

Search for Sterile Neutrinos with the IceCube Neutrino Observatory

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften
der RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von
Diplom-Physiker
Marius Wallraff

aus Aachen.

Berichter:

Universitätsprofessor Dr. rer. nat. Christopher Wiebusch
Universitätsprofessor Dr. rer. nat. Thomas Hebbeker

Tag der mündlichen Prüfung: 28.10.2019

Diese Dissertation ist auf den Internetseiten der Hochschulbibliothek online verfügbar.

知彼知己，百戰不殆

If you know the enemy and know yourself,
you need not fear the result of a hundred battles.

Sun Tzu, The Art of War

Abstract

One of the open questions of particle physics is whether there are more than the three known flavors of neutrinos. The properties of these new particles would have to deviate from the known neutrinos by either being heavier than half the mass of the Z boson, or by not participating in the weak interaction. The latter of these alternatives gained relevance in the recent years from neutrino oscillation experiments that observed anomalies in their measured neutrino rates that could be explained by the introduction of one or more of these *sterile neutrinos* with a squared-mass difference Δm_{42}^2 of roughly 1 eV².

In vacuum, the oscillation probabilities between the known flavors and potential sterile neutrinos have to be small, because otherwise they would have been detected already. However, when neutrinos traverse matter, the differences in the way the flavors can interact with matter induce an additional potential that can lead to resonances that can strongly increase these probabilities. Together with the large energy range, the high density of the Earth's core and mantle causes atmospheric neutrinos to be ideal to search for muon neutrino disappearance signatures of sterile neutrinos in the 1 eV² range.

The IceCube Neutrino Observatory at the Geographic South Pole is capable of recording several ten thousands of atmospheric muon neutrino events per year with a purity of over 99.9%. This analysis uses one year of data from an incomplete detector configuration with 59 of the final 86 strings installed. A likelihood ratio test is performed as scan in two parameters – the sterile mixing angle θ_{24} and the squared mass difference Δm_{42}^2 – on two-dimensional histograms of neutrino events versus reconstructed energy and zenith angle to search for disappearance signatures. Systematic uncertainties in the atmospheric neutrino flux and detector responses have been taken care of, and have been parameterized and implemented into the likelihood function as nuisance parameters where applicable.

The analysis does not find evidence for sterile neutrinos. The best fit lies at $\theta_{24} = 12.08^\circ$, $\Delta m_{42}^2 = 0.0401$ eV² and is compatible with the null hypothesis at a significance of 0.463 σ ; it lies outside the median expected 90% C.L. sensitivity at the lower boundary of the analysis' energy range. The goodness of fit is good, and exclusion contours in the two scan parameters are given. Possible improvements to this analysis method are discussed in the outlook.

Zusammenfassung

Eine der offenen Fragen der Teilchenphysik ist, ob mehr als die drei bekannten Neutrino-Flavor existieren. Die Eigenschaften dieser neuen Teilchen müssten von denen der bekannten Neutrinos abweichen indem sie entweder mindestens halb so schwer sind wie das Z-Boson, oder indem sie nicht an der schwachen Wechselwirkung teilnehmen. Die letztere Alternative hat in den vergangenen Jahren an Relevanz gewonnen aufgrund einiger Experimente, die Anomalien in ihren gemessenen Neutrinoaraten entdeckt haben, die sich durch die Einführung eines oder mehrere *steriler Neutrinos* mit einer Differenz der quadrierten Massen Δm_{42}^2 von etwa 1 eV^2 erklären lassen könnten.

Die Oszillationswahrscheinlichkeiten in Vakuum zwischen den bekannten und potentiellen sterilen Neutrinos müssen klein sein, da sie sonst schon entdeckt worden wären. Wenn Neutrinos hingegen Materie durchqueren, induzieren die unterschiedlichen Arten der Flavor, mit Materie zu interagieren, ein zusätzliches Potential, das zu Resonanzen führen kann, die die Wahrscheinlichkeiten stark erhöhen. Zusammen mit ihrem großen Energiebereich machen die hohe Dichte des Erdkerns und -mantels atmosphärische Neutrinos zu idealen Objekten, um nach Disappearance-Signaturen sterile Neutrinos im 1 eV^2 -Bereich zu suchen.

Das IceCube Neutrino Observatory am geographischen Südpol ist dazu in der Lage, jährlich etliche Zehntausend Myonneutrino-Ereignisse aufzuzeichnen mit einer Reinheit von über 99.9%. Diese Analyse nutzt Daten eines Jahres, die mit einer unvollständigen Detektorkonfiguration von 59 von 86 Strings des fertigen Detektors aufgenommen wurden. Ein Likelihood-Ratio-Test wird als Scan in zwei Parametern durchgeführt – dem sterilen Mischungswinkel θ_{24} und der Differenz der quadrierten Massen Δm_{42}^2 – auf zweidimensionalen Histogrammen der Neutrino-Ereignisse gegen rekonstruierten Energie und Zenithwinkel, um nach Disappearance-Signaturen zu suchen. Systematische Unsicherheiten im Fluss atmosphärischer Neutrinos und im Ansprechverhalten des Detektors wurden berücksichtigt, und wenn möglich parameterisiert und als Störparameter in die Likelihood-Funktion eingebunden.

Die Analyse findet keine Hinweise auf sterile Neutrinos. Die beste Übereinstimmung wird für $\theta_{24} = 12.08^\circ$ und $\Delta m_{42}^2 = 0.0401 \text{ eV}^2$ gefunden, und das Ergebnis ist mit der Nullhypothese kompatibel bei einer Signifikanz von 0.463σ ; es liegt außerhalb der im Median erwarteten Signifikanz im 90%-Konfidenzintervall am unteren Ende des Energiebereichs der Analyse. Die Anpassungsgüte ist gut, und Ausschlussgrenzen in Abhängigkeit der beiden Scan-Parameter werden angegeben. Mögliche Verbesserungen der Analyse-Methode werden im Ausblick diskutiert.

Contents

Abstract	i
Zusammenfassung	iii
Contents	v
List of Figures	viii
List of Tables	xiv
1 Neutrinos	1
1.1 Sources of Neutrinos	2
1.2 Properties of Atmospheric Neutrinos	5
1.2.1 Energy Dependence	5
1.2.2 Angular Spectrum	10
1.2.3 Flavor Composition	11
1.2.4 Antineutrino-Neutrino Ratio	13
1.2.5 Prompt Neutrinos	14
1.3 Standard Neutrino Oscillations	16
1.3.1 Vacuum Oscillation	17
1.3.2 The Induced Matter Potential	22
1.3.3 Matter Effects	25
1.4 Additional Sterile Flavors	28
1.4.1 Vacuum Oscillation	31
1.4.2 Matter Effects	34
1.5 Current State of Research	36
2 The IceCube Neutrino Observatory	37
2.1 Detection Principle	37
2.1.1 Čerenkov Radiation	37
2.1.2 Neutrino Interaction Products	40
2.1.3 Lepton Propagation	45
2.2 Design and Hardware	49

2.2.1	Digital Optical Modules	51
2.2.2	Ice Properties and Models	53
2.3	Processing and Event Reconstruction	56
2.3.1	Triggering, Low-Level Reconstruction and Filtering	57
2.3.2	High-Level Reconstruction	58
2.4	Neutrino and Muon Simulations	59
3	Event Selection	61
3.1	Baseline Event Sample	62
3.2	Adaption to this Analysis	65
4	Signatures of Sterile Neutrinos in IceCube	69
4.1	Four-Flavor Models	69
4.1.1	$\theta_{14} = \theta_{34} = 0$	69
4.1.2	$\theta_{34} = \theta_{24}$, $\theta_{14} = 0$	70
4.1.3	$\theta_{34} \neq \theta_{24}$, $\theta_{14} = 0$	73
4.1.4	$\theta_{14} \neq 0$	77
4.2	Models with More Flavors	77
4.3	Impact of Detector Resolution	79
4.3.1	Angular and Energy Resolutions	79
4.3.2	Effect on the Signatures	82
5	Analysis Strategy	85
5.1	Binning and Probability Calculation	85
5.2	Likelihood Function	91
5.3	Likelihood Ratio Test	92
5.4	Nuisance Parameters	92
5.5	Software Implementation	94
6	Handling of Systematic Uncertainties	97
6.1	Unparameterizable Uncertainties	97
6.1.1	Atmospheric Models	97
6.1.2	Seasonal Variations	99
6.1.3	Ice Models	101
6.2	Parameterized Uncertainties	101
6.2.1	Total Flux Normalization	101
6.2.2	Cosmic-Ray Spectral Index	102
6.2.3	Kaon-Pion Ratio	102
6.2.4	Antineutrino-Neutrino Ratio	103
6.2.5	DOM Efficiency	104
6.2.6	Ice Model Parameters	106
6.3	Prior Values	110

7 Sensitivity and Results	113
7.1 Sensitivity and Full-Circle Tests	113
7.1.1 Asimov Sensitivity Estimates	113
7.1.2 Ensembles of Pseudo Experiments	114
7.2 Best Fit, Resulting Limits, and Checks	116
7.2.1 Pull Distribution	121
7.2.2 Nuisance Parameters and $n - 1$ Plots	124
7.3 Development	126
7.4 Previous Publication of Results	130
8 Conclusions and Outlook	133
8.1 Conclusions	133
8.2 Outlook	134
8.2.1 Starting Events	135
8.2.2 Low-Energy Sample	135
A CP Violation	139
B Neutrino Mass Hierarchy	143
C Calculation of Oscillation Probabilities with nuCraft	145
C.1 Factorization of Oscillations and Non-Coherent Interactions	147
D Locally Weight-Preserving Gaussian Blurring	149
E Kernel Density Estimation	151
F 3×3 Bin Manipulation Study	153
G Brazilian Flag Plots	157
Acknowledgments	I
Declaration	III
Eidesstattliche Erklärung	V
Bibliography	VII

List of Figures

1.1	Expected flux of solar neutrinos versus energy.	3
1.2	All-particle cosmic-ray spectrum vs. energy per nucleus, and spectra of its four components.	5
1.3	Mean free paths and mean energy loss distances for protons interacting with photons of the CMB.	6
1.4	Total cross-sections of hadronic and photonic processes versus invariant mass.	8
1.5	Relative contributions of pion and kaon decays to the muon and muon neutrino fluxes versus particle energy.	9
1.6	To-scale sketch of the Earth with 20 km of atmosphere, showing the path length of particles horizontally arriving at IceCube.	10
1.7	Expected energy and angular distributions of conventional atmospheric, prompt, and extra-terrestrial neutrinos from an analysis of six years of IceCube data.	11
1.8	Flavor ratio of muon neutrinos and antineutrinos compared to electron neutrinos and antineutrinos at ground level according to the HKKMS2006 model.	12
1.9	Measurement of the numbers of muon antineutrinos and neutrinos of two independent samples by the MINOS experiment.	13
1.10	Uncertainties in the electron and muon neutrino-antineutrino ratios and the electron-muon flavor ratio as functions of energy and zenith angle.	14
1.11	Antineutrino-neutrino ratio of muon and electron flavors at ground level according to the HKKMS2006 model.	15
1.12	Sketch of an atmospheric neutrino interacting inside an underground detector, illustrating the relation between angle $\tilde{\theta} = \pi - \theta$ and baseline L	19
1.13	Simplified visualization of (1.13) under assumption of a fixed interaction height of 20 km above ground level.	20
1.14	Survival probability of atmospheric muon (anti-)neutrinos without taking into account matter effects, as a function of neutrino energy and zenith angle.	21
1.15	Feynman graphs of neutral-current elastic scattering of conventional neutrinos and antineutrinos on normal matter.	22
1.16	Feynman graphs of charged-current elastic scattering of neutrinos and antineutrinos on normal matter.	23

1.17	Earth's density as predicted by the Preliminary Reference Earth Model (PREM)	24
1.18	Survival probability of atmospheric muon neutrinos and muon antineutrinos as a function of neutrino energy and zenith angle, taking into account matter effects.	27
1.19	Appearance probability of atmospheric electron neutrinos as a function of neutrino energy and zenith angle, without and with matter effects.	28
1.20	Hadronic cross-section at the Z^0 boson resonance in e^+e^- annihilations at the LEP particle collider.	30
1.21	Survival probability of atmospheric muon neutrinos with an additional flavor of sterile neutrinos as a benchmark, without matter effects.	33
1.22	Survival probability of atmospheric muon neutrinos and muon antineutrinos with an additional flavor of sterile neutrinos, taking into account matter effects.	35
2.1	Sketch of a charged particle traveling faster than the phase speed of light in the medium, producing Čerenkov radiation.	38
2.2	The refractive index n_r of water ice at -7°C , and the relative spectral acceptance of IceCube's Digital Optical Modules for vertically arriving photons. .	39
2.3	Feynman graph of neutral-current non-elastic interactions of conventional neutrinos and antineutrinos on matter.	40
2.4	Feynman graph of charged-current non-elastic interactions of conventional neutrinos and antineutrinos on matter. Equivalent interactions with sea quarks are omitted.	41
2.5	Calculated total neutrino and antineutrino cross-sections for various interactions with electrons or nucleons at high energies.	41
2.6	Predicted and measured total neutrino and antineutrino cross-sections for charged-current interactions with isoscalar nucleons.	43
2.7	Predicted average of the kinematic inelasticity parameter Bjorken-Y for CC and NC neutrino and antineutrino interactions as a function of the incident neutrino energy.	44
2.8	Energy fraction of the electromagnetic component of a hadronic cascade f_{em} in lead calorimeters as simulated by FLUKA90.	44
2.9	Radiation length and photon mean free path in water ice modified by the LPM effect in units of the unmodified radiation length X_0	46
2.10	Parameterization of the energy loss of a muon in water ice due to ionization and three radiative processes.	47
2.11	Distributions of the planar scattering angle ψ and of the offset y due to scattering of 10 TeV muons down to 1 GeV.	48
2.12	Footprint of standard IceCube strings, DeepCore strings, and IceTop tanks, with seasons of deployment.	50
2.13	Layout of an IceCube Digital Optical Module (DOM).	51
2.14	Diagram of the DOM digitizer launch logic, with two adjacent HLC launches.	52

