where $Y_e = N_e/(N_p + N_n)$ is the electron fraction in the medium and is usually close to $\frac{1}{2}$. We can find the resonant energy where the MSW resonance occurs for a given matter density by combining Equations 3.28 and 3.29 into

$$E_{\text{res}} = \frac{\Delta m^2 \cos 2\vartheta}{2V} \approx \frac{(\Delta m^2 \text{eV}^{-2}) \cos 2\vartheta}{Y_e \left[\frac{\rho}{\text{g cm}^{-3}}\right]} \times 6.6 \times 10^{12}\,\text{eV}\ . \tag{3.30}$$

For three-flavor oscillations observed experimentally, the MSW resonance can enhance mixing due to the relatively small mixing angle $\theta_{13} = (8.60 \pm 0.13)°$[21]. This occurs at $\approx 7\,\text{GeV}$ for muon neutrinos traveling through the mantle of the Earth, where the density is $\rho_{\text{mantle}} \approx 4.5\,\text{g cm}^{-3}$ and where the observed oscillations are due to the larger mass splitting $\Delta m^2_{32} \approx 2.5 \times 10^{-3}\,\text{eV}^2$. For hypothetical sterile neutrinos with a mass splitting of $\mathcal{O}(1\,\text{eV})$ passing through the core of the Earth with $\rho \approx 13\,\text{g cm}^{-3}$ and $Y_e \approx \frac{1}{2}$, this gives a resonant energy of close to $1\,\text{TeV}$ assuming a small mixing angle between active and sterile flavors. This effect has been used in previous IceCube studies[29] to search for traces of sterile neutrinos using TeV scale atmospheric neutrinos, but it is not relevant to the analysis presented in this work.

3.2.3 Parametric resonance

Another way in which the presence of matter potentials can enhance neutrino oscillation amplitudes is via parametric resonance. This resonance occurs when the matter potential changes abruptly at just the right frequency to create a situation that is analogous to constructive interference, even if the density is far from the MSW resonance condition. The reason why this can happen is that the quantum state of the neutrino receives a phase shift at every interface between regions of different densities. If such an interface is located at the oscillation maximum and the difference in densities has the right magnitude, the phase is reset to zero and and the flavor transition probability will increase again in the next region, instead of returning to zero. This situation can arise in particular at the transition points between the mantle and the core when neutrinos pass through the Earth as shown in Figure 3.4. The black solid line shows the simplified three-layer matter potential that is used to approximate the density of the Earth in the oscillation calculation. At each interface (dashed lines), the phase shift causes the oscillation pattern to begin again at the rising flank. After the neutrino propagates through the Earth, the oscillation probability has increased to close to one. The condition for maximum parametric resonance is delicately dependent on the energy, mass splitting, mixing angle, and the parameters of the matter potential and are derived in more detail in [30]. In atmospheric neutrino oscillations, the parametric enhancement of the $\nu_\mu \rightarrow \nu_e$ oscillation channel is responsible for the distortions of the oscillation pattern at a few GeV.

3.2.4 Neutrino Mass Ordering

In vacuum, neutrino oscillations are only sensitive to the absolute differences between (squared) neutrino masses, not their sign. As a

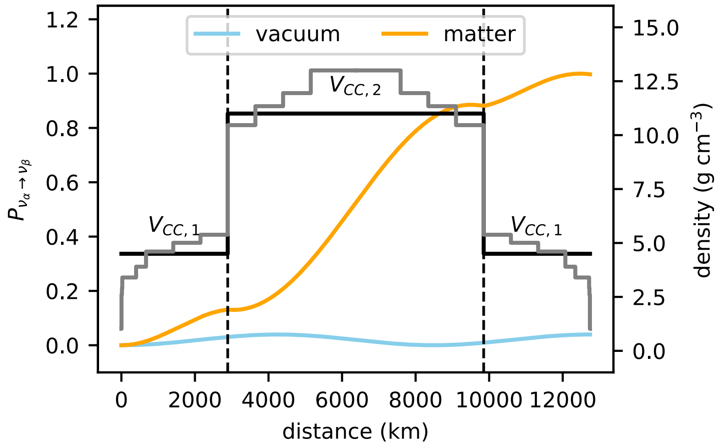

Figure 3.4: Parametric resonance for a simplified mantle-core-mantle propagation with two flavors. The assumed mass-splitting is $\Delta m^2 = 2.5 \times 10^{-3}\,\mathrm{eV}^2$, the mixing angle is $\theta = \arcsin(0.1)$, and the energy is $3.4\,\mathrm{GeV}$. The matter potential used for the calculation. It is an approximation of the density profile of Earth in the 12-layer PREM model, shown in gray.

result, there are two degenerate possibilities for arranging the mass eigenstates that produce the same vacuum oscillation effect. The case where $m_1^2 < m_2^2 < m_3^2$ is referred to as *normal ordering*. In this case, the value of the mass splitting Δm_{ij}^2 is always positive for $i > j$. For *inverted ordering*, we have $m_3^2 < m_1^2 < m_2^2$. In that case, the sign of Δm_{32}^2 is negative. While the ordering has no observable effect in vacuum, it does affect the propagation in the presence of matter effects. This can be seen readily in the MSW resonance condition in Equation 3.26, where the sign of Δm^2 enters the denominator. The resonance condition for neutrinos can only be fulfilled if the sign of Δm^2 is the same as $\cos 2\theta$ for neutrinos, or if the signs of Δm^2 and $\cos 2\theta$ are different for antineutrinos. Therefore, the neutrino mass ordering can be determined experimentally by measuring the differences in matter effects between neutrinos and antineutrinos.

3.3 Standard three-flavor oscillations

The simplest extension to the Standard Model that explains the phenomenon of neutrino oscillations is that of three oscillating flavor eigenstates and three mass eigenstates that result from pure Dirac masses as described in Section 2.1.4. In this picture, the mixing between mass eigenstates is governed by the 3×3 PMNS matrix that is commonly parametrized by three mixing angles and one complex phase,

$$
\begin{aligned}
U &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \\
&= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix},
\end{aligned}
\tag{3.31}
$$

where s_{ij} and c_{ij} are the sine and cosine of the mixing angle θ_{ij}, respectively. The complex phase δ leads to CP-violation of the oscillations. Furthermore, there are two mass splitting values between the three mass eigenstates, Δm_{21}^2 and Δm_{31}^2, which govern the frequency of the oscillations. The mixing angles that parameterize the PMNS matrix and the mass splittings are typically associated with the types of experiments that are most sensitive to them based on the oscillation channel they observe, their energy range, and the distance between the source and the detector, also referred to as the *baseline* of the experiment. In order for an experiment to resolve oscillations, the oscillation phase should be

$$\frac{\Delta m^2 L}{4E} \approx 1.27 \frac{(\Delta m^2_{31}/\mathrm{eV}^2)(L/\mathrm{km})}{(E/\mathrm{GeV})} \sim \mathcal{O}(1) \; . \tag{3.32}$$

This condition leads to two different categories of experiments, the first being sensitive to the very slow oscillations due to $\Delta m^2_{21} \sim \mathcal{O}\left(10^{-5}\,\mathrm{eV}^2\right)$ and the second one probing the faster oscillations due to $\Delta m^2_{31} \sim \mathcal{O}\left(10^{-3}\,\mathrm{eV}^2\right)$. Because the slow oscillation channel was first observed in neutrinos originating from the Sun, Δm^2_{21} is also referred to as the "solar" mass splitting. The larger mass splitting Δm^2_{31} was measured for the first time in neutrinos that are produced in the atmosphere of the Earth and is also called the "atmospheric" mass splitting. Figure 3.5 gives an overview over different oscillation experiments, both operational and planned, with their respective energy ranges and baselines. Experiments on the line marked $L_\odot$ probe the solar mass splitting Δm^2_{21}, while those on the line marked L_atm measure oscillations from the atmospheric mass splitting of Δm^2_{21}.

Similarly, the mixing angles are also associated with those experiments that are most sensitive to their value. Since the angle θ_{13} is small, the PMNS matrix in Equation 3.31 can be decomposed to a first approximation into the upper left and lower right 2×2 blocks. The oscillation probability in each block can be approximated by oscillations between only two flavors with one mass splitting. The upper left block governs the transitions between the ν_e and ν_μ flavors with the mixing angle θ_{12} and the solar mass splitting Δm^2_{21}. For this reason, the mixing angle θ_{12} is also referred to as the "solar angle". The lower right block of the PMNS matrix controls transitions between ν_μ and ν_τ flavors with the mixing angle θ_{23} and atmospheric mass splitting Δm^2_{32}. Therefore, θ_{23} is often called the "atmospheric" mixing angle, although it can also be measured using neutrinos that are produced in accelerators. The effect of the smaller mixing angle θ_{13} is that of a disappearance of electron neutrinos at short distances. Such transitions can be observed using electron antineutrinos that are produced by fission reactors. Thus, the angle θ_{13} is also called the "reactor" angle.

The most recent global best-fit point[31] of the mixing angles and mass splittings under the assumption of normal neutrino mass ordering is shown in Table 3.1. The table also lists the experimental channels that are primarily used to measure each parameter.

3.3.1 Solar neutrinos

The Sun produces energy in two different nuclear reactions known as the *pp* chain and the CNO cycle. The net result of these reactions is the conversion of four protons and two electrons into a $^4\mathrm{He}$ nucleus and

Table 3.1: Best-fit values of the three-flavor oscillation parameters and the oscillation channel that is primarily involved in their measurement. Values taken from [31].

parameter	global fit	experimental channel
$\theta_{12}/^\circ$	$33.44^{+0.77}_{-0.74}$	$\nu_e \to \nu_e$ (solar)
$\frac{\Delta m^2_{21}}{10^{-5}\,\mathrm{eV}^2}$	$7.42^{+0.21}_{-0.20}$	$\bar{\nu}_e \to \bar{\nu}_e$ (reactor)
$\theta_{23}/^\circ$	$49.2^{+1.0}_{-1.3}$	$\nu_\mu \to \nu_\mu$ (atmospheric, accelerator)
$\frac{\Delta m^2_{31}}{10^{-3}\,\mathrm{eV}^2}$	$2.515^{+0.028}_{-0.028}$	
$\theta_{13}/^\circ$	$8.57^{+0.13}_{-0.12}$	$\bar{\nu}_e \to \bar{\nu}_e$ (reactor)
$\delta_{\mathrm{CP}}/^\circ$	194^{+52}_{-25}	$\nu_\mu \to \nu_e$ (accelerator)

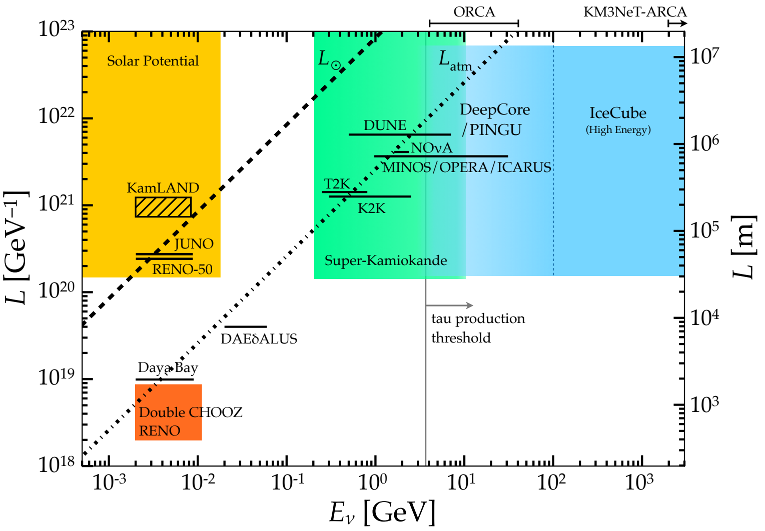

Figure 3.5: Energy ranges and baselines for a selection of neutrino oscillation experiments, operational or planned. The range of each experiment is shown either as a line for experiments with a fixed oscillation distance such as T2K or DUNE, or as a box for experiments with both a variable baseline and energy range such as IceCube. Figure taken from[32].

two neutrinos,

$$4\mathrm{p} + 2\mathrm{e}^- \to^4 \mathrm{He} + 2\nu_\mathrm{e} \; . \tag{3.33}$$

Measurements of the solar neutrino flux are the oldest hints at neutrino oscillations and began at the end of the 1960s with the Homestake[33] experiment. It used the inverse beta decay process $^{37}\mathrm{Cl}+\nu_e \to e^- +^{37}\mathrm{Ar}$ to capture neutrinos with an energy threshold of $0.814\,\mathrm{MeV}$. The observed neutrino flux was only about 30% of the expectation from the Standard Solar Model[34], giving rise to the *solar neutrino problem*. This problem was resolved in 2002, when observations at the Sudbury Neutrino Observatory (SNO) confirmed[5] that this deficit is due to the conversion of electron neutrinos into muon neutrinos, while the total neutrino flux, as measured by the rate of neutral-current interactions, remains constant. Terrestrially, these oscillations can only be observed using very low energy neutrinos from nuclear fission reactors at baselines of hundreds of kilometers. The only experiment of this kind that is currently operational is the KamLAND detector[35].

[34]: Bahcall et al. (2005), *New solar opacities, abundances, helioseismology, and neutrino fluxes*

3.3.2 Reactor neutrinos

Nuclear fission reactors produce primarily electron antineutrinos that can be used to measure the electron-neutrino survival probability at short distances where oscillation effects are dominated by the larger mass splitting Δm^2_{31}. The amplitude of fast $\nu_e \to \nu_e$ oscillations is controlled by the $|U_{e3}|$ element of the PMNS matrix in Equation 3.31 that is proportional to $\sin(\theta_{13})$. The most sensitive experiments measuring θ_{13} are nuclear reactor experiments such as Double Chooz[36], RENO[37], and the Daya Bay Reactor Neutrino Experiment[38]. These experiments feature two or more detectors set up at different distances from a nuclear reactor that measure the neutrino flux from a nuclear reactor. Comparative measurements between these detectors allow the mixing angle θ_{13} to be precisely determined while canceling most of the systematic uncertainties of the experiment.

3.3.3 Atmospheric Neutrino Oscillations

The first experiment that was able to measure the effect of atmospheric neutrino oscillations is the Super-Kamiokande (SuperK) experiment[4], a Cherenkov neutrino detector located in the Kamioka mountains in Japan, which is still operational to this day. Since the measurement presented in this work is also a measurement of atmospheric neutrino oscillations, this kind of measurement is explained in more detail compared to other types of experiment in this section.

Neutrino production in the atmosphere

Neutrinos are produced in the atmosphere of the Earth when highly energetic cosmic radiation interacts with its upper layers. This cosmic radiation is composed mostly of fast protons and helium nuclei and to a smaller fraction of other elements. The composition and spectrum of this radiation have been extensively measured by experiments conducted at high altitudes, such as in weather balloons, satellites, and on the International Space Station. From a few GeV to approximately 100 TeV, it can be well approximated by a power spectrum with a spectral index of $\gamma \approx 2.7$. When a cosmic ray interacts with a nucleus, it creates a shower of hadrons that is mostly composed of Pions and Kaons that subsequently decay into muons, electrons, and neutrinos in the reaction chain

$$\begin{aligned} \mathrm{C.R.} + N &\longrightarrow X + \pi^{\pm}, K^{\pm} \\ \pi^{\pm} &\longrightarrow \mu^{\pm} + \overset{(-)}{\nu}_{\mu} \\ \mu^{\pm} &\longrightarrow e^{\pm} + \overset{(-)}{\nu}_{e} + \overset{(-)}{\nu}_{\mu} \; . \end{aligned} \tag{3.34}$$

At energies of $\mathcal{O}(1\,\mathrm{GeV})$, these reactions lead to a production of muon neutrinos and electron neutrinos at a ratio of $2:1$. However, as the energy increases, muons can increasingly reach the surface of the Earth and interact before they can decay, increasing the relative fraction of muon neutrinos as shown in Figure 3.6b. Between GeV and TeV energies, the ν_{μ} flux has a spectral index that is similar to that of the primary cosmic rays at $\gamma \approx 2.7$ as can be seen in Figure 3.6a, while the spectrum for electron neutrinos falls off more quickly at higher energies. Tau neutrinos can also be produced if a cosmic ray interaction produces charmed mesons, but such interactions are rare and the atmospheric tau neutrino flux is below 0.1% for the energies considered in this work[39] and is therefore neglected.

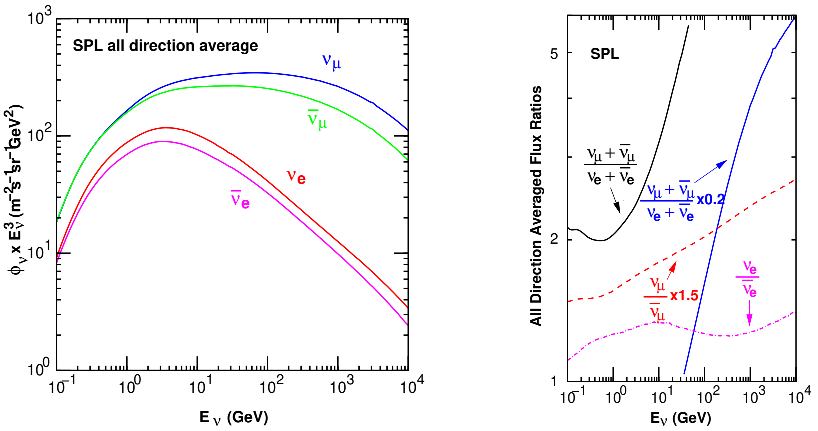

Figure 3.6: Neurino flux at the South Pole, averaged over all directions. Figures taken from [40].

(a) Neutrino flux averaged over all directions

(b) Ratio of muon neutrinos to electron neutrinos

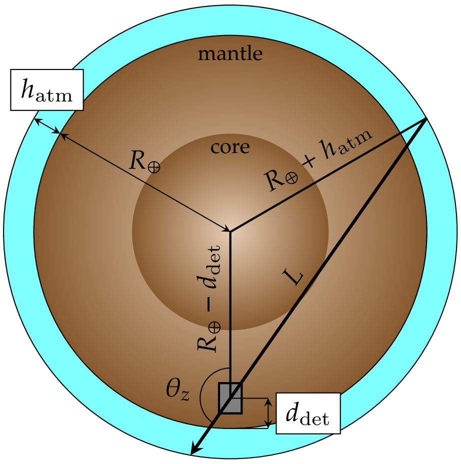

Figure 3.7: Illustration of the geometry of atmospheric oscillation measurements. The detector is shown as the gray box near the bottom and θ_z indicates the observed zenith angle of a neutrino.

Three-flavor oscillations of atmospheric neutrinos

After neutrinos are produced in the atmosphere, they travel through the Earth before being detected at the South Pole. Depending on the zenith angle at which a neutrino is observed, it has traveled a total distance, L, that depends on the radius of the Earth, $R_\oplus$, the depth of the detector, d_{det}, and the production height of the neutrino, h_{atm}, as illustrated in Figure 3.7. From the geometry, L is calculated as

$$L = (d - R_\oplus)\cos(\theta_z) + \sqrt{(R_\oplus + h_{\text{atm}})^2 - (R_\oplus - d)^2 \sin^2(\theta_z)}\,. \quad (3.35)$$

If the height of the atmosphere and detector depth is neglected, Equation 3.35 reduces to

$$L = \begin{cases} 0 & \theta_z \leq \frac{\pi}{2} \\ 2R_\oplus \cos\theta_z & \theta_z > \frac{\pi}{2} \end{cases}\,. \quad (3.36)$$

From the perspective of the IceCube detector, an *up-going event* is one where the zenith angle is $\theta_z > 90°$, that is, $\cos(\theta_z) < 0$. On the other hand, a *down-going* event is one where $\cos(\theta_z) > 0$. In such cases, the neutrino travels only through the atmosphere and the overburden of the detector and therefore matter effects are negligible. The Earth is modeled as a set of concentric shells with different matter den-

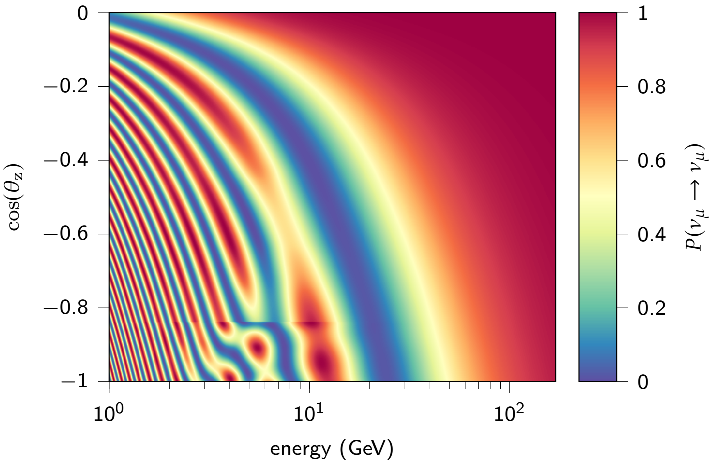

Figure 3.8: Muon-neutrino survival probability calculated at NuFit 4.0[24] global best fit parameters.

sities. The most prominent feature in the density profile is the core region at a depth of close to half an Earth radius, where the density sharply increases from $5\,\mathrm{g\,cm^{-3}}$ to $>10\,\mathrm{g\,cm^{-3}}$[41]. Neutrinos for which $\cos(\theta_z) \lesssim 0.8$ travel partially through the core and generally experience greatly enhanced matter effects. Given the baseline of $\mathcal{O}\left(R_\oplus \sim 10^4\,\mathrm{km}\right)$ and the core energy range of the atmospheric neutrino spectrum between GeV and TeV scales, the mass splitting between mass eigenstates whose oscillation could be observed is $10^{-4}\,\mathrm{eV}^2 < \Delta m^2 < 10^{-1}\,\mathrm{eV}^2$.

Since the primary constituent of the atmospheric neutrino flux are muon (anti-)neutrinos, the most important oscillation channel to be observed is the muon neutrino survival probability ($\overset{(-)}{\nu}_\mu \to \overset{(-)}{\nu}_\mu$) that can be approximated by the vacuum two-flavor oscillation expression

$$\begin{aligned} P(\nu_\mu \to \nu_\mu) &\simeq 1 - 4|U_{\mu 3}|^2(1-|U_{\mu 3}|^2)\sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) \\ &\simeq 1 - \sin^2(2\theta_{23})\sin^2\left(1.27\frac{(\Delta m^2_{31}/\mathrm{eV}^2)(L/\mathrm{km})}{(E/\mathrm{GeV})}\right) , \end{aligned} \tag{3.37}$$

where the second line uses the fact that θ_{13} is small while θ_{23} is close to $\pi/4$. Because most atmospheric neutrinos have energies above the tau production threshold at $\sim 3.5\,\mathrm{GeV}$, it is also possible to directly probe the tau appearance channel $\nu_\mu \to \nu_\tau$ with the approximate transition probability

$$\begin{aligned} P(\nu_\mu \to \nu_\tau) &\simeq 4|U_{\tau 3}|^2|U_{\mu 3}|^2\sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) \\ &\simeq \sin^2(2\theta_{23})\cos^4(\theta_{13})\sin^2\left(1.27\frac{(\Delta m^2_{31}/\mathrm{eV}^2)(L/\mathrm{km})}{(E/\mathrm{GeV})}\right) . \end{aligned} \tag{3.38}$$

Atmospheric neutrino oscillations can therefore directly probe the magnitudes of the $U_{\mu 3}$ and $U_{\tau 3}$ elements of the PMNS matrix.

Matter effects

The distinctive feature of atmospheric neutrino measurements over other terrestrial experimental setups is that the Earth provides a large amount of matter between the point of neutrino production and the detector. As described in Section 3.2, the presence of matter adds an effective potential to the Hamiltonian that is diagonal in the flavor basis. In the case of three-flavor oscillations, this potential affects only electron neutrinos. Ignoring the small mass splitting Δm^2_{21}, all matter effects occur in the transitions between mass eigenstates 1 and 3. The effective mixing angle θ^M_{13} and mass splitting $\Delta m^2_{31,M}$ in matter can then be calculated in the two-flavor picture as described in Section 3.2.2. If the sign of $\Delta m^2_{31}\cos(2\theta_{13})$ is positive, then the resonance condition in Equation 3.26 can be fulfilled for neutrinos given appropriate matter density. If the sign is negative, then it can only be fulfilled for antineutrinos. This means that matter effects can provide sensitivity to the neutrino mass ordering through an enhancement of the mixing angle θ_{13} for neutrinos or antineutrinos. This also applies to parametric enhancement effect described in Section 3.2.3. The CP-violating phase δ_{CP} has only a negligible effect on atmospheric oscillations because it requires sensitivity to Δm^2_{21}. The muon neutrino survival probability for atmospheric neutrinos including matter effects is shown in Figure 3.8 as a function of energy and cosine of the zenith angle. The MSW resonance condition is approximately fulfilled for neutrinos passing through the mantle of the Earth at energies between $6\,\mathrm{GeV}$ and $10\,\mathrm{GeV}$. The distortions for neutrinos passing through the core of the Earth at energies below $10\,\mathrm{GeV}$ are due to parametric resonance.

Current and future atmospheric neutrino experiments

The currently operating dedicated atmospheric neutrino experiments are SuperK, IceCube DeepCore (focus of this work) and ORCA[42]. The ORCA detector is currently under construction in the Mediterranean sea. Once completed, it will consist of 115 detection units with 18 digital optical modules each, instrumenting a volume of approximately $7\,\mathrm{MT}$. The first 6 detection units have been deployed since January 2020 and were able to successfully measure atmospheric neutrino oscillations[43]. Some detectors that are primarily built to measure the neutrino flux from accelerators (see Section 3.3.4), such as MINOS, are also capable of detecting atmospheric neutrinos[44].

3.3.4 Accelerator neutrinos

The $\nu_\mu \rightarrow \nu_\mu$ oscillation channel probed by atmospheric neutrino oscillations is also accessible via long-baseline accelerator neutrino experiments. These experiments use neutrinos that are produced when protons are shot at a stationary target by a particle accelerator. The interactions of the protons with the target produce a hadronic shower consisting mostly of pions that are then focused into a tube where they are allowed to decay into neutrinos and charged leptons. This process is very similar to the production of neutrinos in the atmosphere and leads to a beam that is composed of mostly muon neutrinos and some electron neutrinos. The beam is directed at a detector that is typically located several hundreds of kilometers away from the neutrino production site. An example of such an experimental setup is the MINOS

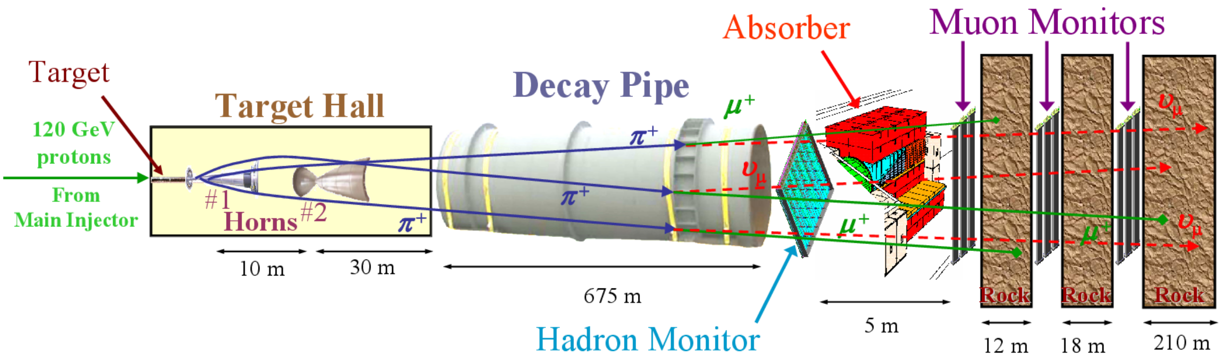

Figure 3.9: Neutrino Main Injector (NuMI) facility producing neutrinos for the MINOS experiment. Figure taken from[48].

experiment[45]. Figure 3.9 shows the setup of the Neutrino Main Injector (NuMI) that produces neutrinos for MINOS at a peak energy of $3\,\mathrm{GeV}$. The neutrino flux is measured at a near and a far detector at baselines of $1\,\mathrm{km}$ and $735\,\mathrm{km}$, respectively, that are both identically constructed magnetized steel scintillator tracking calorimeters. The atmospheric mass splitting and mixing angle can be estimated by comparing the neutrino flux at the near and far detectors, so that most of the inherent systematic uncertainties of the measurement cancel out. Other experiments with a similar setup are T2K[46] and NOνA[47].

3.4 Anomalies in neutrino oscillation measurements

While the three-flavor oscillation picture explains most of the experimental data from reactor, accelerator, atmospheric and solar neutrino experiments fairly well, there are some notable exceptions of anomalous experimental observations. These anomalies are

- Reactor Antineutrino Anomaly: A deficit of $\bar{\nu}_e$ in the neutrino flux of nuclear reactors with respect to the theoretical flux expectation at baselines between $L \sim 10\,\mathrm{m}$ and $L \sim 100\,\mathrm{m}$.
- Gallium anomaly: A deficit of ν_e in the flux of a radioactive $^{51}\mathrm{Cr}$ source. The detector is in direct contact with the source and is of $\mathcal{O}(1\,\mathrm{m})$ in size in each dimension.
- LSND anomaly: An excess of $\bar{\nu}_e$ in the neutrino flux of a proton-on-target (accelerator) source at $L \sim 30\,\mathrm{m}$ and $E \sim 30\,\mathrm{MeV}$.
- MiniBooNE anomaly: Excess of electron-like events in the the flux of neutrinos generated by an accelerator source at $L \sim 500\,\mathrm{m}$ and $E \sim 500\,\mathrm{MeV}$.

This section summarizes each of these measurements and describes how they could be resolved by the introduction of sterile neutrino states.

3.4.1 Reactor neutrino anomaly

The Reactor Antineutrino Anomaly (RAA) was first described in 2011[49] after new corrections were introduced into the theoretical calculations

[49]: Mention et al. (2011), *The Reactor Antineutrino Anomaly*

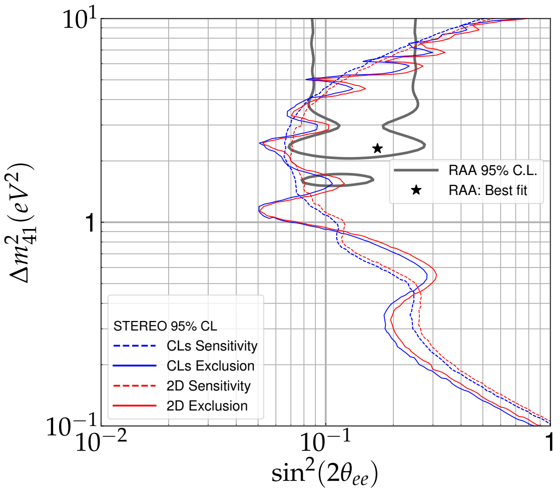

Figure 3.10: Exclusion contours for the sterile mass splitting and mixing angle from the STEREO experiment. The effective mixing angle θ_{ee} is equivalent to θ_{13} in Equation 3.39. The red solid line shows the (FC corrected) exclusion contour, the dashed line shows the expected sensitivity. Blue lines show the sensitivity and exclusion limits using a different statistical method. Gray solid lines show the preferred values of the RAA. Figure taken from [50].

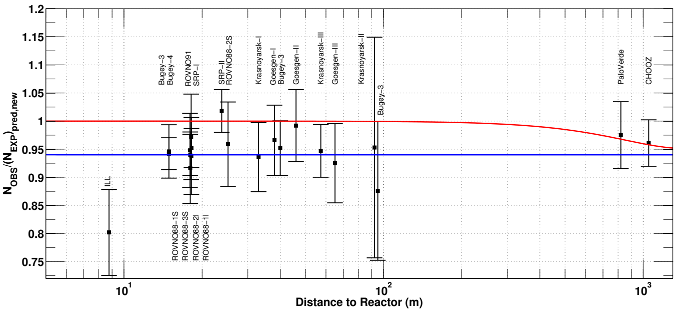

Figure 3.11: Deficit in $\bar{\nu}_e$ flux from commercial nuclear reactors referred to as Reactor Antineutrino Anomaly. Figure taken from [49].

to the predicted neutrino flux from commercial nuclear reactors. The corrections adjusted the predicted flux upwards, leaving a deficit of $\sim 6\%$ in the measured flux with respect to the new prediction. It is suggested in [49] that this deficit could be explained by the introduction of an additional mass eigenstate with a mass splitting, Δm^2_{41}, and mixing angle, θ_{14}. The electron neutrino survival probability can be approximated for distances of $L \lesssim 1\,\mathrm{km}$ as

$$P_{ee} = 1 - \cos^4(\theta_{14}) \sin^2(2\theta_{13}) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) - \sin^2(2\theta_{14}) \sin^2\left(\frac{\Delta m^2_{14} L}{4E}\right). \tag{3.39}$$

The magnitude of the additional mass splitting and mixing angle proposed in [49] is $\left|\Delta m^2_{41}\right| \gg 1\,\mathrm{eV}^2$ and $\sin^2(2\theta_{\mathrm{new}}) = 0.12$, respectively. The oscillations due to this mass eigenstate would be too fast to be resolvable and would only lead to an average deficit that is shown as the blue line in Figure 3.11. The red line shows the flux expectation in the presence of only three active flavors.

Whether or not the RAA truly is an anomaly depends on whether or not one trusts the theoretical flux predictions. The calculations involved require complex nuclear corrections that are not fully understood. The RAA is in fact not the only anomaly in the reactor antineutrino flux. The second anomaly, an excess of $\bar{\nu}_e$ at $\sim 5\,\mathrm{MeV}$, is thought to originate from poorly understood nuclear physics[51]. For this reason, the RAA was put under scrutiny by follow-up experiments that only measure flux differences between different baselines to cancel uncertainties in the absolute flux. Three such experiments are STEREO, PROSPECT, and DANSS. The detectors of STEREO and PROSPECT are segmented between baselines of respectively $9.4\,\mathrm{m}$ to $11.2\,\mathrm{m}$ and $6.7\,\mathrm{m}$ to $9.3\,\mathrm{m}$, while the DANSS detector is positioned on a movable platform that allows for measurements to be taken at baselines of $10.9\,\mathrm{m}$, $11.9\,\mathrm{m}$ and $12.9\,\mathrm{m}$[52]. The results of the measurements of these three experiments show no sign of baseline-dependent flux variations and therefore exclude the oscillation amplitude preferred by the RAA at mass splitting values between $0.1\,\mathrm{eV}^2$ and $10\,\mathrm{eV}^2$. The contours for STEREO[50] are shown in Figure 3.10, those for DANSS and PROSPECT look very similar and can be found in [52]. However, these measurements still allow oscillations due to a mass eigenstate with $\Delta m^2_{41} > 10\,\mathrm{eV}^2$ that cannot be resolved by their detector segmentation. In addition, more recent re-evaluations of the reactor flux model in light of new data from the research reactor at the Kurchatov Institute of the National Research Center (KI) have again increased the predicted flux and thus resolved the anomaly almost entirely[53].

[52]: Licciardi (2021), *Results of STEREO and PROSPECT, and status of sterile neutrino searches*

[50]: Almazán et al. (2023), *STEREO neutrino spectrum of ^{235}U fission rejects sterile neutrino hypothesis*

[53]: Giunti et al. (2022), *Reactor antineutrino anomaly in light of recent flux model refinements*

3.4.2 Gallium anomaly

The Gallium anomaly is a relatively large ($\sim 20\%$) deficit of electron neutrinos observed in measurements of the flux from a $^{51}\mathrm{Cr}$ source at very short distances that was first observed in the GALLEX experiment[55] and that has been confirmed more recently by the SAGE and BEST experiments[54, 56].The BEST experiment consists of an inner and outer detector volume surrounding the radioactive source, as shown in the diagram in Figure 3.12, such that the flux can be compared at two different baselines. The measured ratios in the outer and inner volumes are, respectively, $R_{\mathrm{out}} = 0.77 \pm 0.05$ and $R_{\mathrm{in}} = 0.79 \pm 0.05$. These results do not show any baseline-dependent effects, but they are consistent with the overall deficit observed by SAGE and GALLEX. The weighted average ratio between the expected and predicted flux of all three experiments is $R = 0.80 \pm 0.05$, which brings the significance of the Gallium anomaly to 4σ. In order to explain this ratio with a two-flavor $\nu_e \longrightarrow \nu_s$ oscillation model, the mass splitting would have to be $\Delta m^2 > 1\,\mathrm{eV}^2$ such that the oscillations would average out between the inner and outer volumes of the BEST experiment, and the mixing angle would have to be rather large at $\sin^2(2\theta) \approx 0.4$. This large amount of mixing puts the Gallium anomaly in tension with the reactor anomaly, in particular in light of the more recent flux calculations. The 2σ contours for the effective mixing parameters of both anomalies are shown in Figure 3.13. The reactor anomaly results are shown for different flux models described in [53]. Notably, the contour using the KI and EF flux models is not closed and is therefore compatible with the null hypothesis of no sterile mixing. The gallium anomaly is shown for different cross-section models. The contours are all closed, suggesting that uncertainties in the cross-section calculation cannot resolve the gallium anomaly.

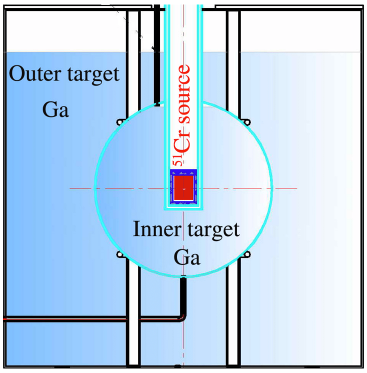

Figure 3.12: Experimental setup of the BEST experiment. The diameters of the inner and outer volumes are $133.5\,\mathrm{cm}$ and $218\,\mathrm{cm}$, respectively. Figure taken from [54].

3.4.3 LSND and MiniBooNE Anomalies

Both the LSND and MiniBooNE anomalies are unexplained excesses in the $\bar{\nu}_e$ flux of accelerator-generated neutrinos. The LSND measurement

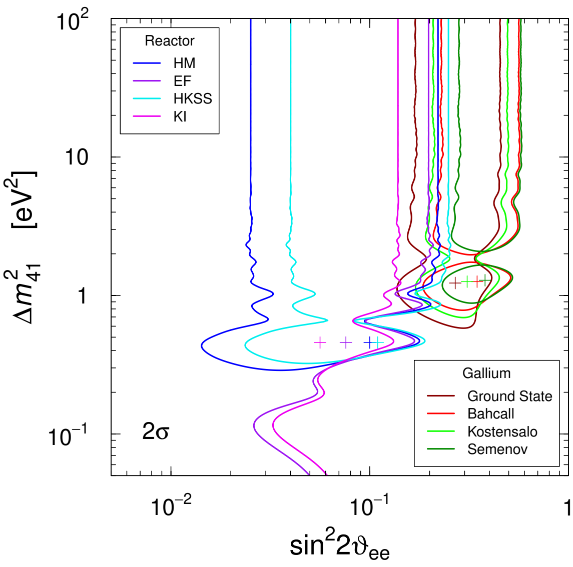

Figure 3.13: Contours delimiting the 2σ allowed regions of the Gallium anomaly and the reactor anomaly in mass splitting and the effective mixing angle $\sin^2(2\theta_{ee}) = 4(1-|U_{e4}|^2)|U_{e4}|^2$. The reactor contours are shown for different flux predictions and the Gallium anomaly contours are shown under different neutrino cross-section models. Figure taken from [53].

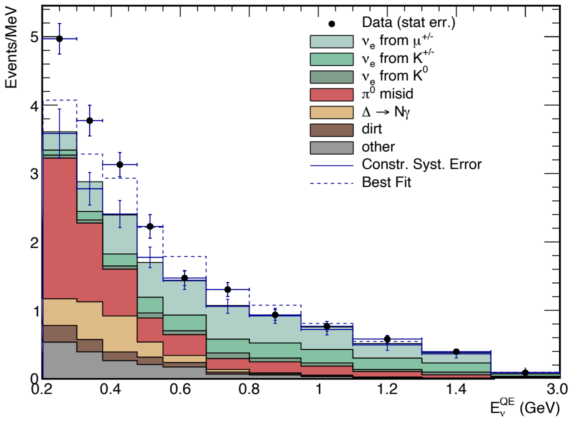

Figure 3.14: MiniBooNE neutrino mode E_ν^{QE} for ν_e CCQE data (points with error bars) and background (histograms). The dashed line shows the best fit of a two-flavor oscillation hypothesis. Figure taken from[58].

found a 3.8σ excess at a baseline of $L \sim 30\,\mathrm{m}$ and neutrino energy $E \sim 30\,\mathrm{MeV}$[57]. The MiniBooNE experiment was set up specifically to probe the parameter space of the LSND result at a very different baseline and energy, where the systematic uncertainties would be different. It found a 4.5σ excess at [58] at a baseline of $L \sim 500\,\mathrm{m}$ and neutrino energy of $E \sim 500\,\mathrm{MeV}$. The preferred regions for the mass splitting and mixing angle of both experiments are mostly compatible, as shown in Figure 3.15. However, the two-flavor oscillation hypothesis does not seem to fit the energy distribution of the observed excess in the lowest energy bins very well, as can be seen in the histogram in Figure 3.14. The mismatch suggests that the simple 3+1 model may not be enough to resolve the anomaly and that more complicated models such as sterile neutrinos with decay[59] or additional "dark sector" interactions[60] might be necessary.

The MiniBooNE excess has been measured on top of a large background dominated by photon-producing neutrino-nucleus interactions that are subject to substantial theoretical uncertainties. A more recent re-analysis of the MiniBooNE data found that a more conservative treatment of all these uncertainties somewhat reduces the magnitude of the observed excess, but leaves the significance of the sterile oscillation hypothesis still at 3.6σ[61]. A followup experiment, MicroBooNE[62], uses a Liquid Argon Time Projection Chamber (LiArTPC) to detect neutrino interactions. In contrast to MiniBooNE, MicroBooNE can separate interactions with electrons in the final state (that is, quasi-elastic charged-current interactions of ν_e) from neutral-current interactions that produce a single photon. Although recent results from MicroBooNE could exclude some of the photon-producing states suspected as the culprit of the anomalous excess[63], they were also unable to confirm the anomaly in the electron neutrino channel[64].