LIST OF FIGURES

2.15	Inverse effective scattering length and inverse absorption length for ice at the South Pole vs. depth for different wavelengths as measured by AMANDA.	54
2.16	Depths of ice layers along the average direction of the ice tilt gradient, and average charge ratio between experimental data and simulated data without taking into account ice anisotropy.	55
3.1	Distributions of the atmospheric muon background in the split stream before the final cuts.	67
3.2	Distributions of the atmospheric muon background in the unsplit stream before the final cuts.	68
3.3	Distributions of the TopologicalTrigger decision before the final cuts, and ratio versus energy.	68
4.1	Change of muon antineutrino disappearance signatures for varying values of Δm_{42}^2 and θ_{24}	71
4.2	Muon antineutrino survival probabilities for vertically upgoing particles for various values of θ_{24}	72
4.3	Muon antineutrino disappearance signatures for large values of θ_{24}	72
4.4	Muon antineutrino disappearance signature for the $\theta_{34} = \theta_{24}$ scheme, and corresponding tau antineutrino appearance probability.	73
4.5	Differences in the muon antineutrino and muon neutrino disappearance probabilities for the model of figure 4.4 compared to the central model of figure 4.1.	74
4.6	Muon antineutrino survival probabilities for vertically upgoing particles for various values of θ_{34}	75
4.7	Integral muon (anti-)neutrino survival probabilities for vertically upgoing particles for various values of θ_{34} above 100 GeV, assuming an $E^{-3.7}$ spectrum, and corresponding muon neutrino disappearance signature at $\theta_{34} = 22.8^\circ$	76
4.8	Differences in the muon antineutrino and muon neutrino disappearance probabilities for $\theta_{14} = \theta_{24} = 7.1^\circ$ compared to the central model of figure 4.1.	78
4.9	Differences in the muon antineutrino and muon neutrino disappearance probabilities analog to figure 4.8, but for a large value of θ_{14}	78
4.10	Differences in the muon antineutrino and muon neutrino disappearance probabilities analog to figure 4.8, but comparing zero and non-zero θ_{14} for the $\theta_{34} = \theta_{24} = 7.1^\circ$ baseline model.	78
4.11	Muon antineutrino survival probabilities for vertically upgoing tracks for a $3 + 2$ model compared to those for the corresponding $3 + 1$ models and to the combination of those $3 + 1$ models.	80
4.12	IceCube angular resolution for muon neutrino track reconstruction as a function of energy, and absolute and relative changes of the baseline L of atmospheric neutrinos due to a 1° zenith angle misreconstruction.	81

4.13	Reconstructed muon neutrino energy according to the muEx algorithm as a function of true neutrino energy, and relative spread of the reconstructed values as function of true neutrino energy.	83
4.14	Muon neutrino and antineutrino survival probabilities sampled over all final-level MC events versus true MC zenith angle and neutrino energy.	84
4.15	Muon neutrino and antineutrino survival probabilities sampled over all final-level MC events versus reconstructed zenith angle and neutrino energy.	84
5.1	Event numbers of the final-level event sample versus reconstructed zenith angle and energy in the analysis binning.	86
5.2	Effective number of events n_{eff} per bin for the baseline MC dataset at final level (with fitted nuisance parameters).	87
5.3	Percentual ratio $100 \cdot n_{\text{exp}} / n_{\text{eff}}$ of experimental event numbers to effective event numbers for the baseline MC dataset at final level, i.e., ratio of figure 5.1 to figure 5.2 in percent.	87
5.4	Fraction of the maximum event weight among the summed event weights per bin in percent for the baseline MC dataset at final level at a four times finer binning.	88
5.5	ν_μ and $\bar{\nu}_\mu$ survival probabilities vs. true and reconstructed energy and zenith angle before and after the smoothing and histogramming procedure described in the text.	89
5.6	Standard deviations of the individual $(\bar{\nu})_\mu$ survival probabilities of all MC events per bin for the lower four plots of figure 5.5.	90
6.1	Effect of the time-averaged seasonal variations on the neutrino event rate as a function of reconstructed energy and zenith angle.	100
6.2	Antineutrino-neutrino ratio uncertainty parameterization, showing both parameters and example lines.	104
6.3	Effects of the DOM efficiency on the distribution of true neutrino energy and reconstructed energy.	105
6.4	Examples for the per-event DOM efficiency parameterizations in nine different regions of energy and zenith angle.	107
6.5	χ^2 / n_{DoF} distribution of the DOM efficiency parameterization, and example parameterizations for ten bins.	108
6.6	Ratio plot of reconstructed energy and zenith angle for the two datasets used for the bulk ice scattering parameterization.	109
6.7	Ratio plots of reconstructed energy and zenith angle for the three datasets used for the bulk ice absorption parameterization.	110
6.8	Ratio plots of reconstructed energy and zenith angle for the three datasets used for the hole ice scattering parameterization.	111
7.1	Asimov estimate for the analysis' sensitivity.	114

LIST OF FIGURES

7.2	Distributions of the test statistic values and best-fit positions of 300 pseudo experiments for the null hypothesis.	115
7.3	Distributions of $2\Lambda_{\text{sig}}$ and best-fit positions of 100 pseudo experiments for a weak signal hypothesis.	117
7.4	Distributions of $2\Lambda_{\text{sig}}$ and best-fit positions of 100 pseudo experiments for a intermediate signal hypothesis.	117
7.5	Distributions of $2\Lambda_{\text{sig}}$ and best-fit positions of 100 pseudo experiments for a strong signal hypothesis.	117
7.6	Likelihood landscape and exclusion contours for the experimental data with the extended parameter space.	118
7.7	Energy and zenith angle distributions of the experimental data compared to the best-fit MC predictions for the null hypothesis and the best fit, including numbers for goodness of fit.	119
7.8	Experimental contour and distribution of contours from null hypothesis pseudo experiments.	120
7.9	Pulls of the experimental data compared to the null hypothesis and to the best fit.	122
7.10	Distribution of the number of rectangular bin clusters exceeding a given Poissonian pull value vs. pull thresholds.	123
7.11	Positions of the clusters above the 3.0σ pull threshold in experimental data.	123
7.12	Number of rectangular bin clusters exceeding a given Poissonian pull vs. pull thresholds, and positions of clusters exceeding 3.0σ	123
7.13	Values of all nuisance parameters over the full scanned parameter space.	125
7.14	Likelihood landscapes with one nuisance parameter each fixed to the default value (called $n - 1$ plots).	127
7.15	Likelihood landscape of the initial analysis result, see text.	128
7.16	Likelihood landscape of a MC sample with ice tilt fitted with a similar sample without.	129
7.17	Likelihood landscape and distribution of pseudo experiment contours of an overly conservative result suggested by parts of the IceCube collaboration, see text.	131
7.18	Contours and sensitivities of the IC-86 analysis.	131
8.1	Estimated sensitivity and likelihood landscape of a simple study with an IC-79 low-energy sample.	136
8.2	Exclusion limits of another IceCube analysis searching for sterile neutrinos, using a low-energy sample.	137
A.1	Muon neutrino and antineutrino survival probabilities for vertically upgoing particles for different values of δ_{24} or δ_{34} vs. neutrino energy.	140
A.2	Muon neutrino survival probabilities for vertically upgoing particles for different values of δ_{13} or δ_{23} vs. neutrino energy.	141

C.1	Expectation value and central 68% interval of the path length distributions of neutrinos in Earth’s atmosphere as a function of the zenith angle.	146
C.2	Comparison of nuCraft default computations vs. common approximations. .	147
C.3	Difference between the fully simulated coherent and non-coherent interactions using ν -SQuIDS and the factorization approach using nuCraft and the IceCube MC framework for nearly maximal mixing angles.	148
F.1	Likelihood landscapes of the original Asimov expectation and manipulated ones.	155
F.2	Continuation of figure F.1.	156
G.1	Contours of twenty (2×10) randomly selected pseudo experiments for the null hypothesis, with their points of best fit.	158
G.2	Brazilian-flag plots of the final result, created by the 1d algorithm.	159
G.3	Brazilian-flag plots of the final result, created by the final 2d algorithm. .	159

List of Tables

1.1	Global-fit oscillation parameters for normal mass hierarchy	21
3.1	Cuts performed to get the baseline event sample, starting from the more generic <i>muon Level 3</i> sample	63
3.2	Simulated atmospheric muon neutrino background events before final cuts	66
6.1	Values of nuisance parameter priors, pseudo experiment randomization parameters, and allowed ranges	111
F.1	Number of added/subtracted events and resulting significance for the 3×3 bin manipulation study	154

Chapter 1

Neutrinos

Neutrinos are electrically neutral particles that were proposed in 1930 by Wolfgang Pauli to solve the problem of the allegedly violated conservations of energy, momentum, and angular momentum in beta decay. Until then, beta decay was perceived as a two-body process, for which energy and momentum in the rest frame assume fixed values that can easily be computed analytically, but measurements in 1914 showed that this is not the case [2]. The introduction of the electron neutrino turned the process into a three-body decay that could correctly explain the continuous distribution of the electron energy as it had been observed:

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

The proposed electron antineutrino $\bar{\nu}_e$ has been experimentally discovered in 1956 [3] through the observation of beta capture of electron antineutrinos generated by beta decay processes inside a nuclear reactor – so-called reactor neutrinos. Six years later, a second flavor, the muon neutrino ν_μ , was discovered [4], after having been proposed in the 1940s to explain the continuous energy and momentum spectra of muon decay in cosmic rays:

$$\mu^- \rightarrow \nu_\mu + e^- + \bar{\nu}_e$$

The third and latest neutrino flavor was discovered 2000 by the dedicated DONUT experiment, after its existence had been proposed 1975 after the discovery of the tau lepton [5].

In the late 1960s, the first measurement of neutrinos produced in nuclear fusion processes inside the sun – so-called solar neutrinos – was successful. However, the measured fluxes were too low by a factor of three compared to the predictions. This disagreement was only conclusively solved about thirty years later by results of Super-Kamiokande and the Sudbury Neutrino Observatory (SNO), which saw experimental evidence for neutrino oscillations [6].

The first sections of this chapter give an overview over various neutrino sources and the properties of the neutrinos that are employed in this work, the later sections explain the phenomenon of neutrino oscillations.

1.1 Sources of Neutrinos

Distinguished by their source, neutrinos can be categorized into several populations that differ in flavor compositions and energy spectra. The neutrinos that are used in this analysis are so-called atmospheric neutrinos. They originate in particle showers that are induced when cosmic rays hit the Earth's atmosphere. They feature a high rate, a large energy range from MeV up to over 100 TeV and come almost exclusively in electron and muon flavors both as particles and antiparticles. As their characteristics are very important for this work, they are discussed in more detail in the next section.

A second source of neutrinos that would be useful to study neutrino oscillations are neutrino beams. These beams can be produced at particle accelerators by directing a proton beam at a fixed target to create charged pions that are then directed into a decay tunnel. Over 99.98% of those charged pions produce muon neutrinos or antineutrinos, which can be selected by filtering pions with magnetic fields depending on their charge [7]:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \quad \pi^- \rightarrow \mu^- + \bar{\nu}_\mu \quad (1.1)$$

Since the processes have two-body final states, the neutrino energy is determined by the pion energy and can therefore be controlled well. Together with their potential for high statistics, this makes them suitable for precision experiments. Unfortunately, due to the large distances to existing accelerators and the steep and therefore expensive decay tunnels that would be required, there are no neutrino beams available to IceCube.

Another population of neutrinos that can be used for oscillation measurements are electron neutrinos that are produced in fusion processes inside the core of the sun. As shown in figure 1.1, they feature both discrete and continuous energy spectra because there are both processes with two and with more particles in their final states [8]. Solar neutrinos have been measured by multiple experiments and were of prime importance for the theoretical prediction and experimental discovery of neutrino oscillations, but at energies up to 18 MeV, they lie far below the energy range of this analysis.

Similarly stemming from nuclear processes are geoneutrinos that originate from radioactive decay of heavy elements inside the Earth [9], as well as reactor neutrinos that originate from fission processes in man-made nuclear reactors. They, too, are below the energy threshold of this analysis, but especially reactor neutrinos are very valuable for

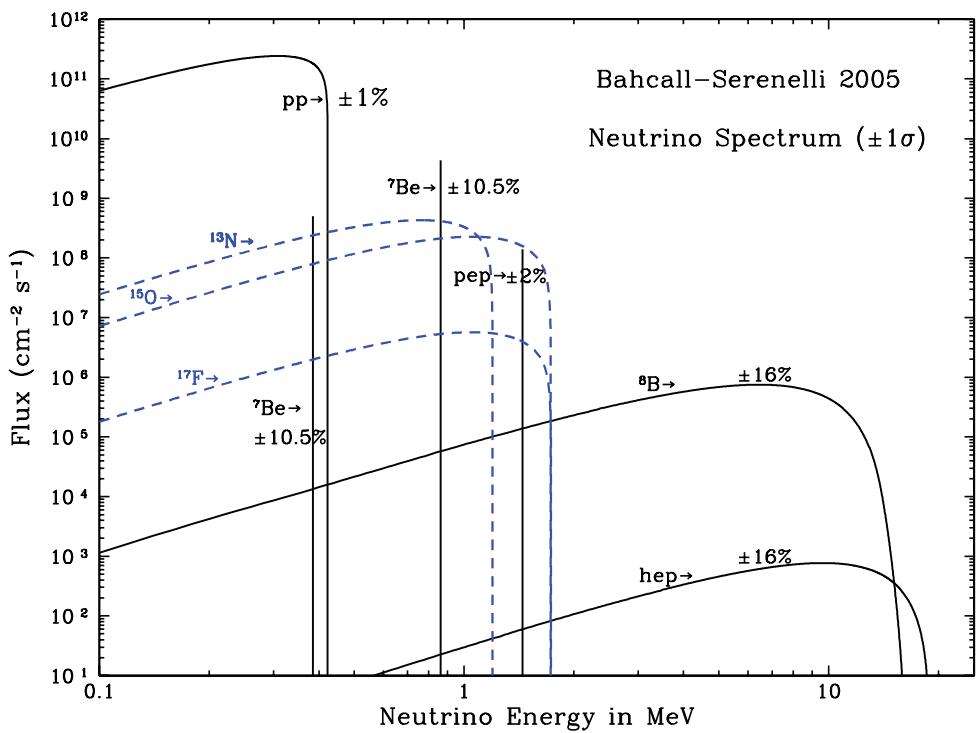


Figure 1.1: Expected flux of solar neutrinos versus energy. For continuous spectra, fluxes are given differentially ($\text{cm}^{-2}\text{s}^{-1}\text{MeV}^{-1}$) [7].

neutrino precision experiments in general as their point of origin is relatively well known.¹ The vast majority of those particles are electron antineutrinos from β^- decay, but there may also be very rare electron neutrinos from β^+ decay or electron capture [10].

Neutrinos from nearby core-collapse supernovæ (~ 10 MeV) would be very valuable because they could be used to improve limits on various neutrino properties such as the magnetic dipole moment and neutrino-neutrino cross-sections, and to probe the supernova explosion mechanism, but they are not very useful for high-precision oscillation studies because their initial flavor distribution (after leaving the supernova remnant) is model-dependent [11].

Neutrinos from the source regions of primary cosmic rays are very interesting from an astroparticle physics point of view because they provide complimentary information compared with cosmic rays and photons of various spectra. However, they are not useful to this analysis because even if their point of origin is known such as the recently identified blazar TXS 0506+056 [12], the exact source distance and the initial flavor composition are not known well enough for precision measurements; for instance, at a redshift of $z = 0.3365 \pm 0.0010$, the uncertainty in the (luminosity) distance to TXS 0506+056 amounts to over 5 Mpc. With high statistics they could still be used to probe quantum decoherence effects (see section 1.3.1). Their contribution quickly becomes negligible below about 300 TeV [13][14], so they can be easily suppressed by removing high-energy events and are not relevant to this work.

A hypothetical neutrino population are decay or annihilation products of yet undiscovered exotic particles. A prime candidate for such particles are WIMPs. Their energy spectrum highly depends on the particle mass and production process [15], and they would most likely only be valuable for oscillation studies if they are recorded with high statistics and if the process is well understood.

The last population in this list are primordial neutrinos that form the cosmic neutrino background (CNB). Analogous to their photon equivalents that form the better known cosmic microwave background (CMB), they froze out of the thermal equilibrium with other particles shortly after the big bang [16]. Due to the expansion of the universe, their energy today is expected to lie well below 1 meV, and they have not yet been directly detected [17].

¹A reactor core is on a scale of 10 m, but often there are multiple cores and experiments can usually not distinguish between those.

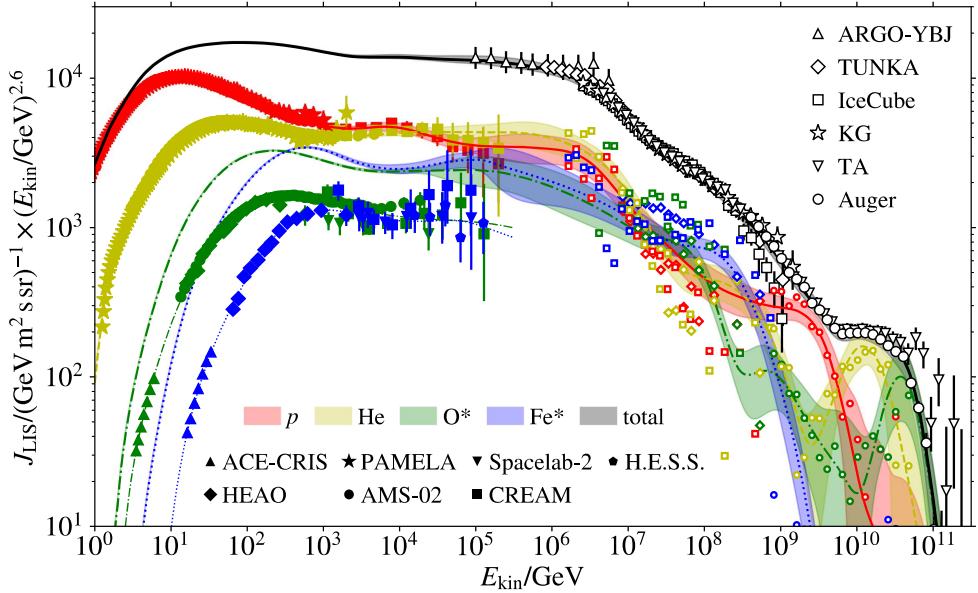


Figure 1.2: All-particle cosmic-ray spectrum vs. energy per nucleus, weighted by $E^{2.6}$ for better readability, and spectra of its four components. O^* and Fe^* indicate groups of multiple elements, contributions of O and Fe alone are shown up to 10^5 GeV [18]; updated from ref. [19].

1.2 Properties of Atmospheric Neutrinos

1.2.1 Energy Dependence

Cosmic-ray Energy Spectrum

As stated above, atmospheric neutrinos originate in particle showers induced by cosmic rays in the Earth's atmosphere. Because of this, the energy spectrum of cosmic rays is fundamental for the energy dependence of atmospheric neutrinos. While the origin of ultra-high-energy cosmic rays and of some of the features of the spectrum have not yet been fully understood to this date, the spectrum has been measured well by many independent experiment. Above an energy of 20 GeV, below which cosmic rays are produced and heavily influenced by the sun, the spectrum can be described by a broken power law, i.e., by a flux $\phi = \phi_0 E^{-\gamma}$, with the spectral index γ changing slightly at certain energies. The three most important changes have been dubbed knee (around 4 PeV), second knee (around 200 PeV), and ankle (around 5 EeV), see figure 1.2. Between 20 GeV and the knee, the spectral index has been measured to a value of $\gamma = 2.65 \pm 0.05$ [20]. Above the knee, the spectral index softens to about 2.95 ± 0.02 [21], up to the second knee, where it softens further to about 3.24 ± 0.08 [22]. At the ankle, it hardens to about 2.71 ± 0.10 [23].

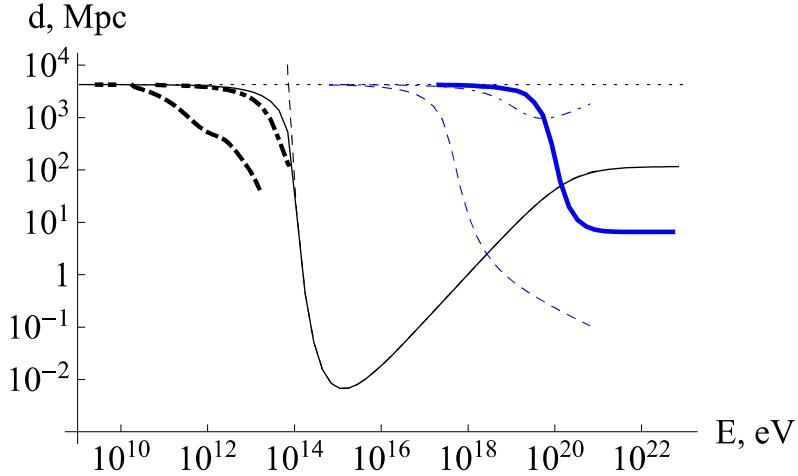


Figure 1.3: Mean free paths for protons interacting with photons of the CMB to produce pions (blue thick line) and electron pairs (blue dashed line). The blue dash-dotted line shows the mean energy loss distance for electron pair production. The horizontal line indicates the cosmological horizon, the remaining black lines deal with photons. For more information, see ref. [24].

At about 50 EeV, the flux falls off steeply, which was predicted as GZK cutoff, named after the theoretical physicists Greisen, Zatsepin, and Kuzmin. At this energy, cosmic-ray protons that interact with photons of the cosmic microwave background (CMB) have enough invariant mass to form proton- or neutron-pion pairs, which often occurs via the Delta resonance: $p + \gamma_{\text{CMB}} \rightarrow \Delta^+ \rightarrow p + \pi^0 \vee n + \pi^\pm$. The mean free path of this process depends on the proton energy and levels out at about 10 Mpc above 500 EeV, see figure 1.3 [24]. In one interaction, the proton loses at least 20% of its energy in the single-pion cases mentioned above, and on average 60% in cases where multiple pions are produced [25].

Similarly, protons that interact with CMB photons can undergo the Bethe-Heitler process and produce electron pairs: $p + \gamma_{\text{CMB}} \rightarrow p + e^+ + e^-$ [24]. The threshold energy for this process is smaller (figure 1.3), but unlike pion production, a single pair production only causes small changes in energy and direction of the proton. Because of this, one can define a mean energy loss distance after which a proton would have lost on average $1/e$ of its initial energy. This distance is shown in the same figure, and is comparable to the mean free path of the pion production.

For heavier nuclei, the GZK cutoff is expected to occur at slightly higher energies proportional to their nucleon number A , but they additionally undergo photo-disintegration, leading to a similar cutoff [18].

The analysis described in this work discards neutrinos with an estimated energy above 50 TeV. Even though there is no upper bound on the energy of a cosmic-ray primary that can induce a neutrino of a given energy, the overall steeply falling energy spectrum of the cosmic rays causes the uncertainties of the high-energy features to be unimportant for this analysis.

Resulting Spectrum of Atmospheric Neutrinos

The three most important processes for neutrino generation are the decays of charged pions, of kaons, and of muons. Mesons can either decay or interact non- elastically with particles of the atmosphere or the ground. If they interact, they do not directly produce a neutrino, but instead further the hadronic particle shower with their energy distributed among the outgoing particles, which themselves again have a chance to decay and produce neutrinos of lower energies, depending on their type. The probability of a particle to interact instead of decaying depends on its energy and on its surroundings. Firstly, the interaction cross-sections are energy-dependent, with cross-sections rising towards higher energies. Figure 1.4 shows total cross-sections of protons with pions and kaons up to 2.6 EeV. It can be seen that in the energy range of about 10^2 to 10^5 GeV, the cross-sections increase by a factor of two, which is a significant but small effect.

A more important effect is the relativistic time dilation as observed in the atmosphere reference frame.² It increases the mean lifetime of a particle with mass m and energy E to

$$\tau = \frac{E}{m} \tau_0, \quad (1.2)$$

which can be orders of magnitude above its mean lifetime at rest τ_0 . In first order, an increased lifetime will inversely proportionally lower the decay probability of the particle, so this effect can have a large impact for particles with non-negligible interaction probabilities.

For charged pions, $\tau_0 = 26.03$ ns. With a mass of 139.57 MeV, their typical decay length at 1 TeV (ignoring energy loss) is about 56 km. Long-lived kaons have similar lifetimes of 12.38 ns (K^\pm) and 51.16 ns (K_L^0). However, their higher masses of $m_{K^\pm} = 493.7$ MeV and $m_{K^0} = 497.6$ MeV lead to typical decay lengths of 8 km (K^\pm) and 31 km (K_L^0) at 1 TeV. With the pion and kaon cross-sections being very similar, this is the reason why kaons become more important for the production of atmospheric neutrinos at high energies that can be seen in figure 1.5.

The short-lived kaon CP-eigenstate K_S^0 has a much shorter lifetime of $\tau_0 = 89.5$ ps which greatly increases its chances to decay, but it decays almost exclusively hadronically into pions; its semi-leptonical branching ratio that also yields neutrinos is only 0.17% [7]. The mean lifetime of muons at rest is $2.197\ \mu$ s and their mass 105.66 MeV [7], so at 10 GeV, their typical decay length is 62.3 km.³ As the usual interaction height of vertically incoming cosmic-ray particles in the Earth's atmosphere is on the order of 15 km [27], and

²For other reference frames, the effect has to be attributed to time dilation, length contraction, or a combination thereof, but the outcome is independent of the reference frame.

³Note that this estimation does not take into account the energy loss of the muon in the atmosphere (2 MeV/cm² for a minimum ionizing particle), because it is a subdominant effect that can be neglected for the qualitative explanation of the production mechanism.

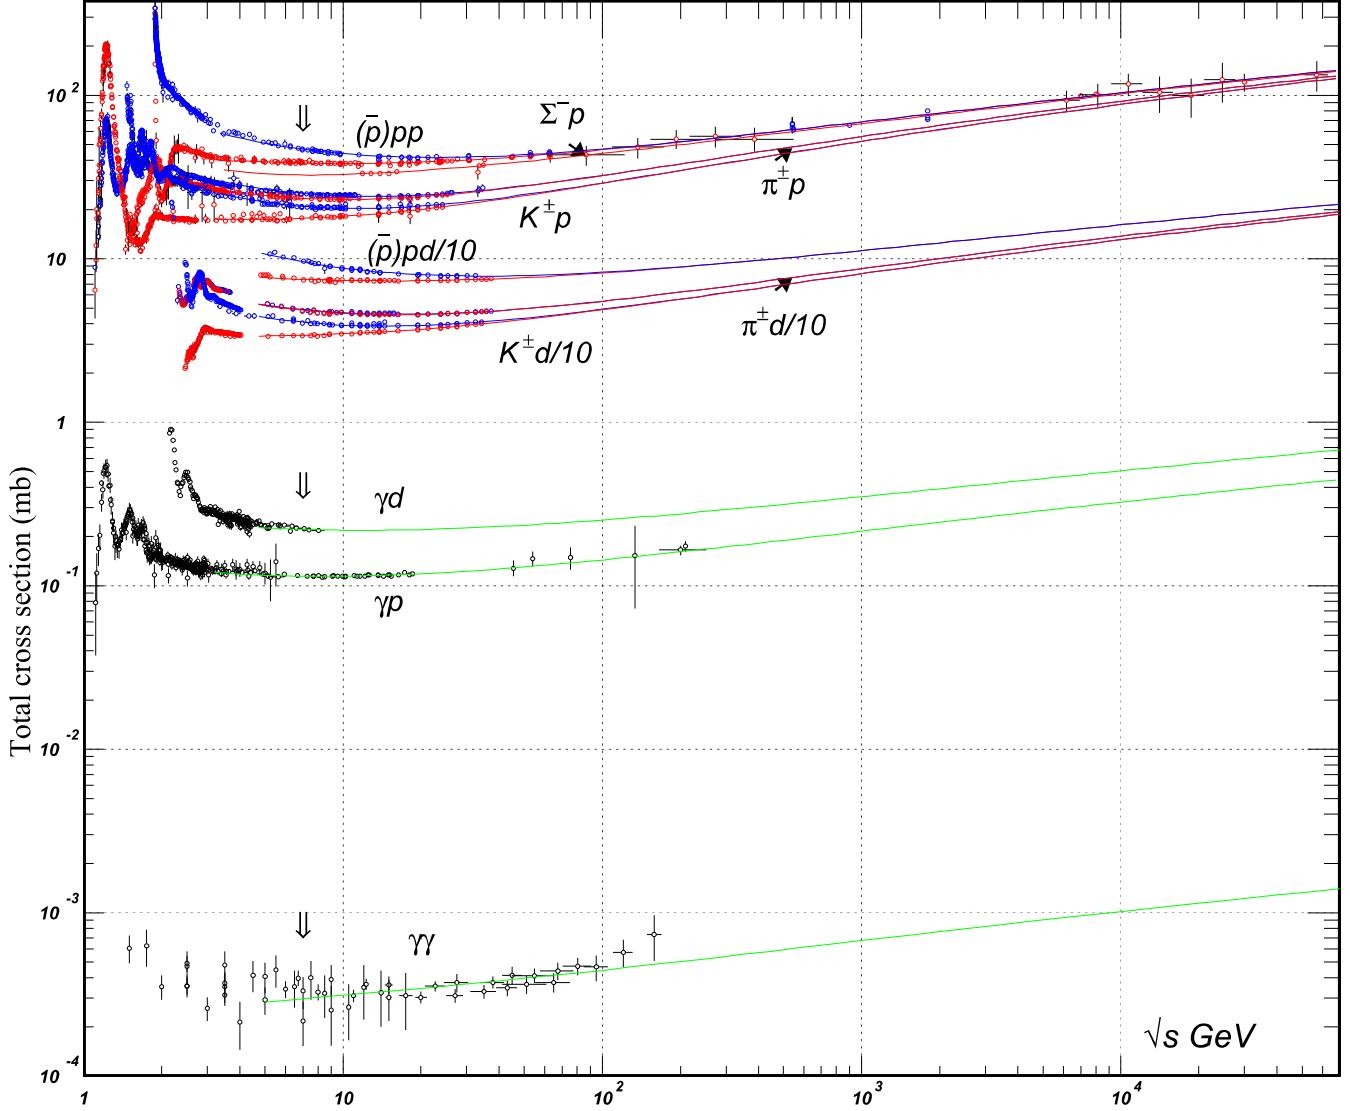


Figure 1.4: Total cross-sections of hadronic and photonic processes versus invariant mass. Invariant mass of fixed-target arrangements such as cosmic rays impacting the atmosphere computes as $\sqrt{s} = \sqrt{2Em + 2m^2}$, so for $m = m_p$, $10 \text{ GeV} \doteq 52.4 \text{ GeV}$, $100 \text{ GeV} \doteq 5330 \text{ GeV}$, and $1000 \text{ GeV} \doteq 533 \text{ TeV}$ [7].