3.4.4 Global picture of oscillation anomalies

While some of the observed anomalies, in particular the LSND and Gallium anomalies, could be confirmed to a high significance by later experiments, it is challenging to reconcile them all in a cohesive global neutrino oscillation model. The simplest model that can be invoked in order to explain the anomalous appearances and disappearances of electron neutrinos is the addition of a fourth neutrinos mass eigenstate, ν_4, and an unobservable sterile flavor eigenstate, ν_s, and to allow the active neutrino flavors to oscillate into and out of the sterile state. Theoretically, such an addition is possible by adding a Majorana mass term as described in Section 2.1.4. The flavor and mass eigenstate would then mix via an extended ”3+1” mixing matrix

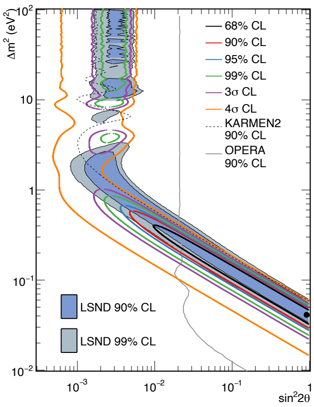

Figure 3.15: MiniBooNE allowed regions for a combined neutrino and antineutrino dataset within a two-neutrino oscillation model. The shaded areas show the 90% and 99% C.L. LSND $\bar{\nu}_\mu \to \bar{\nu}_e$ allowed regions. Figure taken from [58].

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \\ \nu_s \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & U_{\mu 4} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & U_{\tau 4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \end{pmatrix} \tag{3.40}$$

and the fourth mass eigenstate would be much heavier than the first three states with a mass splitting of $\Delta m^2_{41} \gg \Delta m^2_{31}$. Because the oscillations due to this additional mass eigenstate would be much faster than those of the active flavors, the transition probability between flavors α and β can be described in the two-flavor approximation

$$P_{\alpha\beta} \simeq \delta_{\alpha\beta} - 4\left|U_{\alpha\beta}\right|^2 (\delta_{\alpha\beta} - \left|U_{\alpha\beta}\right|^2) \sin^2\left(\frac{\Delta m^2_{41} L}{4E}\right) . \tag{3.41}$$

It is convenient to define the effective mixing angles for each oscillation channel that is probed by different experiments, as listed in Table 3.2. Since there are three measurable mixing angles that depend on only two matrix elements, the model can be over-constrained by measurements of $\theta_{\mu 3}$, θ_{ee} and $\theta_{\mu\mu}$. Attempts to perform a global fit using all three relevant oscillation channels lead to a strong (4.7σ) tension between the

channel	mixing angle definition	experiments
$\nu_\mu \to \nu_e$	$\sin^2(2\theta_{\mu e}) \equiv 4\left\|U_{\mu 4}\right\|^2 \|U_{e4}\|^2$	LSND, MiniBooNE, OPERA, ...
$\nu_e \to \nu_e$	$\sin^2(2\theta_{ee}) \equiv 4\|U_{e4}\|^2(1 - \|U_{e4}\|^2)$	Reactor, solar, Gallium, ...
$\nu_\mu \to \nu_\mu$	$\sin^2(2\theta_{\mu\mu}) \equiv 4\left\|U_{\mu 4}\right\|^2(1 - \left\|U_{\mu 4}\right\|^2)$	MiniBooNE, MINOS, IceCube, ...

Table 3.2: Definitions of effective mixing angles in the ”3+1” oscillation model for oscillation channels relevant to neutrino oscillation anomalies.

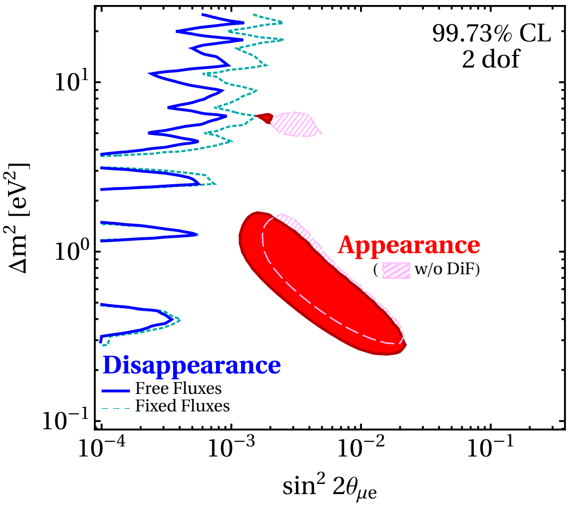

Figure 3.16: Contours delimiting the 99.73% C.L. (3σ) regions in mass splitting and the effective mixing angle $\sin^2(2\theta_{\nu e}) \equiv 4|U_{\mu 4}|^2|U_{e4}|^2$ from the appearance ($\nu_\mu \to \nu_e$) channel and the disappearance ($\nu_e \to \nu_e$) channel. Figure taken from [65].

appearance channel ($\nu_\mu \to \nu_e$) and the combined *disappearance* channels ($\nu_e \to \nu_e$ and $\nu_\mu \to \nu_\mu$)[65]. This can be seen when the 3σ contours in mass splitting and the mixing angle $\theta_{\mu e}$ of the appearance and disappearance datasets are plotted together, as shown in Figure 3.16. The appearance results are combined fits of the LSND and MiniBooNE ν_e excess. The disappearance results combine ν_e results from reactor, solar, and gallium measurements and ν_μ measurements from accelerator and atmospheric sources. Even though the Gallium anomaly requires a sizable magnitude for $|U_{e4}|^2$, the fact that there are no anomalies in the $\nu_\mu \to \nu_\mu$ channel constrains the product $|U_{\mu 4}|^2|U_{e4}|^2$ to a small value. In order to attribute the observed phenomena to sterile neutrino oscillations, the model would require additional degrees of freedom. One possibility is to give the heavy neutrino state the ability to decay, which can reduce the tension between appearance and disappearance results to the level of 3.2σ[66]. A detailed discussion on global fits of the neutrino anomalies can be found in [66].

[66]: Diaz et al. (2020), *Where are we with light sterile neutrinos?*

3.4.5 Other constraints on sterile neutrinos

While the neutrino oscillation anomalies described in Section 3.4.4 hint at the possibility of an additional neutrino mass eigenstate with a mass splitting o $\Delta m^2_{41} \gtrsim \mathcal{O}(1\,\mathrm{eV}^2)$, there are other constraints on the masses and number of neutrino flavor states that need to be considered.

Unitarity constraints on the PMNS matrix

If the neutrino mixing matrix is extended by adding mass eigenstates, the 3×3 block matrix for the three active neutrino flavors is no longer unitary. Therefore, such extensions can be constrained in a generic way by testing the unitarity of the PMNS matrix. Unitarity of the mixing

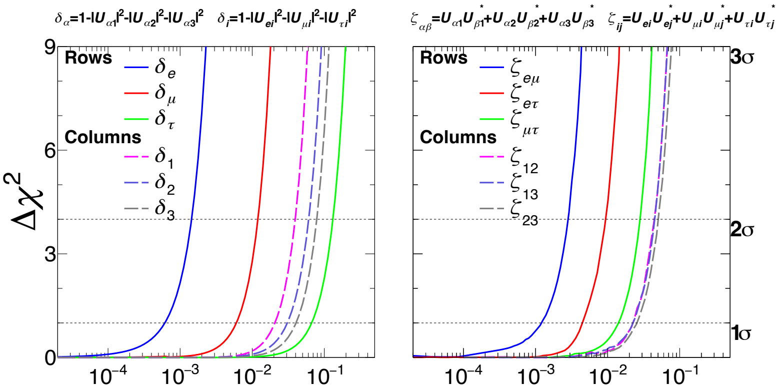

Figure 3.17: Test statistic for the magnitudes of unitary condition violations calculated in [67]. The shown quantities correspond to those in Equations 3.42 to 3.45.

matrix imposes four conditions on the elements of the matrix:

$$|U^{3\nu}_{\alpha1}|^2 + |U^{3\nu}_{\alpha2}|^2 + |U^{3\nu}_{\alpha3}|^2 - 1 = \delta_\alpha = 0, \quad \alpha = e, \mu, \tau, \tag{3.42}$$

$$|U^{3\nu}_{ei}|^2 + |U^{3\nu}_{\mu i}|^2 + |U^{3\nu}_{\tau i}|^2 - 1 = \delta_i = 0, \quad i = 1, 2, 3, \tag{3.43}$$

$$U^{3\nu}_{\alpha1}U^{3\nu,*}_{\beta1} + U^{3\nu}_{\alpha2}U^{3\nu,*}_{\beta2} + U^{3\nu}_{\alpha3}U^{3\nu,*}_{\beta3} = \zeta_{\alpha\beta} = 0, \quad \alpha, \beta = e, \mu, \tau, \quad \alpha \neq \beta, \tag{3.44}$$

$$U^{3\nu}_{ei}U^{3\nu,*}_{ej} + U^{3\nu}_{\mu i}U^{3\nu,*}_{\mu j} + U^{3\nu}_{\tau i}U^{3\nu,*}_{\tau j} = \zeta_{ij} = 0, \quad i, j = 1, 2, 3, \quad i \neq j. \tag{3.45}$$

The first two conditions in Equations 3.42 and 3.43 are the normalization of the rows and columns. The third and fourth conditions in Equations 3.44 and 3.45 are the triangle closure conditions. A global analysis constraining the degree to which each of these conditions can be violated given the experimental data can be found in [67]. The analysis in [67] combines non-anomalous measurements from reactor, solar, and long-baseline accelerator experiments as well as several published sterile neutrino searches. The reactor data used is exclusively comprised of those experiments that have a near and far detector setup and constraints are derived using only ratios between fluxes at different baselines. The accelerator data consists of non-anomalous measurements from NOνA and T2K. The χ^2 test statistic for the quantities in Equations 3.42 to 3.45 calculated in [67] are shown in Figure 3.17. No evidence for non-unitarity is found. The normalization of the $|U_{ei}|$ row is constrained to $\mathcal{O}(10^{-3})$, which puts the result in strong tension with the sterile neutrino interpretation of the Gallium anomaly.

[67]: Hu et al. (2021), *Global oscillation data analysis on the* 3ν *mixing without unitarity*

Cosmological constraints

Standard model cosmology places constraints on the number of relativistic neutrino species, N_{eff}, and the sum of neutrino masses, $\sum m_\nu$. Both of these parameters can be estimated from observations of the cosmic microwave background (CMB), the large-scale structures (LSS), and the abundance of light elements from Big Bang Nucleosynthesis (BBN). These estimates primarily rely on the contribution of neutrinos to the total energy and matter density of the universe, which affect its expansion history. The expansion history of the universe, in turn, affects the size of anisotropies in the CMB. A summary of the effects of neutrinos on cosmological observations can be found in [68]. Re-

[68]: Lesgourgues et al. (2014), *Neutrino cosmology and Planck*

cent results from the Planck Collaboration for these parameters are $N_{\text{eff}} = 2.99 \pm 0.17$ and $\sum m_\nu < 0.1\,\text{eV}$[17]. These results are compatible with the absence of additional heavy neutrino states and are in severe tension with hypotheses involving a mass eigenstate with $\Delta m^2_{41} \sim \mathcal{O}(1\,\text{eV})$. This tension can only be alleviated by the introduction of further extensions to the standard model.

[17]: Planck Collaboration et al. (2020), *Planck 2018 results - VI. Cosmological parameters*

3.5 Open Questions in Neutrino Oscillation Physics

Neutrino oscillations are the only phenomenon beyond the Standard Model that has been confirmed experimentally to a very high statistical significance. As such, they are of great interest to the study of the foundations of physics and cosmology. Precision measurements of neutrino oscillations can potentially lead to important exclusions for theories that make predictions about their masses and mixing parameters. Current observations of the mixing angles θ_{12}, θ_{13}, and the absolute mass differences $|\Delta m^2_{\ell 1}|$ are largely consistent with one another, and global studies can constrain these parameters with relative errors of $\sim 10\%$.[31].

However, important open questions remain about the values of θ_{23}, δ_{CP} and the mass ordering. It is unknown whether θ_{23} is in the upper octant ($\theta_{23} > \frac{\pi}{2}$) or in the lower octant ($\theta_{23} < \frac{\pi}{2}$). This ambiguity affects the sensitivity of atmospheric neutrino experiments to the neutrino mass ordering (NMO), which is also important for the search for double beta decay ($0\nu\beta\beta$)[69]. The NMO is also important for the sensitivity of accelerator experiments to the value of the CP-violating phase δ_{CP}[70]. The existence of such decays would be strong evidence for the existence of Majorana mass terms. Finally, neutrino masses, the process of mass generation and the amount of CP violation in their mixing have consequences for the expansion history and matter distribution of the universe[68]. Given the interdependence between the neutrino oscillation parameters, it is clear that no single experiment is equipped to answer these questions on its own. A joint effort of several experiments measuring oscillations in different flavor transition channels at various energies, baselines and matter densities is needed to uncover the origin and nature of neutrino masses. The measurement of atmospheric neutrino oscillations presented in this work is an important piece in this global effort.

[69]: Avignone et al. (2008), *Double Beta Decay, Majorana Neutrinos, and Neutrino Mass*

[70]: Blennow et al. (2014), *Quantifying the sensitivity of oscillation experiments to the neutrino mass ordering*

4 Neutrinos in IceCube and DeepCore

4.1 The IceCube in-ice Array and DeepCore 41
4.1.1 The Antarctic Ice 41
4.1.2 In-Ice Array 42
4.1.3 IceTop 44
4.1.4 Digital Optical Modules 44
4.2 Propagation of particles in ice 45
4.2.1 Cherenkov Effect 45
4.2.2 Muons 46
4.2.3 Electromagnetic Showers 46
4.2.4 Hadronic Showers 47
4.3 Particle Signatures in IceCube 48
4.3.1 Neutrinos 48
4.3.2 Atmospheric muons . . . 49

The IceCube Neutrino Observatory is a gigaton-scale Cherenkov detector located at the geographic South Pole in close proximity to the Amundsen-Scott South Pole Station. Constructed over the course of several deployment seasons between 2006 and 2011, it instruments approximately one cubic-kilometer of Antarctic glacier with optical sensors that can detect faint flashes of light that are produced when charged particles travel through the ice, such as those produced by neutrino interactions. The detector serves as both a telescope to study the astrophysical origin of cosmic neutrinos and an instrument to measure their fundamental properties.

This chapter describes the instrumentation and layout of the IceCube detector, the interactions that particles undergo when they interact with the ice, and finally the signals that these interactions produce in the detector.

4.1 The IceCube in-ice Array and DeepCore

The IceCube Neutrino Observatory consists of the so-called *in-ice* array, optimized for astrophysical neutrino observations, the *DeepCore* array, used primarily for the observation of atmospheric neutrinos, and the *IceTop* surface array that can be used to study air showers from cosmic rays.

4.1.1 The Antarctic Ice

The detection medium of the IceCube detector is the Antarctic glacier that has formed from layers of snow being deposited top of each other over the course the past $\sim 100\,000$ years[71]. The weight of the upper layers compresses the lower layers into a dense, crystalline structure. As a result, the optical properties of the ice change mostly in the direction perpendicular to the layers, forming a geological record of the atmospheric conditions of the Earth. The transmission of light

[71]: Price et al. (2000), *Age vs depth of glacial ice at South Pole*

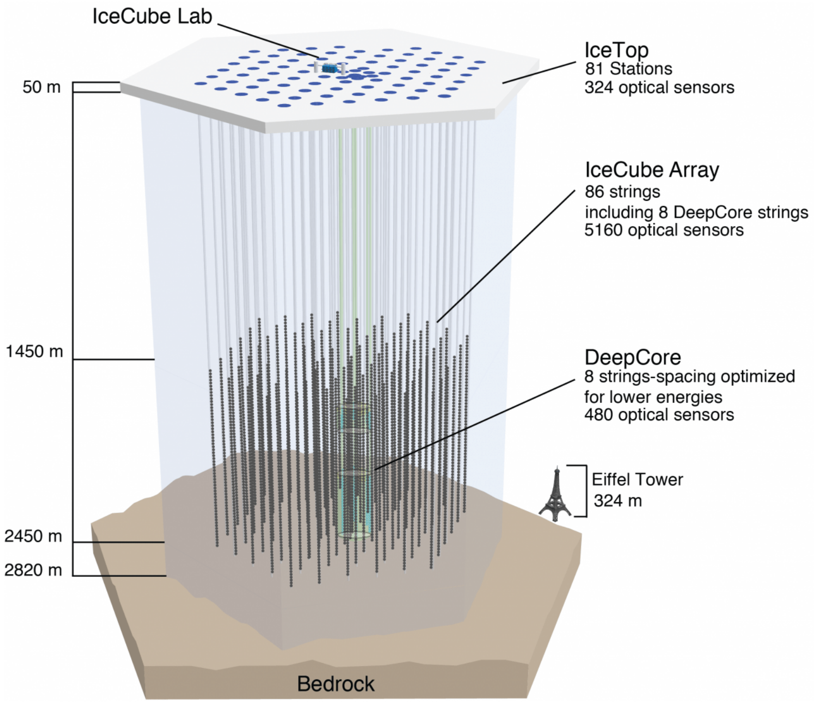
IceCube Lab
IceTop
81 Stations
324 optical sensors
50 m
IceCube Array
86 strings
including 8 DeepCore strings
5160 optical sensors
1450 m
DeepCore
8 strings-spacing optimized
for lower energies
480 optical sensors
Eiffel Tower
324 m
2450 m
2820 m
Bedrock

Figure 4.1: An overview of the IceCube detector

through the ice is primarily characterized by the scattering and absorption length. Within the volume of IceCube, scattering lengths vary between 20 m and 100 m, while absorption lengths range from 100 m to 400 m. Both quantities are highly correlated, such that the absorption length is approximately four times as large as the scattering length[72]. This stratigraphy was traced at millimeter resolution using a laser dust logger deployed down seven IceCube drill holes as described by [73]. The most notable feature of the stratigraphy is the *dust layer* at depths between 2000 m and 2100 m as shown in Figure 4.2. The optical properties of the ice within the dust layer are particularly poor. The ice below the dust layer where the DeepCore fiducial volume is located has the best optical properties of the entire IceCube volume.

[72]: Abbasi et al. (2022), *In-situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory*

4.1.2 In-Ice Array

The 5160 Digital Optical Modules (DOMs) that make up the IceCube in-ice array are distributed over 86 strings. Of these, 78 are arranged on a hexagonal grid spanning an area of approximately one square-kilometer with a horizontal spacing of ~150 m with respect to their closest neighboring strings[74]. Each of these strings holds 60 DOMs at depths between 1450 m and 2450 m with a 17 m vertical spacing. The volume and instrumentation density of this array is optimized for astrophysical neutrinos that are found at energies above 1 TeV[74]. The electric signals measured in each DOM are digitized and sent to the *IceCube lab (ICL)*, where the signal is processed, compressed, and sent North via satellite for offline processing. Figure 4.1 gives an overview of the detector including the IceCube Lab, the ice surface and the bedrock, and a schematic of the layout of the strings is shown in Figure 4.2.

[74]: Aartsen et al. (2017), *The IceCube Neutrino Observatory: instrumentation and online systems*

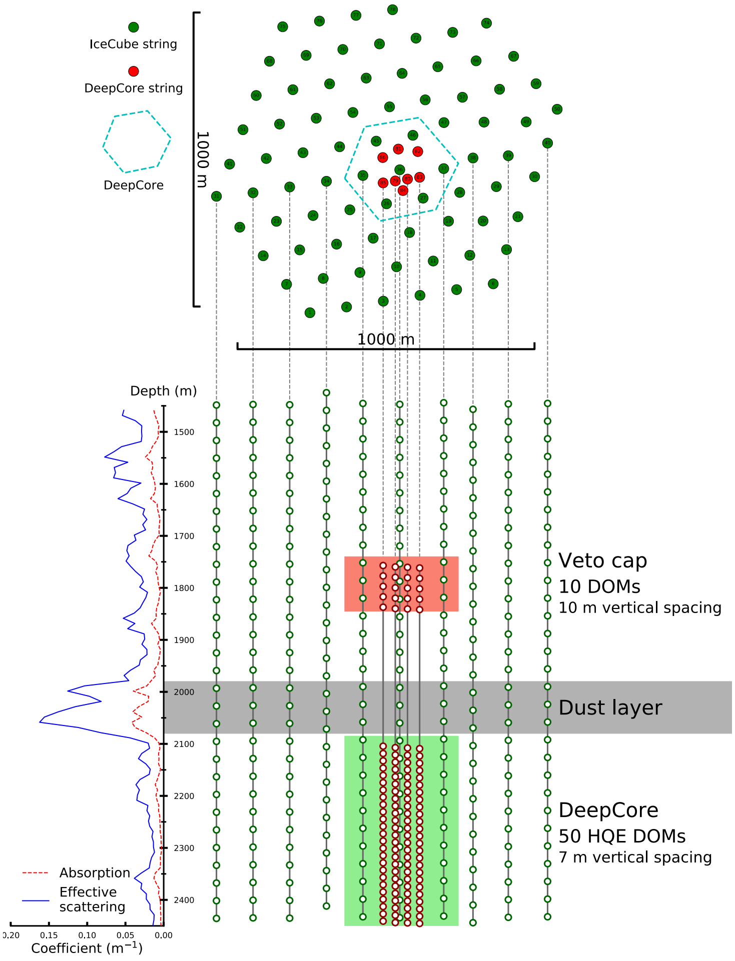
IceCube string
DeepCore string
DeepCore
1000 m
1000 m
Depth (m)
1500
1600
1700
1800
1900
2000
2100
2200
2300
2400
Veto cap
10 DOMs
10 m vertical spacing
Dust layer
DeepCore
50 HQE DOMs
7 m vertical spacing
Absorption
Effective scattering
0.20
0.15
0.10
0.05
0.00
Coefficient (m^{-1})

Figure 4.2: Schematic view of the IceCube detector as seen from the top (upper panel) and the side(lower panel). The DeepCore fiducial volume is indicated by the hexagon in the upper panel and the green shaded area in the bottom panel. The side-band on the lower panel shows the scattering and absorption coefficients as a function of depth.

DeepCore

The remaining 8 strings that are not part of the hexagonal grid are located near the center of the IceCube detector and form the *DeepCore* sub-array[75]. The DOMs on the DeepCore strings have a higher quantum efficiency than those in the rest of the detector and are placed more closely together to lower the minimum energy threshold for neutrino observations to a few GeV. Of the 60 DOMs on each DeepCore string, 50 are placed at depths between 2100 m and 2500 m, where the ice is the most transparent compared to the rest of the IceCube's volume (see also the side band in the bottom panel of Figure 4.2). Together with 7 strings from the in-ice array, the DeepCore strings instrument the DeepCore 20 MT *fiducial volume* as shown in the upper panel of Figure 4.2. The remaining 10 DOMs are located at depths between 1750 m and 1850 m and are used as a veto cap to reject atmospheric muons entering the detector directly from above. In addition, the larger hexagonal IceCube array also serves as a veto for observations inside the DeepCore fiducial volume.

[75]: Abbasi et al. (2012), *The design and performance of IceCube DeepCore*

4.1.3 IceTop

In addition to the in-ice array, IceCube also contains a surface array called *IceTop*, consisting of 81 stations spread across an area of 1 km^2 that is used to detect muons from air showers. It is typically used as a veto against atmospheric muons, but also functions as a detector in its own right measuring the spectrum and composition of cosmic particles. However, it is not relevant to the measurement presented in this thesis.

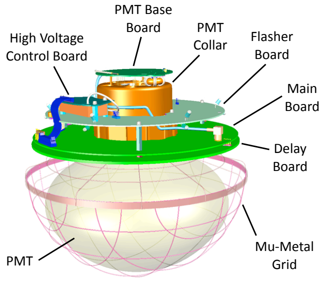

Figure 4.3: Schematic of a DOM, taken from [74].

4.1.4 Digital Optical Modules

The Cherenkov radiation produced by charged particles in the ice is detected and digitized by Digital Optical Modules (DOMs). Each module consists of a photo-multiplier tube (PMT)[76] and electronics housed in a transparent, spherical glass vessel that can withstand the enormous pressure below a water column of 2.5 km[74]. They are each held in place by a harness attached to chains that allows the string cable to pass beside the DOM as shown in Figure 4.4.

[76]: Abbasi et al. (2010), *Calibration and characterization of the IceCube photomultiplier tube*

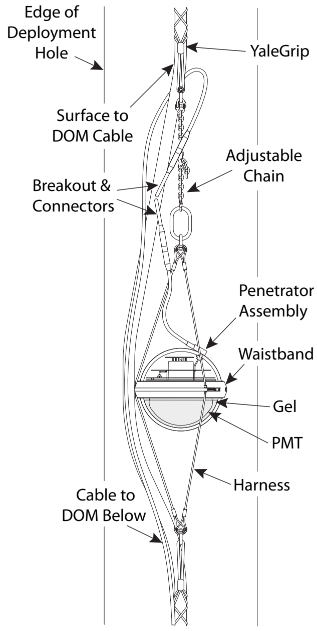

Figure 4.4: Schematic of the cable assembly of a DOM. Figure taken from [74].

The PMTs have a diameter of 10 inches and are sensitive to photons with wavelengths between 300 nm and 650 nm, with a maximum quantum efficiency of about 25% at 390 nm. Inside the DeepCore array, the peak efficiency reaches 34%. They are shielded from external magnetic fields with a mu-metal grid as shown in the schematic in Figure 4.3. The voltage at the PMT is measured and digitized by the on-board electronics[77] of the DOM in two separate readouts that are activated when the measured voltage rises above the equivalent of 0.25 photo-electrons (PE). The first readout is the *fast Analog-Digital Converter (fADC)* and measures the waveform continuously at a rate of 40 MHz. The second readout, the *Analog Transient Waveform Digitizer (ATWD)*, records the PMT voltage at a rate of 300 MHz in three channels with different gain levels to ensure that a large range of voltages can be recorded without saturation of the output. The readout frequency of the ATWD is too high to be directly digitized and sent to the surface. Instead, the ATWD voltage readout is buffered in 128 analog capacitors, corresponding to a readout time of ~420 ns. The buffered voltages are only digitized when at least one of the nearest or next-to-nearest DOMs on the same string also measures a signal within a 1 μs time window, which is referred to as the *hard local coincidence (HLC)* condition. The recorded waveforms are sent to the ICL on the surface, where they are compressed by applying the *wavedeform* algorithm[78]. The output of this algorithm are reconstructed times and charges of single photo-electrons, which are taken as input by all further data processing steps described in section 5.2.

[77]: Abbasi et al. (2009), *The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping*

[78]: Aartsen et al. (2020), *In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes*

The DOMs also contain a flasher board with 12 LEDs that can be used to emit pulsed light for the purpose of *in-situ* detector calibration during special *flasher runs* of the detector. During such runs, the charge and time distributions of the observed pulses in the DOMs in response to the LED flashes are measured. Since the light is emitted at known locations and at known times, the measured distributions allow inference on the absorption and scattering properties of the ice. Because the total amplitude of the emitted light is less well known, this calibration method is less well suited for calibrating the total optical efficiency of the DOMs. Instead, this property of the detector is calibrated more accurately from measurements of minimally-ionizing atmospheric muons, for which the energy loss is well known[79].

[79]: Kulacz (2019), *In Situ Measurement of the IceCube DOM Efficiency Factor Using Atmospheric Minimum Ionizing Muons*

4.2 Propagation of particles in ice

Neutrinos interacting with the ice mostly interact via Deep Inelastic Scattering (DIS), creating muons, electromagnetic showers, and hadronic showers, depending on the flavor of the neutrino and interaction type. The secondary particles produced by those interactions travel through the ice at highly relativistic velocities and lose energy primarily through ionization, bremsstrahlung, pair production and photo-nuclear interactions. The fraction that each of these mechanisms contributes to the total energy loss of the particle depends on the type of particle and its energy. When they are electrically charged, they also give off Cherenkov radiation that is then measured by IceCube.

4.2.1 Cherenkov Effect

The IceCube Neutrino Observatory relies entirely on the Cherenkov effect[80] to detect particle interactions. It is created by any electrically charged particle travelling through a transparent medium with velocities faster than the speed of light in that medium, c/n, where n is the refractive index of the medium. This produces a cone of light moving with the particle similar to a super-sonic shock that is generated by an object travelling through a gas at a velocity above the speed of sound. The effect can be most easily understood according to Huygen's principle as a superposition of spherical light emissions that are produced every time that the particle displaces the charges in the dielectric medium in its closest vicinity, as shown in Figure 4.5. When the particle is over-taking its own light emissions, they overlap coherently and form a conical light front as illustrated in the bottom panel of Figure 4.5.

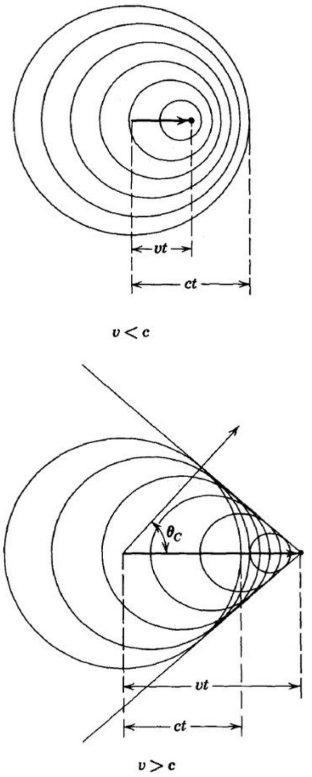

Figure 4.5: An electrically charged particle emitting light while travelling below (upper panel) and above (lower panel) the speed of light in a medium. Image taken from [81].

When the velocity of the particle is very close to the speed of light, as is the case for all (known) particles observable by IceCube, the opening angle of the cone only depends on the refractive index of the medium with

$$\cos(\vartheta_c) = \frac{1}{\beta n} , \tag{4.1}$$

where n is the index of refraction and ϑ_c is the Cherenkov opening angle.

The frequency spectrum of the Cherenkov emissions of highly relativistic particles depends only on the charge of the particle, q, and the (wavelength-dependent) index of refraction, $n(\omega)$, and permeability, $\mu(\omega)$, of the medium. The emitted energy per unit of distance and frequency is given by the Frank-Tamm-Equation[82][83], which simplifies in the case of $v \approx c$ to

$$\frac{\mathrm{d}E}{\mathrm{d}x\mathrm{d}\omega} = \frac{q^2}{4\pi}\mu(\omega)\omega\left(1 - \frac{1}{n^2(\omega)}\right) . \tag{4.2}$$

[82]: Frank et al. (1937), *Coherent visible radiation of fast electrons passing through matter*

[83]: Tamm (1991), *Radiation Emitted by Uniformly Moving Electrons*

The equation shows that the intensity of the Cherenkov emission generally increases with frequency, and indeed the strongest emissions are in the ultraviolet part of the spectrum.

4.2.2 Muons

At energies below 100 GeV, the dominant energy loss for muons is via ionization, and has only a weak dependence on energy. Because the ionization loss is continuous and nearly constant, muons at these energies produce long, track-like signatures in the detector. Above 100 GeV, the losses due to bremsstrahlung, pair production and photo-nuclear interactinos rise quickly in their amplitude and become dominant over ionization at ~1 TeV. The total average energy loss per unit distance, $\langle \mathrm{d}E/\mathrm{d}x \rangle$, can be approximated combining all radiative energy losses (i.e. all losses except for ionization) into one component and adding it to the ionization loss such that

$$\left\langle -\frac{\mathrm{d}E}{\mathrm{d}x} \right\rangle = a_I(E) + b_R(E)E\,, \tag{4.3}$$

where $a_I(E)$ and $b_R(E)$ are slowly changing functions describing the ionization loss and the radiative losses, respectively[84]. For the energy ranges relevant for this work, the energy dependence of $a_I(E)$ and $b_R(E)$ is weak enough such that they can be approximated as constant numbers with $a_I(E) \approx 2\ \mathrm{MeV/cm}$ and $b_R(E) \approx 3.4 \times 10^{-6}\ \mathrm{cm}^{-1}$[84]. In this approximation, we can calculate the average length of a muon track, $\langle L \rangle$, as a function of energy with

$$\langle L \rangle = \frac{1}{b_R} \log\left(\frac{b_R}{a_I} E + 1 \right)\,, \tag{4.4}$$

which gives an average travel distance of 50 m at 10 GeV and 460 m at 100 GeV.

4.2.3 Electromagnetic Showers

In contrast to muons, electrons and positrons lose their energy very quickly by emitting highly energetic photons due to bremsstrahlung. The energy of the emitted photons is high enough that they spontaneously produce pairs of electrons and positrons. This process is repeated until the electrons and positrons reach their critical energy, which is approximately 78 MeV in ice[21]. Below the critical energy, ionization takes over as the predominant mechanism of energy loss, which produces no new shower particles. Another important quantity is the *radiation length*, X_0, defined as the distance at which the energy of an electron is reduced to $1/e$ of its initial energy via bremsstrahlung, which is 36 cm in ice[21]. The radiation length also determines the scale of the longitudinal development of the shower. Expressing distances in units of radiation length as $t = x/X_0$, the shower intensity follows roughly a gamma distribution parametrized as

$$\frac{\mathrm{d}E}{\mathrm{d}t} = E_0 b \frac{(bt)^{a-1} e^{-bt}}{\Gamma(a)}\,, \tag{4.5}$$

[21]: Workman et al. (2022), *Review of Particle Physics*

where the parameters a and b need to be fit empirically[21]. Their values for electrons, positrons and photons interacting in ice have been determined from GEANT4[85] shower simulations in[86] to be

[85]: Agostinelli et al. (2003), *Geant4—a simulation toolkit*

[86]: Rädel et al. (2013), *Calculation of the Cherenkov light yield from electromagnetic cascades in ice with Geant4*

$$a \approx 2.01 + 1.46 \log_{10}(E_0/\mathrm{GeV}),\ b \approx 0.63 \qquad (e^+, e^-), \tag{4.6a}$$

$$a \approx 2.84 + 1.34 \log_{10}(E_0/\mathrm{GeV}),\ b \approx 0.65 \qquad (\gamma). \tag{4.6b}$$

The shower reaches its maximum intensity at a distance of

$$t_{\max} = \frac{a-1}{b}\,, \tag{4.7}$$

which corresponds to a logarithmic growth of the size of the cascade according to Equations 4.6a and 4.6b. The electrically charged components of the electromagnetic shower produce Cherenkov light, where the emissions peak at the Cherenkov angle since the secondary particles are emitted very close to the forward direction as shown in Figure 4.6.

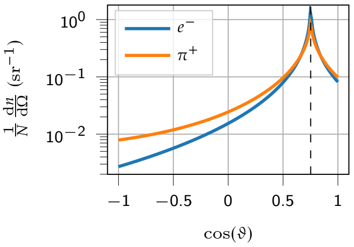

Figure 4.6: Angular profile of the Cherenkov emission of an electromagnetic cascade (e^-) and a hadronic cascade (π^+) using the parametrization from [86].

4.2.4 Hadronic Showers

As discussed in section 2.3.3, neutrino interactions above 10 GeV happen almost exclusively via Deep-Inelastic Scattering (DIS). These interactions always produce a hadronic cascade in addition to any leptons in the final state. Hadronic cascades are also the only visible part of the final state of neutral-current interactions. Hadrons (mostly Pions) that are produced in neutrino-nucleon interactions interact strongly with the surrounding ice to create secondary particles and then decay to form additional photons and leptons. Charged secondary particles produce Cherenkov radiation, while neutral secondary particles are invisible to the detector. Because part of the energy deposited in a hadronic shower is not measurable, the inherent uncertainty on the true energy of the primary particle that initiated the interaction is larger. The average visible electromagnetic fraction of a hadronic shower can be parametrized[86] as a function of the initial energy with

$$F(E_0) = 1 - (1 - f_0)\left(\frac{E_0}{E_s}\right)^{-m} \tag{4.8}$$

with a variance of

$$\sigma_F(E_0) = \sigma_0 \log(E_0)^{-\gamma}\,. \tag{4.9}$$

The parameters f_0, E_s, m, σ_0, and γ are fit to GEANT4 simulation results for hadronic showers induced by different primary particles in[86]. The Cherenkov emissions from the charged components of the shower still peak around the Cherenkov angle as they do for electromagnetic showers, but the emission profile is more smeared out due to the larger variations in particle kinematics as can be seen in figure 4.6 for the example of a shower induced by a pion.

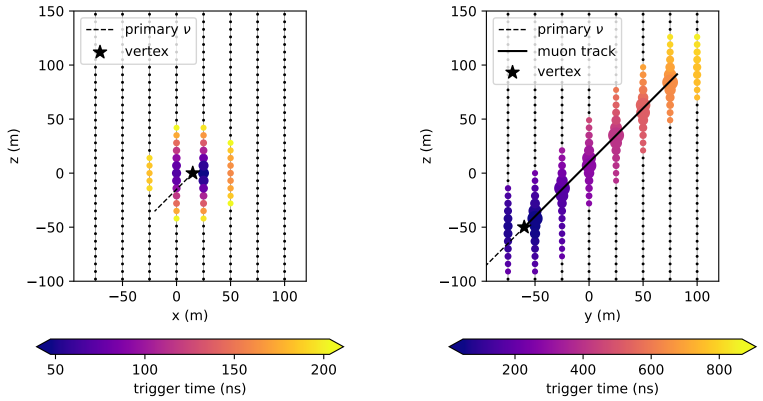

Figure 4.7: An idealized cascade event (left) and starting track event (right) seen from the side. Each DOM that has received light is highlighted with a colored bubble, where the size is proportional the total charge seen by the DOM and the color indicates the time of the hits relative to the time at which the neutrino interaction happened.

4.3 Particle Signatures in IceCube

All particle signatures in IceCube can be approximated as being combinations of compact *cascades* that are produced by hadronic and electromagnetic showers (see section 4.2.4 and 4.2.3), and elongated *tracks* that are only produced by muons travelling through the detector.

4.3.1 Neutrinos

At energies above 10 GeV, nearly all neutrino interactions are due to Deep Inelastic Scattering (DIS) and therefore always produce at least a hadronic cascade originating at the interaction vertex. In Neutral-Current (NC) interactions, this hadronic cascade is the only visible part of the interaction. Charged-current (CC) interactions also produce a lepton of the same flavor as the primary neutrino. For electron-neutrinos, this creates an electromagnetic (EM) shower that originates at the interaction vertex. While the direction of the EM shower and the hadronic shower might not be exactly the same, they are in practice not distinguishable by the detector and can therefore be approximated as a single cascade-like signature. Since the directions of the particles that make up a shower are randomly distributed with a strong bias towards the direction of the primary neutrino, the angular profile of the light emission of a cascade follows a smeared Cherenkov emission profile. At distances of several scattering lengths ($L_s \approx 25\ \mathrm{m}$), this emission profile averages out and the cascade can be approximated as a single point emitting light uniformly in all directions. The left panel of figure 4.7 shows the detector response of such an idealized, perfectly symmetric cascade event. As described in Section 2.3.2, the only distinction between the interactions of neutrinos and antineutrinos is a difference in their cross-section as a function of the inelasticity that is due to their spin configuration. It is therefore impossible to distinguish between the two signatures on an event-by-event basis, even though statistical inferences about the relative distributions of neutrinos and antineutrinos can be made given a sufficiently large population[87].

For muon-neutrinos, a muon is produced in the CC interaction which can then travel a significant distance through the ice beyond the extent of the initial hadronic shower, creating a track-like signature that sticks

[87]: Halve (2018), *Measurement of the Atmospheric Neutrino to Antineutrino Ratio above 100GeV with IceCube*

interaction	secondary particles	signature
ν_e CC	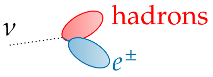	cascade
ν_μ CC	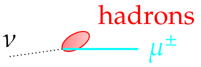	cascade + track
ν_τ CC (83% BR)	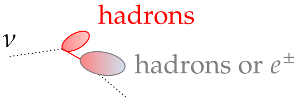	cascade
ν_τ CC (17% BR)	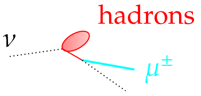	cascade + track
ν NC	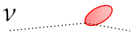	cascade

Table 4.1: Secondary particles and signatures produced by each type of neutrino interaction.

out of the cascade. An idealized version of such an event is shown in the right panel of figure 4.7.

Charged-current interactions of tau-neutrinos produce a tauon that decays after a short distance, creating a second EM or hadronic shower at the point of its decay. At TeV-scale energies, the distance covered by the tauon before its decay can be large enough to make the separation between the two showers resolvable, creating a *double-bang* signature consisting of two cascades. At energies below 100 GeV that are more relevant to this work, however, the two cascades are too close together to be cleanly separable and they are effectively approximated as a single cascade as well. About 17% of tauons produce a muon upon decay, creating a track-like signature as well[21].

A summary of the secondary particles and corresponding signatures for each type of neutrino interaction is given in Table 4.1.

4.3.2 Atmospheric muons

Atmospheric muons are a significant background for oscillation measurements in DeepCore. Not only do they make up the majority of events that pass the initial DeepCore filter, but they also produce a track-like signature in the detector that is challenging to separate from that of a charged-current muon neutrino interaction. For this reason, most of the data filtering steps described in Section 5.2.3 are devoted to rejecting atmospheric muons while keeping as many muon neutrinos in the sample as possible. At energies below 100 GeV, the dominant energy loss for muons is ionization, which creates a continuous energy loss that can pass through the entirety of the instrumented volume of IceCube. As energies increase above 100 GeV, radiative energy losses become more relevant that create a series of stochastically distributed cascades along the muon's trajectory. The fraction of total energy loss that is concentrated in these cascades is referred to as *stochasticity*. Atmospheric muons also often arrive in bundles originating from a single interaction of cosmic rays with the upper atmosphere. Within such a bundle, stochastic energy losses of individual muons may average out over its trajectory, such that the bundle as a whole can be approximated as a single long track with a relatively low stochasticity.

Simulation and Data Processing 5

5.1 Event Simulation

5.1 Event Simulation 51
5.1.1 Neutrino Interactions . . 51
5.1.2 Atmospheric muons . . . 54
5.1.3 Photon Propagation . . . 56
5.1.4 Simulation of Detector Response 56
5.2 Data Processing 58
5.2.1 Trigger 58
5.2.2 Online Filter 58
5.2.3 Offline Filter 59
5.2.4 Event Reconstruction . . 62
5.2.5 Signature Classification . 67
5.2.6 Final Sample Selection and Binning 69
5.2.7 Seasonal Stability 72
5.3 Implementation of systematic uncertainties 73
5.3.1 Variation of Detector Properties 73
5.3.2 Variation of the Atmospheric Neutrino Flux 76

The method by which all of the measurements presented in this thesis are performed is that of *Monte-Carlo (MC) forward folding*. In a nutshell, this method involves producing a large set of simulated signal and background events that are then re-weighted in such a way that their distribution matches that of the observed data events as closely as possible. To give reliable results, an accurate simulation of all particle interactions described in Section 4.2 as well as the detector electronics described in section 4.1.4 is required. The simulated and observed events are then passed through the same data processing chain described in section 5.2. The resulting MC simulated dataset and the observed dataset are then histogrammed in the same binning, and the weights of the MC events are adjusted to give the best match between the histograms according to a loss function as defined in section 6.1.1.

The simulation chain for neutrinos and atmospheric muons can generally be divided into three steps that are described in this chapter:

1. Simulation of particle interactions
2. Photon propagation in ice
3. Response of detector DAQ systems

A special case is the simulation of detector noise, for which no particle production or photon propagation is necessary.

5.1.1 Neutrino Interactions

Because of the inherently low interaction rate of neutrinos, it would be impractical to simulate a constant flux of neutrinos from any particular direction, the vast majority of which would simply pass through the detector without producing any signal at all. Instead, every simulated neutrino is forced to interact within a given volume, and the event is given a weight corresponding to the inverse of the simulated fluence,

$$w = \frac{1}{F_{\mathrm{sim}}} \frac{1}{N_{\mathrm{sim}}} \ . \tag{5.1}$$

Here, N_{sim} is the number of simulated events and F_{sim} is the number of neutrino events per area, solid angle, energy, and time in the simulation. This weight, when multiplied with the flux of a given physics model and a live time, gives the expected number of events that this simulated event corresponds to. The baseline neutrino flux model used in this work is that proposed by Honda *et. al*[40] that is specifically computed for the South Pole[1].

Under the assumption that neutrino absorption is negligible and that the material consists of isoscalar targets, the simulated fluence is given by the chosen probability density in the direction and energy, $\phi_\Omega \times \phi_E$,

[40]: Honda et al. (2015), *Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model*

[1] Variations on this flux model and how they are propagated into the analysis are described in Section 5.3.2.

Table 5.1: Table of generation volumes used for Genie neutrino simulation. The cylinder is centered in DeepCore in all cases.

Flavor	Energy (GeV)	Radius (m)	Length (m)
$\nu_e + \bar{\nu}_e$	1-4	250	500
	4-12	250	500
	12-100	350	600
	100-10000	550	1000
$\nu_\mu + \bar{\nu}_\mu$	1-5	250	500
	5-80	400	900
	80-1000	450	1500
	1000-10000	550	1500
$\nu_\tau + \bar{\nu}_\tau$	1-4	250	500
	4-10	250	500
	10-50	350	600
	50-1000	450	800
	1000-10000	550	1500

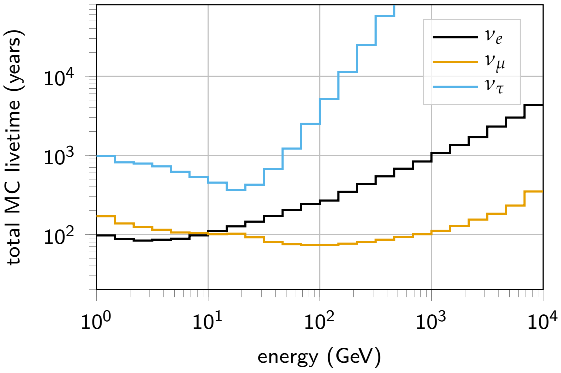

Figure 5.1: Simulated MC livetime as a function of energy, calculated using the HKKM[40] model flux with Nu-Fit 2.2[88] oscillation parameters.

the size of the interaction volume, V, the cross-section of the interaction, σ, and the density of the material, ρ, by

$$F_{\text{sim}}^{-1} = V \times \rho \times N_A \times 1\frac{\text{mol}}{\text{g}} \times \sigma \times \frac{1}{\phi_\Omega} \times \frac{1}{\phi_E}\,, \tag{5.2}$$

where N_A is Avogadro's number. The volume in which neutrino interactions are simulated is a cylinder centered in DeepCore, with a height and radius chosen such that all events that have a chance of producing a signal in DeepCore should be contained in it, depending on the neutrino flavor and energy (see also table 5.1). Neutrino directions are isotropically distributed in zenith and azimuth, implying $\phi_\Omega = \frac{1}{4\pi}$. The neutrino energies are sampled from a power law with $\phi_e \propto E^{-2}$. The simulated live time corresponding to a single simulated event is $T_{\text{sim}} = F_{\text{sim}}/\Phi$, where Φ is the expected neutrino flux including neutrino oscillations at global best-fit parameters. The amount of simulation generated for each neutrino flavor is chosen such that the total simulated live time is > 70 years over the entire energy range. Neutrinos and anti-neutrinos are produced in ratios of 70% and 30%, respectively. The simulated live time as a function of energy is shown in Figure 5.1. The livetime for electron neutrinos increases with energy because the simulated spectrum is harder than the real spectrum. The livetime for tau neutrinos is much higher than that of other flavors because the contribution of tau neutrinos to the expected neutrino flux is very small.

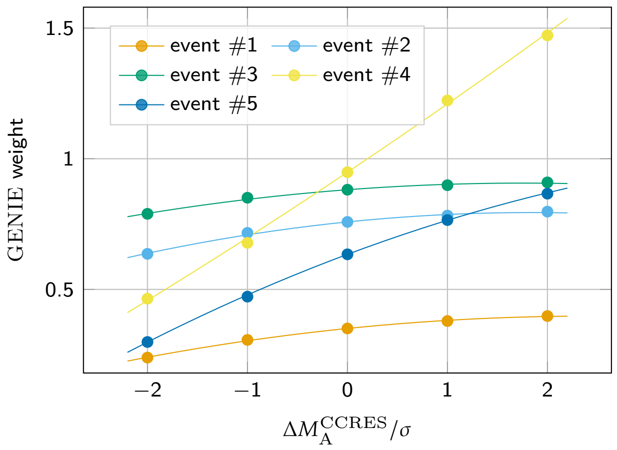

Figure 5.2: GENIE interaction weights as a function of the pull of the axial mass term M_A^{CCRES}, for five ν_e CC events produced via resonance interactions. Each dot represents a discrete point for which the event's cross section is computed in GENIE. The line represents the quadratic fit made used to interpolate the weight value over the continuous range allowed for the systematic parameter.

After sampling the parameters of the primary neutrino, the GENIE[89] software is used to simulate its interaction with the ice and the production of secondary particles and to calculate the cross-section of the interaction. The propagation and Cherenkov light production of any muon that is produced in these interactions is simulated with PROPOSAL[90]. The light output of secondary electrons, positrons, and gamma rays above 100 MeV, and that of hadronic showers above 30 GeV, are generated using analytic approximations from [86] as described in sections 4.2.3 and 4.2.4. At lower energies, the full GEANT4 simulation of the shower development is run to produce Cherenkov photons.

[89]: Andreopoulos et al. (2015), *The GENIE Neutrino Monte Carlo Generator: Physics and User Manual*

[90]: Koehne et al. (2013), *PROPOSAL: A tool for propagation of charged leptons*

Cross-section uncertainties

Two systematic parameters are included to account for uncertainties in the form factors of charged-current quasi-elastic (M_A^{CCQE}) events and charged-current resonant (M_A^{CCRES}) events. Both these form factors have a dependency on the momentum transfer, Q^2, of the form:

$$F(Q^2) \propto \frac{1}{(1-(Q^2/M_A^2)^2} \tag{5.3}$$

Where M_A is called the *axial mass*, and can be measured experimentally. The differential cross-section of each event is computed with GENIE at five discrete points, that is, the nominal mass and -2σ,-1σ,1σ and 2σ away from the nominal mass, where σ is a fractional uncertainty of 20%. This uncertainty approximates the recommendation of the GENIE collaboration, which suggests an asymmetric error of -15% and +25% for M_A^{CCQE} and a symmetric error of ±20% for M_A^{CCRES}[89]. In order to apply a continuous variation of that systematic parameter over the course of a minimization, a quadratic function is fit to interpolate between these discrete points. Figure 5.2 shows the GENIE weights of a handful of ν_e CC events from resonance production, across the allowed range of axial masses, along with their fitted quadratic dependence. The upper panel of Figure 5.3 illustrates an example of the varying M_A^{RES} on the final level sample.

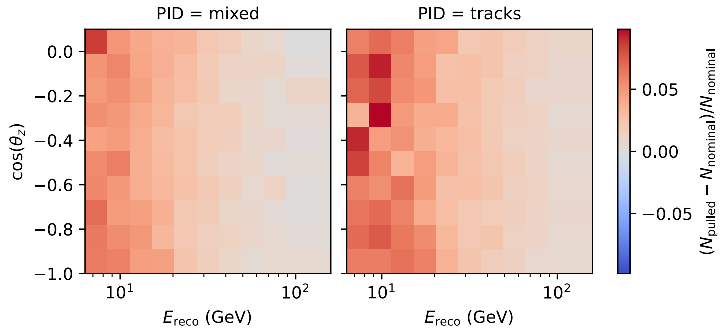

(a) GENIE M_A^{RES}

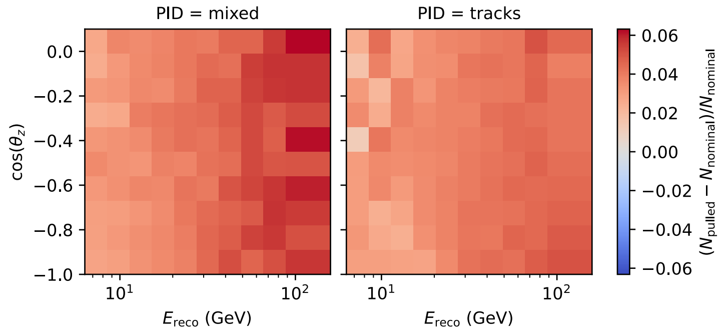

(b) DIS CSMS

Figure 5.3: Fractional difference in event rates between (top)M_A^{RES} (bottom) dis_csms at 1σ and at nominal value for both PID bins.

The uncertainty on the DIS cross-section is primarily given by the disagreement in DIS calculation between CSMS[91] and GENIE[89] cross-sections at energies above 100 GeV. This analysis includes a parameter that interpolates between these two calculations with a linear extrapolation to energies below 100 GeV. The bottom panel of Figure 5.3 illustrates an example of the varying this parameter, DIS, on the final level sample. As expected, the impact of the parameter is largest in the highest energy bins. There is an additional uncertainty of 20% on the normalization of NC events to account for uncertainties of the hadronization process and the Weinberg angle in line with previous oscillation studies[92].

5.1.2 Atmospheric muons

The offline filter steps described in section 5.2.3 decrease the rate of atmospheric muons by several orders of magnitude as events pass through each of its stages. This makes it challenging to produce a sufficiently large amount of simulated muon events to accurately estimate the expected background at the final level. To overcome this challenge, two separate muon simulation sets are produced, one of which is used to tune the lower level (up to L4) offline filters and the other is used to estimate muon background at levels L5 and above.

For both sets, atmospheric muons are generated on the surface of a cylinder encompassing the entire IceCube detector with a radius of 800 m and a height of 1600 m. Positions and directions of muons interacting in the detector are sampled using parametrized tables based on the approach described in [93]. These tables are tuned to approximate the output of a detailed CORSIKA[94] simulation of cosmic

[93]: Becherini et al. (2006), *A parameterisation of single and multiple muons in the deep water or ice*

[94]: Heck et al. (1998), *CORSIKA: A Monte Carlo code to simulate extensive air showers*

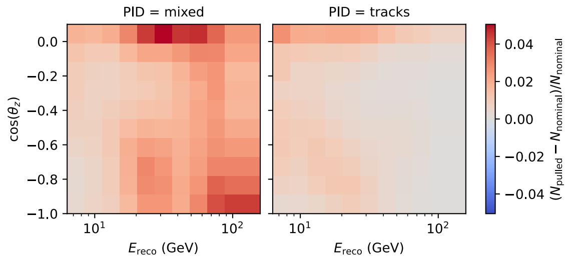

Figure 5.4: Impact on the final histograms when the muon normalization is increased by 50%. The largest impact is seen above the horizon in the mixed PID channel with a change in bin count of 5%.

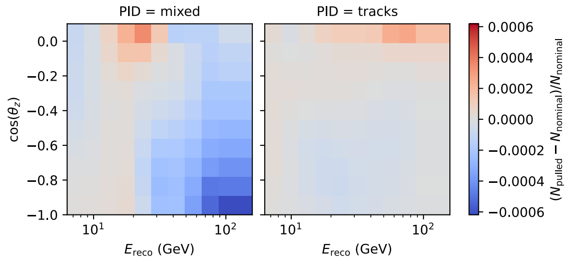

Figure 5.5: Impact on the final histograms when the muon spectral index is increased by 1σ.

ray interactions and subsequent shower production using the cosmic ray flux model described in [95] and the SIBYLL 2.1[96] hadronic interaction model. This flux is also used to weight simulated muon events and is distinct from the flux model used to weight neutrino events. For the simulations used to tune the lower selection levels, the muon energy is sampled from a power law with a spectral index of -3 and all events are accepted to cover the entire IceCube array. To produce the simulation that is used starting at the L5 trigger level, muons are only accepted if they intersect an inner cylinder centered in the DeepCore fiducial volume with a radius of 180 m and a height of 400 m. Furthermore, muons are rejected based on a KDE estimate of the muon density in energy and zenith angle at the L5 filter level. In this way, the sampling preferably produces such muon events that have a higher chance of passing the offline filtering up to L5, which greatly improves the efficiency of the simulation production.

After the position, direction and energy for a muon has been sampled, its propagation and photon production is simulated using PROPOSAL in just the same way as any muon that is produced in neutrino interaction would be.

Muon Uncertainty

Because the muon background contamination is cut to only ~2% at the final level of the event selection (see Section 5.2.6), the impact of muon systematic uncertainties is generally small. Only the over-all scale is left as a free parameter in the analysis, its impact is shown in Figure 5.4. This scale also largely absorbs the effects of DOM efficiency uncertainties, since, to first order, an increase in DOM efficiency leads to a better muon rejection. The spectral index of the muon flux has a very small effect far below the percent-level as shown in Figure 5.5 and is therefore not accounted for in the analysis.

[95]: Gaisser (2012), *Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio*

[96]: Engel et al. (2017), *The hadronic interaction model Sibyll – past, present and future*

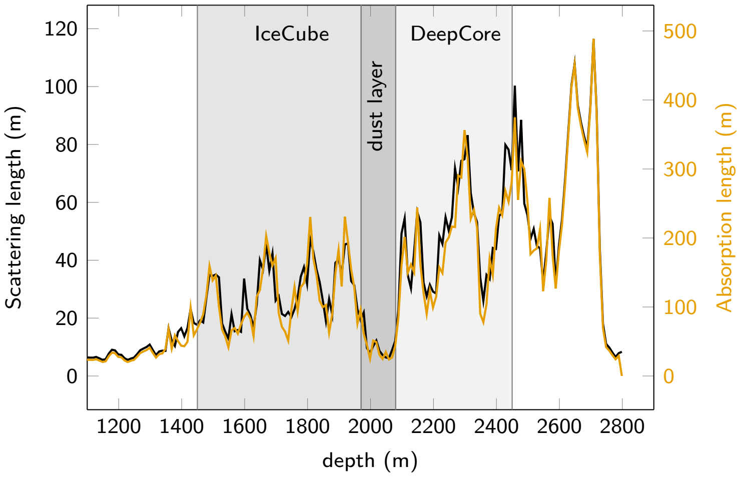

Figure 5.6: Scattering and absorption lengths as a function of depth in the South Pole Ice (SPICE) model that is used to produce the simulation for this work.

5.1.3 Photon Propagation

Photons are individually traced through the ice using the CLSIM[97] package, which is a GPU-accelerated OpenCL re-implementation of the Photon-Propagation Code[98]. The ice is modeled as 10 m thick layers with individual scattering and absorption lengths that are shown in Figure 5.6. The ice model used for the simulation in this work, also referred to as *South Pole ICE (SPICE)*[99], incorporates the fact that the ice layers are slightly tilted with respect to the vertical axis, and that scattering and absorption strengths are not uniform as a function of azimuth. For every photon, CLSIM first samples the absorption length from an exponential distribution whose expectation value is the absorption length of the current layer. It then propagates all photons in parallel steps, where every step corresponds to one scattering event and the step length is sampled from an exponential distribution where the expectation value is the scattering length of the current layer. The scattering angle is then sampled from a mixture of a Henyey-Greenstein distribution and a simplified Mie scattering distribution, where the shape parameters of these distributions have previously been calibrated using the in-situ LED calibration system[99]. Each photon stops when it has either reached its total absorption length or if it has intersected a DOM. After all photons have either been absorbed or reached a sensor, the simulations stops and passes the photons that reached a sensor on to the next step simulating the detector response.

[98]: Chirkin et al. (2019), *Photon Propagation using GPUs by the IceCube Neutrino Observatory*

[99]: Aartsen et al. (2013), *Measurement of South Pole ice transparency with the IceCube LED calibration system*

5.1.4 Simulation of Detector Response

After the photons have reached the surface of the optical sensors, the simulation determines for each one if it is converted into a Monte-Carlo photo-electron (MCPE). The probability that this occurs depends on the wavelength-dependent sensitivity of the DOM, as well as the angular acceptance. The angular acceptance not only depends on the geometry of the DOM itself, but also incorporates the effect of the re-frozen column of ice at the center of each bore hole. If a photon is accepted and converted into an MCPE, the next step is to simulate how much charge would be measured by the PMT inside the DOM as a response. The charge is drawn from a combination of a normal

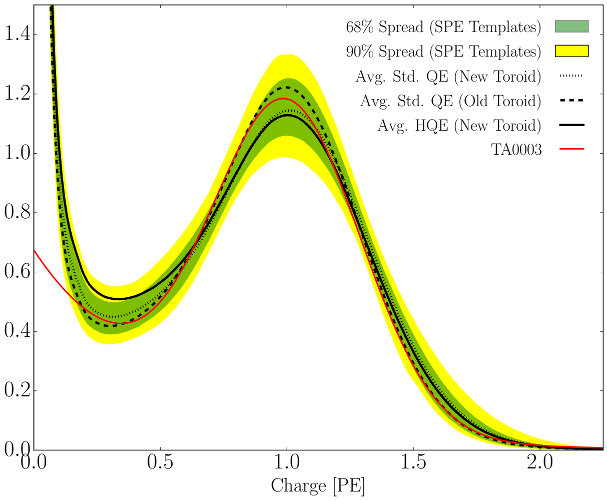

Figure 5.7: The green (yellow) regions show the 68% (90%) spread in the SPE charge templates for a given charge. Superimposed are the average SPE charge templates for the variety of hardware configurations shown in the black dotted, dashed, and solid lines. The TA0003 distribution, shown in red, originates from laboratory measurements. Figure taken from [78].

distribution and two exponential distributions whose parameters have been calibrated *in-situ* to match the observed charge distribution in each individual DOM[78]. This distribution, also referred to as the Single Photo-Electron (SPE) template, is shown in Figure 5.7. The MCPEs with the samples charge are then converted into simulated waveforms for the ATWD and fADC readouts which are then passed into the data processing chain starting from the *wavedeform* algorithm described in Section 4.1.4. From there, the simulated events pass through all the same trigger and filter steps that are described in Section 5.2.

[78]: Aartsen et al. (2020), *In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes*

Detector Noise

In addition to Cherenkov photons induced by relativistic charged particles in the ice, IceCube detects photons from radioactive decays inside the glass housing of the DOMs and PMTs that are simulated using the *Vuvuzela* module[100][101]. These "noise" MCPEs are simulated parametrically by sampling their times from distributions that take both thermal and non-thermal noise components into account. The thermal component comes from uncorrelated photons and PMT dark noise and is modeled as a Poisson process with a constant rate. The non-thermal component comes from correlated bursts of photons that are produced by radioactive decays. To simulate it, decay times are first drawn from a Poisson process with a constant rate, and the number of photons produced in each decay is sampled from a Poisson distribution. The time differences between the non-thermal MCPEs produced by each decay are then sampled from a Log-Gaussian distribution. This simulation method has five free parameters listed in Table 5.2 that are calibrated *in-situ* for every DOM. All thermal and non-thermal MCPEs are injected into each simulated event together with the MCPEs from photons and passed into the rest of the simulation chain.

Table 5.2: Parameters used in the noise simulation. Typical values taken from [100], actual values are fit for each DOM individually.

Parameter	Typical value
Therm. rate	$\lambda_{\mathrm{th}} \approx 20\,\mathrm{Hz}$
Decay rate	$\lambda_{\mathrm{dec}} \approx 250\,\mathrm{Hz}$
Decay hits	$\eta \approx 8$
Decay μ	$\log_{10}(\frac{\mu}{\mathrm{ns}}) \approx -6$
Decay σ	$\log_{10}(\frac{\sigma}{\mathrm{ns}}) \approx 2.7$

[100]: Larson (2013), *Simulation and Identification of Non-Poissonian Noise Triggers in the IceCube Neutrino Detector*

[101]: Larson (2018), *A Search for Tau Neutrino Appearance with IceCube-DeepCore*

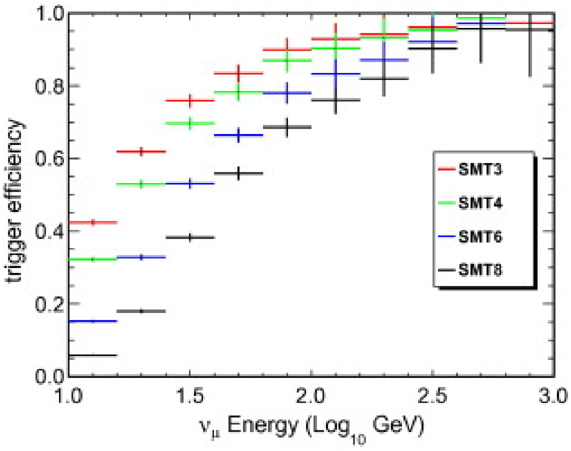

Figure 5.8: Efficiency of the IceCube and DeepCore triggers as a function of the primary neutrino energy. Figure taken from [75].

5.2 Data Processing

5.2.1 Trigger

As described in section 4.1.4, the high-frequency ATWD waveform digitization in each DOM is triggered when it and its adjacent or next-to-adjacent neighbors on the same string record a voltage corresponding to at least 0.25 PE-equivalent within a ±1 μs time window, which is referred to as the Hard Local Coincidence (HLC) condition. Data acquisition for DeepCore is triggered when this condition is fulfilled for at least three DOMs inside the DeepCore fiducial volume within a ±2.5 μs window. If this condition is met, the waveforms for all DOMs that have observed voltages of at least 0.25 PE within a ±10 μs time window centered around the trigger time are recorded. This trigger is referred to as the "SMT3" trigger and is distinct from the so-called "SMT8" trigger that is used to activate the data acquisition of the larger IceCube array, which requires eight DOMs to fulfill the HLC condition in a ±5 μs window. A DOM that has recorded PEs within the readout window but for which the HLC condition has not been met is said to fulfill the *Soft Local Coincidence* (SLC) condition. The DeepCore SMT3 trigger rate is less than 10 Hz while accepting ~70% of ν_μ events at 10 GeV and >90% of ν_μ events at 100 GeV[75]. The trigger efficiency for atmospheric muon neutrinos as a function of the primary neutrino energy is shown in Figure 5.8 for several different triggers that are used in IceCube.

5.2.2 Online Filter

Once the trigger condition is met, the recorded waveforms within the trigger window are converted into reconstructed pulses as described in Section 4.1.4 and are then passed into a set of *online* filters (i.e. filters running on hardware at the Pole). These filters are each designed to select events that are relevant to different physics measurements that are performed within the IceCube collaboration. For the purposes of the analysis presented in this thesis, events are selected using the *DeepCore filter*[75]. This filter is designed to select events that start inside the DeepCore fiducial volume and to reject those that are consistent

[75]: Abbasi et al. (2012), *The design and performance of IceCube DeepCore*

with muons entering the detector from the outside. The filter splits the observed series of hits between those hits that fall within the DeepCore fiducial volume and those outside of it. It then estimates the "center of gravity" (COG) in space and time of the hits inside the fiducial volume and then calculates the velocity that a signal would have to travel from each hit occurring outside the fiducial volume to coincide with the COG. If this velocity is close to the speed of light (between $0.25\ \mathrm{ns/s}$ and $0.4\ \mathrm{ns/s}$) for at least one hit, the event is rejected because it is consistent with a muon traveling through the veto region and entering DeepCore. Figure 5.9 shows an example of an event that would be rejected by the online filter. Only events that pass the trigger and filter conditions are sent north via satellite for further *offline* filtering.

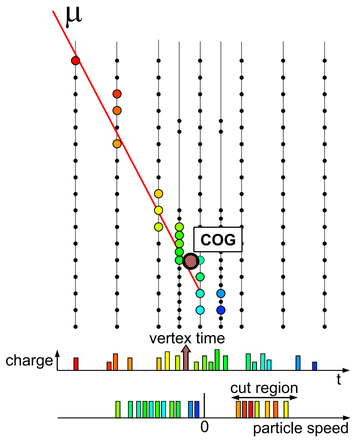

Figure 5.9: Example of an event that would be rejected by the online filter algorithm. DOMs that have observed light are highlighted in color depending on time from red (early hits) to blue (late hits). DOMs that have not observed any light are shown as black dots. Figure taken from [75].

5.2.3 Offline Filter

The offline filter is separated into subsequently applied *levels*, referred to as L3, L4 and L5, where each level reduces the amount of background (atmospheric muons and noise) by approximately an order of magnitude while keeping most of the DeepCore starting events that are the target of the selection.

Level 3

At the lowest offline filter level, L3, cuts are applied to simple variables that remove the most easily identifiable background events while using only few computational resources. The types of relevant types background events at this level of the event selection are pure noise, atmospheric muons and events with several coincident muons.

If a muon enters the detector after the data acquisition has already been triggered, it will create a series of pulses that extends much longer in time than what would be expected from a single particle interaction. Because the MC simulation only simulates single particles, however, these events cause a significant disagreement between data and MC. The time length of the observed hit series after noise cleaning is shown in the left panel of Figure 5.10, where this disagreement at large times is apparent. A cut at $5000\,\mathrm{ns}$, also shown in the figure, removes such events.

To identify noise events, the observed series of hits is first passed into a cleaning algorithm that uses time window coincident conditions between DOMs to remove hits that are likely to originate from pure noise. Only when at least six hits remain in the series after this cleaning procedure, the event is kept. Another algorithm checks whether the observed hits show any sign of directionality and only accepts the event if that is the case. Finally, an event should have more than two hits within a $300\,\mathrm{ns}$ sliding window.

The cuts aimed at removing muons consist of conditions on the number of hits in the veto region and conditions on the vertical position of the first HLC hit. One of these variables is the z-position of the first hit DOM for which the HLC condition was fulfilled and its distribution is shown in the right panel of Figure 5.10 A cut at $-120\,\mathrm{m}$ from the

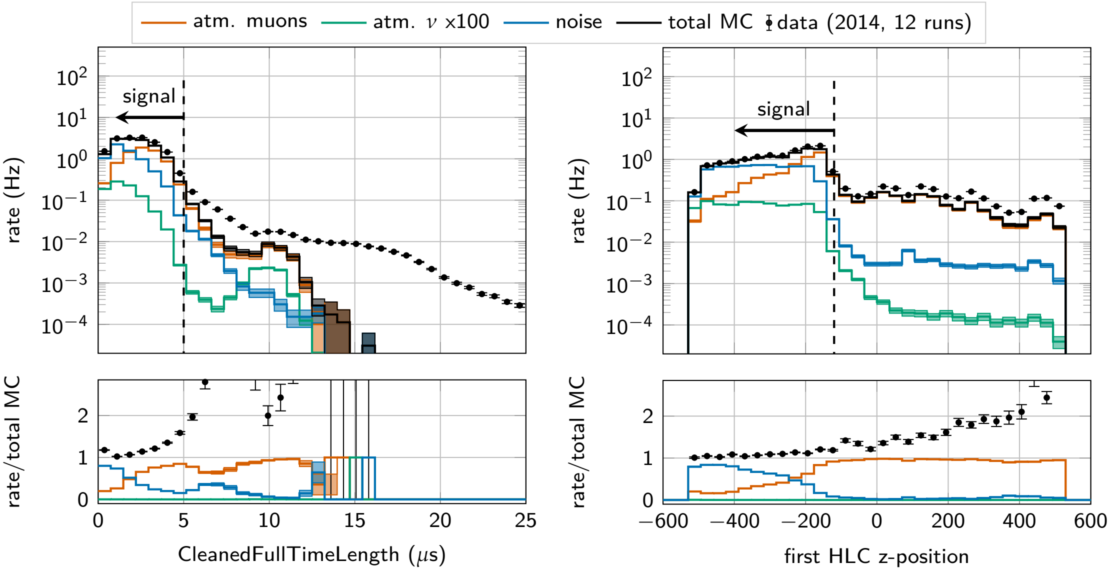

Figure 5.10: Distribution of one of the variables used in the L3 offline filter, the time between the last hit and the first hit after noise cleaning (left) and the z-position of the first HLC hit (right). Histograms show the distributions in simulated data separated by event type, data points with error bars show the distribution of real data. The bottom panel shows the ratio between data and simulation. Events falling on the "signal" side of the histogram are passed to the next filter level.

origin of the IceCube coordinate system (corresponding to a depth of $2068\,\mathrm{m}$ from the surface[2]) removes events that are likely to originate from atmospheric muons since they begin above the fiducial volume of DeepCore. The overall event rate after all L3 cuts have been applied is below $1\,\mathrm{Hz}$.

Level 4

In the next level, L4, more advanced selections based on the output of Boosted Decision Trees (BDTs) are applied, with a separately trained BDT for noise and muon rejection, respectively. The output of each BDT is a probability score between zero (background-like) and one (signal-like) and is shown in Figure 5.11. The first BDT to be evaluated is the one used to reject pure noise events. Its inputs consist of five variables:

- Cleaned N_{ch}: Number of hits in the noise cleaned hit series that was also used at L3.
- STW m3500p4000 DTW200: Slide a $200\,\mathrm{ns}$ time window over all pulses occurring from $-3.5\,\mu\mathrm{s}$ to $4\,\mu\mathrm{s}$ around the trigger time and take the largest number of pulses to fall within the window.
- The speed that is returned by the *LineFit* algorithm[102] on the observed hits
- The "fill ratio", that is, the fraction of DOMs that have recorded any hits inside a sphere centered around the first HLC hit. This variable effectively measures how compactly the hits are distributed around the starting point of the event and has been used in the past to identify cascades in IceCube[103].
- The ratio between the total duration of the cleaned hit series and the uncleaned hit series.

[2] The origin of the IceCube coordinate system is at a depth of $1948.07\,\mathrm{m}$ from the surface with the z-axis oriented upwards. The depth of a given z-coordinate is therefore $d = 1948.07\,\mathrm{m} - z$[74].

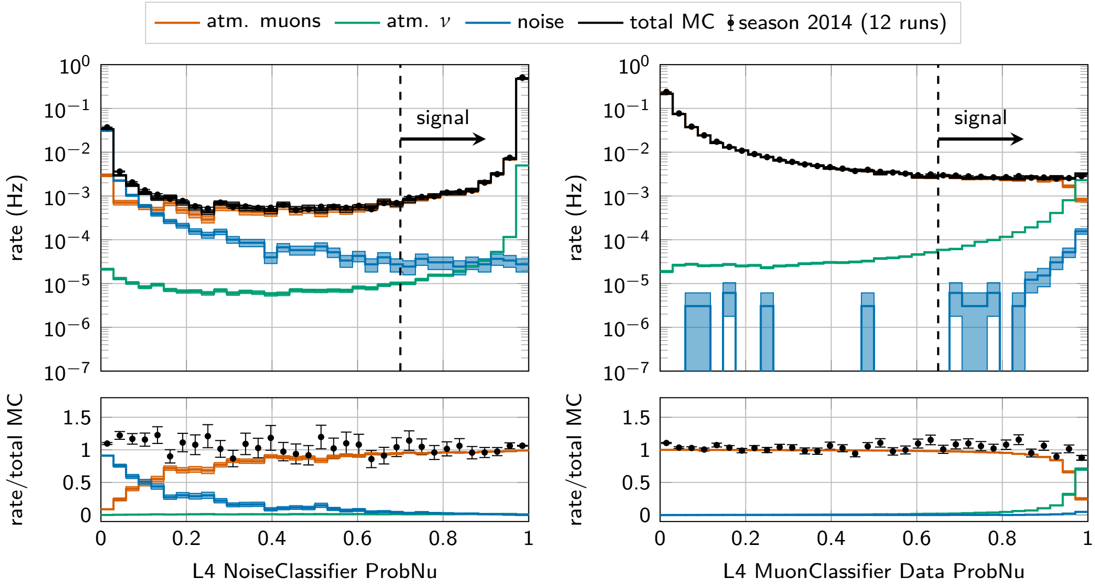

Figure 5.11: Distribution scores for the noise (left) and muon (right) BDT. The distributions of the muon classifier are shown for events where the score of the noise BDT is greater than 0.7. Histograms show the distributions in simulated data separated by event type, data points with error bars show the distribution of real data. The bottom panel shows the ratio between data and simulation. Events falling on the "signal" side of the histogram are passed to the next filter level.

The BDT is trained using simulated pure noise and neutrino events. An event passes the noise filter if the BDT score is above 0.7, which reduces the number of pure noise events by two orders of magnitude from 36.6 mHz to approximately 0.3 mHz. Passing events are passed into the second BDT that is used to reject atmospheric muons. This BDT takes in a larger number of variables that can be summarized as belonging to three different categories:

- Several N_{ch}-like variables counting the number of hits in different veto regions of the detector.
- The time to reach 75% of an event's charge in the cleaned pulse series.
- Several variables that characterize the spacial distribution of hits in z-coordinate and the radius with respect to string 36 (roughly the center of DeepCore).

In contrast to the noise BDT, however, the muon BDT is trained using real data and simulated neutrino events, with the goal of rejecting data events. This is possible because the data sample consists to 99% of atmospheric muons at this stage of the event selection. Events pass the L4 muon cut if the output score of the muon BDT is greater than 0.65, removing 94% of all muon events while keeping 87% of all neutrinos. These thresholds are shown alongside the distribution of the BDT outputs in Figure 5.11.

Level 5

The final filter that is applied before the event reconstruction step is L5. This filter searches specifically for hits occurring in un-instrumented *corridors* within the IceCube array through which an atmospheric muon can sneak into the DeepCore volume while evading previous veto

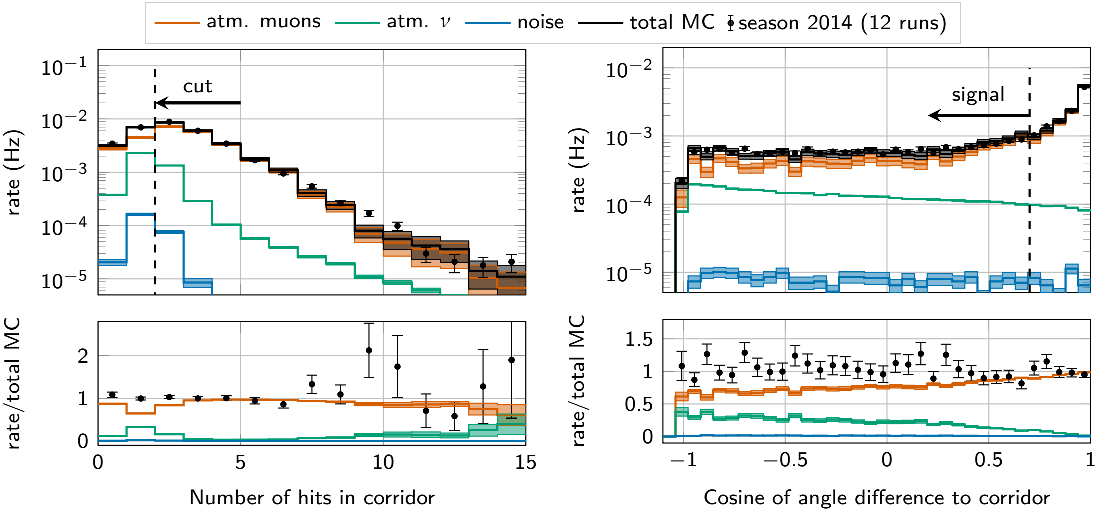

Figure 5.12: Distributions for two of the L5 corridor cut variables. Histograms show the distributions in simulated data separated by event type, data points with error bars show the distribution of real data. The bottom panel shows the ratio between data and simulation. Events falling on the "signal" side of the histogram (or, equivalently, opposite to the "cut" side of the histogram) are passed to the next filter level.

Table 5.3: Summary of the rates obtained after each level of selection. Neutrinos are weighted to an atmospheric spectrum with oscillations included.

	rate (μHz)				
Event type	DeepCore filter	L3	L4	L5	Eff. (%)
Atm. μ	7273	505	28.1	0.93	0.012
Pure noise	6621	36.6	0.28	0.07	0.001
Atm. ν_e CC	1.61	0.95	0.84	0.48	29.8
Atm. ν_μ CC	6.16	3.77	3.11	1.39	22.5
Atm. ν_τ CC	0.19	0.13	0.12	0.07	36.8
Atm. ν NC	0.86	0.53	0.46	0.23	26.7

cuts. In addition, events with more than seven hits in the outermost strings of the IceCube array or that have a down-going pattern of hits in the uppermost region of the detector are vetoed to remove events containing atmospheric muons entering the detector coincidentally with neutrinos. The distribution for one of the corridor variables and one of the muon rejection variables are shown in Figure 5.12. Table 5.3 shows the rates of each event type expected at each level of the selection up to L5 together with the efficiency of the filter at the final level.

5.2.4 Event Reconstruction

After the L5 selection, the rate of muons is reduced enough so that the majority of the total sample is expected to consist of atmospheric neutrinos, and it is at this point that the event reconstruction and signature classification are run. For the measurement presented in this thesis, three reconstructed quantities are required: the zenith angle, the energy, and a proxy score determining the flavor of a neutrino. As described in Section 4.3, all neutrino events in DeepCore can be effectively approximated as a cascade (ν_e CC events, all NC events and 83% of ν_τ CC events) or a combination of a cascade at the neutrino interaction point with an outgoing muon track (ν_μ CC events and 17% of ν_τ CC events). The zenith angle can be most accurately reconstructed for track-like events due to their elongated, highly directional signature.

For cascades, the reconstruction of the direction is more difficult because of their more compact and diffuse light distribution. The energy of a neutrino event is reconstructed by comparing the expected light output of a combined track and cascade hypothesis with the observed hits. Finally, the flavor proxy is calculated using variables that characterize the elongation of the observed hit signature and the goodness of fit of a combined track and cascade hypothesis compared to that of a cascade-only hypothesis. The resulting score allows the separation of muon neutrino interactions from other interactions, which is ideally suitable to observe the muon neutrino disappearance oscillation channel.

Zenith angle reconstruction

The zenith angle is reconstructed using the Single-string Antares-inspired Analysis (SANTA)[104]. It is an older algorithm aimed at reconstructing the direction of muon tracks that was originally developed for use in the ANTARES neutrino telescope [105]. It has since been refurbished and improved in IceCube, as described in detail in [106].

[104]: Garza (2014), *Measurement of neutrino oscillations in atmospheric neutrinos with the IceCube DeepCore detector*

[105]: Aguilar et al. (2011), *A fast algorithm for muon track reconstruction and its application to the ANTARES neutrino telescope*

[106]: Abbasi et al. (2022), *Low Energy Event Reconstruction in IceCube DeepCore*

The reconstructed pulse series in every DOM is summarized by the time of the first pulse and the sum of charges of all pulses. This time and charge are the only information used by the reconstruction and are referred to as a *hit* in the following. The first step of the reconstruction algorithm is a cleaning routine that removes hits produced by photons that have been scattered many times as they traveled through the ice, leaving only hits from photons that have traveled in approximately straight lines based on the time difference between hits on the same string. It calculates the signal speed between hits on the same string, and removes a hit if this velocity is below the speed of light in ice. This is a simplification of the algorithm described in [104], where the effective signal velocity was updated during the selection process. The selection is run separately for each string, and if fewer than three hits remain on a string, all hits on the string are discarded. In total, it is necessary for at least five hits to remain in an event in order to run the directional reconstruction. If only hits on one string remain after the selection, the event is referred to as a *single-string* event, otherwise it is a *multi-string* event. The reconstruction is generally more accurate for multi-string events, because the spacing between strings provides a long lever arm to constrain the direction of a track. In addition, the azimuth angle of the track can only be reconstructed for *multi-string* events due to the rotational symmetry of a single string.

The directional reconstruction itself is a regression that minimizes a modified χ^2 loss that is defined as

$$L(\vec{\theta}) = \sum_{i=1}^{N} \phi(r_i^2(\vec{\theta})) + \frac{1}{\tilde{q}} \sum_{i=1}^{N} \tilde{q}_i \frac{d_{\gamma,i}}{d_0} \,. \tag{5.4}$$

The first term in Equation 5.4 is a χ^2 loss that is modified to be robust against outliers. The second term is a regularization term that penalizes solutions where large charges are observed at large distances. Here, r_i^2 is the chi-square residual for each observed hit, i, between the observed

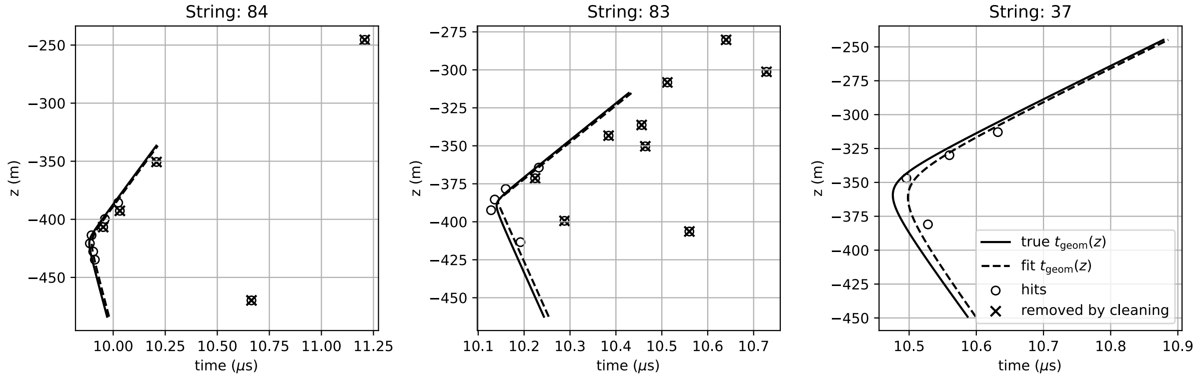

Figure 5.13: Example of a $\nu_{\mu,\text{CC}}$event reconstructed with SANTA with hits on several strings. Strings 84, 83 and 37 are spaced $\sim 80\,\text{m}$ apart from each other and form a highly obtuse triangle.

time, $t_{\text{obs},i}$ and the geometric arrival time, $t_{\text{geom},i}(\vec{\theta})$,

$$r_i^2(\vec{\theta}) = \left(\frac{t_{\text{geom},i}(\vec{\theta}) - t_{\text{obs},i}}{\sigma_t} \right)^2 . \tag{5.5}$$

The residual is wrapped in a *robust loss function*,

$$\phi(r_i^2) = \log\left(1 + r_i^2/C^2\right) C^2 , \tag{5.6}$$

which grows much more slowly than r_i^2 for values of r_i greater than the "soft cutoff", C, while behaves very similarly to r_i^2 for values of r_i smaller than C. Effectively, this robust loss reduces the influence of hits that pass the hit selection procedure despite having undergone a significant amount of scattering. The uncertainty in the pulse-time measurement is approximately $\sigma_t = 3\,\text{ns}$, corresponding to the readout rate of the modules [77].