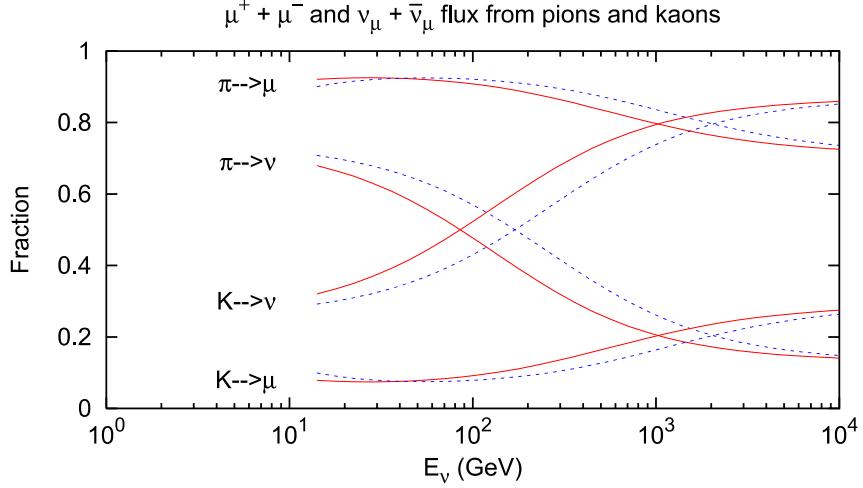


Figure 1.5: Relative contributions of pion and kaon decays to the muon and muon neutrino fluxes versus particle energy. The solid red lines indicate fractions for vertical tracks, the dashed blue lines were computed for a zenith angle of 60° [26].

the lifetime scales linearly with energy, the production of neutrinos from muon decays gets increasingly less important at energies above a few dozen GeV, since muons relatively quickly lose their energy when they reach the ground.

Overall, the energy dependence of these two processes disfavor atmospheric neutrino production at high energies and causes the energy spectrum of atmospheric neutrinos to be steeper, i.e., to feature fewer high-energy particles, than the energy spectrum of the cosmic rays that induce the particle showers. Ignoring muon decay, a parameterization of these effects is given by the Gaisser equation [28]

$$\frac{dN_\nu}{dE_\nu} = \phi_0 E_\nu^{-\gamma} \left(\frac{\mathcal{A}_{\pi\nu}}{1 + \mathcal{B}_\pi \cos(\theta^*) E_\nu / \epsilon_\pi} + \frac{\mathcal{A}_{K\nu}}{1 + \mathcal{B}_K \cos(\theta^*) E_\nu / \epsilon_K} \right). \quad (1.3)$$

The term $\phi_0 E_\nu^{-\gamma}$ is a simple parameterization of the primary cosmic-ray flux with spectral index γ ; $\mathcal{A}_{\pi\nu}$ and $\mathcal{A}_{K\nu}$ are constants that depend on shower kinematics and branching ratios, and \mathcal{B}_π and \mathcal{B}_K are similarly constants that depend on shower kinematics and the meson attenuation in the atmosphere. The energies $\epsilon_\pi \approx 125$ GeV and $\epsilon_K \approx 850$ GeV are so-called critical energies; far below these energies, the respective mesons are expected to always decay, while far above, their probability for decay decreases with their energy as described by (1.2) [28]. The critical energies depend on the mean free path of the mesons in the atmosphere, which depends on the atmospheric density and therefore on the local temperature.

The term $\cos(\theta^*)$ is the cosine of an effective, averaged zenith angle θ^* : The zenith angle θ is not measured relative to the center of the Earth, but relative to the position of the detector which is close to the surface. Because of this, especially when close to the horizon, the curvature of the Earth becomes relevant and causes the zenith angle between

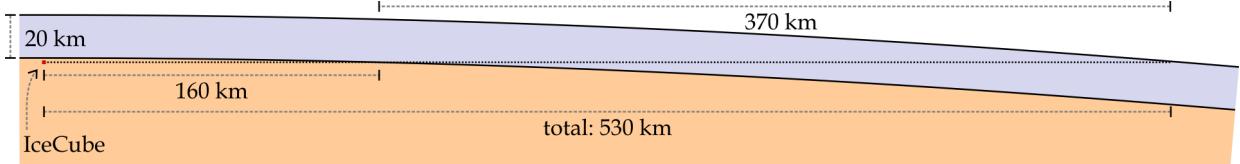


Figure 1.6: To-scale sketch of the Earth with 20 km of atmosphere. The IceCube detector is positioned 2 km below the ground, and the lines indicate path lengths in the atmosphere and in the ground for horizontally arriving particles.

the particle track and the local atmosphere to change along the track. This is illustrated in figure 1.6, which shows the track of a particle arriving horizontally at a detector that is 2 km below the ground. The angle with which it passes the atmosphere at a height of 20 km is 4.76° , the angle at ground level is only 1.44° . The effective angle θ^* has been determined as function of θ by using Monte-Carlo methods to convolute the local zenith angles with the interaction probabilities at the corresponding atmospheric heights [27].

1.2.2 Angular Spectrum

As indicated by (1.3), the flux of atmospheric neutrinos does not only depend on the neutrino energy, but also on the zenith angle, which is defined as 0° for straight down-going and 180° for straight up-going particles. On the zeroth order, the angular distribution of atmospheric neutrinos is given by the arrival directions of the cosmic rays that induce the particle showers. The directions of neutrinos emerging from the resulting particle showers do not deviate significantly from the initial cosmic-ray directions at energies above a few dozen GeV. For example, for neutrinos from pion decay, the mean angle between the cosmic-ray primary and the neutrino can be estimated to be

$$\langle \psi_{\text{CR},\nu} \rangle \approx \frac{\langle p_\pi^\perp \rangle}{p_\pi} \approx \frac{0.3 \text{ GeV}}{4E_\nu} \approx 4.3^\circ \frac{\text{GeV}}{E_\nu},$$

where p_π^\perp is the momentum of the pion perpendicular to the primary, p_π its total momentum, and E_ν the neutrino energy [29].

At energies above about 100 GeV, at which the influence of solar magnetic fields becomes weak, the arrival directions of cosmic rays can be approximated well to be isotropic, although statistically significant deviations on a scale of up to 0.1% have been observed by multiple experiments and for many energies [30]. A much more important correction on the angular distribution of atmospheric neutrinos are the effects of the particle shower propagation in the atmosphere. The change of air density per path length depends on the angle at which the primary particle enters the atmosphere. If the angle is flat, the shower particles have a better chance to decay and produce neutrinos instead of interacting because they spend more time in a less dense environment than particles in showers

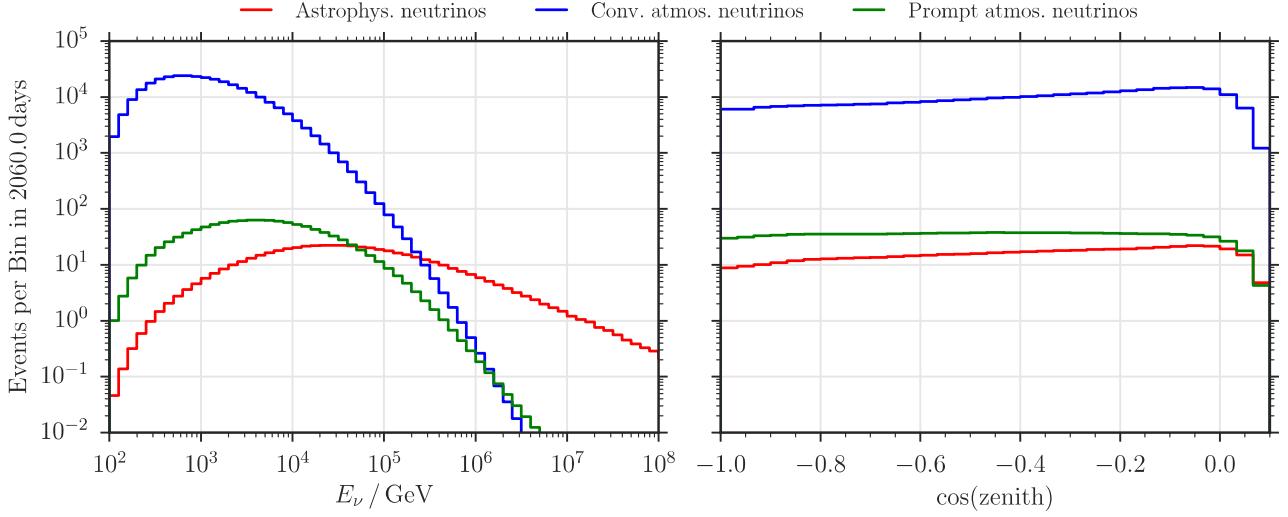


Figure 1.7: Expected energy (left) and angular (right) distributions of conventional atmospheric, prompt, and extra-terrestrial neutrinos from an analysis of six years of IceCube data. For details on the assumed fluxes, see ref. [13].

perpendicular to the atmosphere. As the decay probability depends on the energy as explained in section 1.2.1, the angular spectrum is correlated to the neutrino energy.

Because of the approximate symmetry of the Earth along its rotational axis, only the zenith angle matters, while the azimuthal distribution can be expected to be uniform. The angular distributions of atmospheric and extra-terrestrial neutrinos as measured by the IceCube detector are shown in figure 1.7. The plot also features the distribution of so-called prompt neutrinos from decays of heavier mesons, which are described in section 1.2.5. Where necessary for clarity, the atmospheric neutrinos from light mesons and muons are named conventional.

1.2.3 Flavor Composition

To predict the flavor composition, one has to differentiate between muon, pion, and kaon decays. As all of these particles are lighter than tau leptons ($m_\tau = 1.777 \text{ GeV}$), tau neutrinos are not produced in their decays. To make the discussion less verbose, this section will not differentiate between particles and antiparticles.

Muons produce electron and muon neutrino flavors in a ratio of 1:1 as shown in this chapter’s introduction, but they seldom decay at high energies as explained above. Pions stay relevant up to much higher energies than muons. As cited above, they predominantly produce muon neutrinos (see (1.1)). The decay of charged pions into electrons and electron neutrinos is suppressed strongly because the smaller mass of electrons compared to muons

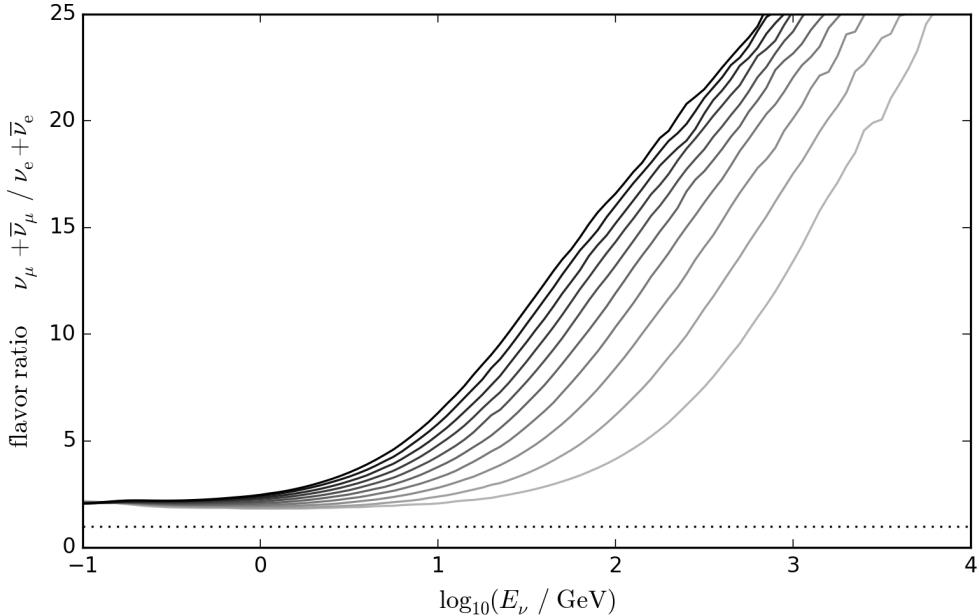


Figure 1.8: Flavor ratio of muon neutrinos and antineutrinos compared to electron neutrinos and antineutrinos at ground level according to the HKKMS2006 model [32][33]. The model does not include detector effects and has to be extrapolated towards high energies and ratios. The lines are equally distributed in the cosine of the zenith angle θ between vertical (black) and horizontal (light gray) trajectories, with the model being symmetric between up- and downgoing neutrinos.

makes them more difficult to be boosted to achieve helicity conservation [31], leading to an electron branching ratio of only about 0.012%.

Kaons have a higher mass and therefore more diverse decay options. As stated above, the “short” eigenstate of the neutral kaon, K_S^0 , decays mostly hadronically into pions. Only 0.17% of the decays produce neutrinos semi-leptonically with an expected flavor ratio of 2:1 [7]. K_L^0 produce neutrinos 67.6% of the time, with a flavor ratio of 3:2 [7]. The branching ratio of charged kaons into products that contain neutrinos is about 75.0% with a flavor ratio of roughly 1:13, thereby highly favoring muon neutrinos.

For the overall flavor ratio, not only the branching ratios are relevant, but also the relative abundances and decay probabilities of the shower particles. Figure 1.8 shows the flavor ratio of conventional atmospheric neutrinos at generation as function of the neutrino energy as predicted by the HKKMS2006 model, which computes atmospheric neutrino fluxes based on measured cosmic-ray fluxes and shower models [32][33]. As expected from the branching ratios and the muon lifetime, muon neutrinos are over five times more abundant than electron neutrinos for vertical trajectories above 10 GeV, and the flavor distribution for horizontal trajectories is shifted by slightly more than one magnitude towards higher energies.