In the second term of Equation 5.4, $\tilde{q}_i$ is the total observed charge in DOM i divided by the effective area of the DOM at the angle of incidence of the photon, and $\bar{q}$ is the average over all $\tilde{q}_i$. For every DOM, the charge per effective area is multiplied by the distance traveled by the photon, $d_{\gamma,i}$, and divided by a typical scaling distance, d_0. The distance d_0 determines the strength of the regularization term and has been optimized empirically to achieve the optimal resolution of the reconstruction to a value of $7\,\text{m}$.

The expected arrival time for unscattered Cherenkov photons is calculated geometrically under the assumption of an infinitely long track characterized by a normalized direction vector $\vec{u} = (u_x, u_y, u_z)$, an anchor point $\vec{q} = (q_x, q_y, q_z)$ and a time t_0 at which the particle passes through $\vec{q}$. The velocity is fixed to the vacuum speed of light, c. Since the reconstruction ignores DOMs that have not recorded any pulses, the fact that the true track length is finite only makes a negligible difference. The arrangement of these vectors is shown in Figure 5.14. Without scattering, all Cherenkov photons lie on a cone with an opening angle θ_c whose tip is in the position of the particle at the time $\vec{p}(t)$. The opening angle satisfies $\cos(\theta_c) = 1/n_{\text{ph}}$, where n_{ph} is the phase index of refraction of the ice. Assuming that a photon has traveled in a straight line at the group velocity in ice, the geometric arrival time,

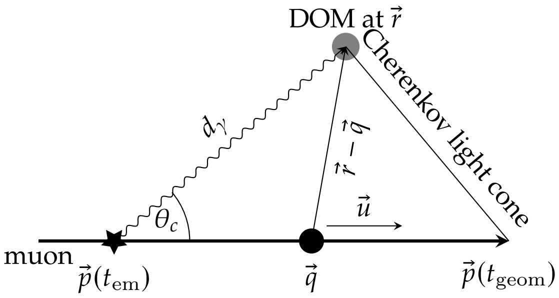

Figure 5.14: Detailed geometry of a light cone created by a track. $\vec{q}$ is the position of the anchor point and $\vec{r}$ is the position of the optical module. $\vec{p}(t_{\mathrm{em}})$ and $\vec{p}(t_{\mathrm{geom}})$ are the positions of the muon at the time the photon is emitted and when it is geometrically expected to arrive, respectively.

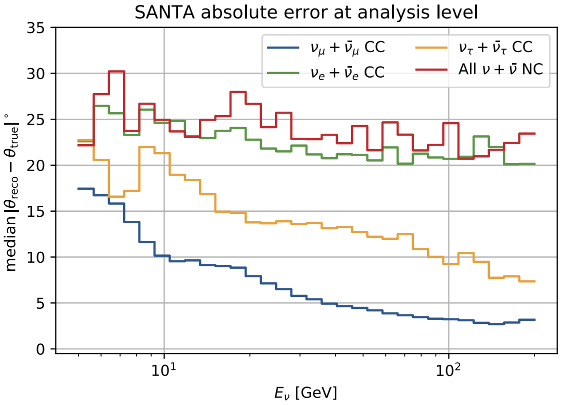

Figure 5.15: Median error on the reconstructed zenith angle at the final level of the sample selection as a function of the true simulated neutrino energy.

t_{geom}, for a DOM at position $\vec{r}$ is

$$t_{\mathrm{geom}} = t_0 + \frac{1}{c}\left((\vec{r}-\vec{q})\cdot\vec{u} + \frac{d_\gamma}{n_{\mathrm{ph}}}\left(n_{\mathrm{ph}}n_{\mathrm{gr}} - 1\right)\right) \tag{5.7}$$

where the distance traveled by the photon d_γ is

$$d_\gamma = n_{\mathrm{ph}}\sqrt{\frac{1}{n_{\mathrm{ph}}^2 - 1}\left(\vec{u}\times(\vec{r}-\vec{q})\right)^2}\,. \tag{5.8}$$

The group and phase indices of refraction depend on the wavelength, but for this reconstruction the value for a wavelength of $\lambda = 400$ nm [i] is used, where $n_{\mathrm{gr}} = 1.356$ and $n_{\mathrm{ph}} = 1.319$ from [107]. An example of a simulated event reconstructed with SANTA is shown in Figure 5.13. The solid and dashed lines show the geometric arrival time calculated according to equation 5.7 using reconstructed and true track parameters, respectively. The circles indicate hits in DOMs, and those hits that have been removed by the hit cleaning procedure are crossed out. The median error on the zenith angle reconstructed using SANTA is shown in Figure 5.15, split by neutrino interaction type. As expected, the error is the smallest for ν_μ-CC interactions, since those produce track signatures that most closely resemble the infinite track hypothesis underlying the SANTA reconstruction algorithm. The worst resolution is achieved for interactions that only produce electronic or hadronic showers, since they produce cascade signatures hardly resembling the infinite track assumption. It is also apparent that the median resolution for ν_τ-CC events lies between that of ν_μ-CC events and pure

[107]: Price et al. (2001), *Role of group and phase velocity in high-energy neutrino observatories*

---

[i] 400 nm is near the wavelength of the highest acceptance of the optical modules.[74]

cascade events. This is readily explained by the fact that 17% of these interactions also produce a muon in their final state.

In addition to the zenith angle reconstruction, SANTA can also be used to fit a simplified cascade hypothesis to the observed hits. For this purpose, it is assumed that light is emitted uniformly in all directions originating from the interaction vertex as shown in the left panel in Figure 4.7. With this assumption of perfect rotational symmetry, it is not possible to reconstruct a direction, and the cascade is fully characterized by the position of the vertex and the interaction time. The ratio of the χ^2 of the infinite-track regression and the χ^2 of this *cascade-only* regression is used as a proxy for the neutrino flavor in this analysis. If it is smaller than one, the infinite-track hypothesis achieves a better fit to the data than the cascade-only hypothesis.

Energy reconstruction

The energy reconstruction runs as a separate step after the zenith angle reconstruction. In contrast to SANTA, the Low-Energy Energy Reconstruction Algorithm (LEERA)[108] fits a combined hypothesis consisting of a cascade and a finite-length track originating at the same point of the cascade. Both the cascade and the track are constrained to move only along the infinite track that has been fit in the zenith reconstruction. This means that the model fit in the energy reconstruction is fully characterized by the shift of the vertex along the infinite track, the length of the finite track (which is linearly related to the track energy), and the energy of the cascade. Given these parameters, the expected light yield for all DOMs is calculated using the so-called *photonics tables*. The tables consist of B-spline coefficients that have been fit to simulated photon propagation for cascades and 3 m long tracks segments at different depths and directions inside the IceCube array to give a (time-dependent) expectation value for the photon count at arbitrary positions inside the detector. The expectation of an arbitrarily long track is calculated by chaining the 3 m segments together that fully cover the desired track length, and scaling the amplitude of the last segment by the remainder of the division of the desired length by the length of the segments. Given these expectation values as a function of event parameters, $\lambda_i(\theta)$, for every DOM, i, a simple Poisson "hit vs. no-hit" log-likelihood is calculated as

$$\log(\mathscr{L}) = \sum_{i \in \text{DOMs without hits}} e^{-\lambda_i(\theta)} + \sum_{i \in \text{DOMs with hits}} \left(1 - e^{-\lambda_i(\theta)}\right) . \quad (5.9)$$

This likelihood is maximized under the hypothesis that the shift, track length, and cascade energy are all free parameters, and under the alternative hypothesis where the track length is fixed to zero, the latter of which corresponds to a cascade-only hypothesis. The difference between these two log-likelihoods provides a measure of the degree to which the combined track and cascade hypothesis fits the observed data better than a track-only hypothesis. It is one of the inputs that is used in a BDT to calculate an overall score of how track-like an observed event signature is. The median relative error in the reconstructed total energy (that is, the sum of the track energy and the cascade energy) is shown in Figure 5.16. As with the zenith angle reconstruction, the relative

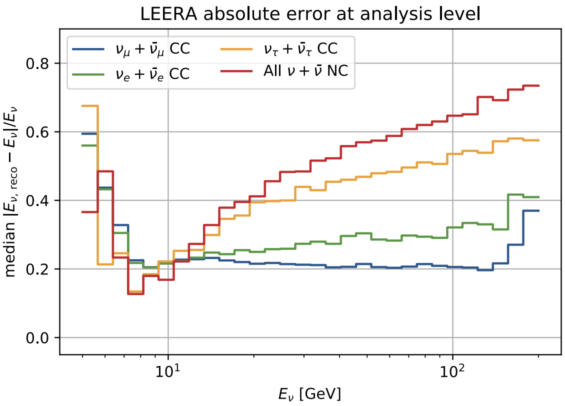

Figure 5.16: Median fractional error on the reconstructed energy at the final level of the sample selection as a function of neutrino energy.

error is smallest for $\nu_{\mu,\,\mathrm{CC}}$-events. This is expected, since these events fit the hypothesis of an initial cascade combined with a finite track the best. The second-best resolution is achieved for $\nu_{e,\mathrm{CC}}$-events, while it is poorest for $\nu_{\tau,\mathrm{CC}}$ and neutral-current events. This is explained by the fact that the expected light yield that is put in Equation 5.9 is based on the assumption that all particles that are produced in the interaction are visible to the detector. While this assumption is a good approximation for $\nu_{e,\mathrm{CC}}$-events, it does not hold for hadronic cascades that contain some neutral components as discussed in chapter 4.2.4. The true energy of the primary particles that produce hadronic cascades is therefore systematically under-estimated and has a larger uncertainty. This additional uncertainty is fundamentally irreducible, because it is not possible to distinguish the signatures of hadronic and electromagnetic showers.

5.2.5 Signature Classification

In addition to the energy and zenith angle, the measurement presented in this thesis requires a score that separates track-like $\nu_{\mu,\,\mathrm{CC}}$-events from other types of interaction. While previous analyses used only single variables such as the reconstructed track length to differentiate between tracks and cascades [109–111], the analysis presented in this thesis uses several variables as input into a Boosted Decision-Tree (BDT) to compute a score for how track-like the observed signature is. The BDT classifier is taken from the scikit-learn[112] package and trained to classify between tracks and cascades using the following input variables:

- SANTA χ^2-ratio, defined as $\frac{(\chi^2/\mathrm{d.o.f.})_{\mathrm{track}}}{(\chi^2/\mathrm{d.o.f.})_{\mathrm{cascade}}}$, i.e. the ratio of goodness-of-fit metrics from each fit hypothesis in the directional reconstruction (see section 5.2.4)
- ΔLLH from energy reconstruction, defined as $\mathrm{LLH}_{\mathrm{track}} - \mathrm{LLH}_{\mathrm{cascade}}$, i.e. the best-fit LLH value from each hypothesis
- Reconstructed muon track length, L_μ
- Radial distance of the reconstructed interaction vertex from string 36[ii], ρ^{36}_{vertex}
- Radial distance of the end-point from string 36, ρ^{36}_{stop}

[ii] String 36 is approximately at the center of the array, and near to the densest region of DeepCore (see Figure 4.2).

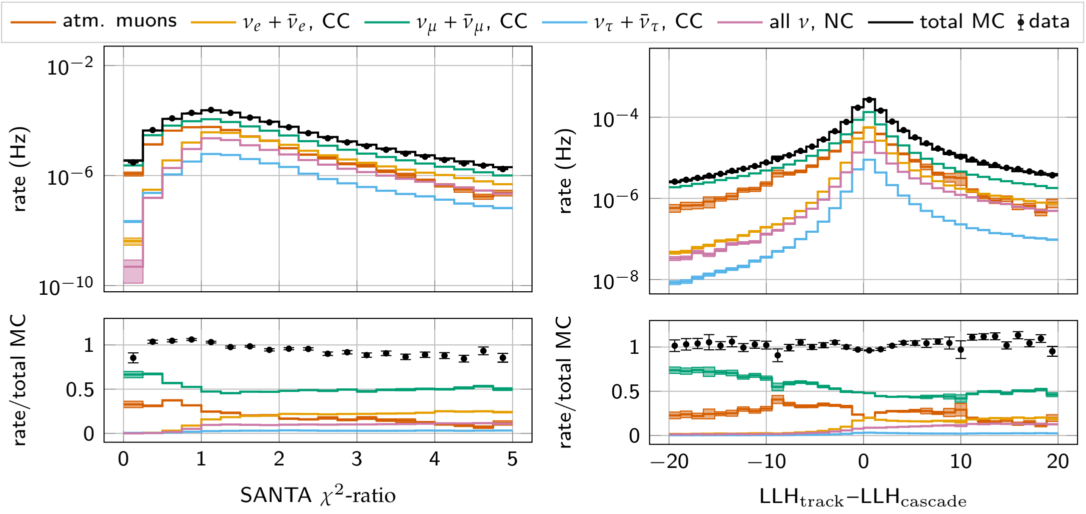

Figure 5.17: Distribution and data/MC comparison for the two most important input variables into the classification BDT.

- Depth of the interaction vertex, z_{vertex}
- Depth of the end point, z_{stop}

Of these variables, the SANTA χ^2-ratio and ΔLLH contribute the most to the final score. Their distributions and comparison between data and simulation can be seen in Figure 5.17 at the L5 selection level, where neutrinos are weighted with NuFit 4.0[24] global fit parameters. The training data consists of simulated ν_e-CC interactions and neutral-current interactions representing cascades, and ν_μ-CC interactions representing tracks. Tau neutrino interactions are not included in the training data in order to avoid confusion due to the 17% of ν_τ-CC interactions that produce track-like signatures. The training samples are weighted to approximate the neutrino flux expected from the HKKM model [40] *without* oscillations. This is done to avoid imprinting the event distributions at certain values of the oscillation parameters into the trained classifier. The distributions for these variables for tracks and cascades as they were used in training can be found in the Appendix B.1.2. Only half of the available simulation is used for training, while the other half is held out to validate that the classifier generalized to events that it has not seen during training. The output score of the classifier is referred to as *particle-ID* (PID) and ranges from zero (very cascade-like) to one (very track-like). The distribution of the PID score for simulated neutrino interactions is shown in Figure 5.18, broken down by flavor and interaction type. The distributions are individually normalized to help visualize the shape differences between the different neutrino interactions. The distributions for all interaction types show a large peak around a probability score of 0.5, suggesting that the event signature cannot be clearly classified for the majority of events. A second peak exists only in the distribution of ν_μ-CC events close to a score of one, meaning that there exists a population of these events that can be very clearly classified as being track-like. There also exists some excess of high PID values in the distribution of ν_τ-CC events corresponding to those events where the decay of the tauon produces a muon. Notably, there is no population of events that can be cleanly classified as a cascade event, i.e., there are no PID scores close to zero. The reason for this is that the two classes are nested hypotheses, one

[40]: Honda et al. (2015), *Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model*

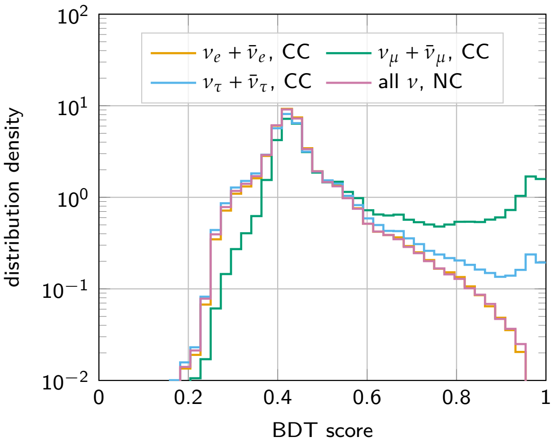

Figure 5.18: PID score distribution for simulated neutrino events at the final level of the event selection, weighted according to the HKKM flux model [40] and neutrino oscillations with NuFit 4.0[24] global fit parameters, normalized to unity. The BDT score ranges from the most cascade-like at 0 to the most track-like event signature at 1.

containing only a cascade and the other containing a combination of a cascade and a track, and it is never possible to prove that a cascade-like signature does not contain at least a short track segment.

5.2.6 Final Sample Selection and Binning

After the reconstruction and classification step, several final cut variables are applied to reduce the background of atmospheric muons to only a few percent, and to remove a small number of events from data containing coincident muons. These cuts are:

- The reconstruction of energy and zenith angle has to be successful. This requires, in particular, that at least five hits remain after the hit cleaning procedure described in Section 5.2.4.
- The reconstructed energy should be in the range between $6\,\mathrm{GeV}$ and $156\,\mathrm{GeV}$.
- Require a minimum PID score (see Section 5.2.5) of 0.55 to only include at least somewhat track-like events.
- Reconstructed $\cos(\theta_z) < 0.1$ to remove events that enter the detector from above the horizon.
- Require a minimum goodness-of-fit of the zenith reconstruction with $\chi^2_{\mathrm{mod}}/\mathrm{d.o.f.} < 50$.
- A tighter cut on the L4 muon BDT score (see Section 2) of $P_\nu > 0.97$.
- Fewer than eight hits in the outermost strings of the IceCube array, and a positive "z-travel" value for hits in the uppermost 15 layers of DOMs in the (non-DeepCore) IceCube array. The "z-travel" value for a given sequence of hits is calculated by subtracting the mean value of the z-coordinate of the first quartile of hits from the mean z-coordinate of all hits.

The cuts on energy, zenith angle, and PID define the range of the binning that will be used in the analysis. The cut on the zenith angle in particular is applied not only to reduce the background of atmospheric muons, but also to remove the phase space of neutrino events where muons that are produced in the same air shower that also produced the neutrino cause it to be vetoed by the muon filter cuts. This effect is referred to as the "self-veto" effect and would lead to a disagreement

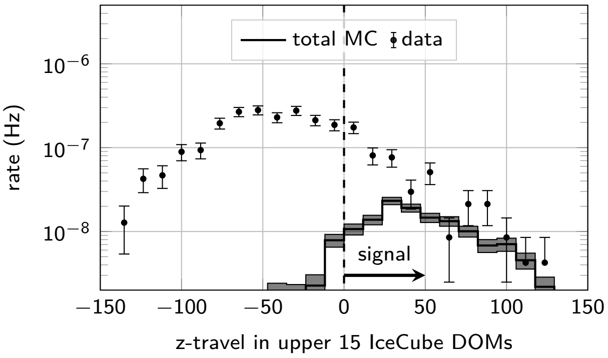

Figure 5.19: Distribution of the "z-travel" variable calculated for the uppermost 15 layers of IceCube DOMs. Only events with at least 4 hits in the uppermost 15 layers of DOMs are included in the histogram.

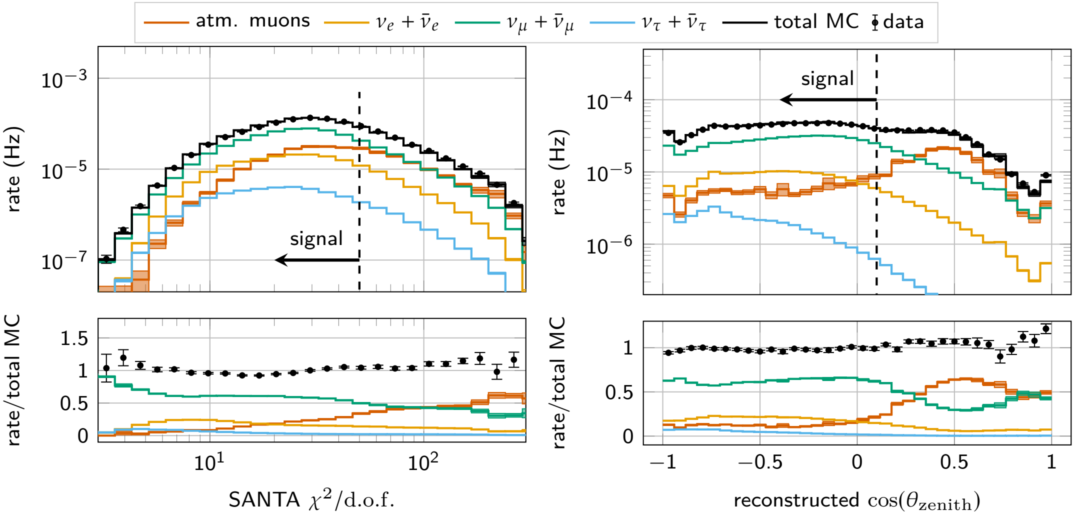

Figure 5.20: Distribution of the SANTA goodness-of-fit variable and the reconstructed zenith angle at L5 of the event selection process.

between data and simulation since coincident muons are never simulated. The distribution of the cosine of the reconstructed zenith angle is shown in the right panel of Figure 5.20, and it is apparent from the distributions that atmospheric muons dominate in the region of down-going events.

The requirement on the SANTA goodness-of-fit not only ensures that the included events are well-reconstructed, but also reduces the fraction of muons in the sample, as can be seen from the distributions shown in the left panel of Figure 5.20. The number of hits on the outermost strings and the "z-travel" variable calculated for hits in the uppermost 15 layers of IceCube DOMs are indicators of muons that hit the detector within the trigger window of a neutrino event. Such coincidences are entirely absent in simulation, which becomes especially apparent in the distribution of the "z-travel" variable shown in Figure 5.19, where a negative value indicates a down-going signal. After the application of all these cuts, the data sample consists of 21,914 well-reconstructed, track-like events with an expected background from atmospheric muons of only ~ 2% as shown in Table 5.4. For the purpose of the oscillation measurements presented in this work, both data and simulation sets are binned in reconstructed energy (E_{reco}), cosine of the reconstructed zenith angle ($\cos(\theta_z)$), and PID as follows:

- E_{reco}: 11 bins spanning the range from 6.31 GeV to 158.49 GeV,

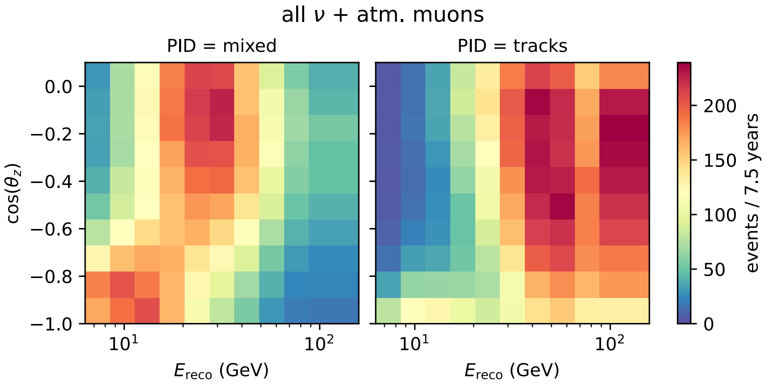

Figure 5.21: Expected event counts in 7.5 years of live time assuming no sterile mixing and NuFit 4.0 [24] global best fit parameters at Normal Ordering.

Table 5.4: Successively applied cuts on the data sample. The bottom row corresponds to the final rates in the sample after all cuts have been applied. The total rate of the data and simulation does not match, which is expected since there is a large amount of uncertainty in the total normalization. Muon contamination is the muon rate divided by the total event rate. Numbers calculated at the NuFit 4.0 global best-fit point.

	rate (μHz)			
condition	ν (sim)	μ (sim)	data	μ fraction
has SANTA reconstruction	957	314	1183	24.7 %
$6\,\mathrm{GeV} < E_{\mathrm{reco}} < 156\,\mathrm{GeV}$	862	311	1095	26.5 %
BDT score > 0.55	232	117	336	33.5 %
$\cos(\theta_{\mathrm{reco}}) < 0.1$	175	17	177	8.8 %
SANTA $\chi^2/\mathrm{d.o.f.} < 50$	164	12	161	6.6 %
L4 muon ν prob > 0.97	101	2	93	2.1 %
coinc. μ cuts (final rate)	**101**	**2**	**93**	**2.1 %**

the two bins with the highest energy are merged.

- $\cos(\theta_z)$: 10 bins spanning the range from -1 to 0.1
- PID: One bin between 0.55 and 0.75, and one bin between 0.75 and 1.0.

The lower PID bin between 0.55 and 0.75 consists to 69% (pre-fit MC estimate) of charged-current $\nu_\mu+\bar{\nu}_\mu$ events and is referred to as the *mixed* channel, while the higher PID channel between 0.75 and 1.0 consists to 94% of charged-current $\nu_\mu + \bar{\nu}_\mu$ events and is referred to as the *tracks* channel. The expectation values of the histogram in both PID channels is shown in Figure 5.21 at current global best-fit parameters for standard three-flavor oscillations. The expectation values are calculated from Monte-Carlo (MC) simulation that is described in detail in Section 5.1. The detailed breakdown of event counts in the final data sample by particle type and PID channel is given in Table 5.5.

Muon Smearing

After all the filtering steps described in section 5.2, the muon contamination of the data sample is reduced to ~ 2% of the sample. This reduces the statistics of muon simulation so much, that the resulting histograms become very sparse as shown in figure 5.22a. Such sparse histograms, in which single MC events have to serve as a stand-in for several real data events, are a poor template for what can be expected in data. To produce a more realistic expectation of the bin counts, the muon histograms are smeared using KDEs as shown in figure 5.22b. Since the KDE operates on events on the entire zenith and energy range, including events that fall outside the analysis binning, some

Table 5.5: Expected event rate with 8 years livetime broken down in event types and PID bins, calculated at Nu-Fit 4.0 global best fit parameters.

Type	PID	Event Count	Rate (μHz)
All MC	mixed	11428	48.3
All MC	tracks	12238	51.7
$\nu_{all} + \bar{\nu}_{all}$ NC	mixed	943	4.0
$\nu_e + \bar{\nu}_e$ CC	mixed	1704	7.2
$\nu_\mu + \bar{\nu}_\mu$ CC	mixed	7901	33.4
$\nu_\tau + \bar{\nu}_\tau$ CC	mixed	470	2.0
muons	mixed	410	1.7
$\nu_{all} + \bar{\nu}_{all}$ NC	tracks	171	0.7
$\nu_e + \bar{\nu}_e$ CC	tracks	294	1.2
$\nu_\mu + \bar{\nu}_\mu$ CC	tracks	11517	48.7
$\nu_\tau + \bar{\nu}_\tau$ CC	tracks	162	0.7
muons	tracks	93	0.4

events bleed into the highest $\cos(\theta_z)$ bin from further above the horizon. The KDE kernel is mirrored at $\cos(\theta_z) = -1$ to avoid spurious disappearance of events at the edge. The smeared muon histogram is added to the expectation values from the neutrino MC simulation to estimate the total expectation value in every bin shown in Figure 5.21.

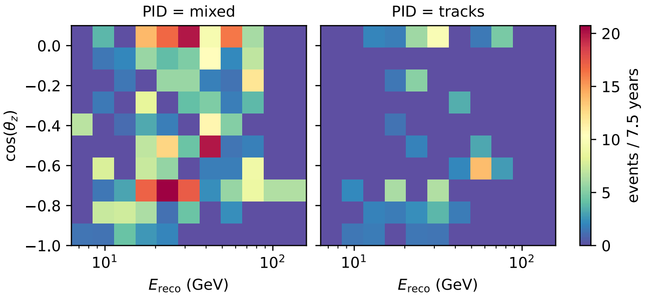

(a) Without KDE smoothing

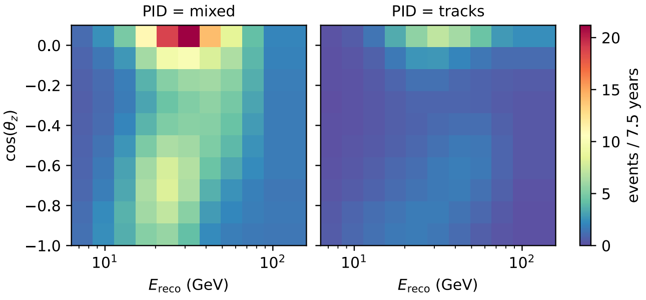

(b) With KDE smoothing

Figure 5.22: Muon template before (top) and after (bottom) the application of KDE smoothing. The shown values are the average of 20 KDE evaluations on different bootstrap samples.

5.2.7 Seasonal Stability

As the operating conditions of the DOMs are very stable after their deployment, the calibration of the DOM response described in Section 5.1.4 is performed only once per year. Such a re-calibration usually also coincides with the release of a new IceCube software package and a new *season* of data taking. To ensure that the data sample is stable under these re-calibrations and software updates, the distributions of the reconstructed energy, zenith angle and PID as well as some control variables are compared between seasons using a Kolmogorov-Smirnov (KS) test[113]. The test calculates the p-value of the largest difference between the cumulative distributions of two samples under the null hypothesis that the samples are drawn from the same distribution. The results for every pair of seasons included in the data sample is shown in heat maps in Figure B.1 in Appendix Section B.1.1. The results show good agreement between the distributions of the different seasons in the sample.

[113]: Kolmogorov (1933), *Sulla determinazione empirica di una legge di distribuzione*

5.3 Implementation of systematic uncertainties

5.3.1 Variation of Detector Properties

Systematic uncertainties on the detector properties that need to be taken into account are the overall optical efficiency of the DOMs as well as the properties of the surrounding ice. The parametrization and priors of each of these properties are informed by IceCube calibration studies.

- DOM efficiency: A factor that scales the probability that a photon hitting the PMT of a DOM will produce a photo-electron that is measured by the electronics. Nominal value is 1, prior standard deviation is 10%.
- Hole ice: Two parameters describe the effect of the optical properties of the column of re-frozen ice within the bore holes in which the strings have been deployed. The details of this parametrization is described below.
- Bulk ice: The over-all absorption and scattering coefficients of all ice layers are multiplied by a scaling factor. The nominal value for ice absorption is 1.0 with a prior standard deviation of 5%. The nominal value for ice scattering is 1.05 with a prior standard deviation of 10%.

In total, the uncertainties on the detector properties are modeled by five parameters, one for the DOM efficiency, two for the hole ice model and two for the bulk ice uncertainty. To model the effect of these parameters on the analysis histogram, several MC sets at different variations of DOM efficiency, hole ice, and bulk ice parameters are produced. These MC sets are used to find a parametrization that will model how the distribution of events in energy, zenith and PID will change as a function of these parameters.

DOM efficiency calibration

As described in Section 4.1.4, the DOMs contain LEDs that are used to calibrate the detector *in-situ*. However, these LEDs are not calibrated with respect to their absolute brightness and therefore are not suited for the calibration of the optical efficiency of the DOMs. Instead, minimally ionizing muons that are produced in air showers are used as a light source with a known brightness. The calibration is performed using a sample of events that pass the *Minimum Bias Trigger*[74] and in which the reconstructed muon track stops within the instrumented volume of IceCube. The DOM efficiency is estimated by comparing the observed charges in the DOMs and the light expectation from the reconstructed muons. Multiple such calibration studies have been run [79][114] and found variations in the optical efficiency of approximately 10%, which is used as a prior for the measurements presented in this work.

[79]: Kulacz (2019), *In Situ Measurement of the IceCube DOM Efficiency Factor Using Atmospheric Minimum Ionizing Muons*

[114]: Feintzeig (2014), *Searches for point-like sources of astrophysical neutrinos with the IceCube Neutrino Observatory*

Hole Ice Parametrization

The bore holes in which IceCube's strings have been deployed were drilled using hot water to melt a column of ice into which the strings with their attached optical sensors could be lowered. This water column re-froze after deployment to form what is referred to as *hole ice*[115]. Camera observations of this re-freezing process suggest that the hole ice is transparent near the edges of be hole and contains a bubble column in its center[116]. The bubble column has a much shorter scattering length than the surrounding bulk ice and therefore decreases the probability of a photon entering a DOM directly from below. The effect of the re-frozen ice column surrounding the strings can be modeled as a modification to the optical efficiency of the DOMs as a function of the incident angle of incoming photons. In the past, many different angular acceptance curves have been produced from *in-situ* calibration measurements[99], the best fit results of previous DeepCore oscillation analyses, in addition to the laboratory measurements that have been made in water tanks before the deployment of IceCube. For the analysis presented in this work, a two-dimensional parametrization was developed that can approximate any of these hole ice models such that it can be used as a *unified* hole ice model. To do this, all previous angular acceptance curves are evaluated as a function of the cosine of the photon incidence angle, $\cos(\eta)$, at 100 points over the entire valid domain between -1 and 1, where 1 represents a photon entering a DOM directly from below. The curves are furthermore normalized to an area of 1 to avoid affecting the total observed charge. Using Principal Component Analysis[117], the variations between the different models are decomposed into a mean and the most important components that explain the variance between models. It was found that the two most important components, p_0 and p_1, describe all known hole ice models adequately. Their effect is shown in Figure 5.23 as variations with the acceptance curve that is used as the baseline in this analysis. The right panel of Figure 5.23 also shows where the older hole ice models are located in the space spanned by p_0 and p_1. The laboratory measurement, which did not include any hole ice effects, notably lies far outside of the region of all other hole ice models that are all produced *in-situ*.

[115]: Fiedlschuster (2019), *The Effect of Hole Ice on the Propagation and Detection of Light in IceCube*

[116]: Rongen (2016), *Measuring the optical properties of IceCube drill holes*

[117]: F.R.S. (1901), *LIII. On lines and planes of closest fit to systems of points in space*

Depth-dependent ice properties

In the parametrization of the uncertainties of the detector properties described in Section 5.3.1, variations of the scattering and absorption coefficients are only described by global, depth-independent scaling factors. In principle, the error on the properties of the ice can also change as a function of depth. Such variations are expected because

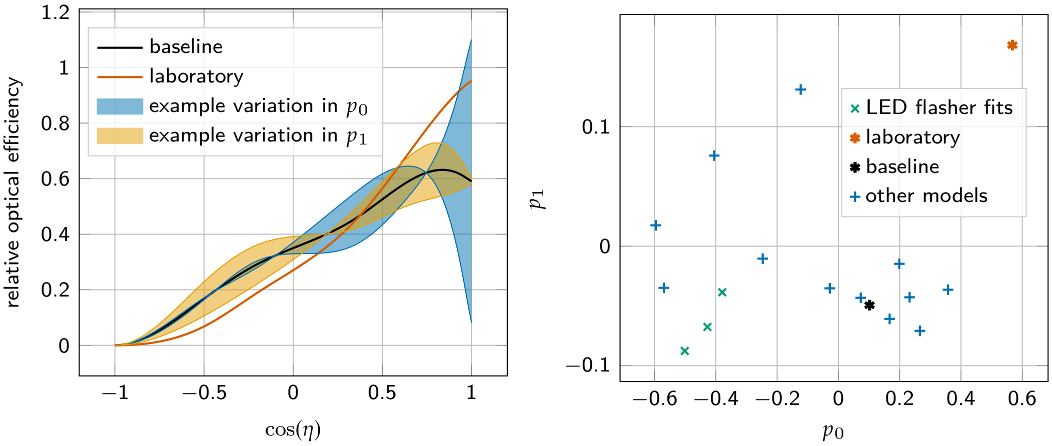

Figure 5.23: Two parameter model used to parametrize the optical efficiency in this analysis (left) and the positions in this two-dimensional space where older hole ice models are located (right). The relative optical efficiency curves are normalized to have the same area, which can lead to acceptance values greater than 1.

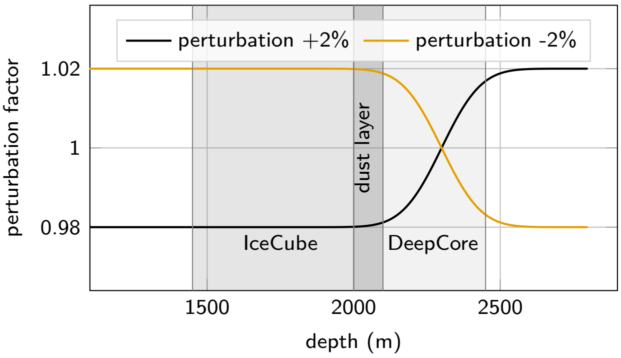

Figure 5.24: Perturbation of the scattering and absorption coefficients with respect to the nominal ice model applied in additional MC sets.

regions of higher absorption and scattering coefficients will also absorb and scatter the light from the LED flashers that is used to do the calibration. Higher uncertainties are also expected near the edges of the detector since there are no more calibration light sources outside of the instrumented volume. Of particular interest for the analysis presented in this work are variations of the ice properties at length scales of the DeepCore fiducial volume located within DeepCore. Variations at much longer scales would be indistinguishable from uniform variations given the size of the event signatures observed below 100 GeV, while variations at much shorter scales are expected to average out. To test how significantly such a variation would impact the final level histograms, two MC sets are produced in which the scattering and absorption coefficients vary following a sigmoid function centered in DeepCore with an amplitude of $\pm 2\%$ in opposing directions as shown in Figure 5.24. The size of this variation corresponds approximately a 1σ-allowed variation according to flasher calibration data. For every bin in the final analysis histogram, a linear regression is fit to the bin counts of the nominal MC set and the two variations. By comparing the χ^2 test statistic resulting from the regression with the free fit and a regression where the slope is fixed to zero, a p-value can be calculated for every bin, where the null hypothesis is that the step-function variation has no effect. The p-values for all analysis bins are shown in Figure 5.25 and are consistent with random fluctuations. Therefore, it was concluded that the effect of a depth-dependent ice model variation

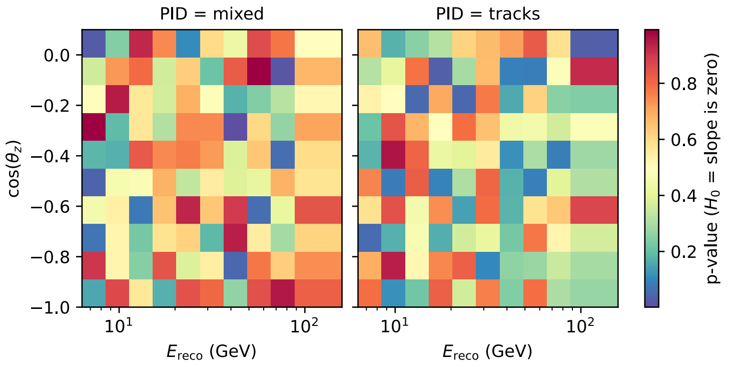

Figure 5.25: Bin-wise p-value of the fitted slopes as a function of the step-function ice model variation.

is well within the statistical uncertainty of the simulation and need not be included in the measurement.

5.3.2 Variation of the Atmospheric Neutrino Flux

The atmospheric neutrino flux can vary depending on the choice of primary cosmic ray (CR) model, assumed meson yield, hadronic interaction (HI) model and atmospheric density model that are used in the calculation. The nominal flux, Φ_{nom}, is modified to a systematic flux, Φ_{sys}, so that

$$\Phi_{\text{sys}}(E) = \Phi_{\text{nom}} \cdot \left(\frac{E}{E_{\text{pivot}}}\right)^{\Delta\gamma} + \sum_{i=1}^{N_{\text{Barr}}} B_i \cdot \frac{\mathrm{d}\Phi_{\text{nom}}}{\mathrm{d}B_i} \tag{5.10}$$

The $\Delta\gamma$ in Equation 5.10 is due to the CR flux uncertainty and corresponds to shifting the spectral index of the neutrino flux, with a pivot point at $E_{\text{pivot}} = 24$ GeV. The second term describes the uncertainty of the Pion and Kaon production yield, where each B_i corresponds to the variation in one *Barr block* (further described below). The gradients with respect to these variations, $\frac{\mathrm{d}\Phi_{\text{nom}}}{\mathrm{d}B_i}$, are calculated using the MCEq[39, 118, 119] flux calculator.

Uncertainty on Meson Production

[120]: Barr et al. (2006), *Uncertainties in Atmospheric Neutrino Fluxes*

The Barr scheme[120] entails dividing the phase space of incident parent particle E_{i} and the outgoing secondary particle E_{s} (or, equivalently, $x_{\text{LAB}} = E_{\text{s}}/E_{\text{i}}$) into regions that are each denoted by a Barr variable. There are eight regions/variables that define the uncertainty on π^+ production, and four regions that define the K^+ production, as shown in Figure 5.26. For every region, a different relative uncertainty is assigned based on the experimental constraints in that region, as shown in Figure 5.27. For primary particle energies > 500 GeV, an additional energy-dependent term is added to the uncertainty to account for the fact that no accelerator measurements are available at these energies to constrain the meson yield. As the pion ratio is well-measured, the uncertainty on π^- is defined by the uncertainty on π^+ combined with the uncertainty on the pion ratio. The uncertainty on K^- production is parametrized separately from the K^+ production. The only modification to the original Barr scheme used in this analysis is that the

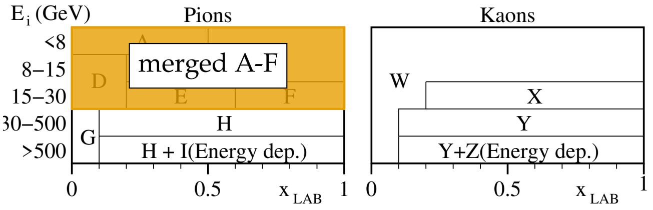
E$_i$ (GeV)
Pions
Kaons
<8
8–15
15–30
30–500
>500
A
merged A-F
D
E
F
G
H
H + I(Energy dep.)
W
X
Y
Y+Z(Energy dep.)
0
0.5
x$_{LAB}$
1
0
0.5
x$_{LAB}$
1

Figure 5.26: Fully correlated regions of uncertainties in the hadronic interaction model. Figure taken from [120].