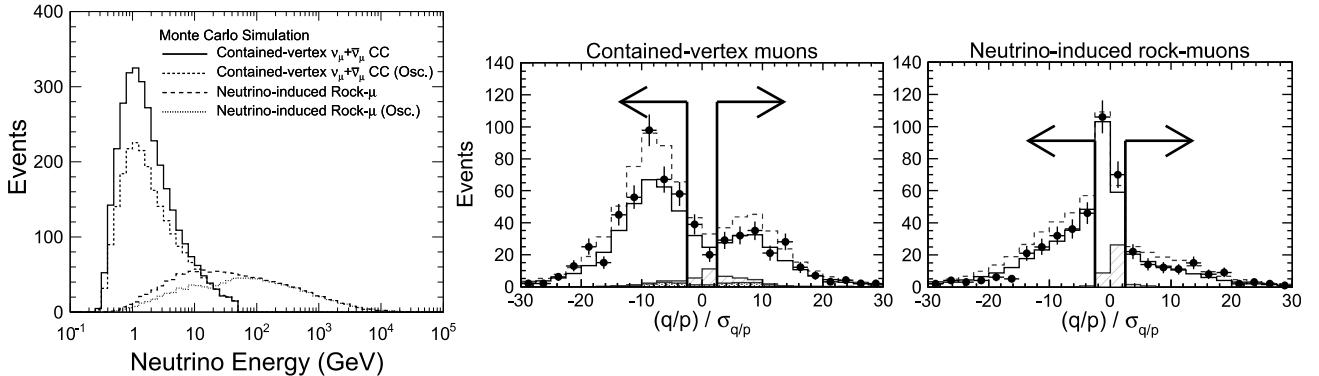


Figure 1.9: Measurement of the numbers of muon antineutrinos and neutrinos of two independent samples by the MINOS experiment [35]. The first plot gives the simulated energy distributions of the two samples, the other plots show the measured charge distributions of the neutrino-induced muons. In the latter plots, solid lines indicate MC predictions with oscillation taken into account, dashed lines are predictions without oscillation, striped histograms indicate misidentified charge, and dotted histograms indicate atmospheric muons.

1.2.4 Antineutrino-Neutrino Ratio

The antineutrino-neutrino ratio at production is relevant for this analysis because oscillation effects differ between particles and antiparticles. Unfortunately, the IceCube detector can not differentiate between them, which is why a good theoretical prediction is important.

The ratio is determined mostly by the composition of the incoming cosmic rays and by the cross-sections of the various shower particles. The hadronic showers in which atmospheric neutrinos are produced almost always originate from positive nuclei as hadronic matter is far more abundant than hadronic antimatter (at the very least on galaxy-cluster scales [34]). Therefore, average showers have a positive net charge. The positive mesons decay (semi-)leptonically into positive (i.e., anti-)muons and muon neutrinos.

The muon-flavor antineutrino-neutrino ratio is experimentally accessible to some experiments such as MINOS and soon NO ν A, which feature magnetic fields strong enough to discern neutrino-induced muons from antimuons [35]. One such measurement is illustrated in figure 1.9, yielding a ratio of $0.63^{+0.09}_{-0.08}$ at a median energy of about 70 GeV. Unfortunately, direct measurements only cover low energies, because either the experiments capable of distinction between antimuons and muons are too small to sample enough events at the lower fluxes prevalent at higher energies, or the experiments lose this capability because the muons' gyroradii become too large.

The ratio can also be calculated indirectly but with higher precision from the charge ratio of atmospheric (i.e., not neutrino-induced) muons. A CMS measurement indicates that muon neutrinos are 1.277 ± 0.005 times as abundant as muon antineutrinos between muon

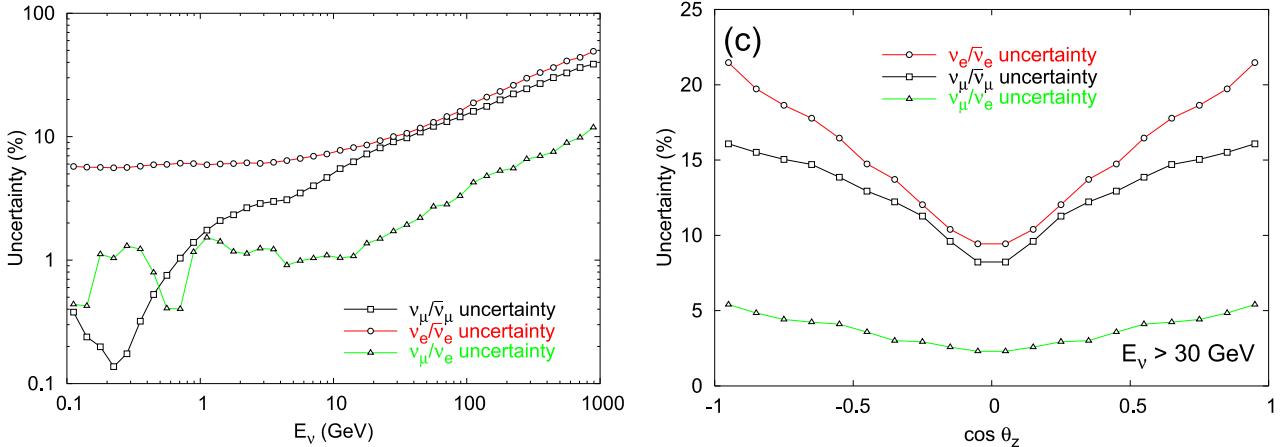


Figure 1.10: Uncertainties in the electron and muon neutrino-antineutrino ratios and the electron-muon flavor ratio as function of energy (**left**) and as function of zenith angle (**right**, for $E_\nu > 30$ GeV) according to ref. [38].

energies of 10 and 100 GeV, increasing to over 1.32 above 400 GeV [36]. However, even though the precision is higher, despite the two-body kinematics a neutrino originating from pion decay can have anywhere between 74.5% and almost down to 0% of the energy of the associated atmospheric muon, depending on the angle of emission in the pion rest frame [37]. Because of this, the distribution is smeared out with regard to the neutrino energy, and the precision of such measurements does not directly translate to the uncertainty on the antineutrino-neutrino ratio.

The uncertainty estimates used as reference in this work are shown in figure 1.10. They are based on calculations starting from cosmic-ray flux measurements. The discussion and implementation of these uncertainties can be found in section 6.2.4. Finally, figure 1.11 shows the prediction for the ratio by the aforementioned HKKMS2006 model as a function of neutrino energy and zenith angle.

1.2.5 Prompt Neutrinos

A population of atmospheric neutrinos that was neglected so far are the so-called prompt neutrinos. They originate from heavier mesons that also contain charm or bottom quarks. These mesons are known to decay much faster than their light-quark counterparts, which causes their expected energy and angular spectra to be closer to the primary cosmic-ray spectrum. Although it is known from accelerator experiments that these mesons exist and that their decays can produce neutrinos, they have not yet been discovered experimentally in cosmic-ray induced air showers because of the strong background of light-meson atmospheric neutrinos, which are often referred to as conventional (atmospheric) neutrinos.

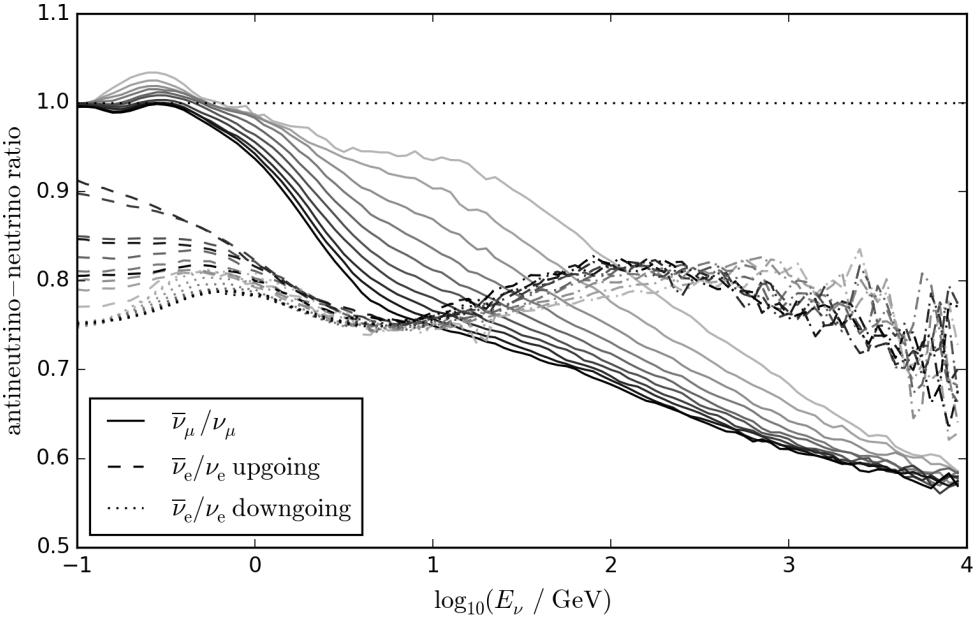


Figure 1.11: Antineutrino-neutrino ratio of muon and electron flavors at ground level according to the HKKMS2006 model, analog to figure 1.8. Again, black depicts vertical and light gray horizontal trajectories. The model differentiates between up- and downgoing electron neutrinos below 10 GeV, but is symmetric otherwise.

Evidence for prompt neutrinos could be provided by unambiguously measuring a population of neutrinos with their predicted energy spectrum and angular distribution (see figure 1.7). The significance of such measurements could be improved drastically by the detection of tau neutrinos, as pions and kaons cannot decay into tau leptons and associated tau neutrinos because their mass is too low.⁴ Tau neutrinos can however be produced as decay products of charm or bottom mesons in air showers [7]. Their detection alone is not sufficient as evidence for prompt neutrinos because they can be produced in neutrino oscillations of conventional atmospheric or (more likely) extraterrestrial neutrinos as described in section 1.3, but together with energy and angular information, they could provide strong evidence.

Prompt neutrinos are part of some of the atmospheric models used in this analysis (section 6.1.1), but they have not been fully incorporated in the fitting scheme because their flux is unknown and so far undetectable even by well-suited analyses [13][14]. The correct implementation of their flux will however become vital in future multi-year searches, where statistics in high-energy regions will become strong enough to notice the sub-dominant contribution of prompt neutrinos.

⁴The same is true for all but the highest-energetic excited states of the light-quark mesons, and the latter ones decay exclusively through strong interactions, not producing any neutrinos.

1.3 Standard Neutrino Oscillations

According to our current knowledge, neutrinos are elementary particles. They are electrically neutral and – as they are leptons – do not carry a color charge, so they do not participate in electromagnetic or strong interactions. They only interact via the weak interaction and via gravity.

Often, it is stated that since neutrinos are produced in weak interactions, their initial states are pure flavor eigenstates that are non-trivial linear combinations of their mass eigenstates. Formally, this is wrong; as the mass eigenstates have distinct masses, such a linear combination can not be an energy and momentum eigenstate at the same time. The correct description of the production process yields a mixed state in which the neutrino is entangled with the other interaction products. However, determining the density matrix that describes the mixed state and projecting out the other states reveals that the pure-state approach is a good approximation that yields the same results as shown in ref. [39], so it will be used in the following.

In this approach, the neutrino flavor eigenstates $|\nu_\alpha\rangle$, $\alpha \in \{e, \mu, \tau\}$ can be expressed through the neutrino mass eigenstates $|\nu_j\rangle$, $j \in \{1, 2, 3\}$; being non-trivial linear combinations, the states are linked with each other by the non-diagonal unitary Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix U :

$$|\nu_\alpha\rangle = \sum_j U_{\alpha j}^* |\nu_j\rangle,$$

where $*$ denotes the complex conjugate, which has to be dropped for antineutrinos. U can be parameterized as the product of three rotation matrices⁵ with additional complex phases. As these matrices do not commute in general, their ordering is important and the usual convention is

$$U = R_{23}(\theta_{23})R_{13}(\theta_{13}, \delta_{13})R_{12}(\theta_{12})J \quad (1.4)$$
$$:= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{13}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{13}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\alpha_1/2} & 0 \\ 0 & 0 & e^{i\alpha_2/2} \end{pmatrix},$$

where $c_{ij} := \cos(\theta_{ij})$ and $s_{ij} := \sin(\theta_{ij})$. The three parameters θ_{12} , θ_{13} , and θ_{23} are called mixing angles, δ_{13} is the Dirac-type phase that governs CP violation, and in case of neutrinos being Majorana particles, the Majorana-type phases α_1 and α_2 in the diagonal matrix J can be different from zero. Note that Dirac-type phases can also occur if neutrinos are Majorana particles [40].

Using δ_{13} as Dirac-type phase is merely a convention. In general, for n flavors, there are

⁵Actually, those matrices are not standard-convention rotation matrices, but matrices generated using Gell-Mann matrices, so the non-diagonal signs of R_{13} are inverted.

$n(n-1)/2$ mixing angles and the same number of phases, of which $(n-1)(n-2)/2$ are of Dirac type; additional Dirac-type phases are redundant as they could be absorbed by redefinition of the fermion fields without observable effects.

1.3.1 Vacuum Oscillation

The vacuum propagation of a neutrino is described by an appropriate relativistic wave equation. Even though neutrinos are spin-1/2-particles and their Dirac/Majorana nature has not yet been determined, it can be shown that for all those cases the appropriate wave equations converge to one that looks like a Schrödinger equation [41]:

$$i\frac{d}{dt}|\nu_j\rangle = \widetilde{H}_0|\nu_j\rangle, \quad (1.5)$$

$$\widetilde{H}_0 = \text{diag}(E_1, E_2, E_3), \quad (1.6)$$

using natural units with $c = \hbar = 1$.

Terms that are proportional to the identity matrix only produce a global phase, but do not cause differences between the phases of different neutrino states. Thus, they do not affect the composition of a neutrino state and can be dropped. For the same reason, Majorana-type phases are irrelevant for neutrino oscillation phenomena and are therefore set to 0 for this work, such that J in (1.4) becomes identity and can be dropped for simplicity. Using this, the Hamiltonian (1.6) can be simplified to

$$\widetilde{H}_0 := \text{diag}(0, E_2 - E_1, E_3 - E_1) \quad (1.7)$$

without changing the results of (1.5).

In the ultra-relativistic limit, $v = c$ holds, so the equation can be easily translated from describing the time dependence to describing the experimentally more intuitive positional dependence:

$$\frac{d}{dt}|\nu_j\rangle \equiv \frac{d}{ds}|\nu_j\rangle$$

In the same limit, the momenta $p_j \gg m_j$ of the various states can be assumed to be identical. Therefore, the energy terms can be expressed through a common neutrino energy $E_\nu \approx p_1 = p_2 = p_3$ by

$$E_j = \sqrt{p_j^2 + m_j^2} \approx p_j + \frac{m_j^2}{2p_j} \approx E_\nu + \frac{m_j^2}{2E_\nu}. \quad (1.8)$$

This approximation holds very well for all neutrinos above about 1 MeV, including the atmospheric neutrinos that are the scope of this work.

Using (1.8), (1.7) can be rewritten as

$$\widetilde{H}_0 := \frac{1}{2E_\nu} \text{diag}(0, \Delta m_{21}^2, \Delta m_{31}^2),$$

where $\Delta m_{ji}^2 = m_j^2 - m_i^2$. From this equation it can be seen that the absolute neutrino mass scale does not influence oscillation phenomena; only the differences between the values are relevant.