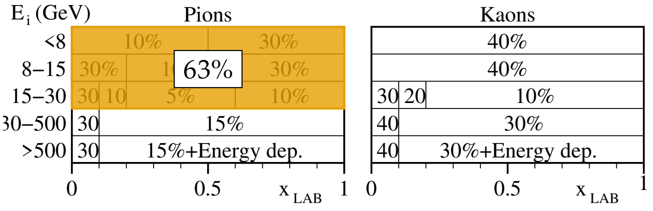
E$_i$ (GeV)
Pions
Kaons
<8
8–15
15–30
30–500
>500
10%
30%
30%
63%
30%
30
10
5%
10%
30
15%
30
15%+Energy dep.
40%
40%
30
20
10%
40
30%
40
30%+Energy dep.
0
0.5
x$_{LAB}$
1
0
0.5
x$_{LAB}$
1

Figure 5.27: Relative uncertainty assigned to each region of hadron phase space in percent. Figure taken from [120].

low-energy π^+ Barr variables A-F are summarized to a single variable with a relative uncertainty of 63%, because their impact was found to be highly correlated. Thus the uncertainty from meson production is described by $N_{\mathrm{Barr}} = 17$ Barr variables that enter Equation 5.10.

Atmospheric density

The development of particle showers in the atmosphere is governed by competing processes of decay and interactions with the surrounding air. The density of the atmosphere can therefore influence the rate of neutrino production and could potentially contribute a systematic uncertainty to oscillation measurements. The size of the effect of atmospheric density uncertainty on the analysis presented in this work is estimated using the same procedure as described in [29]. This is done by obtaining a variation of atmospheric density profile by perturbing the Earth's atmospheric temperature within a prior range given by the NASA Atmospheric InfraRed Sounder (AIRS) satellite [121] temperature data. The resulting atmospheric density profile are injected into MCEQ to calculate new fluxes. This is performed for a variety of CR models and hadronic interaction models available in MCEQ. The resulting fluctuations of the neutrino flux observed at the detector were found to be consistently below 1% for the energy ranges most relevant to DeepCore measurements and is therefore not included as a systematic uncertainty in this work.

Three-flavor oscillation measurement

6

The first measurement made using the data sample described in this work is the measurement of the atmospheric mixing angle θ_{23} and the mass splitting Δm^2_{32}. The experimental setup of DeepCore is ideally suited for this measurement, because the first valley of maximum disappearance for muon neutrinos passing through the entire diameter of the Earth is expected to lie between 20 GeV and 30 GeV as shown in Figure 3.8 in Section 3.3.3. The parameter Δm^2_{32} changes the position of the oscillation valley in energy, while θ_{23} changes its amplitude. In the analysis histogram, this muon neutrino disappearance effect is evident even by just a look in the PID channel for highly track-like events, as shown in Figure 5.21. For this measurement, oscillation probabilities are calculated in the three-flavor oscillation scheme including matter effects. The matter profile of Earth is modeled as concentric shells of a constant density following the Preliminary Reference Earth Model (PREM)[41]. The Monte-Carlo simulated events are weighted in a staged procedure where each stage updates the event weights according to flux, cross-sections and oscillation probabilities[122]. The oscillation probabilities are calculated using a PYTHON implementation of the Barger et al.[123] calculation.

[41]: Dziewonski et al. (1981), *Preliminary reference Earth model*

[122]: Aartsen et al. (2020), *Computational techniques for the analysis of small signals in high-statistics neutrino oscillation experiments*

[123]: Barger et al. (1980), *Matter effects on three-neutrino oscillations*

6.1 Statistical Analysis

6.1.1 Definition of test statistic

To make a measurement, the discrepancy between the histograms of the weighted MC events and the observed data events has to be measured by an appropriate test statistic. This measurement uses a modified χ^2 test statistic defined as

$$\chi^2_{\mathrm{mod}} = \sum_{i \in \mathrm{bins}} \frac{(N_i^{\nu} + N_i^{\mu} - N_i^{\mathrm{obs}})^2}{N_i^{\nu} + N_i^{\mu} + (\sigma_i^{\nu})^2 + (\sigma_i^{\mu})^2} + \sum_{j \in \mathrm{syst}} \frac{(s_j - \hat{s}_j)^2}{\sigma^2_{s_j}}, \qquad (6.1)$$

where N_i^{ν} and N_i^{μ} are the expectation values for neutrinos and atmospheric muons, respectively, and N_i^{obs} is the number of observed events. The expectation value for neutrinos within a bin is calculated as the sum of the neutrino MC event weights $N_i^{\nu} = \sum_i^{\nu\,\mathrm{evts}} w_i$, with the statistical uncertainty due to finite simulation statistics $(\sigma_i^{\nu})^2 = \sum_i^{\mathrm{evts}} w_i^2$. The expectation value for muons, N_i^{μ}, is taken from the KDE-smoothed template shown in Figure 5.22. The variance of the KDE estimate, $(\sigma_i^{\mu})^2$, is calculated from a heuristic described below. The error term due to Poisson fluctuations of the data is calculated with the total MC expectation for muons and neutrinos. The second term in equation 6.1 is included as a penalty term to account for prior knowledge of some systematic parameters.

Muon KDE error estimates

For the three-flavor oscillation analysis, the variance of the muon KDE estimate is calculated using a heuristic based on the theoretical upper bound of the variance of a KDE estimate with a fixed kernel. If the KDE estimate of the density at the point x_0 given n i.i.d. samples from the true PDF p is $\hat{p}_n(x_0)$, the upper bound on the variance is

$$\mathrm{Var}(\hat{p}_n(x_0)) \leq \frac{1}{nh} p(x_0)\sigma_K^2\,, \tag{6.2}$$

where h is the bandwidth and $\sigma_K^2 = \int K^2(y)\,\mathrm{d}y$ with K being the kernel function[124]. The true PDF can be approximated with the estimated PDF to obtain an approximate error estimate. However, the KDE implementation used in this analysis uses a dynamic bandwidth and does not readily give access to σ_K^2 either. For this reason, an ad hoc heuristic is used following

[124]: Li et al. (2007), *Nonparametric Econometrics: Theory and Practice*

$$(\sigma_i^{\mu})^2 = \mathrm{Var}(\hat{p}_n(x_0)) \approx C\frac{\hat{p}_n(x_0)}{n}\,, \tag{6.3}$$

where C is an overall scaling factor that has to be tuned "by hand". To do this, a least-squares fit is run on the systematic MC sets and the value of C is chosen such that the average χ^2/dof from the least-squares fits over all bins is ≈ 1 when n is the number *unweighted* events in the muon MC sample. This heuristic encodes the fact that the variance of the KDE estimate should be proportional to the density and inversely proportional to the number of MC events, where the proportionality factor is tuned to the statistical fluctuations between independent MC sets.

6.1.2 Modeling of Detector Response

For the standard three-flavor fit, the method adopted to model systematic detector uncertainties is an extension to the same method that has been used in previous IceCube oscillation studies[111]. The expectation value of the analysis histogram is calculated using each of the discrete MC sets with perturbed detector properties, and the expectation values in each bin are divided by the expectation given by the nominal MC set. Then, a linear least-squares regression is performed for every bin in the histogram to model the expectation value as a function of all five parameters, resulting in five gradients and one intercept value for every bin. The result of such a fit for an arbitrarily chosen bin is shown in Figure 6.1. A side effect of this process is that the intercept of the linear regression has a much smaller statistical uncertainty than the statistical uncertainty of the nominal MC set alone, which reduces the statistical MC uncertainty of the analysis overall.

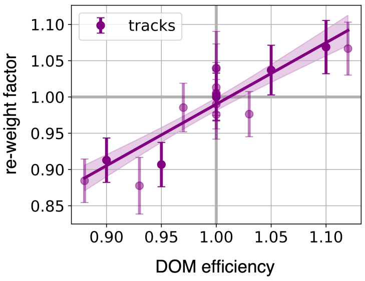

Figure 6.1: Example of a linear regression in one bin of the analysis projected onto the dimension of the DOM efficiency. Data points with translucent error bars originate from MC sets where one or more parameters besides DOM efficiency are at off-nominal points and are projected along the fitted surface to the nominal point.

A fundamental weakness of the linear fit treatment is that the fitted parameters are only valid for the choice of oscillation and flux parameters that they have been fitted with. In particular, there is a strong interaction between the DOM efficiency parameter and the Δm^2_{31} oscillation parameter. As can be seen in Figure 6.2, the gradients of the relative bin counts with respect to DOM efficiency show a distinct imprint

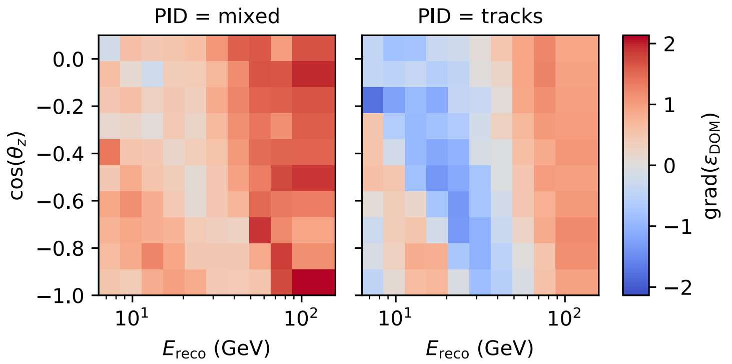

Figure 6.2: Gradient of the relative bin count with respect to DOM efficiency.

of the oscillation minimum. However, the location of this minimum depends on the current value of Δm^2_{31}, which causes a considerable bias if the value of Δm^2_{31} at which the gradients are to be evaluated is different from the value at which they have been fitted. This is mitigated by running the fits at several values of Δm^2_{31} covering the entire plausible range of this parameter, and interpolating all fit parameters with a piece-wise linear function between those points. As a result, all slopes and intercept values change as a function of the mass splitting. Although the interactions between detector systematic uncertainties and analysis parameters are not limited to the mass splitting, the bias produced as a result of the choice of other parameters was found to be much smaller and is therefore neglected in the three-flavor analysis.

6.1.3 Selection of Free Parameters

The systematic uncertainties of the oscillation measurement consists of uncertainties on the properties of the detector, the neutrino flux, neutrino cross-sections, and atmospheric muon background. The effects of each of these uncertainties on the expectation values of the analysis histogram are parametrized by several parameters that are included as nuisance parameters during the fit with Gaussian priors when external constraints are available. The systematic uncertainties of the detector are described by five parameters, namely DOM efficiency, two hole ice parameters and bulk ice scattering and absorption as described in Section 5.3.1. Variations of the atmospheric neutrino flux are parametrized by varying the flux contributions of Pions and Kaons in each "Barr block" (see Section 5.3.2) separately. Neutrino cross-section variations are modeled by varying the axial masses for resonant and quasi-elastic interactions, and by interpolating between the GENIE and CSMS cross-sections for DIS interactions as described in Section 1, with an additional 10% error on the neutral-current contribution to reflect the uncertainty in the hadronization process. For the atmospheric muon background, both the total normalization and the spectral index of the flux can be varied.

When including all plausible sources of systematic uncertainties described above, the test statistic from Equation 6.1 would have to be optimized with respect to 28 nuisance parameters. Together with the two physics parameters, this would require an optimization in 30 dimensions to run the analysis. To reduce this computational burden, the potential bias and its significance that could plausibly be produced

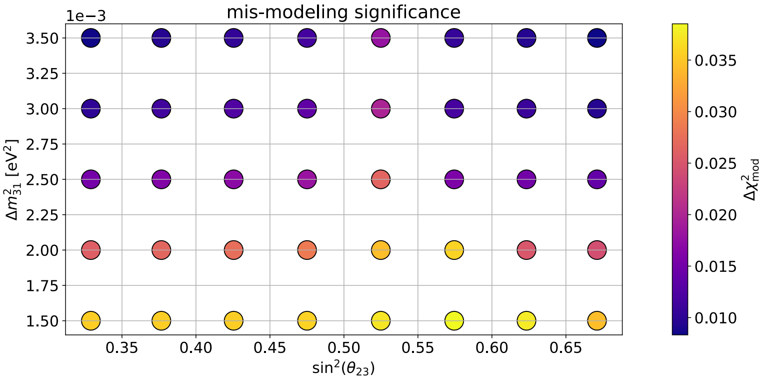

Figure 6.3: Grid of $\Delta\chi^2_{\mathrm{mod}}$ values showing the impact of the axial mass M_A^{CCQE} being pulled by 1σ.

by each parameter is assessed, and the value of parameters that are found to have a negligible impact is fixed to its global best-fit value. The impact of each parameter is tested as follows: First, pseudo-data *without* statistical fluctuations is produced from simulation where the value of the parameter to be tested is increased by 1σ if it has a Gaussian prior, or half-way to its upper boundary if it does not have a prior. The histograms are then fit back while keeping the parameter to be tested fixed at its nominal value. This fit is done once with the physics parameters (θ_{23} and Δm^2_{31}) fixed at the value that was used to create the pseudo-data, and once with the physics parameters left free. The difference in the test statistic χ^2_{mod} between the free fit and the fit with physics parameters fixed to the truth, $\Delta\chi^2_{\mathrm{mod}}$, is referred to as *mis-modeling*. The p-value of the mis-modeling, calculated under the assumption that it should follow a χ^2-distribution with two degrees of freedom, can be interpreted as the significance with which the analysis would have rejected the true physics value *solely* due to the exclusion of the parameter in question. This test neglects any global offset to the test statistic, since it would not affect the estimate of the confidence limits for the physics parameters. The test described above is repeated for a grid of true values of θ_{23} and Δm^2_{31} that span the entire range of values that is not strongly excluded by other measurements, producing a value for $\Delta\chi^2_{\mathrm{mod}}$ at each point in the grid, as shown in Figure 6.3. The largest value of $\Delta\chi^2_{\mathrm{mod}}$ of the entire grid produced for one parameter represents the *maximum mis-modeling* for that parameter. Taking the maximum mismodeling for all parameters, one can produce a ranking of the impacts of all parameters of the analysis as shown in Figure 6.4. Parameters for which the maximum mis-modeling lies below a conservatively chosen value of $\Delta\chi^2_{\mathrm{mod}} < 5 \times 10^{-3}$ are fixed to their global best-fit value in the analysis, reducing the total number of free parameters to 17. The complete list of all parameters of the analysis, their priors and their allowed ranges can be found in Table C.1 in the appendix.

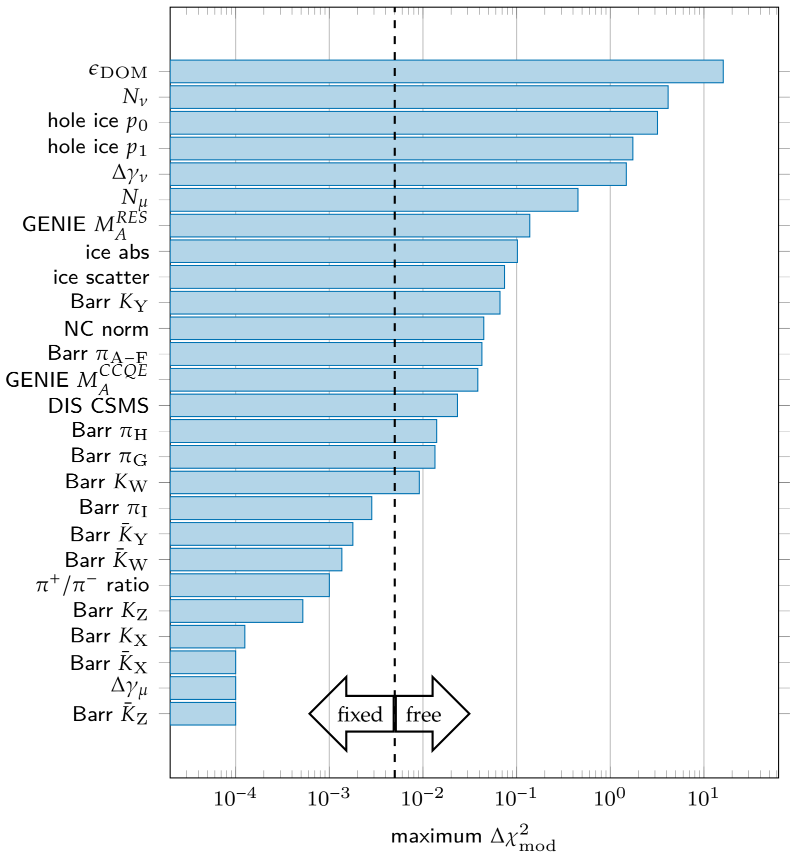

Figure 6.4: Ranking of $\Delta\chi^2_{\mathrm{mod}}$ values for all nuisance parameters considered for the three-flavor oscillation analysis.

6.2 Analysis Checks

Before running the analysis on real data, its robustness is assessed on pseudo-data produced with simulated MC data sets. Once the robustness on pseudo-data has been established, the analysis is first run *blindly*, that is, without showing the analyzer the results of the physics parameters and only revealing a set of goodness-of-fit variables that has been chosen in advance. The actual fit values of the measured parameters are only revealed when the values of these variables lie within the plausible range that can be expected from purely statistical fluctuations.

6.2.1 Robustness of the minimization

The free fit of the physics parameters θ_{23} and Δm^2_{31} is run separately once for the lower octant ($\theta_{23} < 45°$) and once for the upper octant ($\theta_{23} > 45°$) to break the degeneracy between the octants. Each fit uses the SCIPY[125] implementation of the L-BFGS-B algorithm[126] to find the parameter values that minimize the χ^2_{mod} test statistic. To ensure that the minimization will always converge to the global optimum for any true value of the physics parameters, pseudo-data without statistical fluctuations (also referred to as an *Asimov* test set) is produced on a grid spanning all values that are not strongly excluded by other experiments and a fit is run for each grid point. Since there are no statistical fluctuations, the fit is expected to always converge exactly to

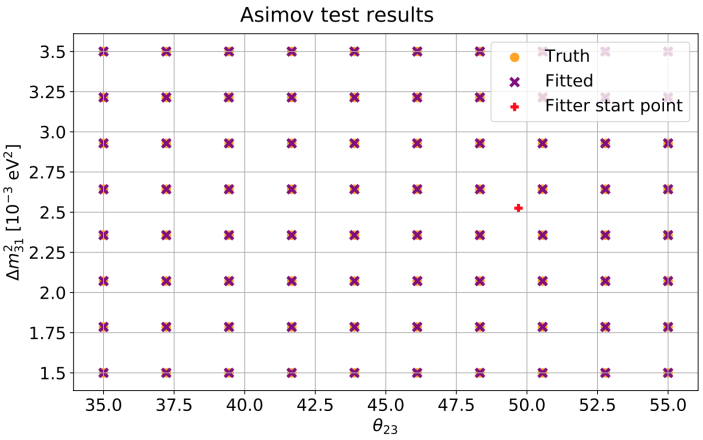

Figure 6.5: Asimov inject/recover test result for the three-flavor oscillation analysis.

the injected true value. As can be seen in the result shown in Figure 6.5, the convergence of the minimizer is robust everywhere.

6.2.2 Ensemble tests

To get expected distributions of the test statistic and parameter fluctuations, the analysis is run on an ensemble of fluctuated pseudodata. For every trial of the ensemble, the expectation value in every analysis bin is first drawn from a normal distribution centered on the MC expectation with a standard deviation corresponding to the MC uncertainty. Using these sampled expectation values, the bin count is drawn from a Poisson distribution independently in every bin. This sampling scheme ensures that the fluctuations reflect both the MC uncertainty and the Poisson fluctuations expected in data. A free fit is run on every trial, producing a set of best-fit parameters, a value for the total $\Delta\chi^2_{\mathrm{mod}}$ test statistic, as well as the chi-square value in each individual bin of the analysis histogram.

Goodness of Fit

Before looking at the best fit parameters of the real data fit, the goodness of fit is assessed using the total and bin-wise test statistic distributions. The distribution of the test statistic acquired from the ensemble described in Section 6.2.2 is shown in Figure 6.6 together with the observed test statistic from real data. The observed test statistic is found to lie very well within the expectation with a p-value of 32%. The bin-wise contribution and the test statistic and its expected distribution are shown in Figure 6.7. The histogram shows no apparent regions of particularly bad agreement between the data and the MC expectation, and the distribution of the bin-wise test statistic is in agreement with the distribution expected from pseudo-data trials.

Test for un-physical mixing

If the real data contains an under-fluctuation in the oscillation valley, it is possible that the fit prefers more than maximal ν_μ disappearance,

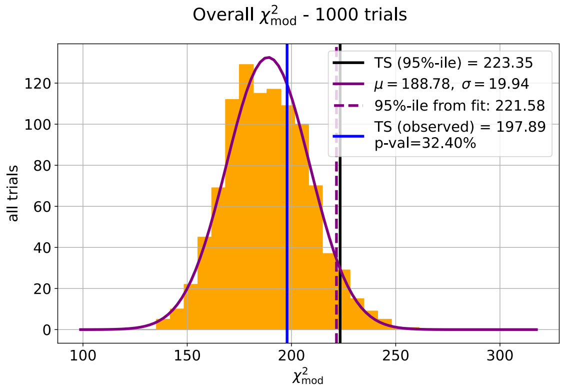

Figure 6.6: Observed test statistic value of the three-flavor oscillation analysis compared to expected distribution from ensemble.

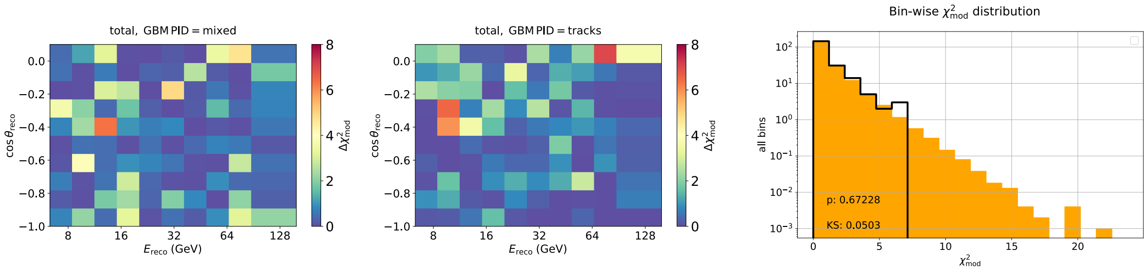

Figure 6.7: Contribution of every bin to the over-all test statistic in the three-flavor analysis (left) and their observed distribution compared to the expected distribution from pseudo-data trials (right).

which is physically not possible. This tendency to fit unphysical magnitudes of ν_μ disappearance is tested by running a fit in which the oscillation probabilities are calculated with a simplified two-flavor model as described in Subsection 3.1.1. Rather than parametrizing amplitudes of the flavor transition probabilities by the mixing angle, however, the scale of the oscillation, $\sin^2(2\theta_{23})$, is replaced with a scaling factor that is allowed to float freely even to unphysical values where $\sin^2(2\theta_{23}) > 1$. If the true mixing angle is $\theta_{23} = 45°$, it is expected that such unphysical best-fit values can occur solely due to random Poisson fluctuations of the data. To quantify this expectation, another ensemble of trials is produced where the injected true mixing is maximal. The two flavor analysis is run on each trial to produce a distribution of expected values that is shown in Figure 6.8. The results show that, while the real data fit does indeed prefer a slightly unphysical ν_μ disappearance, this preference still lies well within the expectation from random fluctuations if the true mixing was assumed to be maximal.

Seasonal Stability

As described in Section 5.2.7, each season of IceCube data taking lasts approximately one year and begins with a re-calibration of the DOM response and, usually, a new IceCube software release. In addition to the stability of the distributions of individual variables shown in that section, the fluctuations of the fit results of the analysis are also examined for signs of inter-season variations.

To this end, pseudo-data trials are generated where the simulated live time in each trial is one year and the three-flavor fit is run on every trial. The trials are produced at the best fit point of the all-season

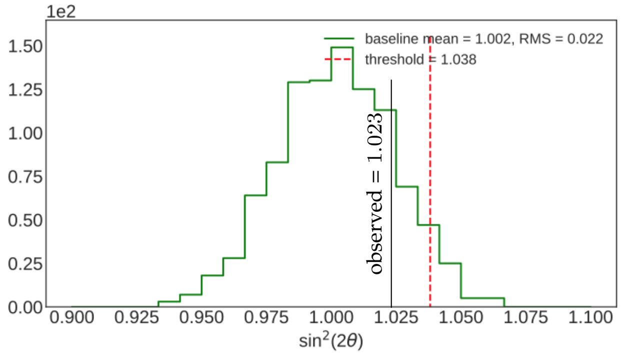

Figure 6.8: Observed best fit values of the two-flavor fit compared to the distribution from pseudo-data trials.

fit in a way that does not reveal the values of these parameters to the analyzers and the fit results of each parameter are placed into a histogram. From these histograms, the p-value for the observed variations of the best fit results for the individual seasons around the all-season fit is calculated under the null hypothesis that they are solely attributable to statistical fluctuations of the data. The test does not show any concerning deviations of any parameter from the all-time best fit result. Heat maps of the associated p-values can be found in Figure B.6 in the appendix.

6.3 Results

To prevent biasing the result due to psychological effects such as the well-known confirmation bias[127], the results of the analysis are revealed by scripts of code that only print out goodness-of-fit variables such as the test statistic and the parameter pulls to the console while hiding the physics result from the analyzers. Only after establishing a good fit of all parameters without any extraordinary pulls, the code reveals the final physics result and likelihood contours that are discussed in this section.

[127]: Nickerson (1998), *Confirmation Bias: A Ubiquitous Phenomenon in Many Guises*

6.3.1 Measured Nuisance Parameter Values

The results for all nuisance parameters are shown in Table 6.1. The pull values show that all parameters fit comfortably within 1σ of their defined priors. This is expected, since the analysis is not itself a statistically powerful measurement of most of the parameters being considered, and therefore the best fit values are pulled towards the center of the prior. The fit prefers a slightly harder cosmic ray spectrum and a larger muon background than initially expected. The optical efficiency of the DOMs fits to a slightly larger value than nominal with 106%, while the ice properties stay very close to their initial values.

The best-fit point of the hole ice parameters p_0 and p_1 is close to the results of the LED flasher fits shown in Figure 5.23. The resulting light acceptance curve is shown in Figure 6.9 compared to the curves produced by *in-situ* LED calibration studies. The result of the three-flavor oscillation fit agrees with the LED calibration result in preferring a lower forward acceptance than the baseline hole ice model.

Table 6.1: Fitted values of all nuisance parameters from the all-season three-flavor fit. The pull of the best fit value is shown for parameters with a defined prior.

category	Parameter	Best Fit Value	Pull (σ)
ν flux	$\Delta\gamma_\nu$	0.065	0.65
	Barr π_{AF}	0.233	0.369
	Barr π_{G}	−0.055	−0.183
	Barr π_{H}	−0.0179	−0.119
	Barr K_{W}	0.0824	0.206
	Barr K_{Y}	0.106	0.355
	Barr $\bar{K}_{\mathrm{W}}$	−0.009	−0.0224
cross-section	M_A^{CCQE}	0.0283	0.0283
	M_A^{CCRES}	0.572	0.572
	DIS CSMS	0.0379	0.0379
	NC norm	1.13	0.633
detector systematics	ϵ_{DOM}	1.06	0.625
	hole ice p_0	−0.269	
	hole ice p_1	−0.041	
	ice absorption	0.974	
	ice scattering	0.989	
norm	N_μ	1.39	
	N_ν	0.824	

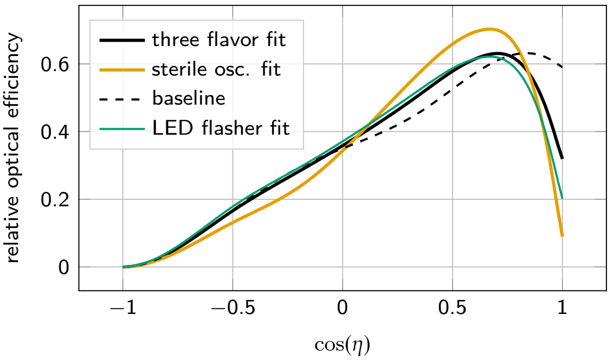

Figure 6.9: Angular acceptance curves corresponding to the best fit point of the three flavor and sterile oscillation fits compared to the results of LED flasher calibration studies.

6.3.2 Oscillation parameters

The fitted values for the three-flavor oscillation parameters are

$$\sin^2\theta_{23} = 0.507^{+0.050}_{-0.053}$$
$$\Delta m^2_{32} = 2.42^{+0.77}_{-0.75} \times 10^{-3}\ \mathrm{eV}^2.$$

The 90% C.L. allowed region for these parameters is shown in Figure 6.10 along with measurements from other experiments. The observed confidence limits for θ_{23} are slightly smaller than what would be expected from a likelihood scan over Asimov pseudo-data that was produced at the best fit point. To make sure that this is compatible with random fluctuations, the likelihood is profiled over $\sin^2(\theta_{23})$ for 100 pseudo-data trials including both Gaussian fluctuations that emulate the MC uncertainty and Poisson fluctuation for the data uncertainty. Figure 6.11 shows the 68% (90%) intervals of the test statistic at each point of the scan over all trials. The observed contour is fully contained in the 68% band, demonstrating that the narrowed 90% range for θ_{23} is fully compatible with expected data fluctuations. The contribution of each category of systematic uncertainties shown in Table 6.1 to the total uncertainty is shown in Table 6.2. The values show that the largest contribution to the total systematic uncertainty of the analysis is by far the uncertainty on the detector properties. Overall, the result provides constraints on the atmospheric mass splitting and mixing angle that are the most stringent in its class of experiments and are competitive with those from accelerator experiments[128–130].

Table 6.2: Contribution of each category of systematic uncertainties to the total error budget in each physics parameter.

	error contrib. (%)	
category	Δm^2_{32}	$\sin^2\theta_{23}$
μ norm	1.8	1.1
ν norm	1.0	0.4
detector	33.6	10.6
ν flux	5.4	1.4
cross-section	6.8	0.3

6.3.3 Post-fit Data/MC agreement

Using the weights of the MC events at the best fit point, the post-fit agreement between data and simulation can be shown for any variable. Of particular interest is the distribution of the argument to neutrino oscillations, L/E, where L is the total distance traveled by a neutrino and E is its energy. Although the true value of the oscillation argument is unknown for data events, it can nevertheless be calculated from the reconstructed energy and zenith angle. The resulting distribution is shown in Figure 6.12 and displays a good agreement between data and simulation, with a reduced χ^2 of close to unity. The comparison between data and simulation for the reconstructed zenith angle and energy gives reduced χ^2 values of 0.782 and 1.19, respectively, when calculated in the same binning.

6.3.4 Likelihood Coverage Test

When drawing the 90% exclusion contour for the oscillation parameters shown in Figure 6.10, it is assumed that Wilks' theorem[132] holds, that is, the distribution of the test statistic follows a χ^2 distribution with two degrees of freedom. If more than 90% of repeated measurements fall below the 90% threshold, the likelihood is said to be *over-covering*. In the reverse case where fewer than 90% of repeated measurements fall below the 90% threshold, the likelihood is said to be *under-covering*. The coverage of the likelihood may change depending on the assumed true parameter values. For this measurement in particular, it is expected that the likelihood should over-cover near maximal mixing, because

[132]: Wilks (1938), *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*

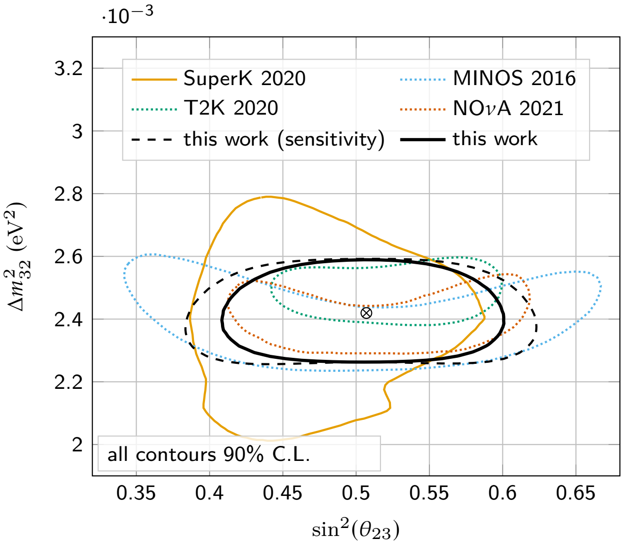

Figure 6.10: Contours showing the 90% C.L. allowed region for the physics parameters of the three-flavor analysis and of other experiments[128–131]. The sensitivity for this work is calculated at the best fit point and the cross shows the best fit value. Other experiments shown in dotted lines are accelerator results, while solid lines are atmospheric oscillation results. All contours shown assume Normal Ordering.

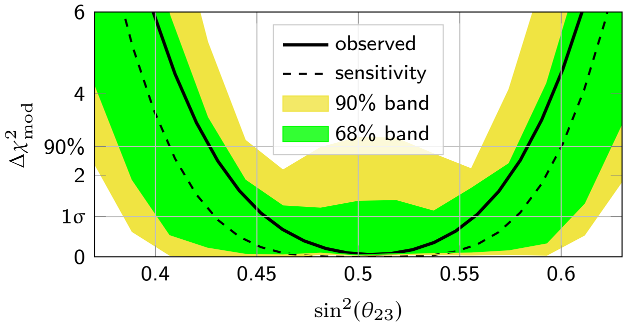

Figure 6.11: Observed contour in $\sin^2(\theta_{23})$ (solid) compared to the Asimov expectation (dashed) and the distribution of 100 pseudo-data trials (yellow and green bands) produced at the best fit point of the three-flavor oscillation analysis. The observed contour is fully contained within 68% of the fluctuations of the trials.

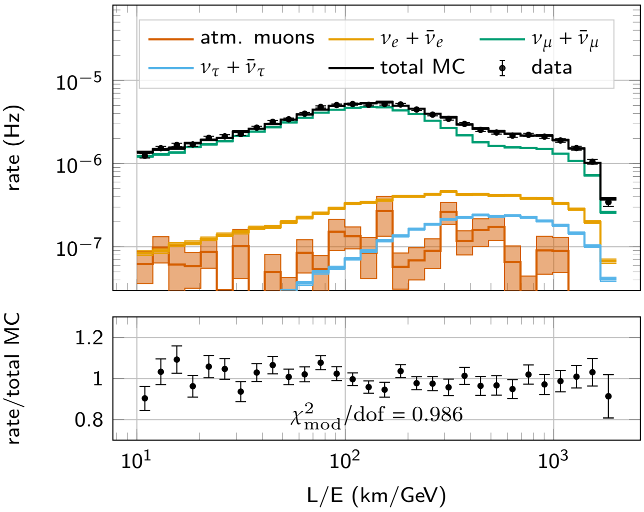

Figure 6.12: Oscillation argument L/E calculated from reconstructed quantities at the best fit point of the three-flavor oscillation analysis.

the mixing angle can no longer provide a full degree of freedom. To test the coverage for particular values of θ_{23} and Δm^2_{31}, pseudo-data is generated where these values are injected as true values. The bin counts of the pseudo-data are Poisson-fluctuated to create an ensemble of trials. Then, one free fit is run and another fit where the physics parameters are fixed to their true values. The coverage is then evaluated by counting the fraction of trials for which $\Delta\chi^2_{\mathrm{mod}}$ between these two fits is smaller than the 90% threshold given by Wilks' theorem. The results are shown in Figure 6.13 for a range of points in mixing angle and mass splitting. As expected, the likelihood is over-covering near maximal mixing, while there is very little dependence of the coverage on the mass splitting. The likelihood is over-covering for all points in mass splitting in the right panel of Figure 6.13, because the injected mixing angle was at the best fit point of the analysis, which is very close to maximal. In conclusion, the 90% contours shown in Figure 6.10 are slightly over-conservative in the region close to maximal mixing.

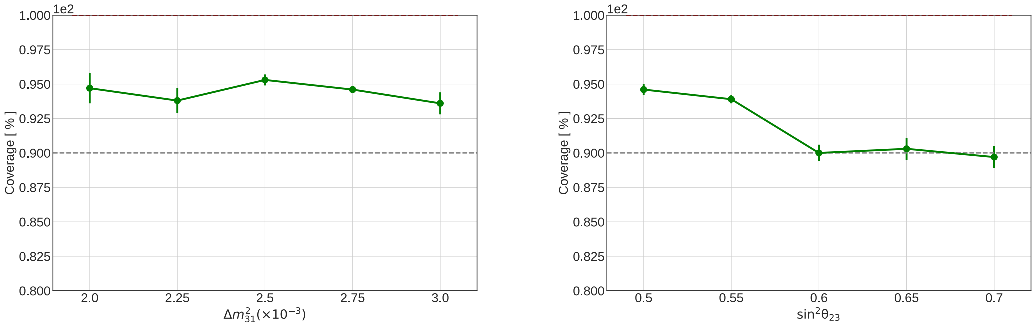

Figure 6.13: Fraction of trials below the 90% threshold expected from Wilks' theorem for a range of points in mass splitting (left) and mixing angle (right).

Search for sterile neutrino mixing

7

7.1 Atmospheric oscillations in the presence of an eV-scale sterile neutrino 91
7.1.1 The 3+1 model 91
7.1.2 Neutrino production height effects 92
7.1.3 Oscillation signal for large mass splittings . . . 92
7.1.4 Oscillation signal for small mass splittings . . 93
7.2 Oscillation Probability Calculation with nuSQuIDS 93
7.2.1 Node placement 94
7.2.2 Production height averaging 94
7.2.3 Low-pass filtering 95
7.3 Nuisance oscillation parameters 97
7.4 Statistical Analysis . . 98
7.4.1 Signal in Analysis Binning 98
7.4.2 Definition of test statistic 99
7.4.3 Modeling of detector response via likelihood-free inference 99
7.4.4 Selection of Free Parameters 104
7.4.5 Likelihood Optimization 105
7.4.6 Analysis Checks 107
7.5 Results 108
7.5.1 Best Fit Parameters 108
7.5.2 Likelihood Scan and Contour 110

This analysis assumes the extended 3+1 neutrino PMNS framework, with the parameters to be constrained being the oscillation parameters θ_{24} and θ_{34}. The three-flavor atmospheric oscillation parameters and the CP-violating phase δ_{24} are treated as nuisance parameters. The mass-splitting of the additional mass eigenstate is fixed at $\Delta m^2_{41} = 1\ \mathrm{eV}^2$. Because the oscillations at Earth-scale baselines occur on much smaller energy scales than can be resolved by DeepCore, the analysis effectively becomes indifferent to the magnitude of Δm^2_{41} and constrains the values of θ_{24} and θ_{34} based on the *averaged* oscillation effect. As a consequence, the constraints calculated based on the assumption that $\Delta m^2_{41} = 1\ \mathrm{eV}^2$ are still valid even if the true mass splitting is much larger. This holds true up to mass splitting values of about $m^2_{41} \gtrapprox 100\ \mathrm{eV}^2$, where the heavy mass eigenstate becomes so much slower than the light eigenstates that it would no longer interfere with them and decohere[25].

7.1 Atmospheric oscillations in the presence of an eV-scale sterile neutrino

7.1.1 The 3+1 model

This analysis probes the ”3+1” oscillation model, in which a fourth ”sterile” (i.e., non-interacting) neutrino flavor eigenstate ν_s and mass eigenstate ν_4 with mass splitting Δm^2_{41} is added to the standard three-flavor model. The PMNS mixing matrix is extended by a fourth row and column, which is parametrized with additional mixing angles θ_{14}, θ_{24}, θ_{34} and CP violating phases δ_{14} and δ_{24} as

$$\begin{aligned} U_{3+1} =& \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & U_{\mu 4} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & U_{\tau 4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{pmatrix} \\ =& R_{34}(\theta_{34})\tilde{R}_{24}(\theta_{24},\delta_{24})\tilde{R}_{14}(\theta_{14},\delta_{14})R_{23}(\theta_{23})\tilde{R}_{13}(\theta_{13},\delta_{13})R_{12}(\theta_{12})\,, \end{aligned} \tag{7.1}$$

where R_{kl} are rotation matrices. The goal of this analysis is to constrain the matrix elements $U_{\mu 4}$ and $U_{\tau 4}$ with magnitude $|U_{\mu 4}|^2 = \sin^2(\theta_{24})$ and $|U_{\tau 4}|^2 = \sin^2(\theta_{34})\cos^2\theta_{24}$, respectively, via the measurement of ν_μ disappearance.

In the presence of an eV-scale sterile neutrino, the standard three-flavor oscillation pattern as a function of neutrino energy and zenith angle is distorted and overlaid with a much faster secondary oscillation pattern. Figure 7.1 shows the muon-neutrino survival probability in the presence of a fourth mass eigenstate with $\Delta m^2_{41} = 1\ \mathrm{eV}^2$ and $\theta_{24} = 15°$ as a function of the energy and cosine of the zenith angle, where $\cos(\theta_z) = -1$ indicates that the neutrino is coming directly from below

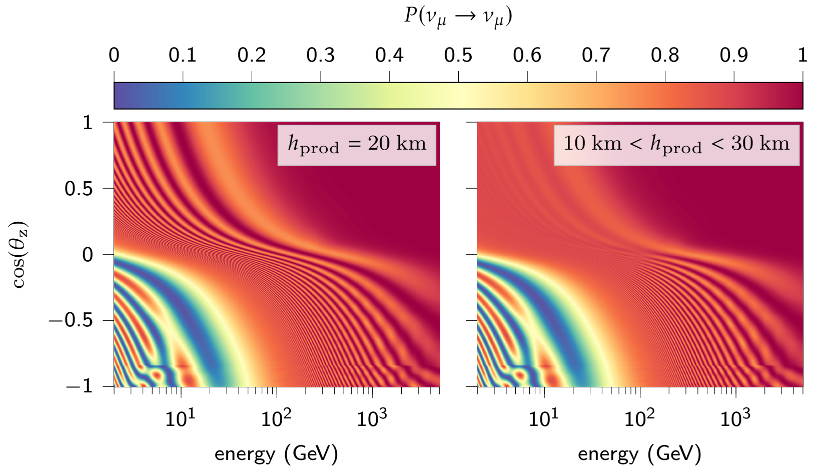

Figure 7.1: Muon neutrino survival probability in the presence of a fourth mass eigenstate with $\Delta m^2_{41} = 1\ \mathrm{eV}^2$ and $\theta_{24} = 15°$ with a fixed production height of 20 km (left) and with production heights averaged between 10 km and 30 km (right).

and $\cos(\theta_z) = 1$ directly from above. For up-going neutrinos, the oscillation pattern induced by Δm^2_{41} is only resolved at energies of $> \mathcal{O}(1\ \mathrm{TeV})$. Below ~200 GeV, the oscillation length reaches values of as low as $\mathcal{O}(\mathrm{km})$, which makes them unresolvable below the horizon where baselines are of $\mathcal{O}(10000\ \mathrm{km})$.