The Schrödinger equation (1.5) can be translated into the flavor basis, which is more useful as it is the basis in which neutrinos partake at weak interactions that are responsible for both their creations and possible experimental detections:

$$\frac{d}{dx} |\nu_\alpha\rangle = \frac{-i}{2E_\nu} H_0 |\nu_\alpha\rangle \quad (1.9)$$

$$H_0 = U \text{diag}(0, \Delta m_{21}^2, \Delta m_{31}^2) U^\dagger. \quad (1.10)$$

The dagger symbol † indicates the adjoint matrix, i.e., the complex conjugate of the transposed matrix.

The probability of a neutrino of initial state $|\nu_\alpha\rangle$ to be in state $|\nu_\beta\rangle$ after traveling a distance L can then be calculated as

$$P_{\alpha \rightarrow \beta} = |\langle \nu_\beta(L) | \nu_\alpha \rangle|^2.$$

The vacuum solution to equation (1.9) is the well-known plane-wave solution

$$P_{\alpha \rightarrow \beta} = \sum_{j,i} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* e^{-i \frac{\Delta m_{ij}^2 L}{2E_\nu}} \quad (1.11)$$

$$\begin{aligned} &= \delta_{\alpha\beta} - 4 \sum_{j, i > j} \Re(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E_\nu}\right) \\ &\quad + 2 \sum_{j, i > j} \Im(U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin\left(\frac{\Delta m_{ij}^2 L}{2E_\nu}\right) \end{aligned} \quad (1.12)$$

with Kronecker delta $\delta_{\alpha\beta}$. This vacuum solution is not an adequate approximation for this work as explained below, but it can still be used to demonstrate the phenomenological behavior of neutrino oscillations. For a more detailed derivation see ref. [42].

Thus, neutrino oscillations in vacuum can be described as a superposition of various sine functions. The second sum in (1.12) vanishes if all CP-violating phases are zero. Since CP violation is not relevant for this analysis, the terms are dropped for now (see appendix A). The sine function of the remaining terms can be translated from natural units to SI units:

$$\sin^2\left(\frac{\Delta m_{ij}^2 L}{4E_\nu}\right) = \sin^2\left(\frac{\Delta m_{ij}^2 L c^3}{4\hbar E_\nu}\right) = \sin^2\left(1.266933 \frac{\Delta m_{ij}^2 L}{E_\nu} \frac{\text{GeV}}{\text{eV}^2 \text{km}}\right),$$

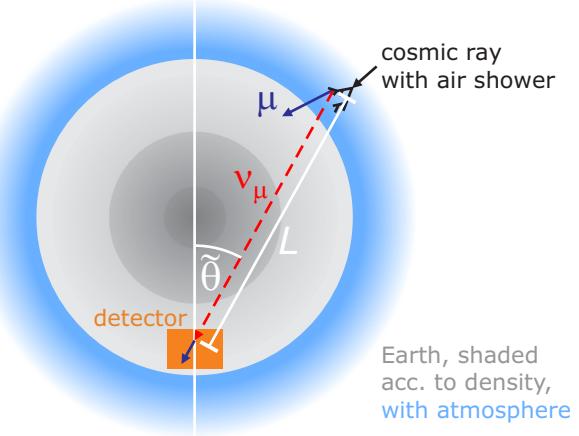


Figure 1.12: Sketch of an atmospheric neutrino interacting inside an underground detector, illustrating the relation between angle $\tilde{\theta} = \pi - \theta$ and baseline L .

where 1.266933 is the approximate value of $\text{GeV fm}/(4\hbar c)$. The two non-constant parameters of the sine's argument are L and E_ν , which in this vacuum solution only appear as a quotient. Therefore, the oscillations become faster, the smaller the neutrino energy E_ν is, and oscillation probabilities are identical for neutrinos with equal values of $\frac{L}{E_\nu}$. The amplitude of the oscillations, i.e., the amount by which the probability for a given flavor change can vary, is given solely by elements of U , and therefore only depends on the mixing angles and phases, but not on the neutrino mass differences [43].

For atmospheric neutrinos measured on Earth, the distance L , which is also known as the oscillation baseline, can be expressed as function of the zenith angle θ at the detector as illustrated in figure 1.12:

$$L = \sqrt{r_{\text{int}}^2 + r_{\text{det}}^2 - 2 r_{\text{int}} r_{\text{det}} \cos\left(\theta - \arcsin\left(\frac{\sin(\pi - \theta)}{r_{\text{int}} r_{\text{det}}}\right)\right)}, \quad (1.13)$$

where r_{int} is the interaction altitude and r_{det} the altitude of the detector, both relative to the center of the Earth. A visualization of this relation is given in figure 1.13. For this work, the Earth is assumed to be spherical and the atmosphere is modeled according to ref. [27]; a more detailed discussion of these aspects can be found in appendix C.

Using this relation, one can compute the expected oscillation probabilities as function of the two detector observables E_ν and θ . For the vacuum propagation discussed in this section, this could be done analytically, but for consistency with the following sections, the probabilities in figure 1.14 have been computed numerically using the tool nuCraft, which is described in appendix C. The figure shows the so-called disappearance probability for muon neutrinos, i.e., the probability for an atmospheric neutrino that was created with muon flavor to be detected as a neutrino of non-muon flavor.⁶ The dominant pattern is

⁶All figures and numbers are assuming normal neutrino mass hierarchy; inverted hierarchy would have little impact and would only complicate the discussion, see appendix B.

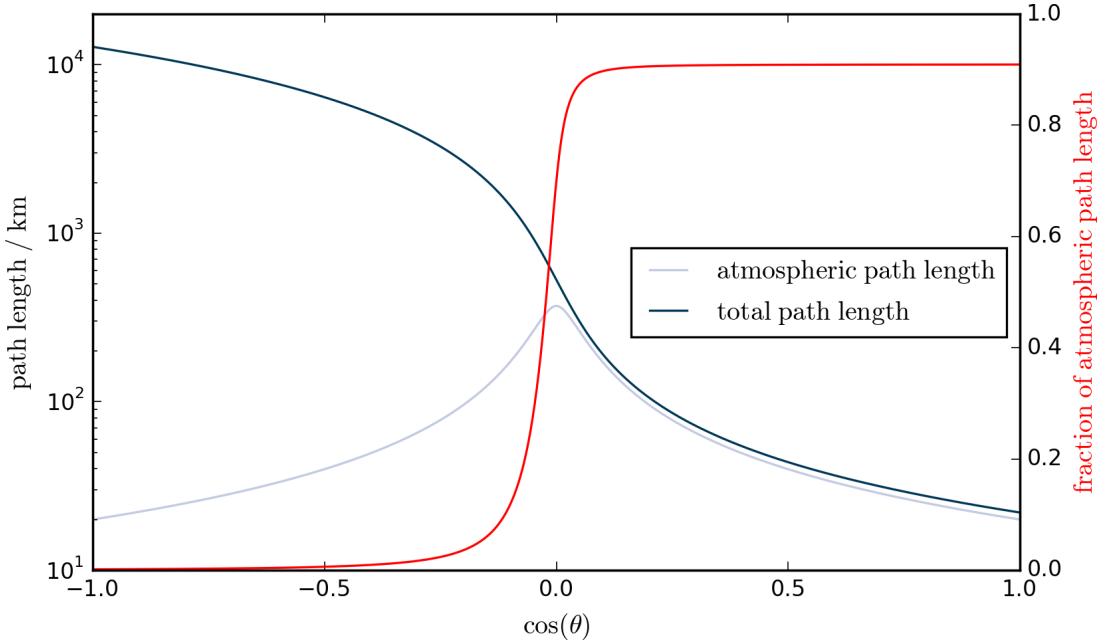


Figure 1.13: Simplified visualization of (1.13) under assumption of a fixed interaction height of 20 km above ground level.

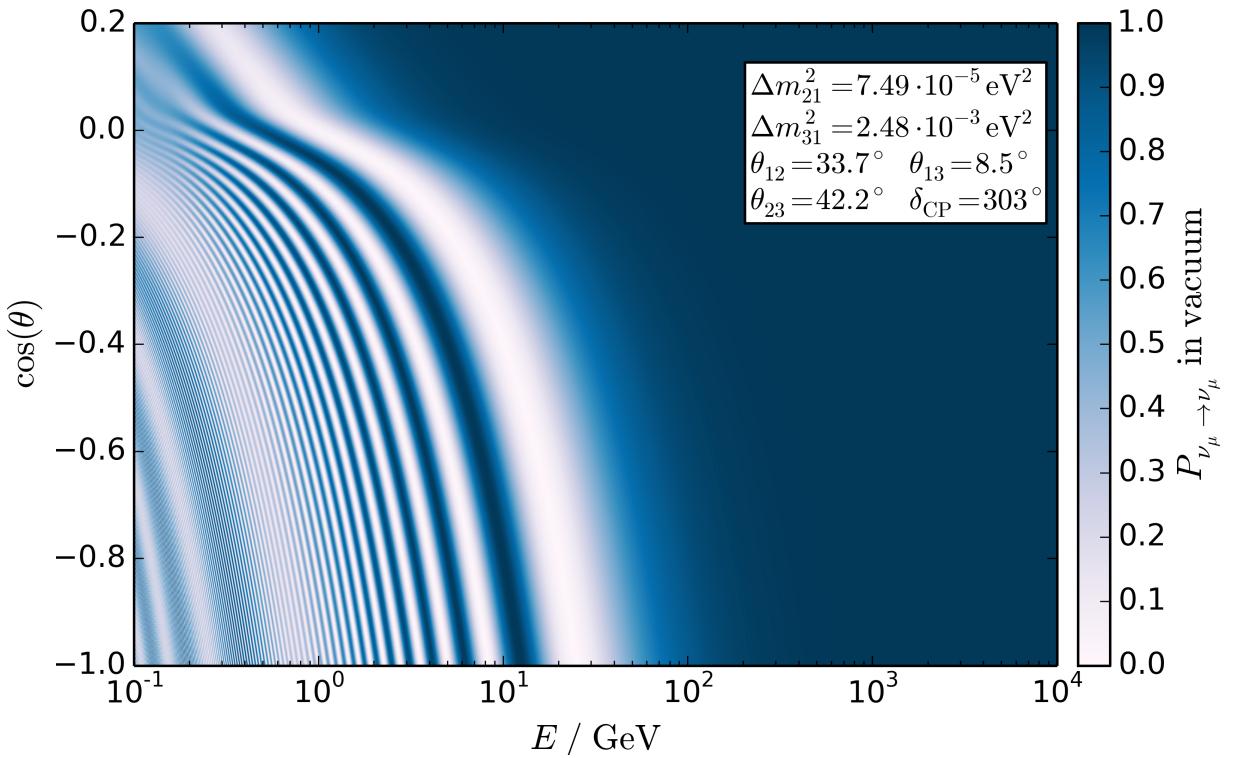
one sine term that corresponds to the oscillation of muon neutrinos to tau neutrinos, with its frequency given by $\Delta m_{32}^2 = 2.40 \cdot 10^{-3}$ eV (see table 1.1). It has its first oscillation minimum at around 25 GeV for vertically upgoing neutrinos and continues towards lower energies. At significantly lower energies around 0.8 GeV, an additional sine pattern has its first oscillation minimum; it belongs to the oscillation of muon neutrinos to electron neutrinos, given by $\Delta m_{21}^2 = 7.49 \cdot 10^{-5}$ eV, which is a factor of $\frac{\Delta m_{32}^2}{\Delta m_{21}^2} = 32.0$ below the first one. The muon-tau oscillations have been measured by a multitude of experiments, but the muon-electron channel would be very difficult to observe with atmospheric neutrinos because of the very low neutrino energies and because of the high energy resolution that would be required. The next section will however demonstrate that the muon-electron oscillations as shown in figure 1.14 do not occur in such regular manner if matter effects are taken into account.

For extremely large distances $\frac{L}{\text{km}} \gg 10000 \frac{E_\nu}{\text{GeV}}$ or small energies, respectively, the oscillations can be assumed to be averaged out: the neutrinos already went through many oscillation periods and for example a small uncertainty in the measurement of the neutrino energy that would only have a small effect on the oscillation probability per period now has a large impact and can make the oscillation probability unpredictable. Similarly, a spatial extent of the neutrino production region that is larger than an oscillation period causes the probability to smear out on average.

The resulting average can be computed easily by taking into account the fact that the sine and cosine functions in (1.11) disappear on average unless the argument of the functions

Table 1.1: Oscillation parameters for normal mass hierarchy according to the global fit [44]

parameter	value	σ_+	σ_-
$\theta_{12}/^\circ$	33.72	0.79	0.76
$\theta_{23}/^\circ$	42.2	2.2	1.4
$\theta_{13}/^\circ$	8.50	0.19	0.20
$\delta_{13}/^\circ$	303	39	50
$\Delta m_{21}^2/10^{-5} \text{ eV}^2$	7.49	0.19	0.17
$\Delta m_{31}^2/10^{-3} \text{ eV}^2$	2.477	0.042	0.042

**Figure 1.14:** Survival probability of atmospheric muon (anti-)neutrinos without taking into account matter effects, as a function of neutrino energy and zenith angle. The oscillation parameters are displayed in the plot and have been taken from the global fit [44] (also shown in table 1.1).

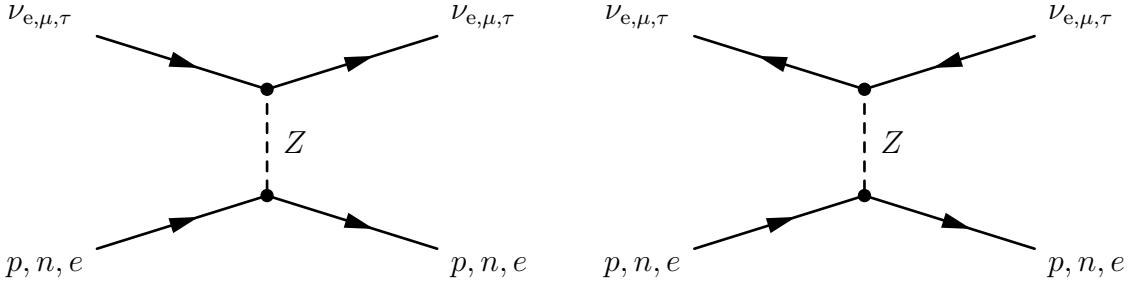


Figure 1.15: Feynman graphs of neutral-current elastic scattering of conventional neutrinos (left) and antineutrinos (right) on normal matter, lowest order only.

is constant, which it is only for $i = j$:

$$\langle P_{\alpha \rightarrow \beta} \rangle = \sum_j U_{\alpha j}^* U_{\beta j} U_{\alpha j} U_{\beta j}^* = \sum_j |U_{\alpha j}|^2 |U_{\beta j}|^2$$

This result is different from the flavor ratio of 1:1:1 that is sometimes assumed universally in literature, even though a source ratio of 1:2:0 as expected from pion and consequent muon decays would come close to 1:1:1 at Earth (namely 1.065:0.977:0.958 using table 1.1). If future measurements of the flavor ratio of astrophysical neutrinos deviate from 1:1:1 at Earth, this would provide interesting clues for the sources of these neutrinos (namely that they are not originating exclusively from uninterrupted pion decay). In addition to that, it could also provide strong limits for quantum decoherence induced by quantum gravity, which is predicted to shift the flavor ratio towards 1:1:1 independent of oscillation phenomena, but has not been observed so far [45].

1.3.2 The Induced Matter Potential

Even though neutrinos interact only weakly with matter, the vacuum propagation equation (1.9) established in the previous section needs to be modified to take into account the effects that matter has on neutrino oscillation probabilities. For a neutrino to be detectable, it needs to be conserved until it reaches the detector, so only elastic forward scattering processes are taken into account. In the neutral-current channel, the possible processes are elastic interactions of neutrinos of all three flavors with protons, neutrons, and electrons of the matter they traverse (t -channel, figure 1.15). Since the energy transfer has to be negligible, interactions are coherent which means that the neutrinos can not resolve the individual scattering partners. Accordingly, the scattering amplitudes have to be summed up, and cross-sections become relatively large. However, as only relative phases between different neutrino states cause oscillation effects, only matter interactions that differ between the states are relevant. Therefore, neutral-current interactions can be ignored in the three-flavor case [41].

In the charged-current channel, only the electron flavor is capable of elastic scattering, because normal matter does not contain muons or taus. The possible processes are shown

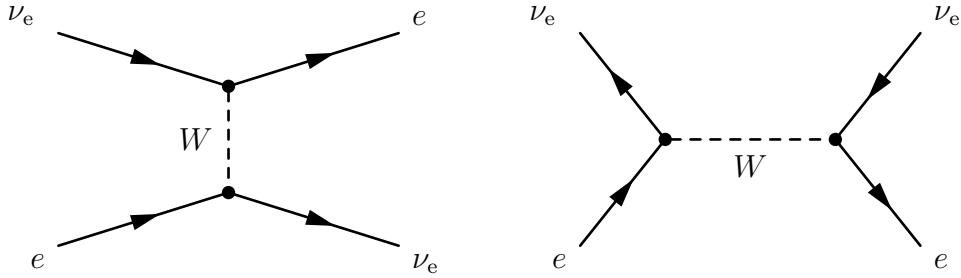


Figure 1.16: Feynman graphs of charged-current elastic scattering of neutrinos (**left**) and antineutrinos (**right**) on normal matter, lowest order only.

in figure 1.16: Electron neutrinos participate in this process using the t -channel, while electron antineutrinos use the s -channel [46]. However, the cross-section for the s -channel process is strongly suppressed by the large mass of the W^- , so this process is usually neglected.

The charged-current interactions are exclusive to one flavor and therefore induce a matter potential A_{CC} for the electron flavor that acts as an effective squared-mass difference:

$$\frac{d}{dx}|\nu_\alpha\rangle = \frac{-i}{2E_\nu}(H_0 + A)|\nu_\alpha\rangle \quad (1.14)$$

$$A = \frac{1}{2} \text{diag}(A_{\text{CC}}, 0, 0) \quad (1.15)$$

$$A_{\text{CC}} = \pm 2\sqrt{2}G_F \frac{\rho}{m_N} Y_e E_\nu \quad (1.16)$$

$$Y_e = \frac{n_e}{n_p + n_N}, \quad (1.17)$$

where G_F is Fermi's constant, ρ the local mass density, m_N the mean nucleon mass, and Y_e the electron fraction, i.e., the number of electrons per number of nucleons. For the electrically neutral matter that is discussed in this work, Y_e is related to the analogously defined proton and neutron fractions by $Y_e = Y_p = 1 - Y_n$.

For a detailed derivation starting with the Feynman diagram, see ref. [47]. In this derivation, it also becomes apparent that the positive sign applies to neutrinos, while the negative sign applies to antineutrinos.

With $G_F = 1.166\,38 \cdot 10^{-5} \text{ GeV}^{-2}$, $m_N = 0.939\,565 \text{ GeV}$ [7], and expressed in SI units, this yields

$$A_{\text{CC}} = \pm 15.1337 \cdot 10^{-5} \text{ eV}^2 Y_e \frac{\rho}{\text{kg/dm}^3} \frac{E_\nu}{\text{GeV}}. \quad (1.18)$$

The density ρ of the Earth can be calculated well by using seismologic measurements in combination with basic astronomical and geodetic constraints. The most well-established density model for the Earth is the so-called Preliminary Earth Reference Model (PREM),

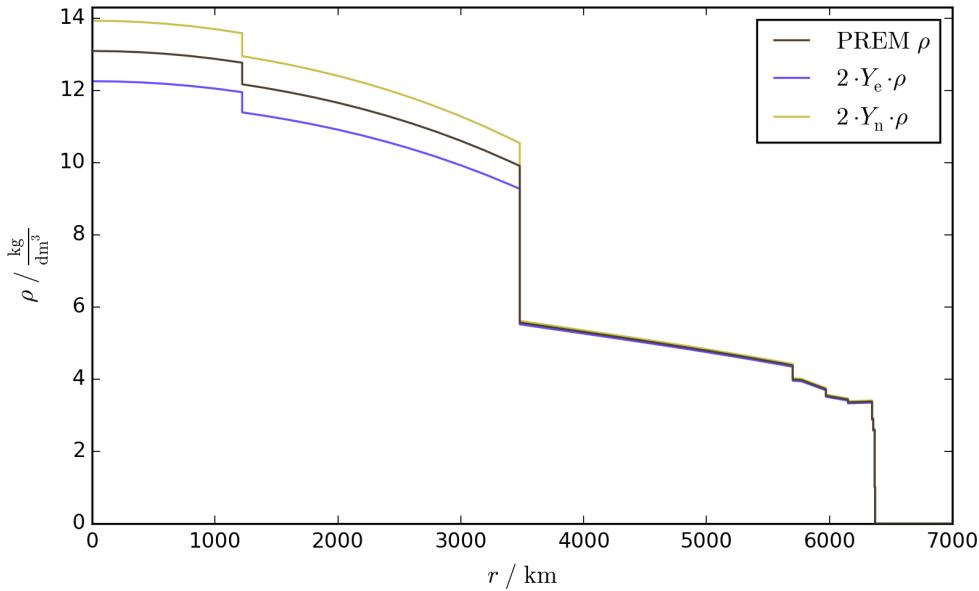


Figure 1.17: Earth's density as predicted by the Preliminary Reference Earth Model (PREM). The other two lines demonstrate the effects of the electron and neutron fractions $Y_e = 1 - Y_n \neq 0.5$, which can be regarded as corrections to the matter density as their deviations from 0.5 are often neglected in other works.

which has been published in 1981 [48]. Among other data, it provides a one-dimensional parameterization for the density of the Earth in form of a non-continuous spline of order three, as shown in figure 1.17.

The electron fraction has to be estimated separately for the two main components of the Earth: the mantle and the core. The crust is relatively thin (35 km on average) and composed of similarly light elements as the mantle, mainly oxygen, silicon, aluminium, and magnesium. While the mantle contains much more magnesium than the crust, the only abundant heavy element for both is iron at about 6%, so the crust can be modeled as part of the mantle. Using the elemental abundancies from [49] and the list of isotope abundancies from [50], one can compute $Y_e^{\text{mantle}} = 0.4957$. The outer core that according to PREM begins at a radius of 3630 km is composed mostly of iron with small amounts of nickel, oxygen, and sulfur, leading to $Y_e^{\text{outer core}} = 0.4680$. Finally, the inner core, starting at 1221.5 km, is composed almost exclusively of iron and nickel, with $Y_e^{\text{inner core}} = 0.4646$. Note that it is difficult to find error estimates for these numbers, but rough comparisons to other models show that the numbers seem to be agreed upon widely on the one-percent level in the fields of geophysics. Figure 1.17 shows how large the impact of Y_e or Y_n is if they are regarded as corrections to the matter density, with the reasoning that their deviations from 0.5 are often neglected in the context of neutrino oscillations.

Because of the wide ranges of atmospheric neutrino energies and density values prevalent on Earth ($1 \lesssim \frac{\rho}{\text{kg/dm}^3} \lesssim 13$), the induced matter potential A_{CC} can assume a relatively

wide range of numerical values and can therefore significantly influence the oscillation behavior as shown in the next section.

1.3.3 Matter Effects

With the induced matter potential A_{CC} that has been introduced in section 1.3.2, the Schrödinger equation (1.14) does not have a constant Hamiltonian anymore. Instead, $H_0 + A$ depends on the local matter density and therefore on the location of the neutrino. This was first predicted and analyzed independently by Wolfenstein and by Mikheyev and Smirnov in the 1970s and 80s, and subsequently named MSW effect.⁷

For illustrative purposes, the effect of the matter potential can best be demonstrated in a two-flavor approximation with electron and muon neutrinos:

$$\begin{aligned}
 \frac{d}{dx} |\nu_\alpha\rangle &= \frac{-i}{2E_\nu} (H_0 + A) |\nu_\alpha\rangle = \frac{-i}{2E_\nu} (U \text{diag}(0, \Delta m_{21}^2) U^\dagger + A) |\nu_\alpha\rangle \\
 &= \frac{-i}{2E_\nu} \left(\begin{pmatrix} \cos(\theta_{12}) & \sin(\theta_{12}) \\ -\sin(\theta_{12}) & \cos(\theta_{12}) \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & \Delta m_{21}^2 \end{pmatrix} \begin{pmatrix} \cos(\theta_{12}) & -\sin(\theta_{12}) \\ \sin(\theta_{12}) & \cos(\theta_{12}) \end{pmatrix} \right. \\
 &\quad \left. + \begin{pmatrix} \frac{1}{2}A_{\text{CC}} & 0 \\ 0 & 0 \end{pmatrix} \right) |\nu_\alpha\rangle \\
 &= \frac{-i}{2E_\nu} \begin{pmatrix} \sin^2(\theta_{12})\Delta m_{21}^2 + \frac{1}{2}A_{\text{CC}} & \sin(\theta_{12})\cos(\theta_{12})\Delta m_{21}^2 \\ \sin(\theta_{12})\cos(\theta_{12})\Delta m_{21}^2 & \cos^2(\theta_{12})\Delta m_{21}^2 \end{pmatrix} |\nu_\alpha\rangle \\
 &= \frac{-i}{4E_\nu} \begin{pmatrix} A_{\text{CC}} - \cos(2\theta_{12})\Delta m_{21}^2 & \sin(2\theta_{12})\Delta m_{21}^2 \\ \sin(2\theta_{12})\Delta m_{21}^2 & \cos(2\theta_{12})\Delta m_{21}^2 \end{pmatrix} |\nu_\alpha\rangle, \tag{1.19}
 \end{aligned}$$

again dropping terms proportional to the identity matrix in the last step. This Schrödinger equation is not fundamentally different from the vacuum formulation (which is in fact just a special case with $A_{\text{CC}} = 0$), so the neutrino propagation can again be described as oscillation, but the oscillation amplitude is different from the vacuum solution. As explained in section 1.3.1, in vacuum, the amplitude of the transformation between the flavors only depends on the mixing angles; in that case, a pure electron state would always remain in electron flavor by at least

$$\min(P_{\nu_e \rightarrow \nu_e}^{\text{vac}}) = \frac{(\cos(2\theta_{12})\Delta m_{21}^2)^2}{(\cos(2\theta_{12})\Delta m_{21}^2)^2 + (\sin(2\theta_{12})\Delta m_{21}^2)^2} = \cos^2(2\theta_{12}).$$

When matter is present, i.e., when $A_{\text{CC}} \neq 0$, the quotient can not be simplified in the same manner:

$$\min(P_{\nu_e \rightarrow \nu_e}) = \frac{(A_{\text{CC}} - \cos(2\theta_{12})\Delta m_{21}^2)^2}{(A_{\text{CC}} - \cos(2\theta_{12})\Delta m_{21}^2)^2 + (\sin(2\theta_{12})\Delta m_{21}^2)^2} =: \cos^2(2\theta_{12}^{\text{mat}}) \tag{1.20}$$

⁷Sometimes, this name is reserved for matter effects in very slowly changing electron densities such as inside the sun, called the adiabatic regime, but this distinction is not made in this work.

The amplitude now not only depends on the mixing angle and the induced mass potential, but also on the squared-mass difference Δm_{21}^2 . Most importantly, if $A_{CC} - \cos(2\theta_{12})\Delta m_{21}^2 = 0$, the mixing between the two flavors is maximal, even if θ_{12} is very small. Because of the wide numerical range of the induced matter potential as mentioned in section 1.3.2, this condition can always be met inside Earth and is called MSW resonance. Note that because of the sign of A_{CC} switching for antineutrinos, the resonance condition between electron and muon flavors is only met in the neutrino channel, but the antineutrino channel still features weaker, non-resonant modifications due to the presence of matter. Also note that the resonance condition alone does not guarantee a strong oscillation effect; for that, the phase has to be correct as well.

For phenomenological discussions, it is convenient to define effective mixing angles such as θ_{12}^{mat} in (1.20). With these angles and effective squared-mass differences defined like

$$\Delta m_{21}^{2, \text{mat}} := \sqrt{(A_{CC} - \cos(2\theta_{12})\Delta m_{21}^2)^2 + (\sin(2\theta_{12})\Delta m_{21}^2)^2}, \quad (1.21)$$

the solutions to the Schrödinger equation visually have the same shape as the vacuum solutions [51]. This makes the equations easier to read and often easier to understand intuitively, but of course the solutions themselves do not change and retain their dependencies on the density. Numerically, these quantities are not of much use because they are expensive to compute and are usually defined using pairwise two-flavor approximations. For this reason, they will not be used further in this work. They are useful however to demonstrate that not only the oscillation amplitude is affected by matter effects, but also the oscillation frequency, which can therefore differ between neutrinos and antineutrinos.

One such resonance can be seen in figure 1.18, where it manifests as a strong disturbance in the neutrino channel just below 10 GeV compared to the regular pattern of figure 1.14. To show the matter effects more clearly, figure 1.19 features the appearance probabilities for electron neutrinos with and without matter effects side by side. As the figures illustrate realistic three-flavor probabilities, the appearance probability for ν_e at the resonance does not reach 100% because the resonance conditions for $\nu_e - \nu_\mu$ and $\nu_e - \nu_\tau$ oscillations differ by a factor of $\cos(2\theta_{13})\Delta m_{31}^2 / \cos(2\theta_{12})\Delta m_{21}^2 = 82.4$ in energy. The distortion just below 10 GeV marks the $\nu_e - \nu_\tau$ resonance, while the deviation from the $\nu_e - \nu_\mu$ vacuum pattern below 1 GeV is caused by the $\nu_e - \nu_\mu$ resonance.