7.1.2 Neutrino production height effects

Above the horizon, the distance from the upper layers of the atmosphere where the neutrinos are produced to the detector is small enough to create a distinct oscillation pattern for a fixed production height, as can be seen in the left panel of Figure 7.1. Because neutrino production heights vary over a range of tens of kilometers, it is necessary to average the oscillation probability over production heights to obtain a more realistic expectation. This is done analytically in nuSQuIDS[133, 134] by calculating the averaged vacuum oscillation probability over a uniform distribution. The right panel of Figure 7.1 shows the oscillation probability with production heights averaged between 10 km and 30 km as describe in Section 7.2.2. The oscillation pattern above the horizon is no longer clearly resolvable, but an average disappearance effect for neutrino energies below 20 GeV remains. The uniform distribution that is assumed to calculate the averaged oscillation probabilities is of course not entirely realistic. For this reason, only events that arrive at most slightly above the horizon ($\cos(\theta_z) < 0.1$) are included in this analysis.

7.1.3 Oscillation signal for large mass splittings

In the mass-splitting range where $\Delta m_{41} \approx \mathcal{O}(1\,\mathrm{eV}^2)$ and in the energy range of the event sample (< 150 GeV), the presence of a sterile neutrino produces rapid oscillations overlaid on the standard three-flavor oscillation pattern as well as distortions to that pattern itself as shown in Figure 7.2 for a mixing angle of $\theta_{24} = 15°$. The oscillation frequency in energy is too large to be resolved by DeepCore, but the average effect still allows us to constrain the magnitudes of $U_{\mu 4}$ and $U_{\tau 4}$. The precise

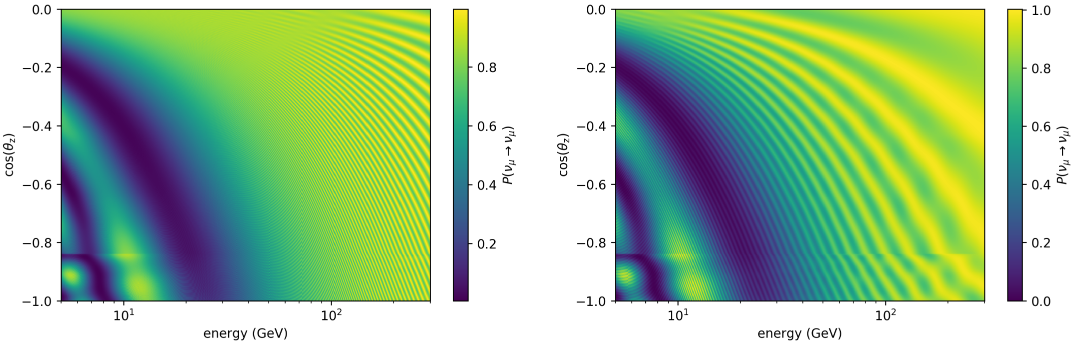

Figure 7.2: ν_μ survival probability at $\Delta m^2_{41} = 1\ \mathrm{eV}^2$ (left) and at $\Delta m^2_{41} = 0.1\ \mathrm{eV}^2$ (right). The mixing angle is $\theta_{24} = 15°$ in both cases

value of Δm^2_{41} has very little influence on the average amplitude of the oscillations and therefore cannot be recovered in this mass splitting regime. The oscillation averages stay approximately constant up to mass splitting values of well above $100\ \mathrm{eV}^2$ where decoherence effects begin to play a role[25].

7.1.4 Oscillation signal for small mass splittings

For mass-splitting values of Δm^2_{41} well below $1\ \mathrm{eV}^2$, the oscillation pattern is no longer completely averaged out. The right panel of Figure 7.2 shows the muon-neutrino survival probability in the presence of a sterile neutrino state with mass splitting $\Delta m^2_{41} = 0.1\ \mathrm{eV}^2$ and mixing angle $\theta_{24} = 15°$. The highest oscillation minimum in energy and cosine of the zenith angle (upper right corner of the figure) is large enough to be resolvable with DeepCore. This makes it possible in principle to produce constraints of the mixing matrix elements as a function of the mass-splitting Δm^2_{41}, although this is beyond the scope of the analysis presented in this thesis.

7.2 Oscillation Probability Calculation with nuSQuIDS

This analysis uses as customized version of the neutrino Simple Quantum Integro-Differential Solver (nuSQuIDS)[133, 134] package to calculate oscillation probabilities. The basic principle behind nuSQuIDS is to calculate the state transition probabilities in the Interaction (Dirac) Picture of quantum mechanics, where the Hamiltonian is divided into the time-independent vacuum oscillation part H_0 and the time-dependent interaction part $H_1(t)$:

$$H(t) = H_0 + H_1(t) \tag{7.2}$$

In this picture, the operators evolve with H_0 as

$$\bar{O}_I(t) = e^{iH_0 t} O_S e^{-iH_0 t}\ , \tag{7.3}$$

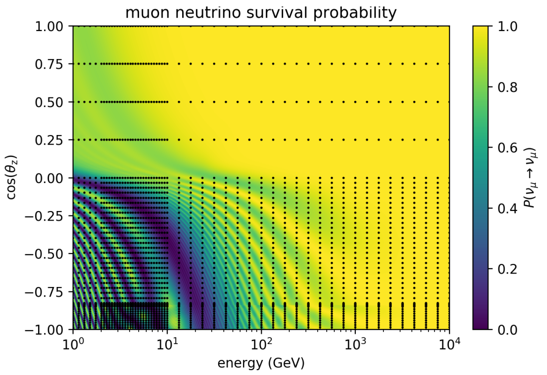

Figure 7.3: Optimized placement of nuSQuIDS nodes (black dots) with extra dense node spacing in the three critical regions. The injected value for Δm^2_{41} is 0.1 eV^2.

while state densities evolve with the interaction Hamiltonian $H_1(t)$

$$\partial_t \bar{\rho}_I(t) = -i[\bar{H}_{1,I}(t), \bar{\rho}_I(t)]\,. \tag{7.4}$$

This state evolution is solved by numerical integration in nuSQuIDS, which is computationally expensive. However, because fast oscillations within H_0 only play a sub-leading role, this difficult calculation does not have to be performed at every point in the analysis space. It is sufficient to calculate state densities at a selected set of points, referred to as *nodes*, and to then interpolate the densities between them. The fast oscillations between the nuSQuIDS nodes are recovered when the probabilities for each flavor, i, are projected out with the trace operation on the state density with the (time-evolved) projection operator for that state:

$$p_i(t) = \mathrm{Tr}(\underbrace{\bar{\Pi}^{(i)}(t)}_{\mathrm{proj.\,op.}}\; \bar{\rho}_I(t)) \tag{7.5}$$

7.2.1 Node placement

The nodes where the difficult state integration is calculated do not need to be spaced densely enough in the analysis space to resolve the fast oscillations due to sterile neutrinos, but they do need to resolve matter effects. The state densities change most rapidly (as function of energy) for neutrinos that traverse a lot of matter at low energies. Additionally, there is a sharp break at $\cos(\theta_{\mathrm{zenith}}) = -0.84$ where neutrinos begin to pass through the core. For this reason, the nuSQuIDS nodes are concentrated in three places:

- the energy region between 2 GeV and 10 GeV
- within a small interval around $\cos(\theta_{\mathrm{zenith}}) = -0.84$
- the region below $\cos(\theta_{\mathrm{zenith}}) = -0.84$

Figure 7.3 shows the optimized placement of nuSQuIDS nodes as black dots.

7.2.2 Production height averaging

At eV-scale mass splittings, oscillations are fast enough that significant oscillations occur even at 10-km-scale distances. If the production height is assumed to be fixed at an exact position, a strong oscillation

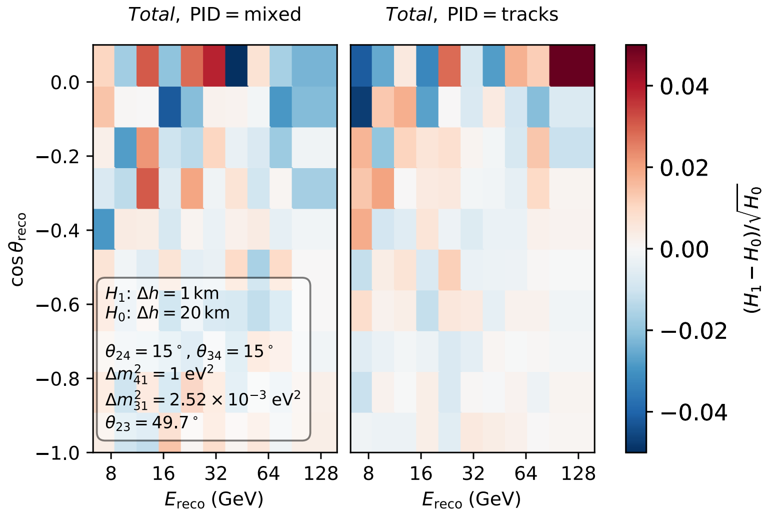

Figure 7.4: Pull introduced into each bin of the analysis histogram when changing the distance over which the production height is averaged from 20 km down to 1 km at sterile mixing angles of $\theta_{24} = \theta_{34} = 15°$.

pattern appears above the horizon that could falsely produce a very high sensitivity in the analysis, driven entirely by events above the horizon. In reality, production heights can vary within the atmosphere, which smears out the oscillations. For this analysis, an analytical averaging method was implemented in nuSQuIDS that assumes a uniform distribution of propagation distances between two points. The start and end point depends on the zenith angle and corresponds to the intersection of the neutrino path with a height of 10 km and 30 km, respectively. Because the bulk of events entering the detector from above the horizon are excluded from the analysis, the precise choice of the lower and upper bound of the height averaging is not very consequential. Figure 7.4 shows the pull in units of the standard deviation of the Poisson fluctuations of the data in each bin that would be expected as a consequence of changing the distance over which the production height is averaged from 20 km down to 1 km. As expected, bins near the horizon are most affected, but even there bin-wise pulls do not rise above 0.05σ.

7.2.3 Low-pass filtering

The fast oscillations in the presence of sterile neutrinos are filtered with a low-pass filter during both state evolution and the calculation of probabilities. This step dramatically increases the speed with which oscillation probabilities can be evaluated.

Vacuum oscillations enter the differential equation governing the state evolution via the time evolution of the interaction Hamiltonian $\bar{H}_{1,I}(t)$. At low energies, they cause tiny but extremely fast oscillations of the time derivative that the numerical integrator has to keep track of by drastically reducing the step size, slowing down the calculation. To mitigate this problem and increase performance, a low-pass filter is applied when calculating the RHS of the differential equation.

In the presence of sterile oscillations, transition probabilities usually have to be calculated for every single simulated event to average them out in the analysis binning. Because the MC set has millions of events,

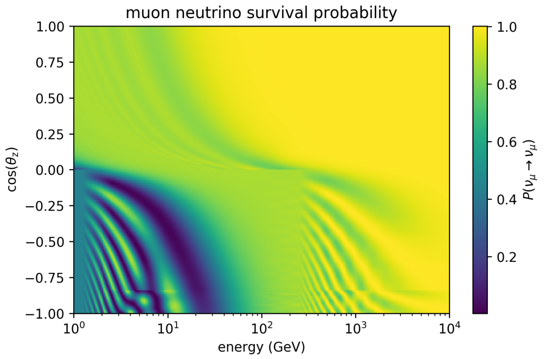

Figure 7.5: Muon neutrino survival probability in the presence of a sterile neutrino after application of both height averaging and low-pass filtering.

doing this would be very expensive even when using nuSQuIDs' state interpolation feature. In this analysis, a low-pass filter as a function of energy is applied when the transition probabilities are projected out from the state densities to average out very fast oscillations. With this filtering, it is possible to calculate oscillation probabilities on a fine binning with 20k bins. One caveat is that this filtering is not appropriate to apply above the horizon because propagation distances there are short enough that oscillation probabilities don't average out completely. Therefore, it is applied only below the horizon as shown in Figure 7.5.

Test of Filter Performance

To ensure that the chosen cut-off points of the low-pass filters do not introduce significant distortions to the oscillation signal, an inject/recover test is performed on a grid in $|U_{\mu 4}|$ and $|U_{\tau 4}|$ where pseudo-data is produced *without* any low-pass filtering and fit back with the filtering enabled. For every grid point, a second fit is run where θ_{24} and θ_{34} are fixed to their true injected values. The p-value of the difference in the test statistic between these two fits, $\Delta\chi^2_{\mathrm{mod}}$, can be interpreted as the significance with which the true value of the mixing angles would be rejected solely due to the mismodeling of the true oscillation probabilities. The results of the mixing parameters fit and the corresponding values of $\Delta\chi^2_{\mathrm{mod}}$ are shown in Figure 7.6. Assuming that $\Delta\chi^2_{\mathrm{mod}}$ should follow a χ^2 distribution with two degrees of freedom, the significance of the mis-modeling is very small. The fit errors for small values of $|U_{\mu 4}|$ and $|U_{\tau 4}|$ are insignificant because the likelihood is very flat in this region of the parameter space. There are a few points for which the fit did not converge to the true global optimum, as indicated by a negative value of $\Delta\chi^2_{\mathrm{mod}}$, but the likelihood difference from the injected true parameters is also very small in these cases.

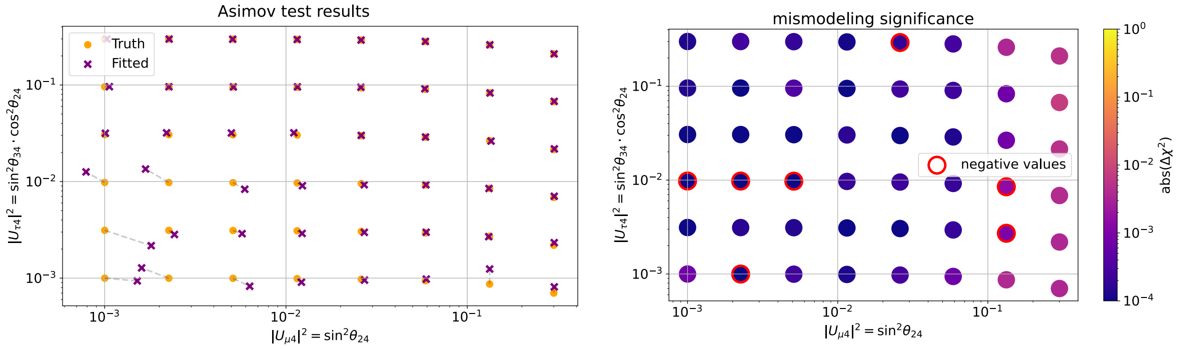

Figure 7.6: Results of the inject/recover test (left) and the corresponding mis-modeling values, $\Delta\chi^2_{\mathrm{mod}}$, attributable to low-pass filtering in NUSQUIDS (right). A negative value of $\Delta\chi^2_{\mathrm{mod}}$ indicates that the fit did not converge to the true global optimum.

7.3 Nuisance oscillation parameters

Besides the physics parameters θ_{24} and θ_{34} to be constrained by the analysis (at a fixed sterile mass splitting Δm^2_{14}), there are 4 additional mixing angles and 3 CP-violating phases in the 3+1 PMNS matrix as well as the mass splittings Δm^2_{12} and Δm^2_{13} that influence the oscillation probability. The solar and reactor angles θ_{12} and θ_{13} as well as the solar mass splitting Δm^2_{12} are constrained by other experiments beyond the sensitivity of this analysis and are fixed at their current global best fit point[24]. The effect of the standard three-flavor CP violating phase $\delta_{\mathrm{CP}} = \delta_{13}$ is negligible and is fixed to zero. The mixing angle θ_{14} and the phase δ_{14} are also fixed to zero, since recent reactor data constrains $|U_{e4}|^2 = \sin^2(\theta_{14})$ to $\mathcal{O}(10^{-3})$, which is well below the sensitivity of this analysis[67]. The only nuisance parameters that remain free are the standard 3-flavor atmospheric oscillation parameters θ_{23} and Δm^2_{13} as well as the sterile CP-violating phase δ_{24}. The effect of δ_{24} is to shift the oscillation pattern of muon neutrinos as shown in Figure 7.7. The effect runs in the opposite direction for *anti*neutrinos. Because neutrinos and antineutrinos are nearly indistinguishable in DeepCore, the combined effect of δ_{24} is a smearing of the oscillation minimum. Additionally, the sign of $\cos(\delta_{24})$ is approximately degenerate with the neutrino mass hierarchy effect. It is therefore expected that the analysis will produce very similar results for NO and IO when δ_{24} is free. Table 7.1 gives an overview over all oscillation parameters in the 3+1 model and their treatment in this analysis.

[24]: Esteban et al. (2019), *Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ23, δCP, and the mass ordering*

[67]: Hu et al. (2021), *Global oscillation data analysis on the 3ν mixing without unitarity*

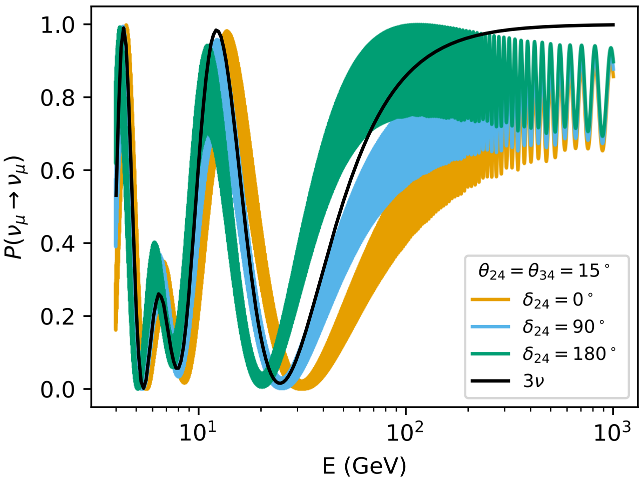

Figure 7.7: Muon neutrino survival probability for a directly up-going neutrino as a function of energy in the presence of sterile neutrinos at different values of the sterile CP violating phase.

Table 7.1: Oscillation parameters of the 3+1 model and their treatment in this analysis.

parameter	nominal value	fixed?	comment
θ_{14}	$0°$	fixed	Constr. by reactor data
θ_{24}	–	free	Physics parameter
θ_{34}	–	free	Physics parameter
θ_{12}	$33.82°$	fixed	Constrained by reactor and solar data
θ_{13}	$8.61°$	fixed	Constrained by reactor and accelerator data
θ_{23}	$49.6°$	free	Atm. mixing angle
δ_{13}	$0°$	fixed	Negligible effect
δ_{14}	$0°$	fixed	No effect when $\theta_{14} = 0°$
δ_{24}	–	free	Smears osc. minimum
Δm^2_{21}	7.39×10^{-5} eV^2	fixed	Constrained by reactor and solar data
Δm^2_{31}	2.525×10^{-3} eV^2	free	Atm. mass splitting
Δm^2_{41}	1 eV^2	fixed	Averaged out above 1 eV^2

7.4 Statistical Analysis

7.4.1 Signal in Analysis Binning

The change in bin counts with respect to the nominal expectation for different combinations of sterile mixing angles at $\Delta m^2_{41} = 1$ eV^2 is shown in Figure 7.8. Although a pull in only θ_{34} by $20°$ has only a very small effect (see the middle panels of Figure 7.8), the combination of both angles can greatly amplify the signal. The CP-violating phase δ_{24} only plays a role when both angles θ_{24} and θ_{34} are non-zero. The sensitivity of this analysis to θ_{34} is entirely due to matter effects experienced by neutrinos passing through the dense core of the Earth. This can be understood from the approximation of the 3+1 oscillation probability under the assumption of very large sterile mass splitting Δm^2_{41} derived in [135] and [136]. To leading order, the vacuum part of the $\mu \to \tau$ oscillation channel does not have any dependence on θ_{34}. When matter effects are considered, however, θ_{34} modifies the effective atmospheric mixing angle in matter, θ^m_{23}, to be further off-maximal.

[135]: Aiello et al. (2021), *Sensitivity to light sterile neutrino mixing parameters with KM3NeT/ORCA*

[136]: Maltoni et al. (2007), *Sterile neutrino oscillations after first MiniBooNE results*

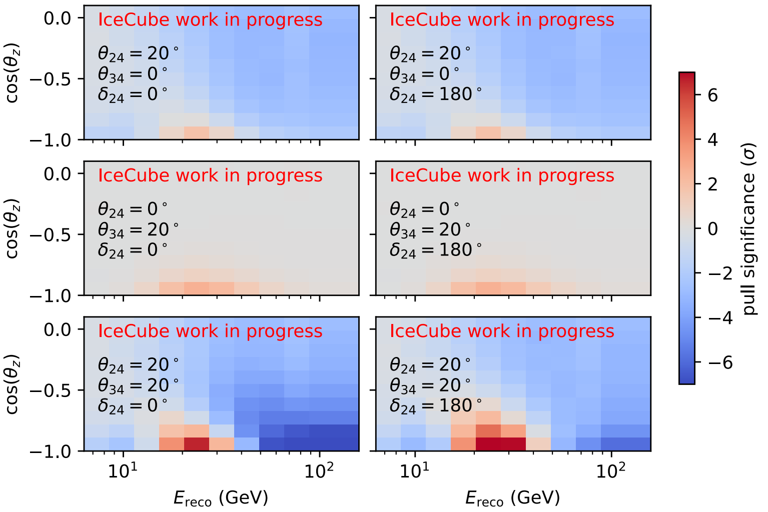

Figure 7.8: Signal in the analysis binning produced by different combinations of θ_{24}, θ_{34}, and δ_{24} as a fraction of the Poisson error in each bin. The mass splitting of the sterile state is $\Delta m^2_{41} = 1\ \mathrm{eV}^2$.

7.4.2 Definition of test statistic

The histograms of the data and the expected values estimated by MC are compared using the same modified χ^2 likelihood described in Section 6.1.1 that is also used for the three-flavor oscillation analysis.

KDE error estimates

Similarly to the three-flavor oscillation analysis, the MC histograms for atmospheric muons are smoothed using a KDE. However, the method of estimating the error on these KDE templates has been changed from the heuristic described in Section 6.1.1 to a more generic bootstrapping method. To estimate the error on the KDE output, 20 bootstrap samples are drawn separately in each PID channel and the KDE is re-evaluated for each trial. The expectation value and the standard deviation of the expectation in each bin of the histogram are the mean and standard deviation of these samples, respectively. The samples are always produced with the same initial random seed to ensure that expectations and errors are reproducible.

7.4.3 Modeling of detector response via likelihood-free inference

As described in Section 6.1.2, the gradients acquired by bin-wise linear regressions in the analysis histogram are only correct for the particular choice of the flux and oscillation parameters at which they have been fitted. For the three-flavor analysis, this can be mitigated by a piecewise interpolation of the gradients as a function of the mass splitting Δm^2_{31}, and the impact of other parameters is small enough that this one-dimensional interpolation is sufficient. However, for the sterile analysis, several additional oscillation parameters have the potential to substantially change the shape of the oscillation pattern and therefore change the detector response in each bin. The grid required to apply the interpolation to all potentially relevant parameters would have to have at least five dimensions and the RAM requirement to hold

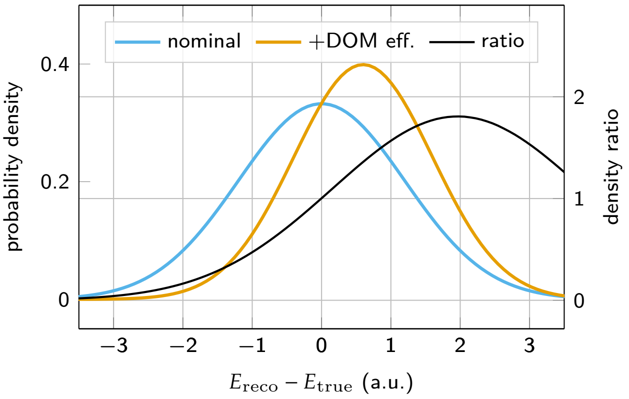

Figure 7.9: Idealized sketch of the effect of DOM efficiency on the distribution of the energy reconstruction error. The blue line shows the distribution of the error in the nominal MC set, the orange line that of events in an MC set with an increased DOM efficiency. The black line shows the ratio between the two PDFs.

all those parameters would have been prohibitive. It was therefore deemed necessary to develop an entirely new method for the treatment of detector systematic uncertainties that completely decouples the detector response from flux and oscillation weights, as described in this section.

Qualitative description

The goal of this treatment of detector systematic effects is to find reweighting factors for every event in the nominal MC set that correspond to how much more or less likely that particular event would be if the detector properties were different. These reweighting factors should be independent from any other event weights, especially flux and oscillation weights. To get an intuition, one may look at the effect of the DOM efficiency on the energy reconstruction error that is illustrated qualitatively in Figure 7.9. The figure shows the probability density function (PDF) of the energy reconstruction error under the assumption of the nominal DOM efficiency (—) and an increased DOM efficiency (—). In the case where the DOM efficiency is higher than nominal, the detector observes more light than is expected from the baseline simulation, and it becomes more likely that the reconstructed energy is larger than the true energy. Given a sample of MC events that drawn from the nominal distribution, the off-nominal distribution can be emulated by re-weighting each event by the ratio (—) of the PDFs. An event where the reconstructed energy is larger than the true energy will get a larger weight, while an event where the reconstructed energy is lower than the true energy will get a smaller weight. Importantly, this relationship will not change depending on the initial flux of the primary neutrino or oscillation effects because the detector reacts only to the final state. For a practical application, more variables than just the energy reconstruction error will influence the PDF of the MC events between different variations of detector properties, and more than one of these properties may change at a time. The qualitative description above can be made more generic by replacing the 1D PDF shown in Figure 7.9 with a multi dimensional one depending on event parameters $\boldsymbol{x}$ that may contain several true and reconstructed variables, and by replacing DOM efficiency with a vector of detector properties, $\boldsymbol{\alpha}_k$, where the index k corresponds to one discrete MC set.

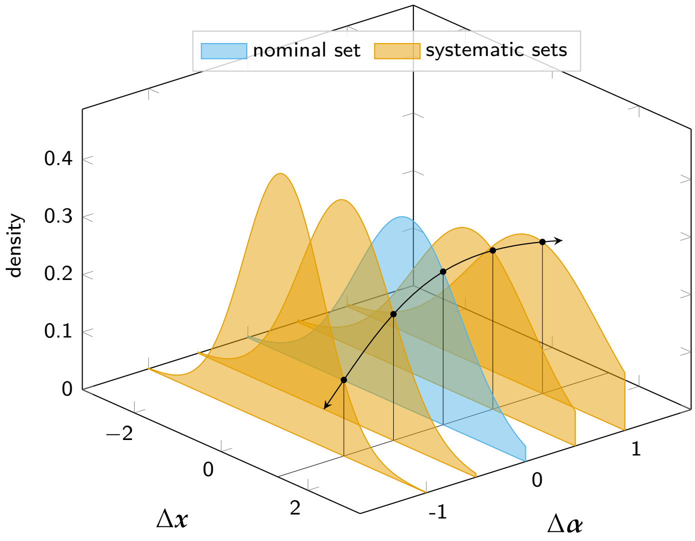

Figure 7.10: Illustration of the re-weighting process used to model changes in the detector response. The axis labeled Δx stands in for the parameters that characterize each individual event such as true reconstructed energy and zenith angle. The axis labeled $\Delta \alpha$ stands in for the detector parameters that vary between systematic sets, such as the DOM efficiency. The distribution of event parameters, x, are shown for the nominal MC set in blue, the off-nominal sets in yellow. The black line shows the function along which an event at a particular x is re-weighted as a function of α.

Quantitative derivation

Using the discrete MC sets with different variations of detector parameters, one needs to find the re-weighting factors, R_{ik}, for each event in the nominal MC set, i, such that re-weighting every event by its weighting factor produces the same distribution in energy, zenith angle and PID as the off-nominal MC set, k. If the true distributions of event parameters in each MC set was known, this weighting factor could be calculated as the ratio

$$R_{ik} = \frac{P(\boldsymbol{x}_i|\boldsymbol{\alpha}_k)}{P(\boldsymbol{x}_i|\boldsymbol{\alpha}_{\text{nominal}})} \frac{P(\boldsymbol{\alpha}_k)}{P(\boldsymbol{\alpha}_{\text{nominal}})} \, , \tag{7.6}$$

where $\boldsymbol{\alpha}_k$ are the detector parameters of the off-nominal MC set, k. The vector $\boldsymbol{x}_i$ contains the true and reconstructed values for energy and zenith angle, as well as the PID of the event. The probability distribution $P(\boldsymbol{x}_i|\boldsymbol{\alpha}_k)$ is the distribution of the event parameters in MC set k, and $\boldsymbol{\alpha}_{\text{nominal}}$ contains the values of the detector parameters of the nominal MC set. The second fraction is the ratio of the total normalization of events under different values of the detector parameters. For example, if an increase in DOM efficiency increases the total number of events by 10%, then $P(\boldsymbol{\alpha})/P(\boldsymbol{\alpha}_{\text{nominal}}) = 1.1$. In practice, of course, the probability distributions $P(\boldsymbol{x}_i|\boldsymbol{\alpha}_k)$ are unknown, but the factors R_{ik} can still be extracted from the available MC sets. The first step is to apply Bayes' theorem to Equation 7.6 to express R_i as the ratio of the posterior probability distribution of the detector parameters given the event parameters,

$$R_{ik} = \frac{P(\boldsymbol{\alpha}_k|\boldsymbol{x}_i)}{P(\boldsymbol{\alpha}_{\text{nominal}}|\boldsymbol{x}_i)} \, . \tag{7.7}$$

The posterior distributions, $P(\boldsymbol{\alpha}_k|\boldsymbol{x}_i)$, can be acquired from a classifier trained to give the posterior probability that an event with parameters $\boldsymbol{x}_i$ belongs to the MC set k. This means that the task of finding the reweighting factors can be translated into a *classification task*. Such an inference method, where probability distributions are learned as a ratio of posteriors from a classifier, is also known as a *likelihood-free inference* method.

K-Neighbors method to calculate posteriors

In principle, any classifier that provides well-calibrated posterior outputs can be plugged into eq. 7.7. For this analysis, the simple and robust *k-neighbors* method is used. The K-Neighbors classifier calculates posterior probabilities by finding the set $\mathcal{N}$ of the N nearest neighbors for every event, i. This set is defined as the set of N events with the smallest Euclidean distance in the event parameters X. Then, the estimate for the posterior for set k is the fraction of the total weight of the neighbors belonging to set k,

$$P(\boldsymbol{\alpha}_k|\boldsymbol{x}_i) = \frac{\sum_{j\in\mathcal{N}\cup k} w_j}{\sum_{j\in\mathcal{N}} w_j} \; . \tag{7.8}$$

The weights w_j are the weighted effective area that every simulated MC event corresponds to and correct for the different amount of MC that was produced for every systematic MC set. They do not include neutrino flux or oscillation effects.

While this method is very robust, it is also prone to over-fitting if the number of neighbors is chosen to be too small. On the other hand, if the number of neighbors is too large, it might blur out important features and under-fit. An increased number of neighbors also causes a systematic bias in the probability estimate due to the fact that the distribution across the selected neighbors is not perfectly uniform. This bias is corrected in linear order by reweighting the events in each neighborhood as described in the Appendix A.1.1. With this bias correction applied, a neighborhood size of 200 per included MC set was found to be a good compromise between bias and overfitting.

Classification variables

The input variables passed into the classifier for each event, $\boldsymbol{x}_i$, need to cover all variables that are used in the binning ($E_{\mathrm{reco}}, \cos(\theta_{\mathrm{reco}}), \mathrm{PID}$) as well as all variables that are used when re-weighting events by flux and oscillation probabilities ($E_{\mathrm{true}}, \cos(\theta_{\mathrm{true}})$). This gives a total of five input variables that are used for classification. Because the K-neighbors classifier calculates euclidean distances between events in these five dimensions to determine which events are neighbors, all dimensions are transformed to be approximately normally distributed as follows and then scaled to have a unit variance:

- Energies are replaced by their logarithm
- The zenith angle is used directly, rather than its cosine
- The PID, which is the probability output of a BDT, is transformed into the log-odds ratio, $\mathrm{LOR} = \log(\mathrm{PID}) - \log(1 - \mathrm{PID})$. This transformed variable turns the pileup of events near a PID value of one into a long tail.

The classifier is fit to the transformed variables separately for each flavor for CC interactions and to the combined set of all NC interactions.

Calculating event-wise gradients

The K-Neighbors calculation produces event weights that can reweight the events of the nominal MC set to imitate the distribution of any other MC set. To be useful in an analysis, however, it is a requirement that this re-weighting can be interpolated to any value of detector parameters between the discrete MC sets. This is accomplished by fitting a vector of gradients, g_i, for every event, minimizing the negative log-likelihood

$$-\log \mathcal{L} = -\sum_k P_{\mathrm{obs}}(\boldsymbol{\alpha}_k|\boldsymbol{x}_i) \log P_{k,\mathrm{pred}}(g_i, \boldsymbol{x}_i) \,, \tag{7.9}$$

where the observed probability is $P_{\mathrm{obs}}(\boldsymbol{\alpha}_k|\boldsymbol{x}_i)$ from the K-neighbors calculation and the predicted probability $P_{k,\mathrm{pred}}(g_i, \boldsymbol{x}_i)$ is

$$P_{k,\mathrm{pred}}(g_i, \boldsymbol{x}_i) = \frac{\exp\left(\sum_j \alpha_{k,j} g_{i,j}\right)}{\sum_l \exp\left(\sum_j \alpha_{l,j} g_j\right)} \,. \tag{7.10}$$

The motivation for the negative log-likelihood loss in Equation 7.9 is that minimizing this quantity is equivalent to minimizing the same cross-entropy between labels P_{obs} and class predictions $P_{k,\mathrm{pred}}$ that is typically used to train neural network classifiers. It has been shown that classifiers that minimize the cross-entropy end up learning posterior distributions[137].

To model nonlinear effects, gradients are fit not only to the five detector parameters, but also to their squared values, for a total of ten gradients per event.

Evaluation

Once the event-wise gradients for all detector uncertainty parameters have been obtained, all events can be easily re-weighted during a fit for any given set of detector parameters, $\boldsymbol{\alpha}$, by multiplying the weight for each event by the ratio

$$R_i(\boldsymbol{\alpha}) = \frac{P(\boldsymbol{\alpha}|\boldsymbol{x}_i)}{P(\boldsymbol{\alpha}_{\mathrm{nominal}}|\boldsymbol{x}_i)} = \exp\left(\sum_j (\alpha_j - \alpha_{j,\mathrm{nominal}}) g_{ij}\right) \,. \tag{7.11}$$

Performance

To verify that the reweighting according to Equation 7.11 gives the expected result, they are used to reproduce each systematic MC set and calculate the bin-wise pulls between the reproduction and the systematic set. When the gradients are correct, the pull between the set and its reproduction,

$$p_n = \frac{N_{\mathrm{reprod},i} - N_{\mathrm{syst},i}}{\sqrt{\sigma^2_{\mathrm{nominal}} + \sigma^2_{\mathrm{syst}}}} \,,$$

should follow a standard normal distribution. Figure 7.11 shows the result of this test for the four MC sets in which only the DOM efficiency is varied between 90% and 110% percent. The spread of the bin-wise

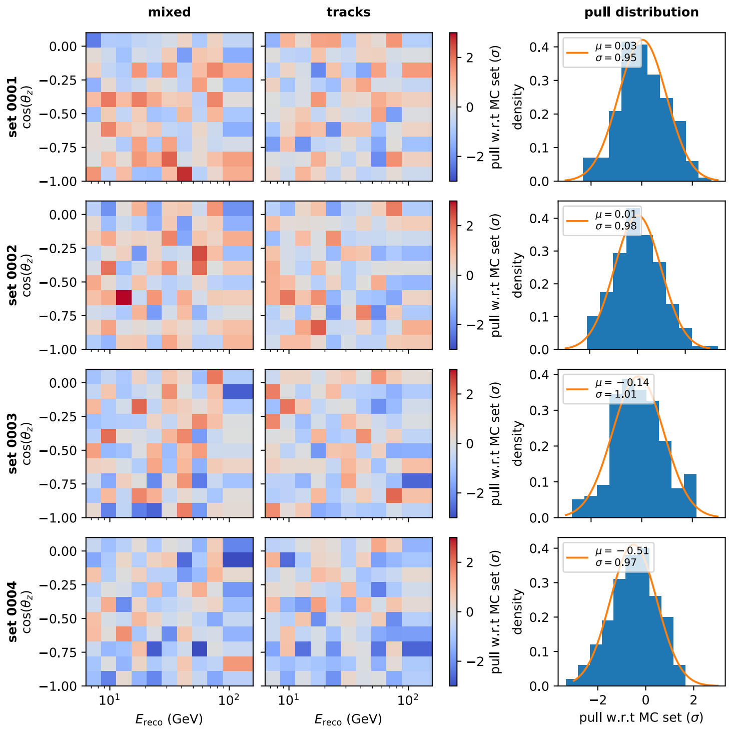

Figure 7.11: Binwise pulls between the nominal set after re-weighting according to eq. 7.11 and the systematic MC sets 0001, 0002, 0003, and 0004 representing DOM efficiency values of 90%, 95%, 105%, and 110%, resepctively. The 1D histogram in each row shows the distribution of the pulls over all bins.

pulls closely follows a normal distribution standard deviation of one, as expected, but the total normalization is slightly under-estimated for the MC set 0004 with DOM efficency of 110%. This is not too concerning for this analysis, since the total normalization is a free parameter without a prior. A similar performance is found for all included MC sets.

Figure 7.12 shows the prediction of the bin count as a function of the DOM efficiency scale, ϵ_{DOM} at the nominal point (left panel) and for different injected values of the mass splitting Δm^2_{31} (right panel) for one arbitrarily chosen bin of the analysis. The prediction matches the shape of the bin count change very well, although it is offset slightly towards lower bin counts. This is expected, since the prediction is based on re-weighting the nominal MC set events without any corrections on the bin count at the nominal point. The error band shown in the figure corresponds to the uncertainty of the nominal set. The right panel of Figure 7.12, demonstrates that the prediction automatically adjusts itself to any injected value of oscillation parameters. This happens despite the fact that the flux and oscillation weights have not been used at all when fitting the event-wise gradients, which demonstrates that the detector response has truly been decoupled from flux and oscillation effects.

7.4.4 Selection of Free Parameters

To reduce the computational cost of optimizing the likelihood in a high-dimensional space, the impact of each nuisance parameter in consideration is assessed and its value fixed if it is found to be negligible. Since the standard three-flavor oscillation model is a nested hypothesis within the 3+1 model, the parameter selection described in Section 6.1.3 is taken as a starting point for the parameter selection of this analysis.

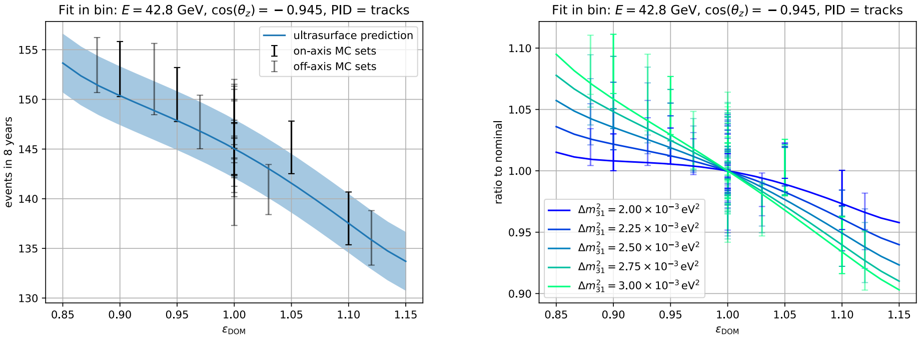

(a) Prediction at best fit point of three-flavor analysis **(b)** Prediction at different mass splitting values.

Figure 7.12: Prediction of bin counts in one bin of the analysis as a function of the DOM efficiency scale, ϵ_{DOM}. The error band in the left panel corresponds to the error on the nominal MC prediction without errors on the event-wise gradients.

Starting from that selection, a test is run to determine if a paraemter is entirely dominated by its prior. For this test, 200 trials are run where all nuisance parameters are sampled randomly according to their prior. Then, Asimov pseudo-data is produced at that point, and this pseudo-data is fit back with the default fit settings. The resulting pairs of true injected parameter values and fitted values for every free parameter are shown in a scatter plot in Figure 7.13. If a parameter is entirely dominated by its prior, it will fit back to its nominal value regardless of the injected value. For parameters where this is the case, their value is fixed to their nominal value during a fit. Parameters to which this applies are framed in red in Figure 7.13. The framed parameters are (from top to bottom and from left to right in the bottom row): Barr K_{W}, Barr $\bar{K}_{\mathrm{W}}$, Barr Π_{G}, $\Pi/\bar{\Pi}$, Barr Π_{H}. The test shown in the figure was run before the priors on $\mathrm{barr\_i\_Pi}$, $\mathrm{barr\_z\_K}$, and $\mathrm{barr\_z\_antiK}$ have been inflated as described in Section 5.3.2. After that change, $\mathrm{barr\_i\_Pi}$ was also added to the set of free parameters of the analysis. The full list of free parameters with their respective ranges and priors is shown in Table 7.2.

7.4.5 Likelihood Optimization

The optimization of the likelihood of the sterile oscillation model is considerably more difficult than that of the standard three-flavor oscillation analysis due to the increased complexity of the parameter interactions. Two additional approximate degeneracies arise in the likelihood space from the sterile mixing angles alone: Firstly, for every best-fit point in $|U_{\mu 4}|^2$ and $|U_{\tau 4}|^2$, there will also be another local optimum where $|U_{\mu 4}|^2$ and $|U_{\tau 4}|^2$ are flipped. In what follows, this is termed the *triangle* degeneracy. Secondly, flipping the sign of $\cos(\delta_{24})$, that is, flipping the *quadrant* of δ_{24} around the angle of $90°$, usually also results in another local minimum. Together with the octant degeneracy of θ_{23}, this means that at least eight local minima have to be checked for every fit. During the development of this analysis, it was found that the likelihood is not entirely convex even within each of these eight distinct segments, leading to a poor performance of local second-order optimizers such as MIGRAD[138]. For this reason, the search for the global optimum in this analysis is done in three steps that are run

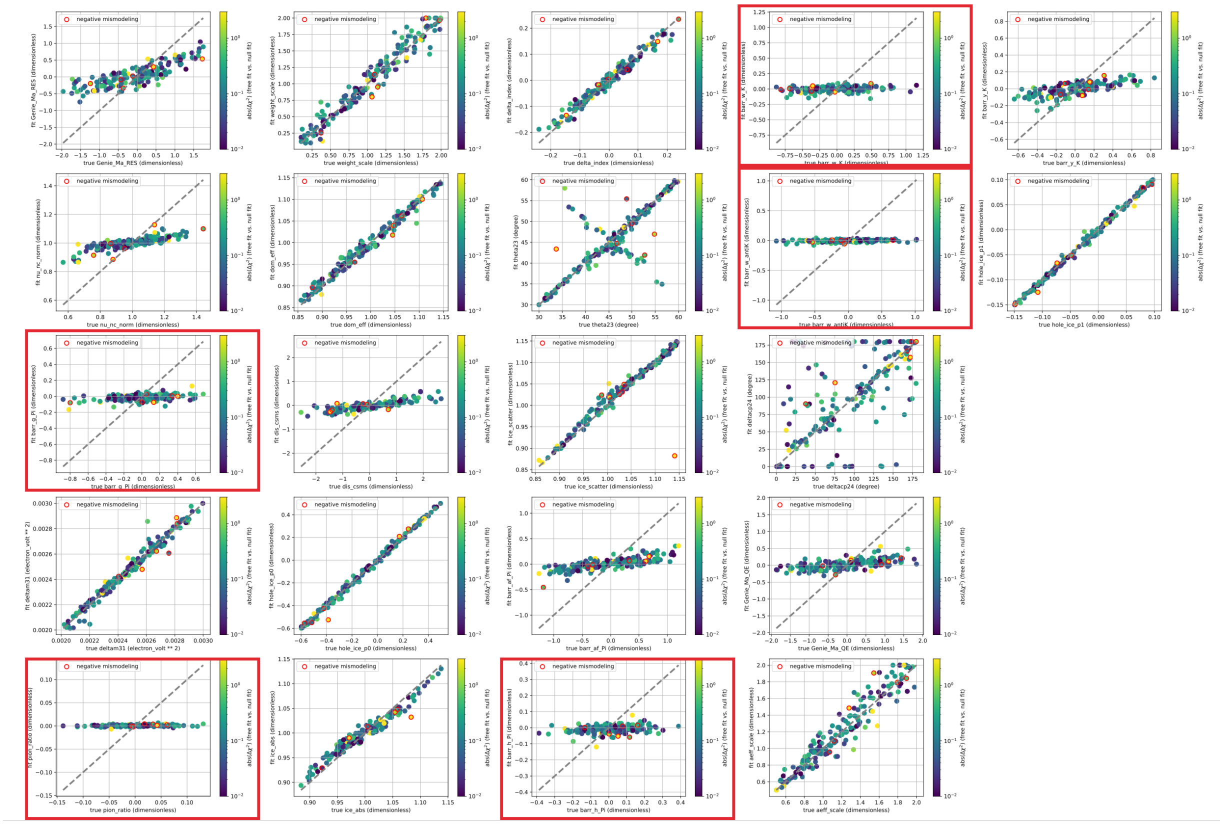

Figure 7.13: Result of the ensemble test with randomly injected nuisance parameters. This test was run before the priors on barr__i__Pi, barr__z__K, and barr__z__antiK have been inflated as described in Section 5.3.2. Parameters framed in red have been deemed to be negligible for the analysis. The color scale shows the likelihood difference between the free fit and a fit in which the physics parameters (θ_{24}, θ_{34} have been fixed to their true value. If this number is negative, the trial is circled in red. In such cases, the minimizer failed to find the correct global optimum in the free fit.

Table 7.2: List of all free parameters in the sterile oscillation analysis with their respective ranges and priors (if applicable).

parameter	nominal value	range	prior
θ_{24}	0°	0° to 45°	uniform
θ_{34}	0°	0° to 45°	uniform
θ_{23}	49.6°	20° to 70°	uniform
δ_{24}	0°	0° to 180°	uniform
Δm^2_{31}	2.525×10^{-3} eV^2	$(2 \text{ to } 3) \times 10^{-3}$ eV^2	uniform
$\Delta\gamma_\nu$	0.0	-3σ to $+3\sigma$	$\sigma = 0.1$
Barr, $a - f_{\pi^+}$	0.0	-3σ to $+3\sigma$	$\sigma = 0.63$
Barr, i_{π^+}	0.0	-3σ to $+3\sigma$	$\sigma = 0.61$
Barr, y_{K^+}	0.0	-3σ to $+3\sigma$	$\sigma = 0.3$
$M_{A,QE}$	0.0	-2σ to $+2\sigma$	$\sigma = 1.0$
$M_{A,res}$	0.0	-2σ to $+2\sigma$	$\sigma = 1.0$
DIS	0.0	-3σ to $+3\sigma$	$\sigma = 1.0$
$N_{\nu,NC}$	1.0	0.5 to 1.5	$\sigma = 0.2$
N_ν	1.0	0.5 to 2.0	uniform
N_μ	1.0	0 to 3	uniform
ϵ_{DOM}	1.0	0.85 to 1.15	$\sigma = 0.1$
ice absorption	1.0	0.85 to 1.15	$\sigma = 0.05$
ice scattering	1.05	0.85 to 1.15	$\sigma = 0.1$
hole ice, p_0	0.101569	-1.1 to 0.5	uniform
hole ice, p_1	-0.049344	-0.15 to 0.1	uniform

separately for every combination of triangle, quadrant and octant. The first step is a global optimization within each segment of the parameter space using the CRS2[139] algorithm with a population of 300 samples. The optimum found by CRS2 is then taken as the starting point for a local, derivative-free optimization using the SbPlx[140] algorithm from the nlopt[141] package. Finally, the best fit point from SbPlx is used as a starting point to "polish" the result using migrad[138].

7.4.6 Analysis Checks

The sterile oscillation analysis follows a procedure similar to that of the three-flavor oscillation analysis, in which the true results for nuisance and physics parameters are only revealed after a good fit quality has been established. Because the data sample used to run the analysis is exactly the same, the checks for post-fit agreement between data and simulation and the seasonal stability need not be repeated. However, the increased complexity of the likelihood space due to the increased number of oscillation parameters requires additional tests to ensure that the true optimum has been reached.

Fit Convergence

The inject/recover test with randomly sampled physics and nuisance parameters described in Section 7.4.4 is used to quantify the reliability with which the correct global optimum is found. To do this, a second fit is run for each trial where the values of θ_{24} and θ_{34} are fixed to their true values. If the test statistic for this fit is smaller than that of the free fit, then the global optimization is proven to have failed to find the correct optimum. Trials for which this is the case are marked in Figure 7.13 with red circles as having negative mis-modeling values. It is apparent from Figure 7.13 that this is the case for a small percentage of trials. The failure rate decreases after the parameters framed in red are fixed. Because it is not possible to generally prove the success of a fit, the failure rate of this test can only provide a lower bound to the true failure rate. The failure rate of the optimization could be reduced by either splitting the likelihood space into even more chunks, or by increasing the number of samples used by the CRS2 search, but only at a significant increase to the already considerable computational cost of the analysis. For the real data fit, the convergence of the likelihood is ensured by running a second fit where the physics parameters are fixed to the best fit point of the free fit. The best fit parameter values and the test statistic of both fits are compared with a script that only prints the differences between them without revealing their absolute values to the analyzer, and the results will only be revealed if they are close to within minimizer precision.

Goodness of Fit

In a similar fashion as for the three-flavor oscillation analysis, the goodness of fit is established using an ensemble of pseudo-data trials with fluctuations taking both the MC uncertainty and the Poisson

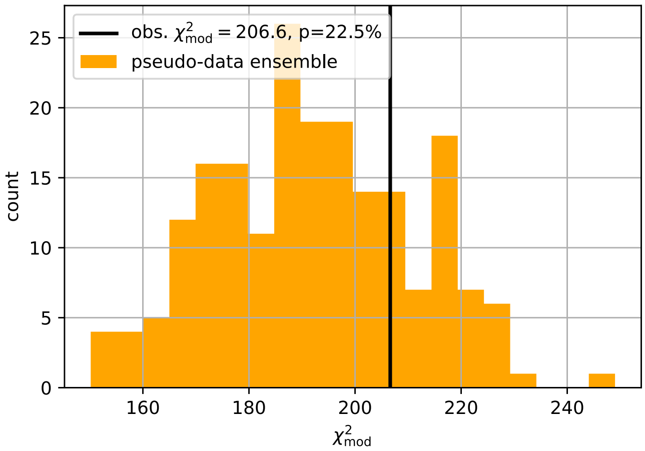

Figure 7.14: Observed test statistic compared to the distribution expected from pseudo-data trials for the sterile oscillation analysis.

fluctuation of the data into account. The pseudo-data expectation is generated by injecting the best fit values of the real data fit as true values. The distribution of test statistic values resulting from this ensemble compared to that of the real data fit is shown in Figure 7.14. The p-value of the observed test statistic is 22.5%, demonstrating that the goodness of fit is well within the expectation.

7.5 Results

7.5.1 Best Fit Parameters

The best fit result for the mixing amplitudes is

$$|U_{\mu 4}|^2 = 0.00449$$
$$|U_{\tau 4}|^2 = 0.00307 \ .$$

Such small mixing amplitudes lie very well below the sensitivity of the analysis and are practically indistinguishable from the null hypothesis, in which there is no sterile neutrino mixing. The value of the sterile CP-violating phase δ_{24} is inconsequential at such low sterile mixing amplitudes. The best fit values of the nuisance parameters are shown in Table 7.3. Both analyses prefer close to maximal mixing in the atmospheric mixing angle θ_{23} and the atmospheric mass splitting fits to a value in the sterile oscillation fit that is only slightly higher than that of the three-flavor fit. Almost all of the remaining nuisance parameters fit into their respective 1σ intervals, except for the scattering coefficient of the bulk ice, which pulls down by 1.5σ of its prior to a value of close to 90%. A comparison to the corresponding best fit parameters of the three-flavor oscillation analysis in Table 6.1 shows that most variables are in agreement in their general trend, but that the magnitude of their deviations from the baseline point is on average larger in the sterile oscillation analysis. For instance, both fits prefer a harder cosmic ray spectrum, a higher DOM efficiency and less forward angular acceptance than nominal (see Figure 6.9), but the effect is more pronounced in the sterile oscillation fit in each case. The best fit point of the sterile analysis also prefers an even higher atmospheric muon background than the three-flavor fit.

Table 7.3: Fitted values of all nuisance parameters from the all-season sterile oscillation fit. The pull of the best fit value is shown for parameters with a defined prior.

category	Parameter	Best Fit Value	Pull (σ)
ν flux	$\Delta\gamma_\nu$	0.091	0.91
	Barr π_{AF}	0.106	0.169
	Barr π_{I}	0.446	0.731
	Barr K_{Y}	−0.0401	−0.134
cross-section	M_A^{CCQE}	−0.05	−0.050
	M_A^{CCRES}	0.0947	0.095
	DIS CSMS	0.301	0.301
	NC norm	1.005	0.024
detector systematics	ϵ_{DOM}	1.08	0.812
	hole ice p_0	−0.589	
	hole ice p_1	0.0408	
	ice absorption	0.988	−0.243
	ice scattering	0.895	−1.546
oscillation	Δm^2_{32}	$2.476 \times 10^{-3}\,\mathrm{eV}^2$	
	$\sin^2\theta_{23}$	0.502	
	δ_{24}	180°	
norm	N_μ	1.93	
	N_ν	0.744	

The oscillation parameters themselves are unlikely to explain these differences, since the sterile oscillation analysis fits very close to the null hypothesis of no sterile mixing. However, there are two major distinctions in the implementation of both analyses in addition to the oscillation model that might be responsible for the observed discrepancies. First, the weighting scheme that implements the effect of systematic uncertainties of detector properties is fundamentally different between both analyses. In the three-flavor analyses, bin-wise linear regressions provided a prediction of the expected bin count for a given set of detector parameters as described in Section 6.1.2. These predictions reduce the statistical uncertainty by incorporating the information from all the MC sets, but depend on the particular choice of parameters at which the linear regressions were performed. In the sterile analysis, each individual MC event is re-weighted using gradients that are fully decoupled from the choice of flux and oscillation parameters that are derived in Section 7.4.3. The downside of this method is that the statistical uncertainty of the MC prediction is increased. It is possible that the coupling of the detector response to the flux and oscillation parameters in the three-flavor fit and the increased statistical uncertainty in the sterile oscillation fit contribute to the different fit outcomes. The second major difference between the fits is the choice of free parameters, in particular with respect to the parameterization of the meson production in the atmosphere with "Barr blocks" described in Section 5.3.2. The sterile analysis allows only a small number of meson production parameters to be varied, but includes the high-energy pion fluctuation in the I block (see Figure 5.26) to vary with an increased prior width. In summary, several factors can contribute to the small discrepancies in the best fit parameter outcomes between the analyses, but none of the observed differences are large enough to be particularly concerning.

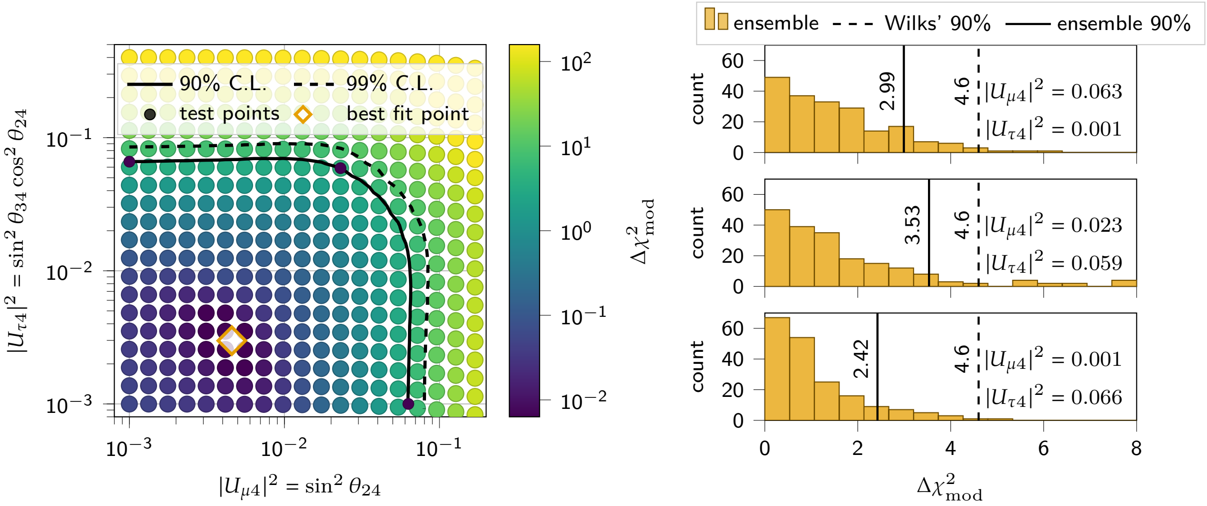

Figure 7.15: Scan of the χ^2_{mod} difference with respect to the best fit point with 90% and 99% C.L. contours assuming Wilks' theorem with two degrees of freedom (left) and ensemble test of the coverage on three points along the 90% C.L. line (right). The test points indicated in the left panel correspond to the points at which the ensemble test were produced in the right panel.

7.5.2 Likelihood Scan and Contour

The likelihood is scanned over $|U_{\mu4}|^2$ and $|U_{\tau4}|^2$ while marginalizing over all other parameters, including the standard three-flavor oscillation parameters and δ_{24}. The degeneracies in θ_{23} and δ_{24} are broken by running four separate fits, one for each possible combination of octant and quadrant. An additional local MIGRAD optimization is run for every grid point of the scan, in which all nuisance parameters start from the global best-fit point. The test statistic value used in the scan is the optimum out of these five fits. The test statistic values on the entire grid and the 90% and 99% C.L. contours assuming Wilks' theorem with two degrees of freedom are shown in Figure 7.15. The marginalized constraints for the matrix elements, calculated assuming that Wilks' theorem holds with one degree of freedom, are

$$\begin{aligned} |U_{\mu4}|^2 &< 0.0534\ (90\%\ \text{C.L.}),\ 0.0752\ (99\%\ \text{C.L.}) \\ |U_{\tau4}|^2 &< 0.0574\ (90\%\ \text{C.L.}),\ 0.0818\ (99\%\ \text{C.L.})\ , \end{aligned} \tag{7.12}$$

setting the strongest limit on $|U_{\tau4}|^2$ to date.

The assumption that Wilks' theorem holds is certainly violated in some parts of the parameter space, especially in regions of very small sterile mixing due to boundary effects and the fact that the CP-violating phase δ_{24} can no longer provide a degree of freedom if either mixing amplitude is zero. The coverage of the assumed χ^2 distribution is tested on three points along the contour to determine if it is over-covering or under-covering using an ensemble of 200 pseudo-data trials, where the point to be tested is injected as the truth. The right panel of Figure 7.15 shows the distribution of the test statistic for each of the points that were tested together with the 90% quantile from the ensemble and the expectation from Wilks' Theorem. In each case, the 90% quantile of the trials lies below the expected value, suggesting that the likelihood is over-covering and that the contour is therefore conservative. As expected, this effect is stronger when one of the mixing amplitudes is close to zero.

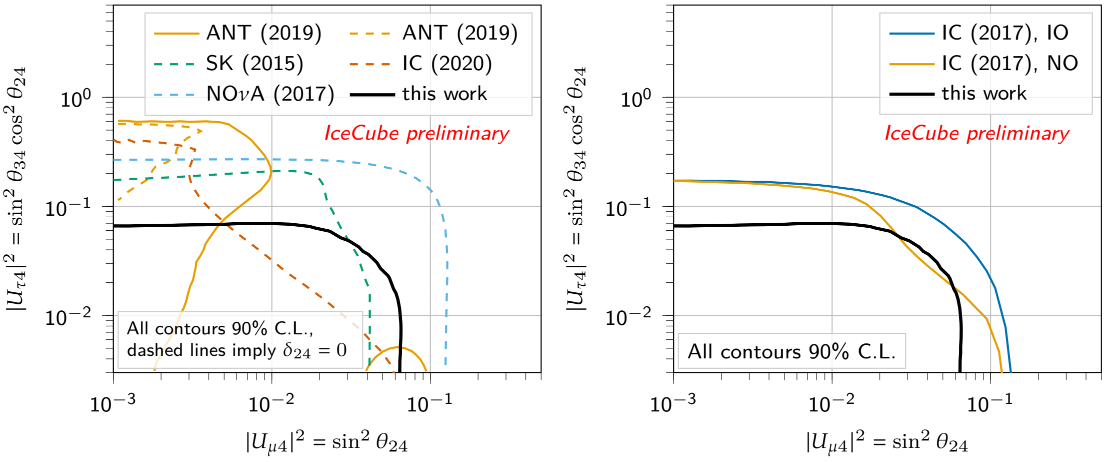

Figure 7.16: Contour of the 90% C.L. of this analysis compared to measurements from the ANTARES[142], Super-Kamiokande[143] and NOνA[144] experiments and a previous high-energy IceCube oscillation study[29] (left), and compared to the previous DeepCore study[109] (right).

Summary and Outlook | 8

8.1 Summary 113
8.1.1 Three-Flavor Oscillation Measurement 113
8.1.2 Sterile Neutrino Search . 114
8.2 Outlook 114
8.2.1 Reconstruction Improvements 115
8.2.2 Ice model 115
8.2.3 Treatment of systematic uncertainties 116
8.2.4 IceCube Upgrade 116

8.1 Summary

This work presented two oscillation measurements of atmospheric neutrinos using the DeepCore sub-array of the IceCube Neutrino Observatory. Both measurements were based on a newly developed data sample of 21 914 highly pure track-like events in the energy range between 5 GeV and 150 GeV. The selection process for these events was described in Chapter 5 and consisted of several filtering steps in which background due to detector noise and atmospheric muons was removed. At the final filter level, the contribution of muon background was reduced to $\sim 2\%$, the noise background was entirely negligible and the sample consisted almost entirely of neutrino interactions. The zenith angle of each event was reconstructed using a modified chi-square fit to the hit times of the observed light under the simplified assumption that they are well described by the Cherenkov light cone of a minimally ionizing muon. This method required a cleaning step prior to the reconstruction that removes hits from photons that are likely to have undergone a significant amount of scattering. Because the zenith reconstruction requires that at least five hits remain after the cleaning procedure, only a subset of all events could be reconstructed this way. However, the stringent hit selection also resulted in a sample of high-quality events.

The energy was estimated with a likelihood that takes into account whether or not a sensor in the array has observed light.

Several quantities produced by the reconstruction algorithms, such as the length of the reconstructed track and the goodness-of-fit, were used as input variables for a Boosted Decision Tree (BDT) that calculates a particle ID (PID) score that estimates the probability of an event to have originated from a charged-current muon neutrino interaction. Data and simulated pseudo-data were binned in zenith angle, energy, and PID. The best-fit neutrino oscillation parameters were then fit by weighting the simulated events to match the histograms of the observed data as closely as possible.

8.1.1 Three-Flavor Oscillation Measurement

The first data analysis shown in this work was a measurement of the atmospheric mixing angle and mass splitting in the three-flavor neutrino oscillation model assuming normal mass ordering. This measurement is complementary to oscillation analyses of accelerator neutrinos and constitutes the most precise measurement using atmospheric neutrinos to date. The result,

$$\begin{aligned} \sin^2\theta_{23} &= 0.507^{+0.050}_{-0.053} \\ \Delta m^2_{32} &= 2.42^{+0.77}_{-0.75} \times 10^{-3}\ \mathrm{eV}^2, \end{aligned}$$

is consistent with previous DeepCore measurements and current global fits.

8.1.2 Sterile Neutrino Search

The second measurement presented in this work was the search for eV-scale sterile neutrinos. The search was performed under the "3+1" model, where the PMNS matrix is extended by an additional row and column to accomodate the mixing of a fourth neutrino mass eigenstate. The measurement used the same data sample as the three-flavor fit and the same likelihood function calculated in an identical binning. The major technical difference between the analyses was that the neutrino oscillation calculation for the sterile neutrino model was done using a customized version of the nuSQuIDS package. This allowed to efficiently calculate flavor transition probabilities in the presence of a heavy fourth mass eigenstate that produces a very fast oscillation pattern. The customizations that were developed specifically for this work were the addition of low-pass filters that can analytically produce oscillation probabilities where the contributions due to that heaviest mass eigenstate are averaged out. Another major technical development that separates the sterile neutrino analysis from the three-flavor fit is the introduction of a novel method of incorporating uncertainties in the detector response in a way that is fully decoupled from neutrino oscillation probabilities.

The analysis constrained the $|U_{\mu 4}|$ and $|U_{\tau 4}|$ elements of the extended PMNS matrix to

$$\begin{aligned} |U_{\mu 4}|^2 &< 0.0534\ (90\%\ \text{C.L.}),\ 0.0752\ (99\%\ \text{C.L.}) \\ |U_{\tau 4}|^2 &< 0.0574\ (90\%\ \text{C.L.}),\ 0.0818\ (99\%\ \text{C.L.})\ , \end{aligned} \tag{8.1}$$

while marginalizing over the CP violating phase δ_{24}. This result is valid for both normal ordering and inverted ordering due to the approximate degeneracy between the mass ordering and the sign of $\cos(\delta_{24})$. Spot-checks of the likelihood distributions showed that these limits err on the conservative side. More stringent limits could be obtained by correcting the critical values of the likelihood according to Feldman and Cousins[145]. However, the computational expense was estimated to be too high and the conservative limits sufficient for the purposes of this work. This result is a substantial improvement over the previous DeepCore result and provides the most stringent limit on $|U_{\tau 4}|^2$ to date. The constraint on $\left|U_{\mu 4}\right|^2$ is competitive with other experiments and has the potential to further increase the tension between appearance and disappearance datasets in global fits of the 3+1 sterile neutrino model described in Section 3.4.4.

8.2 Outlook

The measurement presented in this thesis only used the fraction of the available DeepCore data that could be reconstructed with established reconstruction methods that are optimized for robustness. The purpose of this analysis was not to achieve the highest possible sensitivity, but

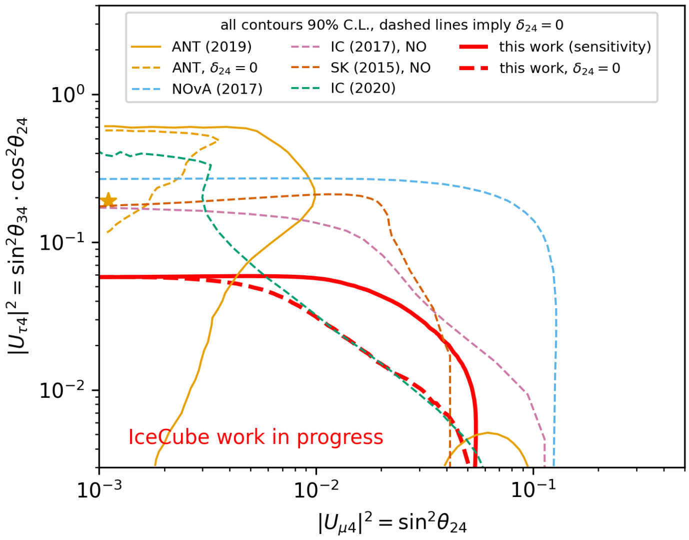

Figure 8.1: Projected sensitivity of the sterile neutrino search when using the table-based reconstruction method.

to verify the integrity of the newly developed methods of data selection and of treating uncertainties in the detector properties. Once the tools that were developed for this analysis are combined with more capable reconstruction methods, this new data sample will provide constraints on oscillation parameters that are more stringent than those of any other atmospheric neutrino oscillation measurement and rival the precision of the most recent accelerator experiments.

8.2.1 Reconstruction Improvements

A substantial increase in the sensitivity of both analyses could be achieved by using the table-based reconstruction method described in [106]. This algorithm can provide an estimate for the energy and zenith angle for nearly all events passing the Level 5 event filter described in Section 5.2 and has a much higher resolution than the reconstruction method used in this analysis, substantially increasing the statistical power of the analysis. The projected sensitivity that could be achieved with this method is shown in Figure 8.1. With the increased statistical power and resolution also comes a larger burden to accurately model the properties of the neutrino flux, particle interactions and detector characteristics. The work to bring data and simulation into agreement for the higher statistics sample is still ongoing at the time of writing this thesis.

8.2.2 Ice model

The major distinction of the IceCube Neutrino Observatory compared to most other neutrino detectors is that the detection medium is the naturally occurring ice at the South Pole, rather than an artificially produced material. The ice consists of many layers with different optical qualities that are the result of varying snow depositions over the last 100 000 years. Understanding the precise properties of the ice at every position in the detector requires complex calibration procedures that are beyond the scope of this work. Recently, calibration efforts within IceCube have led to a new ice model that takes into account the birefringent scattering of light at the boundary of ice crystals[72]. The birefringence bends light rays into the direction of the flow of the glacier, which affects the zenith angle reconstruction for events in DeepCore in particular. Future iterations of DeepCore oscillation studies are expected to include this new ice model in order to achieve a better agreement between data and simulation.

[72]: Abbasi et al. (2022), *In-situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory*

8.2.3 Treatment of systematic uncertainties

Detector uncertainties

In order to perform the sterile oscillation search, a novel method of interpolating between off-nominal MC sets has been developed that allows to weight individual events based on changes in the detector response. The need for this development arose as the preceding uncertainty treatment showed artificial dependencies on the values of physics parameters. In the three-flavor analysis, this problem was addressed by interpolating the gradients of the bin counts with respect to detector parameters on a grid in the mass splitting parameter Δm^2_{31} as described in Section 6.1.2. For the sterile analysis, however, the dimensionality of such a grid made this approach unfeasible. The newly developed method, described in Section 7.4.3, completely decouples the detector response from oscillation parameters and thereby eliminates the need for any interpolation. This property makes this approach universally applicable to any oscillation study probing arbitrarily complex oscillation phenomena. In addition, the method is compatible with unbinned likelihood functions in future analyses, unlike previous methods[146] that are inherently binned. This new method of calculating event-wise weights based on posterior estimates will be expanded upon and has the potential to become the new standard treatment for detector uncertainties for neutrino telescopes.

[146]: Aartsen et al. (2019), *Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube*

Atmospheric flux

The treatment of atmospheric neutrino flux uncertainties used in this work is based on the Barr blocks method that was published in 2006[120]. Since then, new methods of modeling variations of the atmospheric flux have been developed that decrease the overall relative uncertainty by up to 40% and also provide a data-driven parametrization of the flux variations[147]. Incorporating these developments into the oscillation data analysis has the potential to improve the agreement between data and simulation and to increase the sensitivity.

[147]: Fedynitch et al. (2022), *Data-driven hadronic interaction model for atmospheric lepton flux calculations*

8.2.4 IceCube Upgrade

The IceCube Collaboration will deploy seven strings of densely spaced optical sensors within the DeepCore fiducial volume in the near future that will form the *IceCube Upgrade* [148]. The Upgrade will not only be instrumented more densely than the existing DeepCore array, but

[148]: Ishihara (2019), *The IceCube Upgrade – Design and Science Goals*

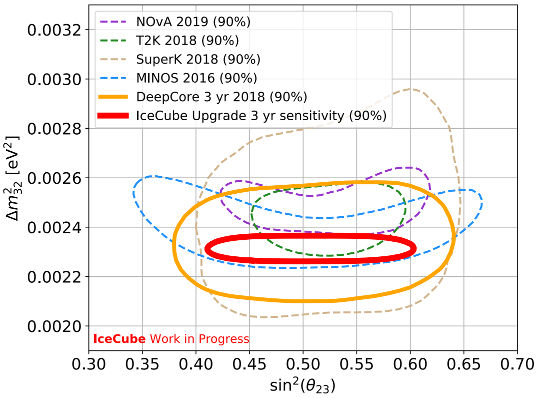

Figure 8.2: Expected sensitivity to the atmospheric neutrino oscillation parameters with three years of Upgrade data compared to recent results from DeepCore and other experiments. Figure taken from [148].

will also contain new types of optical sensors that contain multiple PMTs. This will increase the total photocathode density and add the ability to differentiate the arrival direction of light on each sensor. The new array is expected to lower the energy threshold for the detection of atmospheric neutrinos to $\sim 1\,\mathrm{GeV}$. The Upgrade strings will also incorporate several calibration devices that will further improve the understanding of the ice and detector properties. The sensitivity of the Upgrade array to neutrino oscillation parameters is expected to greatly improve over that of DeepCore as shown in Figure 8.2 for the atmospheric mass splitting and mixing angle.

9

Conclusion

This thesis presented the first results obtained using a newly developed eight-year data sample of neutrinos produced in the atmosphere of the Earth, as detected by the IceCube DeepCore detector. The sample is the result of a collaborative effort to improve the detector calibration and selection process to achieve better agreement between data and simulation. The events used for the measurement were reconstructed using a simple and fast geometric algorithm that reconstructs the direction for very clean, track-like events as well as an energy reconstruction based on a simple hit/no-hit likelihood. With this sample of events, the analysis was able to achieve a good fit to the data and produce a measurement of the standard atmospheric neutrino oscillation parameters θ_{23} and Δm^2_{32}, as well as constraints on sterile neutrino mixing that are the best in this class of experiments.

The analysis also led to important developments of analysis tools and techniques that are directly applicable to studies with higher statistics and improved reconstruction methods. These include a method of interpolating between different Monte Carlo sets in a way that is decoupled from neutrino oscillation effects, as well as several filtering techniques to efficiently calculate neutrino oscillation probabilities in the presence of heavy neutrino states.

The search for sterile neutrinos, which are a potential byproduct of the process that generates neutrino masses and might explain the different mass scale of neutrinos with respect to other SM particles, was performed by probing the ν_μ disappearance probability under the assumption of the minimal "3+1" model with one additional mass splitting of $\mathcal{O}(1\,\mathrm{eV}^2)$. The choice of mass splitting was based on experimental anomalies found in accelerator experiments as well as measurements of the neutrino flux produced in radioactive decay. The null-result of the search further increases the tension between the anomalies observed in the ν_e appearance channel probed in LSND and MiniBooNE and the constraints from ν_e and ν_μ disappearance datasets. Given that this tension already approached the 5σ threshold in global fits, it is unlikely that the simple "3+1" model of sterile neutrino oscillation can explain the anomalies that initially motivated the measurement. The data selection and analysis methods developed during this work will enable future DeepCore measurements with greatly increased statistical power to probe neutrino oscillations for signals of a rich landscape of other possible phenomena beyond the Standard Model.

APPENDIX

A

Mathematical derivations

Detailed mathematical derivations that are too involved for the main text go here.

A.1 Detector systematics via Likelihood-free Inference

A.1.1 Linear correction for KNN bias

The probability estimate from a K-Neighbors classifier with a large number of neighbors shows a systematic bias due to the fact that the distribution of the samples in each neighborhood are not uniform. This bias is corrected by reweighting the samples inside the neighborhood in such a way that their "center of gravity" is located at the point where the KNN is queried. This is done under the assumption that the weight should be a linear function of the coordinate of each sample after subtracting the query coordinate. In one dimension, this leads to the condition

$$\sum_{i \in \mathcal{N}} w_i x_i = \sum_{i \in \mathcal{N}} (1 + g x_i) x_i = 0 \,, \tag{A.1}$$

where the index i runs over all samples in the neighborhood around the query point, w_i is the weight to be assigned to each sample and x_i is the position of the sample relative to the query point. The weight is replaced in the second equality with the linear function $w_i = (1 + g x_i)$ with the gradient g. The condition in Equation A.1 is solved for the gradient as

$$\begin{aligned} & \sum_{i \in \mathcal{N}} (1 + g x_i) x_i = 0 \\ & \Leftrightarrow \sum_{i \in \mathcal{N}} x_i = -g \sum_{i \in \mathcal{N}} x_i^2 \\ \Leftrightarrow & -\frac{\sum_{i \in \mathcal{N}} x_i}{\sum_{i \in \mathcal{N}} x_i^2} = g \,. \end{aligned} \tag{A.2}$$

The gradient g thus found is the ratio of the first and second moments of the distribution of samples in the neighborhood. The non-uniformity of the distribution of events around the query point is corrected to first order by applying the weights $w_i = (1 + g x_i)$ to every event. In higher dimensions, the same calculation is done independently in each dimension and the sample weights for all dimensions are multiplied.

B

Additional Figures

B.1 Event Selection

B.1.1 Seasonal Stability of Variables (KS Tests)

Figure B.1 shows heat maps of the p-value calculated between each pair of seasons used in the data sample underlying this work. The null hypothesis of the test is that the samples of each season are drawn from the same distribution. The top row shows the results for the analysis variables energy, zenith angle and PID, while the bottom row shows the results for three selected control variables.

B.1.2 PID Variables

Histograms showing the distributions of variables used for event signature classification (see section 5.2.5).

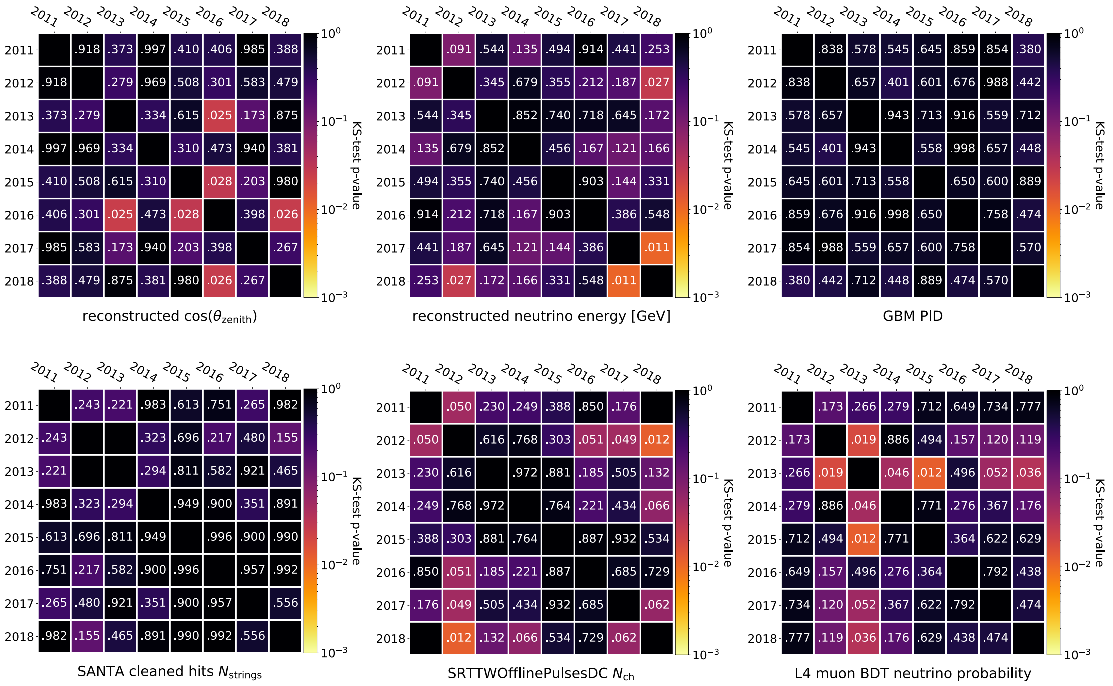

reconstructed cos(θ_{zenith})

	2011	2012	2013	2014	2015	2016	2017	2018
2011		.918	.373	.997	.410	.406	.985	.388
2012	.918		.279	.969	.508	.301	.583	.479
2013	.373	.279		.334	.615	.025	.173	.875
2014	.997	.969	.334		.310	.473	.940	.381
2015	.410	.508	.615	.310		.028	.203	.980
2016	.406	.301	.025	.473	.028		.398	.026
2017	.985	.583	.173	.940	.203	.398		.267
2018	.388	.479	.875	.381	.980	.026	.267	

reconstructed neutrino energy [GeV]

	2011	2012	2013	2014	2015	2016	2017	2018
2011		.091	.544	.135	.494	.914	.441	.253
2012	.091		.345	.679	.355	.212	.187	.027
2013	.544	.345		.852	.740	.718	.645	.172
2014	.135	.679	.852		.456	.167	.121	.166
2015	.494	.355	.740	.456		.903	.144	.331
2016	.914	.212	.718	.167	.903		.386	.548
2017	.441	.187	.645	.121	.144	.386		.011
2018	.253	.027	.172	.166	.331	.548	.011	

GBM PID

	2011	2012	2013	2014	2015	2016	2017	2018
2011		.838	.578	.545	.645	.859	.854	.380
2012	.838		.657	.401	.601	.676	.988	.442
2013	.578	.657		.943	.713	.916	.559	.712
2014	.545	.401	.943		.558	.998	.657	.448
2015	.645	.601	.713	.558		.650	.600	.889
2016	.859	.676	.916	.998	.650		.758	.474
2017	.854	.988	.559	.657	.600	.758		.570
2018	.380	.442	.712	.448	.889	.474	.570	

SANTA cleaned hits N_{strings}

	2011	2012	2013	2014	2015	2016	2017	2018
2011		.243	.221	.983	.613	.751	.265	.982
2012	.243			.323	.696	.217	.480	.155
2013	.221			.294	.811	.582	.921	.465
2014	.983	.323	.294		.949	.900	.351	.891
2015	.613	.696	.811	.949		.996	.900	.990
2016	.751	.217	.582	.900	.996		.957	.992
2017	.265	.480	.921	.351	.900	.957		.556
2018	.982	.155	.465	.891	.990	.992	.556	

SRTTWOfflinePulsesDC N_{ch}

	2011	2012	2013	2014	2015	2016	2017	2018
2011		.050	.230	.249	.388	.850	.176	
2012	.050		.616	.768	.303	.051	.049	.012
2013	.230	.616		.972	.881	.185	.505	.132
2014	.249	.768	.972		.764	.221	.434	.066
2015	.388	.303	.881	.764		.887	.932	.534
2016	.850	.051	.185	.221	.887		.685	.729
2017	.176	.049	.505	.434	.932	.685		.062
2018		.012	.132	.066	.534	.729	.062	

L4 muon BDT neutrino probability

	2011	2012	2013	2014	2015	2016	2017	2018
2011		.173	.266	.279	.712	.649	.734	.777
2012	.173		.019	.886	.494	.157	.120	.119
2013	.266	.019		.046	.012	.496	.052	.036
2014	.279	.886	.046		.771	.276	.367	.176
2015	.712	.494	.012	.771		.364	.622	.629
2016	.649	.157	.496	.276	.364		.792	.438
2017	.734	.120	.052	.367	.622	.792		.474
2018	.777	.119	.036	.176	.629	.438	.474	

Figure B.1: Kolmogorov-Smirnov p-values calculated between each season of data for reconstructed quantities used in the fit (top) and control variables (bottom)

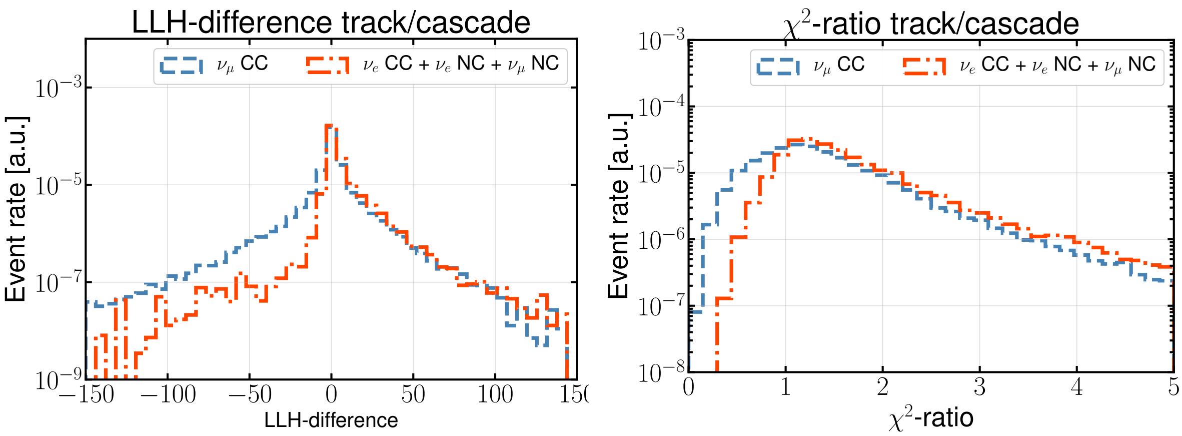

Figure B.2: Histograms of the likelihood score from the energy reconstruction (left) and the goodness-of-fit ratio of the zenith reconstruction (right) in simulation.

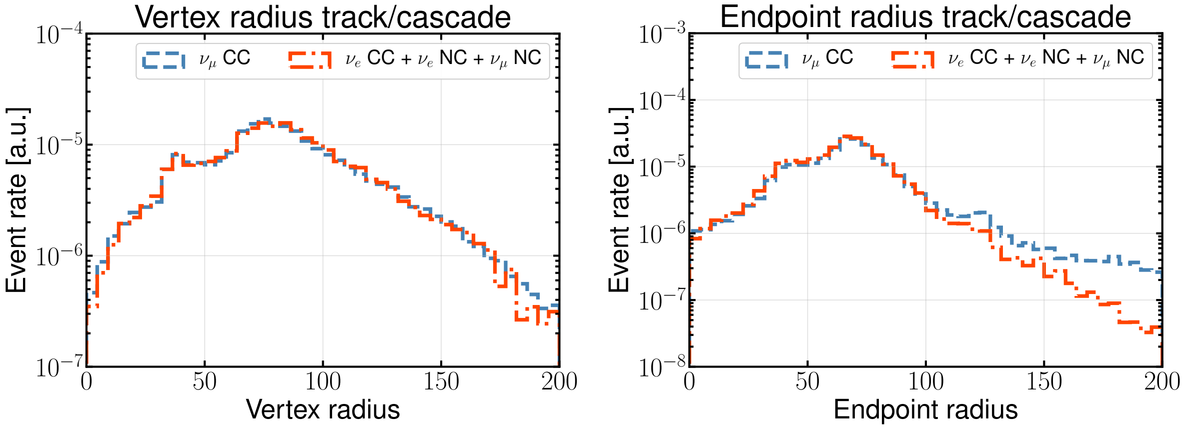

Figure B.3: Histograms of the radius with respect to string 36 of the vertex (left) and the endpoint of the reconstructed track (right) in simulation.

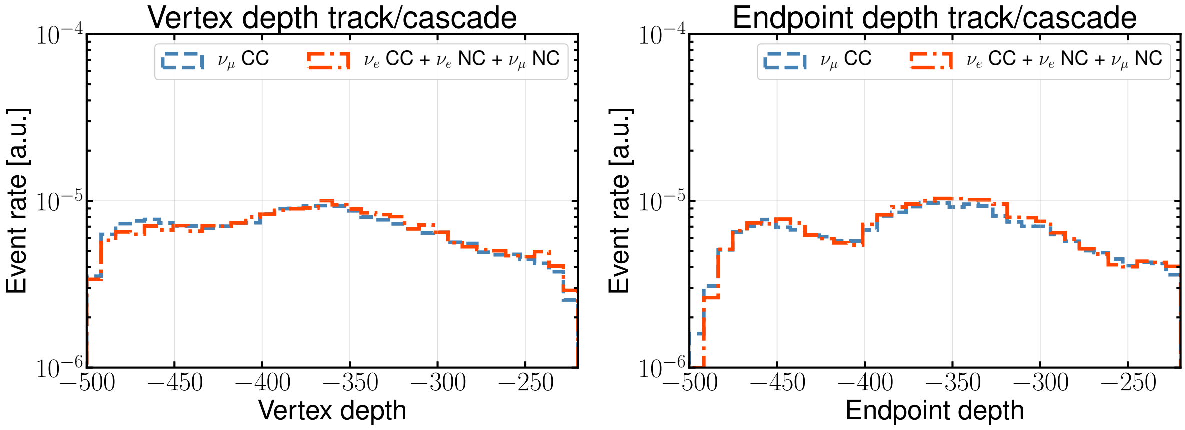

Figure B.4: Histograms of the z-coordinate of the vertex (left) and the endpoint of the reconstructed track (right) in simulation.

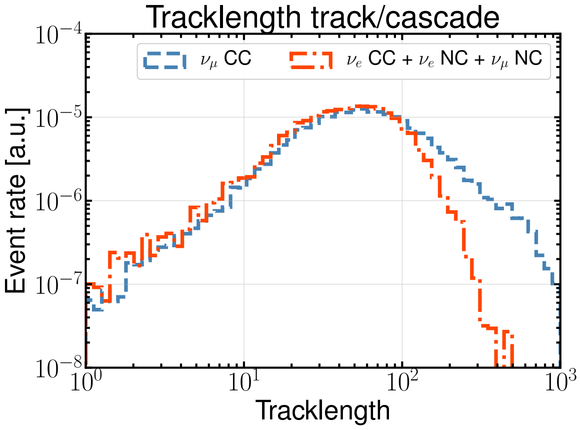

Figure B.5: Histogram of the reconstructed track length in simulation.

B.2 Three-flavor analysis

B.2.1 Seasonal Stability of Nuisance Parameter Fits

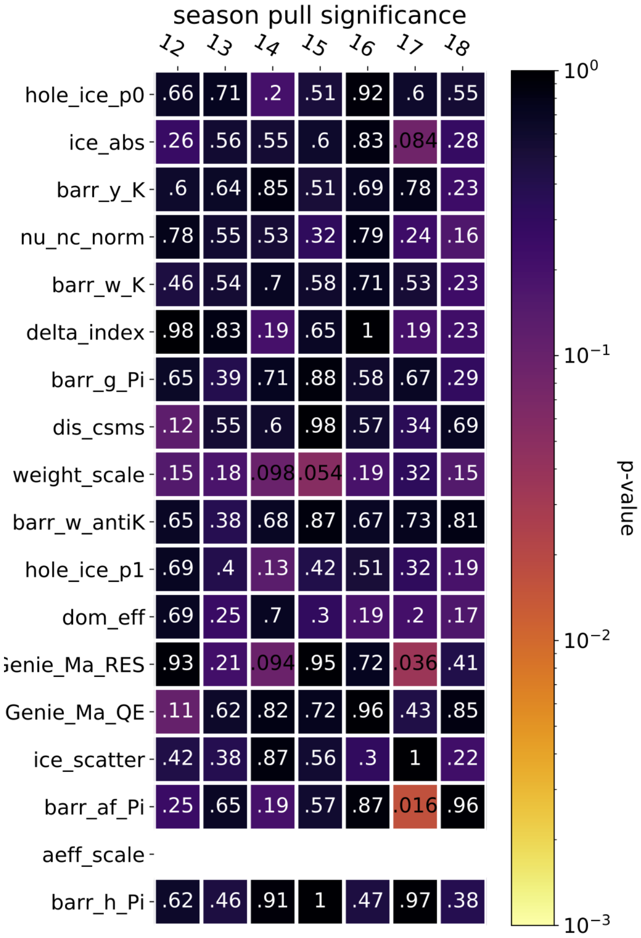

	12	13	14	15	16	17	18
hole_ice_p0	.66	.71	.2	.51	.92	.6	.55
ice_abs	.26	.56	.55	.6	.83	.084	.28
barr_y_K	.6	.64	.85	.51	.69	.78	.23
nu_nc_norm	.78	.55	.53	.32	.79	.24	.16
barr_w_K	.46	.54	.7	.58	.71	.53	.23
delta_index	.98	.83	.19	.65	1	.19	.23
barr_g_Pi	.65	.39	.71	.88	.58	.67	.29
dis_csms	.12	.55	.6	.98	.57	.34	.69
weight_scale	.15	.18	.098	.054	.19	.32	.15
barr_w_antiK	.65	.38	.68	.87	.67	.73	.81
hole_ice_p1	.69	.4	.13	.42	.51	.32	.19
dom_eff	.69	.25	.7	.3	.19	.2	.17
enie_Ma_RES	.93	.21	.094	.95	.72	.036	.41
Genie_Ma_QE	.11	.62	.82	.72	.96	.43	.85
ice_scatter	.42	.38	.87	.56	.3	1	.22
barr_af_Pi	.25	.65	.19	.57	.87	.016	.96
aeff_scale							
barr_h_Pi	.62	.46	.91	1	.47	.97	.38

Figure B.6: P-value of the blind ensemble test to ensure the compatibility of fit results of the three-flavor analysis between seasons. P-values are calculated using histograms of the fit results from an ensemble of pseudo-data where one year of live time is assumed for each trial. Season 2011 is excluded because its live time is smaller.

C

Additional tables

C.1 Three-flavor oscillation measurement

Table C.1: List of systematic parameters considered in the three-flavor oscillation analysis along with their priors and allowed ranges.

Parameter	Nominal Value	Prior	Fixed?
Flux and Cross Section:			
$\Delta\gamma_\nu$	0	±0.1	False
energy pivot	24.09	-	True
pion ratio	0.0	±0.05	True
barr af Pi	0.0	±0.63	False
barr g Pi	0.0	±0.3	False
barr h Pi	0.0	±0.15	False
barr i Pi	0.0	±0.122	True
barr w K	0.0	±0.4	False
barr x K	0.0	±0.1	True
barr y K	0.0	±0.3	False
barr z K	0.0	±0.122	True
barr w antiK	0.0	±0.4	False
barr x antiK	0.0	±0.1	True
barr y antiK	0.0	±0.3	True
barr z antiK	0.0	±0.122	True
A_{eff} scale	1.0	Uniform	False
NC normlisation	1.0	±0.2	False
ν_τ CC norm	1.0	-	True
ν_τ norm	1.0	-	True
M_A^{CCQE} (in σ)	0.0	±1.0	False
M_A^{CCRES} (in σ)	0.0	±1.0	False
dis csms	0.0	±1.0	False
Detector:			
DOM efficiency	1.0	±1.0	False
hole ice p0	0.101569	Uniform	False
hole ice p1	-0.049344	Uniform	False
ice absorption	1.0	Uniform	False
ice scattering	1.05	Uniform	False
Oscillation:			
θ_{12}	$33.82°$	-	True
θ_{13}	$8.61°$	$\pm0.13°$	True
δ_{CP}	0.0	-	True
Δm^2_{21}	$7.39 \times 10^{-5}\ \mathrm{eV}^2$	-	True
Atm. muons:			
$\Delta\gamma_\mu$ (μ spectral index in σ)	0.0	±1.0	True
Weight scale	1.0	Uniform	False

Bibliography

Here are the references in citation order.

[1] C. L. Cowan et al. "Detection of the Free Neutrino: a Confirmation." In: *Science* 124.3212 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103 (cited on page 1).

[2] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos." In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36 (cited on page 1).

[3] K. Kodama et al. "Observation of tau neutrino interactions." In: *Physics Letters B* 504.3 (2001), pp. 218–224. DOI: 10.1016/s0370-2693(01)00307-0 (cited on page 1).

[4] Y. Fukuda et al. "Evidence for Oscillation of Atmospheric Neutrinos." In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562 (cited on pages 1, 3, 29).

[5] Q. R. Ahmad et al. "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory." In: *Phys. Rev. Lett.* 89 (1 June 2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301 (cited on pages 1, 28).

[6] G. Aad et al. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC." In: *Physics Letters B* 716.1 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020 (cited on page 3).

[7] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Mar. 2007 (cited on pages 3, 4, 20).

[8] M. D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2013 (cited on pages 3, 5).

[9] S. L. Glashow, J. Iliopoulos, and L. Maiani. "Weak Interactions with Lepton-Hadron Symmetry." In: *Phys. Rev. D* 2 (7 Oct. 1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285 (cited on page 8).

[10] Z. Maki, M. Nakagawa, and S. Sakata. "Remarks on the Unified Model of Elementary Particles." In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. DOI: 10.1143/PTP.28.870 (cited on page 8).

[11] M. Kobayashi and T. Maskawa. "CP-Violation in the Renormalizable Theory of Weak Interaction." In: *Progress of Theoretical Physics* 49.2 (Feb. 1973), pp. 652–657. DOI: 10.1143/PTP.49.652 (cited on page 8).

[12] E. Majorana. "Teoria simmetrica dell'elettrone e del positrone." In: *Il Nuovo Cimento (1924-1942)* 14.4 (2008), p. 171. DOI: 10.1007/BF02961314 (cited on page 8).

[13] A. M. Abdullahi et al. "The Present and Future Status of Heavy Neutral Leptons." In: *2022 Snowmass Summer Study*. Mar. 2022 (cited on page 9).

[14] R. N. Mohapatra and G. Senjanovi ć. "Neutrino Mass and Spontaneous Parity Nonconservation." In: *Phys. Rev. Lett.* 44 (14 Apr. 1980), pp. 912–915. DOI: 10.1103/PhysRevLett.44.912 (cited on page 10).

[15] Y. Cai et al. "From the Trees to the Forest: A Review of Radiative Neutrino Mass Models." In: *Frontiers in Physics* 5 (2017). DOI: 10.3389/fphy.2017.00063 (cited on page 10).

[16] S. M. Carroll. "The Cosmological constant." In: *Living Rev. Rel.* 4 (2001), p. 1. DOI: 10.12942/lrr-2001-1 (cited on page 11).

[17] Planck Collaboration et al. "Planck 2018 results - VI. Cosmological parameters." In: *A&A* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910 (cited on pages 11, 40).

[18] S. Alam et al. "Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory." In: *Phys. Rev. D* 103 (8 Apr. 2021), p. 083533. DOI: 10.1103/PhysRevD.103.083533 (cited on page 11).

[19] M. Aker et al. "Direct neutrino-mass measurement with sub-electronvolt sensitivity." In: *Nature Physics* 18.2 (2022), pp. 160–166. DOI: 10.1038/s41567-021-01463-1 (cited on page 11).

[20] ALEPH Collaboration et al. "Precision electroweak measurements on the Z resonance." In: *Physics Reports* 427.5 (2006), pp. 257–454. DOI: https://doi.org/10.1016/j.physrep.2005.12.006 (cited on page 11).

[21] R. L. Workman et al. "Review of Particle Physics." In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097 (cited on pages 14, 25, 46, 47, 49).

[22] J. Formaggio and G. Zeller. "From eV to EeV: Neutrino Cross Sections Across Energy Scales." In: *Rev. Mod. Phys.* 84 (2012), p. 1307. DOI: https://doi.org/10.1103/RevModPhys.84.1307 (cited on pages 14, 17).

[23] V. Lyubushkin et al. "A study of quasi-elastic muon neutrino and antineutrino scattering in the NOMAD experiment." In: *The European Physical Journal C* 63.3 (2009), pp. 355–381. DOI: 10.1140/epjc/s10052-009-1113-0 (cited on page 15).

[24] I. Esteban et al. "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of $\theta 23$, δCP, and the mass ordering." In: *Journal of High Energy Physics* 2019.1 (Jan. 2019), p. 106. DOI: 10.1007/JHEP01(2019)106 (cited on pages 20, 21, 31, 68, 69, 71, 97).

[25] B. Jones. "Dynamical pion collapse and the coherence of conventional neutrino beams." In: *Physical Review D* 91.5 (Mar. 2015). DOI: 10.1103/physrevd.91.053002 (cited on pages 21, 22, 91, 93).

[26] J. Linder. *Neutrino matter potentials induced by Earth*. 2005. DOI: 10.48550/ARXIV.HEP-PH/0504264. URL: https://arxiv.org/abs/hep-ph/0504264 (cited on page 22).

[27] L. Wolfenstein. "Neutrino oscillations in matter." In: *Phys. Rev. D* 17 (9 May 1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369 (cited on page 24).

[28] A. Y. Smirnov. "The MSW effect and solar neutrinos." In: *10th International Workshop on Neutrino Telescopes*. May 2003, pp. 23–43 (cited on page 24).

[29] M. G. Aartsen et al. "Searching for eV-scale sterile neutrinos with eight years of atmospheric neutrinos at the IceCube Neutrino Telescope." In: *Phys. Rev. D* 102.5 (2020), p. 052009. DOI: 10.1103/PhysRevD.102.052009 (cited on pages 25, 77, 111).

[30] E. K. Akhmedov. "Parametric resonance in neutrino oscillations in matter." In: *Pramana* 54 (2000). Ed. by A. S. Joshipura and P. B. Pal, pp. 47–63. DOI: 10.1007/s12043-000-0006-4 (cited on page 25).

[31] I. Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations." In: *JHEP* 09 (2020), p. 178. DOI: 10.1007/JHEP09(2020)178 (cited on pages 27, 28, 40).

[32] M. G. Aartsen et al. "PINGU: A Vision for Neutrino and Particle Physics at the South Pole." In: *J. Phys. G* 44.5 (2017), p. 054006. DOI: 10.1088/1361-6471/44/5/054006 (cited on page 28).

[33] R. Davis. "A review of the homestake solar neutrino experiment." In: *Progress in Particle and Nuclear Physics* 32 (1994), pp. 13–32. DOI: https://doi.org/10.1016/0146-6410(94)90004-3 (cited on page 28).

[34] J. N. Bahcall, A. M. Serenelli, and S. Basu. "New solar opacities, abundances, helioseismology, and neutrino fluxes." In: *Astrophys. J. Lett.* 621 (2005), pp. L85–L88. DOI: 10.1086/428929 (cited on page 28).

[35] K. Eguchi et al. "First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance." In: *Phys. Rev. Lett.* 90 (2 Jan. 2003), p. 021802. DOI: 10.1103/PhysRevLett.90.021802 (cited on page 28).

[36] A. Cabrera. "The Double Chooz Experiment." In: *Nuclear Physics B - Proceedings Supplements* 229-232 (2012). Neutrino 2010, pp. 87–91. DOI: https://doi.org/10.1016/j.nuclphysbps.2012.09.014 (cited on page 29).

[37] J. K. Ahn et al. "RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang." In: (Mar. 2010) (cited on page 29).

[38] X. Guo et al. "A Precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya-Bay." In: (Jan. 2007) (cited on page 29).

[39] A. Fedynitch et al. "Calculation of conventional and prompt lepton fluxes at very high energy." In: *EPJ Web of Conferences*. Vol. 99. EDP Sciences. 2015, p. 08001 (cited on pages 29, 76).

[40] M. Honda et al. "Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model." In: *Phys. Rev. D* 92.2 (2015), p. 023004. DOI: 10.1103/PhysRevD.92.023004 (cited on pages 30, 51, 52, 68, 69).

[41] A. M. Dziewonski and D. L. Anderson. "Preliminary reference Earth model." In: *Physics of the Earth and Planetary Interiors* 25.4 (1981), pp. 297–356. DOI: https://doi.org/10.1016/0031-9201(81)90046-7 (cited on pages 31, 79).

[42] S. Adrián-Martínez et al. "Letter of intent for KM3NeT 2.0." In: *Journal of Physics G: Nuclear and Particle Physics* 43.8 (June 2016), p. 084001. DOI: 10.1088/0954-3899/43/8/084001 (cited on page 32).

[43] L. Nauta et al. "First neutrino oscillation measurement in KM3NeT/ORCA." In: *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*. Vol. 395. 2021, p. 1123. DOI: 10.22323/1.395.1123 (cited on page 32).

[44] P. Adamson et al. "First observations of separated atmospheric nu(mu) and anti-nu(mu) events in the MINOS detector." In: *Phys. Rev. D* 73 (2006), p. 072002. DOI: 10.1103/PhysRevD.73.072002 (cited on page 32).

[45] D. Michael et al. "The magnetized steel and scintillator calorimeters of the MINOS experiment." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 596.2 (2008), pp. 190–228. DOI: https://doi.org/10.1016/j.nima.2008.08.003 (cited on page 33).

[46] K. Abe et al. "The T2K Experiment." In: *Nucl. Instrum. Meth. A* 659 (2011), pp. 106–135. DOI: 10.1016/j.nima.2011.06.067 (cited on page 33).

[47] R. B. Patterson. "The NOvA Experiment: Status and Outlook." In: *Nucl. Phys. B Proc. Suppl.* 235-236 (2013). Ed. by T. Kobayashi, M. Nakahata, and T. Nakaya, pp. 151–157. DOI: 10.1016/j.nuclphysbps.2013.04.005 (cited on page 33).

[48] R. M. Zwaska. "Accelerator systems and instrumentation for the NuMI neutrino beam." In: (Dec. 2005). DOI: 10.2172/879065 (cited on page 33).

[49] G. Mention et al. "The Reactor Antineutrino Anomaly." In: *Phys. Rev. D* 83 (2011), p. 073006. DOI: 10.1103/PhysRevD.83.073006 (cited on pages 33, 34).

[50] H. Almazán et al. "STEREO neutrino spectrum of ^{235}U fission rejects sterile neutrino hypothesis." In: *Nature* 613.7943 (2023), pp. 257–261. DOI: 10.1038/s41586-022-05568-2 (cited on pages 34, 35).

[51] P. Huber. "Reactor antineutrino fluxes – Status and challenges." In: *Nuclear Physics B* 908 (2016). Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015, pp. 268–278. DOI: https://doi.org/10.1016/j.nuclphysb.2016.04.012 (cited on page 34).

[52] M. Licciardi. "Results of STEREO and PROSPECT, and status of sterile neutrino searches." In: *55th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. May 2021 (cited on page 35).

[53] C. Giunti et al. "Reactor antineutrino anomaly in light of recent flux model refinements." In: *Physics Letters B* 829 (2022), p. 137054. DOI: 10.1016/j.physletb.2022.137054 (cited on pages 35, 36).

[54] V. Barinov et al. "Results from the Baksan Experiment on Sterile Transitions (BEST)." In: *Physical Review Letters* 128.23 (2022). DOI: 10.1103/physrevlett.128.232501 (cited on page 35).

[55] W. Hampel et al. "Final results of the 51Cr neutrino source experiments in GALLEX." In: *Physics Letters B* 420.1 (1998), pp. 114–126. DOI: https://doi.org/10.1016/S0370-2693(97)01562-1 (cited on page 35).

[56] J. N. Abdurashitov et al. "Measurement of the response of a Ga solar neutrino experiment to neutrinos from a ^{37}Ar source." In: *Physical Review C* 73.4 (2006). DOI: 10.1103/physrevc.73.045805 (cited on page 35).

[57] A. Aguilar-Arevalo et al. "Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam." In: *Phys. Rev. D* 64 (2001), p. 112007. DOI: 10.1103/PhysRevD.64.112007 (cited on page 36).

[58] A. A. Aguilar-Arevalo et al. "Significant Excess of ElectronLike Events in the MiniBooNE Short-Baseline Neutrino Experiment." In: *Phys. Rev. Lett.* 121.22 (2018), p. 221801. DOI: 10.1103/PhysRevLett.121.221801 (cited on pages 36, 37).

[59] O. Fischer, Á. Hernández-Cabezudo, and T. Schwetz. "Explaining the MiniBooNE excess by a decaying sterile neutrino with mass in the 250 MeV range." In: *Physical Review D* 101.7 (2020). DOI: 10.1103/physrevd.101.075045 (cited on page 36).

[60] E. Bertuzzo et al. "Dark Neutrino Portal to Explain MiniBooNE Excess." In: *Phys. Rev. Lett.* 121 (24 Dec. 2018), p. 241801. DOI: 10.1103/PhysRevLett.121.241801 (cited on page 36).

[61] C. Giunti, A. Ioannisian, and G. Ranucci. "A new analysis of the MiniBooNE low-energy excess." In: *JHEP* 11 (2020). [Erratum: JHEP 02, 078 (2021)], p. 146. DOI: 10.1007/JHEP11(2020)146 (cited on page 37).

[62] R. Acciarri et al. "Design and Construction of the MicroBooNE Detector." In: *JINST* 12.02 (2017), P02017. DOI: 10.1088/1748-0221/12/02/P02017 (cited on page 37).

[63] P. Abratenko et al. "Search for Neutrino-Induced Neutral-Current Δ Radiative Decay in MicroBooNE and a First Test of the MiniBooNE Low Energy Excess under a Single-Photon Hypothesis." In: *Phys. Rev. Lett.* 128 (2022), p. 111801. DOI: 10.1103/PhysRevLett.128.111801 (cited on page 37).

[64] P. Abratenko et al. "First constraints on light sterile neutrino oscillations from combined appearance and disappearance searches with the MicroBooNE detector." In: (Oct. 2022) (cited on page 37).

[65] M. Dentler et al. "Updated global analysis of neutrino oscillations in the presence of eV-scale sterile neutrinos." In: *Journal of High Energy Physics* 2018.8 (2018). DOI: 10.1007/jhep08(2018)010 (cited on page 38).

[66] A. Diaz et al. "Where are we with light sterile neutrinos?" In: *Physics Reports* 884 (2020), pp. 1–59. DOI: 10.1016/j.physrep.2020.08.005 (cited on page 38).

[67] Z. Hu et al. "Global oscillation data analysis on the 3ν mixing without unitarity." In: *JHEP* 01 (2021), p. 124. DOI: 10.1007/JHEP01(2021)124 (cited on pages 39, 97).

[68] J. Lesgourgues and S. Pastor. "Neutrino cosmology and Planck." In: *New Journal of Physics* 16.6 (2014), p. 065002. DOI: 10.1088/1367-2630/16/6/065002 (cited on pages 39, 40).

[69] F. T. Avignone III, S. R. Elliott, and J. Engel. "Double Beta Decay, Majorana Neutrinos, and Neutrino Mass." In: *Rev. Mod. Phys.* 80 (2008), pp. 481–516. DOI: 10.1103/RevModPhys.80.481 (cited on page 40).

[70] M. Blennow et al. "Quantifying the sensitivity of oscillation experiments to the neutrino mass ordering." In: *Journal of High Energy Physics* 2014.3 (2014), p. 28. DOI: 10.1007/JHEP03(2014)028 (cited on page 40).

[71] P. B. Price, K. Woschnagg, and D. Chirkin. "Age vs depth of glacial ice at South Pole." In: *Geophysical Research Letters* 27.14 (2000), pp. 2129–2132. DOI: 10.1029/2000GL011351 (cited on page 41).

[72] R. Abbasi et al. "In-situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory." In: *The Cryosphere Discussions* 2022 (2022), pp. 1–48. DOI: 10.5194/tc-2022-174 (cited on pages 42, 116).

[73] M. Aartsen, R. Abbasi, et al. "South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy." In: *Journal of Glaciology* 59.218 (2013), pp. 1117–1128 (cited on page 42).

[74] M. G. Aartsen et al. "The IceCube Neutrino Observatory: instrumentation and online systems." In: *Journal of Instrumentation* 12.3 (Mar. 2017), P03012. DOI: 10.1088/1748-0221/12/03/P03012 (cited on pages 42, 44, 60, 65, 73).

[75] R. Abbasi et al. "The design and performance of IceCube DeepCore." In: *Astroparticle Physics* 35.10 (2012), pp. 615–624 (cited on pages 43, 58, 59).

[76] R. Abbasi et al. "Calibration and characterization of the IceCube photomultiplier tube." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 618.1-3 (2010), pp. 139–152. DOI: 10.1016/j.nima.2010.03.102 (cited on page 44).

[77] R. Abbasi et al. "The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping." In: *Nucl. Instrum. Meth. A* 601 (2009), pp. 294–316. DOI: 10.1016/j.nima.2009.01.001 (cited on pages 44, 64).

[78] M. G. Aartsen et al. "In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes." In: *Journal of Instrumentation* 15.6 (June 2020), P06032. DOI: 10.1088/1748-0221/15/06/P06032 (cited on pages 44, 57).

[79] N. Kulacz. "In Situ Measurement of the IceCube DOM Efficiency Factor Using Atmospheric Minimum Ionizing Muons." MA thesis. University of Alberta, 2019 (cited on pages 45, 74).

[80] P. A. Čerenkov. "Visible Radiation Produced by Electrons Moving in a Medium with Velocities Exceeding that of Light." In: *Phys. Rev.* 52 (4 Aug. 1937), pp. 378–379. DOI: 10.1103/PhysRev.52.378 (cited on page 45).

[81] J. Jackson. *Classical Electrodynamics*. Wiley, 2012 (cited on page 45).

[82] I. M. Frank and I. Tamm. "Coherent visible radiation of fast electrons passing through matter." In: *Physics-Uspekhi* 93 (1937), pp. 388–393 (cited on page 45).

[83] I. Tamm. "Radiation Emitted by Uniformly Moving Electrons." In: *Selected Papers*. Ed. by B. M. Bolotovskii, V. Y. Frenkel, and R. Peierls. Berlin, Heidelberg: Springer Berlin Heidelberg, 1991, pp. 37–53. DOI: 10.1007/978-3-642-74626-0_3 (cited on page 45).

[84] D. E. GROOM, N. V. MOKHOV, and S. I. STRIGANOV. "MUON STOPPING POWER AND RANGE TABLES 10 MeV–100 TeV." In: *Atomic Data and Nuclear Data Tables* 78.2 (2001), pp. 183–356. DOI: https://doi.org/10.1006/adnd.2001.0861 (cited on page 46).

[85] S. Agostinelli et al. "Geant4—a simulation toolkit." In: *Nucl. Instr. Meth. Phys. Res.* 506.3 (July 2003), pp. 250–303. DOI: 10.1016/s0168-9002(03)01368-8 (cited on page 47).

[86] L. Rädel and C. Wiebusch. "Calculation of the Cherenkov light yield from electromagnetic cascades in ice with Geant4." In: *Astroparticle Physics* 44 (2013), pp. 102–113. DOI: https://doi.org/10.1016/j.astropartphys.2013.01.015 (cited on pages 47, 53).

[87] L. Halve. "Measurement of the Atmospheric Neutrino to Antineutrino Ratio above 100GeV with IceCube." MA thesis. RWTH Aachen U., Sept. 2018 (cited on page 48).

[88] M. C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz. "Updated fit to three neutrino mixing: status of leptonic CP violation." In: *Journal of High Energy Physics* 2014.11 (Nov. 2014), p. 52. DOI: 10.1007/JHEP11(2014)052 (cited on page 52).

[89] C. Andreopoulos et al. "The GENIE Neutrino Monte Carlo Generator: Physics and User Manual." In: (2015) (cited on pages 53, 54).

[90] J.-H. Koehne et al. "PROPOSAL: A tool for propagation of charged leptons." In: *Computer Physics Communications* 184.9 (2013), pp. 2070–2090. DOI: https://doi.org/10.1016/j.cpc.2013.04.001 (cited on page 53).

[91] A. Cooper-Sarkar, P. Mertsch, and S. Sarkar. "The high energy neutrino cross-section in the Standard Model and its uncertainty." In: *JHEP* 08 (2011), p. 042. DOI: 10.1007/JHEP08(2011)042 (cited on page 54).

[92] M. G. Aartsen et al. "Measurement of Atmospheric Neutrino Oscillations at 6–56 GeV with IceCube DeepCore." In: *Physical Review Letters* 120.7 (Feb. 2018). DOI: 10.1103/physrevlett.120.071801 (cited on page 54).

[93] Y. Becherini et al. "A parameterisation of single and multiple muons in the deep water or ice." In: *Astroparticle Physics* 25.1 (2006), pp. 1–13. DOI: https://doi.org/10.1016/j.astropartphys.2005.10.005 (cited on page 54).

[94] D. Heck et al. "CORSIKA: A Monte Carlo code to simulate extensive air showers." In: 1998 (cited on page 54).

[95] T. K. Gaisser. "Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio." In: *Astropart. Phys.* 35 (2012), pp. 801–806. DOI: 10.1016/j.astropartphys.2012.02.010 (cited on page 55).

[96] R. Engel et al. "The hadronic interaction model Sibyll – past, present and future." In: *EPJ Web Conf.* 145 (2017). Ed. by B. Pattison, p. 08001. DOI: 10.1051/epjconf/201614508001 (cited on page 55).

[97] C. Kopper et al. https://github.com/claudiok/clsim (cited on page 56).

[98] D. Chirkin et al. "Photon Propagation using GPUs by the IceCube Neutrino Observatory." In: *2019 15th International Conference on eScience (eScience)*. 2019, pp. 388–393. DOI: 10.1109/eScience.2019.00050 (cited on page 56).

[99] M. G. Aartsen et al. "Measurement of South Pole ice transparency with the IceCube LED calibration system." In: *Nucl. Instrum. Meth. A* 711 (2013), pp. 73–89. DOI: 10.1016/j.nima.2013.01.054 (cited on pages 56, 74).

[100] M. Larson. "Simulation and Identification of Non-Poissonian Noise Triggers in the IceCube Neutrino Detector." Available at https://ir.ua.edu/handle/123456789/1927. MA thesis. University of Alabama, Tuscaloosa, AL, USA, 2013 (cited on page 57).

[101] M. Larson. "A Search for Tau Neutrino Appearance with IceCube-DeepCore." available at https://discoverycenter.nbi.ku.dk/teaching/thesis_page/mjlarson_thesis.pdf. PhD thesis. University of Copenhagen, Denmark, 2018 (cited on page 57).

[102] M. Aartsen et al. "Improvement in fast particle track reconstruction with robust statistics." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 736 (2014), pp. 143–149. DOI: 10.1016/j.nima.2013.10.074 (cited on page 60).

[103] R. Abbasi et al. "First search for atmospheric and extraterrestrial neutrino-induced cascades with the IceCube detector." In: *Physical Review D* 84 (Oct. 2011), p. 072001. DOI: 10.1103/PhysRevD.84.072001 (cited on page 60).

[104] J. P. Y. Garza. "Measurement of neutrino oscillations in atmospheric neutrinos with the IceCube DeepCore detector." PhD thesis. Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät I, 2014. DOI: http://dx.doi.org/10.18452/17016 (cited on page 63).

[105] J. A. Aguilar et al. "A fast algorithm for muon track reconstruction and its application to the ANTARES neutrino telescope." In: *Astropart. Phys.* 34 (2011), pp. 652–662. DOI: 10.1016/j.astropartphys.2011.01.003 (cited on page 63).

[106] R. Abbasi et al. "Low Energy Event Reconstruction in IceCube DeepCore." In: (Mar. 2022) (cited on pages 63, 115).

[107] P. Price and K. Woschnagg. "Role of group and phase velocity in high-energy neutrino observatories." In: *Astroparticle Physics* 15.1 (2001), pp. 97–100. DOI: https://doi.org/10.1016/S0927-6505(00)00142-0 (cited on page 65).

[108] A. Terliuk. "Measurement of atmospheric neutrino oscillations and search for sterile neutrino mixing with IceCube DeepCore." PhD thesis. Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät, 2018. DOI: http://dx.doi.org/10.18452/19304 (cited on page 66).

[109] M. G. Aartsen et al. "Search for sterile neutrino mixing using three years of IceCube DeepCore data." In: *Physical Review D* 95.11 (2017). DOI: 10.1103/physrevd.95.112002 (cited on pages 67, 111).

[110] M. G. Aartsen et al. "Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data." In: *Physical Review D* 91.7 (2015). DOI: 10.1103/physrevd.91.072004 (cited on page 67).

[111] M. G. Aartsen et al. "Measurement of Atmospheric Tau Neutrino Appearance with IceCube DeepCore." In: *Phys. Rev. D* 99.3 (2019), p. 032007. DOI: 10.1103/PhysRevD.99.032007 (cited on pages 67, 80).

[112] F. Pedregosa et al. "Scikit-learn: Machine Learning in Python." In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830 (cited on page 67).

[113] A. Kolmogorov. "Sulla determinazione empirica di una legge di distribuzione." In: *Inst. Ital. Attuari, Giorn.* 4 (1933), pp. 83–91 (cited on page 73).

[114] J. Feintzeig. "Searches for point-like sources of astrophysical neutrinos with the IceCube Neutrino Observatory." PhD thesis. Wisconsin U., Madison, 2014 (cited on page 74).

[115] S. Fiedlschuster. "The Effect of Hole Ice on the Propagation and Detection of Light in IceCube." In: (Apr. 2019) (cited on page 74).

[116] M. Rongen. "Measuring the optical properties of IceCube drill holes." In: *EPJ Web of Conferences*. Vol. 116. EDP Sciences. 2016, p. 06011 (cited on page 74).

[117] K. P. F.R.S. "LIII. On lines and planes of closest fit to systems of points in space." In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), pp. 559–572. DOI: 10.1080/14786440109462720 (cited on page 74).

[118] A. Fedynitch et al. https://github.com/afedynitch/MCEq (cited on page 76).

[119] A. Fedynitch, J. B. Tjus, and P. Desiati. "Influence of hadronic interaction models and the cosmic ray spectrum on the high energy atmospheric muon and neutrino flux." In: *Physical Review D* 86.11 (2012), p. 114024 (cited on page 76).

[120] G. D. Barr et al. "Uncertainties in Atmospheric Neutrino Fluxes." In: *Phys. Rev. D* 74 (2006), p. 094009. DOI: 10.1103/PhysRevD.74.094009 (cited on pages 76, 77, 116).

[121] Jet Propulsion Laboratory. "AIRS/AMSU/HSB Version 6 Level 3 Product User Guide." In: *Version 1.2* (Nov. 2014) (cited on page 77).

[122] M. G. Aartsen et al. "Computational techniques for the analysis of small signals in high-statistics neutrino oscillation experiments." In: *Nucl. Instrum. Meth. A* 977 (2020), p. 164332. DOI: 10.1016/j.nima.2020.164332 (cited on page 79).

[123] V. Barger et al. "Matter effects on three-neutrino oscillations." In: *Phys. Rev. D* 22 (11 Dec. 1980), pp. 2718–2726. DOI: 10.1103/PhysRevD.22.2718 (cited on page 79).

[124] Q. Li and J. S. Racine. *Nonparametric Econometrics: Theory and Practice*. Princeton University Press, 2007 (cited on page 80).

[125] P. Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2 (cited on page 83).

[126] R. H. Byrd et al. "A Limited Memory Algorithm for Bound Constrained Optimization." In: *SIAM J. Sci. Comput.* 16 (1995), pp. 1190–1208 (cited on page 83).

[127] R. S. Nickerson. "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises." In: *Review of General Psychology* 2 (1998), pp. 175–220 (cited on page 86).

[128] P. Dunne. *Latest Neutrino Oscillation Results from T2K*. July 2020. DOI: 10.5281/zenodo.4154355. URL: https://doi.org/10.5281/zenodo.4154355 (cited on page 88).

[129] P. Adamson et al. "Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset." In: *Phys. Rev. Lett.* 125.13 (2020), p. 131802. DOI: 10.1103/PhysRevLett.125.131802 (cited on page 88).

[130] M. A. Acero et al. "Improved measurement of neutrino oscillation parameters by the NOvA experiment." In: *Phys. Rev. D* 106.3 (2022), p. 032004. DOI: 10.1103/PhysRevD.106.032004 (cited on page 88).

[131] V. Takhistov. "Review of Atmospheric Neutrino Results from Super-Kamiokande." In: *PoS* ICHEP2020 (2021), p. 181. DOI: 10.22323/1.390.0181 (cited on page 88).

[132] S. S. Wilks. "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses." In: *The Annals of Mathematical Statistics* 9.1 (1938), pp. 60–62. DOI: 10.1214/aoms/1177732360 (cited on page 88).

[133] C. A. Argüelles Delgado, J. Salvado, and C. N. Weaver. "A Simple Quantum Integro-Differential Solver (SQuIDS)." In: *Comput. Phys. Commun.* 196 (2015), pp. 569–591. DOI: 10.1016/j.cpc.2015.06.022 (cited on pages 92, 93).

[134] D. Argüelles et al. https://github.com/arguelles/nuSQuIDS (cited on pages 92, 93).

[135] S. Aiello et al. "Sensitivity to light sterile neutrino mixing parameters with KM3NeT/ORCA." In: *JHEP* 10 (2021), p. 180. DOI: 10.1007/JHEP10(2021)180 (cited on page 98).

[136] M. Maltoni and T. Schwetz. “Sterile neutrino oscillations after first MiniBooNE results.” In: *Phys. Rev. D* 76 (2007), p. 093005. DOI: 10.1103/PhysRevD.76.093005 (cited on page 98).

[137] M. D. Richard and R. P. Lippmann. “Neural Network Classifiers Estimate Bayesian *a posteriori* Probabilities.” In: *Neural Computation* 3.4 (1991), pp. 461–483. DOI: 10.1162/neco.1991.3.4.461 (cited on page 103).

[138] F. James and M. Roos. “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations.” In: *Comput. Phys. Commun.* 10 (1975), pp. 343–367. DOI: 10.1016/0010-4655(75)90039-9 (cited on pages 105, 107).

[139] W. L. Price. “Global optimization by controlled random search.” In: *Journal of Optimization Theory and Applications* 40.3 (July 1983), pp. 333–348. DOI: 10.1007/BF00933504 (cited on page 107).

[140] T. Rowan. “Functional Stability Analysis of Numerical Algorithms.” PhD thesis. University of Texas at Austin, 1990 (cited on page 107).

[141] S. G. Johnson. *The NLopt nonlinear-optimization package*. https://github.com/stevengj/nlopt (cited on page 107).

[142] A. Albert et al. “Measuring the atmospheric neutrino oscillation parameters and constraining the 3+1 neutrino model with ten years of ANTARES data.” In: *JHEP* 06 (2019), p. 113. DOI: 10.1007/JHEP06(2019)113 (cited on page 111).

[143] K. Abe et al. “Limits on sterile neutrino mixing using atmospheric neutrinos in Super-Kamiokande.” In: *Phys. Rev. D* 91 (2015), p. 052019. DOI: 10.1103/PhysRevD.91.052019 (cited on page 111).

[144] S. Edayath et al. “Sterile Neutrino Search in the NOvA Far Detector.” In: *Meeting of the APS Division of Particles and Fields*. Oct. 2017 (cited on page 111).

[145] G. J. Feldman and R. D. Cousins. “Unified approach to the classical statistical analysis of small signals.” In: *Physical Review D* 57.7 (1998), pp. 3873–3889. DOI: 10.1103/physrevd.57.3873 (cited on page 114).

[146] M. Aartsen et al. “Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube.” In: *Journal of Cosmology and Astroparticle Physics* 2019.10 (2019), pp. 048–048. DOI: 10.1088/1475-7516/2019/10/048 (cited on page 116).

[147] A. Fedynitch and M. Huber. “Data-driven hadronic interaction model for atmospheric lepton flux calculations.” In: *Physical Review D* 106.8 (2022). DOI: 10.1103/physrevd.106.083018 (cited on page 116).

[148] A. Ishihara. *The IceCube Upgrade – Design and Science Goals*. 2019. DOI: 10.48550/ARXIV.1908.09441. URL: https://arxiv.org/abs/1908.09441 (cited on pages 116, 117).

Ich erkläre, dass ich die Dissertation selbständig und nur unter Verwendung der von mir gemäß § 7 Abs. 3 der Promotionsordnung der Mathematisch-Naturwissenschaftlichen Fakultät, veröffentlicht im Amtlichen Mitteilungsblatt der Humboldt-Universität zu Berlin Nr. 42/2018 am 11.07.2018 angegebenen Hilfsmittel angefertigt habe.

Alexander Trettin