Figures 1.18 and 1.19 also show a faint line at a zenith angle of $\cos(\theta) = 0.84$; this is the transition between the Earth mantle and outer core, with a significant jump in electron density. The transition between outer core and inner core lies at $\cos(\theta) = 0.98$, but it is much less significant.

Another phenomenon that is also present in these figures is the so-called parametric enhancement. The term was coined mainly by Akhmedov and describes a resonance condition different from the MSW resonance: Parametric enhancement occurs when changes

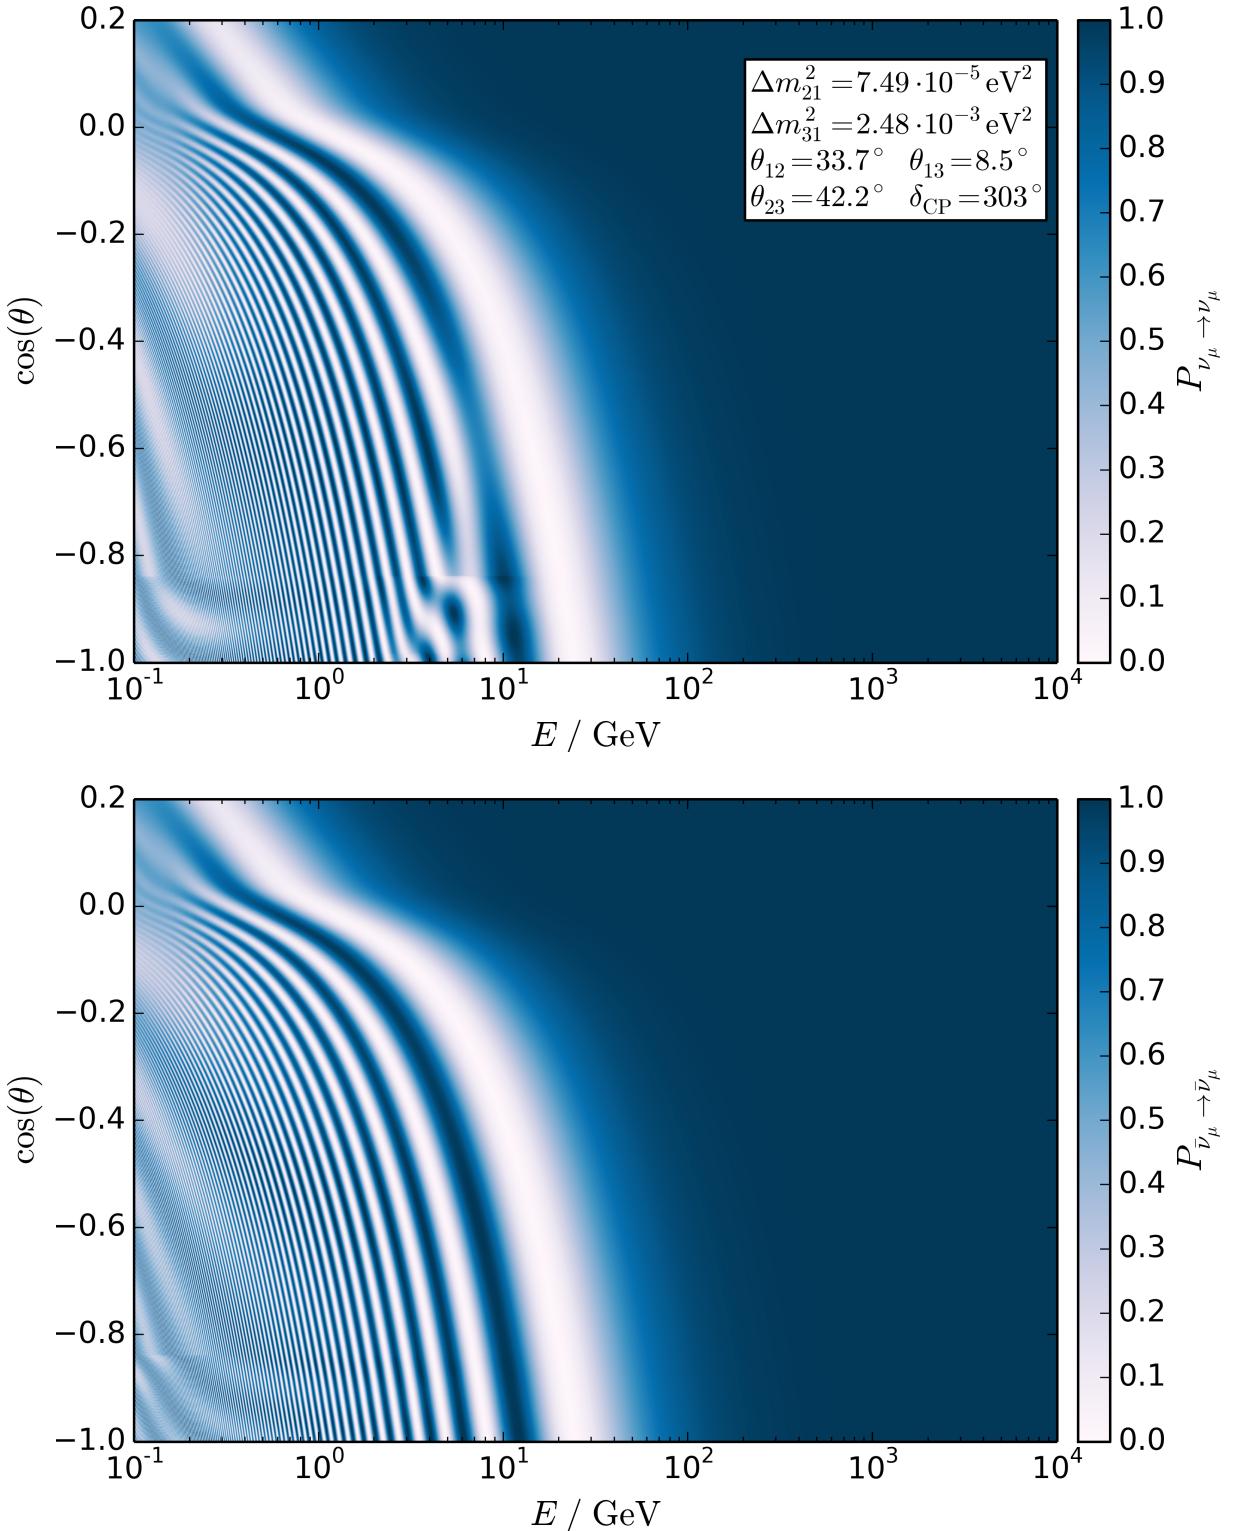


Figure 1.18: Survival probability of atmospheric muon neutrinos (**top**) and muon antineutrinos (**bottom**) as a function of neutrino energy and zenith angle, taking into account matter effects. The oscillation parameters are the same as in figure 1.14.

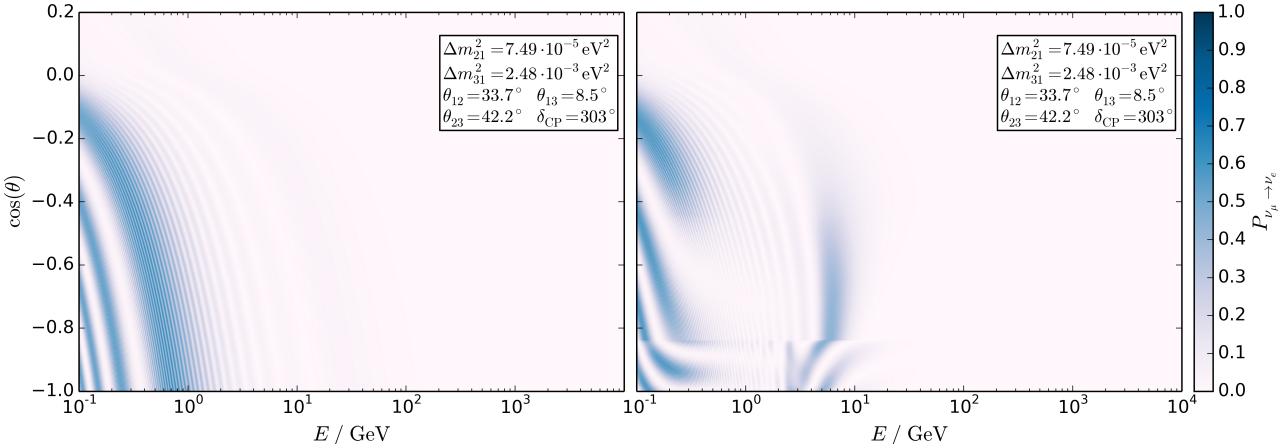


Figure 1.19: Appearance probability of atmospheric electron neutrinos as a function of neutrino energy and zenith angle, without (**left**) and with (**right**) matter effects. This figure is the electron flavor counterpart of figure 1.14 and figure 1.18.

in the density along the baseline are synchronized with the oscillation frequency, such that the phases acquired in regions of different density are odd-integer multiples of π [52]. If the density changes are periodic, phase differences add up constructively over many periods and can lead to strong deviations from the vacuum mixing angles, even if they are small and not strongly modified by the MSW resonance. For atmospheric neutrinos, this condition is fulfilled weakly because traversing the Earth's mantle, core, and again mantle equals 1.5 periods through a periodic profile.

Luckily, assessing the importance of different matter effects is not necessary for this work. Separating matter effects into different categories is only necessary when using approximations to either simplify calculations or discussions of neutrino oscillations in matter, as it is often done in theoretical papers. If one computes oscillations probabilities by numerically solving the Schrödinger equation (1.14) with the induced matter potential A_{CC} and without additional approximations, the different kinds of matter effects can not be distinguished or turned on or off independently.

Since solving the Schrödinger equation is the method of choice for this work as mentioned above and described in more detail in appendix C, the detailed analysis of the impact and phenomenology of parametric enhancements for atmospheric neutrinos will be left to papers such as [53].

1.4 Additional Sterile Flavors

As mentioned during the short introduction of this chapter, there are currently three known flavors of neutrinos: electron, muon, and tau. Neutrinos of all these flavors have

been experimentally detected, and models using these three flavors have successfully described oscillation phenomena in many detectors. Various measurements indicate that there are no additional types of neutrinos with similar properties to those known three flavors. The probably most famous one is the measurement of the decay width of the Z^0 boson resonance in e^+e^- annihilations at the LEP particle collider in CERN [54]. The hadronic cross-section in this process depends strongly on the number of particles into which Z^0 boson can decay. Even though neutrinos are usually not detectable in accelerator experiments but are accounted for as “invisible energy”, a precise measurement of the cross-section in comparison with theoretical predictions allowed to constrain the number of conventional light neutrino flavors to 2.9840 ± 0.0082 [54] as can be seen in figure 1.20. In this context, light means that the neutrinos have to have at most half of the mass of the Z^0 boson, otherwise they can not be produced in this process and do not influence the cross-section. Conventional means that they have to partake in weak interactions for the same reason.

The first condition is easily met by the known three neutrino flavors. Direct measurements of tritium decay product momentum give an upper mass limit of $m_{\bar{\nu}_e} < 2.05$ eV for electron antineutrinos [55][56], and oscillation measurements show that the three known mass eigenstates have to be very close together. The mass up to which the LEP limit holds is $\frac{1}{2}m_Z = (45.5938 \pm 0.0011)$ GeV [7], which is higher by ten orders of magnitude. Therefore, a conventional fourth flavor of neutrinos would have to be much more massive than the known three.

The second condition is the one that is dropped for the hypothetical *sterile* neutrinos that are the subject of this work. To this day, the only known production and detection mechanisms for neutrinos rely on the weak interaction. For this reason it could be possible that there are additional flavors⁸ of neutrinos that so far evaded detection as they do not partake in weak interactions at all. They would still interact gravitationally, but this would be hardly noticeable on non-cosmological scales, where the existence of additional sterile neutrino flavors is not ruled out compellingly (see section 1.5).

Because of the unique sensitivity of the IceCube Neutrino Observatory for direct measurements of light sterile neutrinos which will be explained later, this work focuses on those and keeps condition one intact.

Note that there are various mechanisms by which light sterile neutrinos can come into being. Many of these mechanisms are based on the seesaw mechanism which is often hypothesized to be the reason for the relatively small but non-zero (for at least two mass eigenstates) neutrino masses of the known neutrinos. In the most basic type 1 seesaw mechanism, the mass of the proposed right-handed and therefore sterile neutrino is very high (usually between 10^5 GeV and the GUT scale around 10^{16} GeV), but this model can be modified in various ways that can explain the existence of light sterile neutrinos with masses in the eV or keV range.

⁸*Flavor* could be considered to be a misnomer as it usually indicates a certain charge in relation to the weak interaction which is absent for sterile neutrinos, but it will still be used here to reduce verbosity.

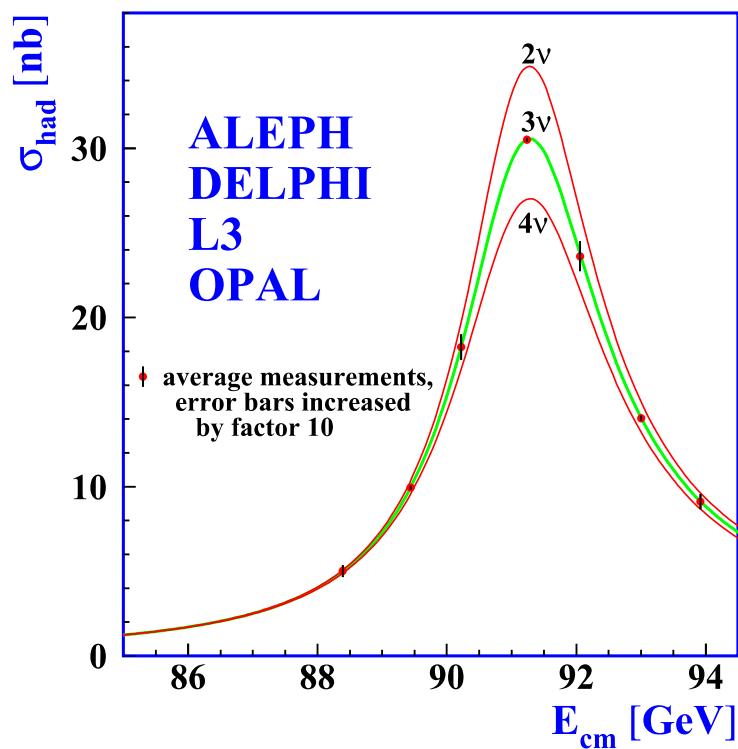


Figure 1.20: Hadronic cross-section at the Z^0 boson resonance in e^+e^- annihilations at the LEP particle collider. Experimental measurements are compared to theoretical predictions for 2, 3, or 4 light non-sterile neutrinos [54].

Because of this and additional, different approaches, the theoretical background of sterile neutrinos is vast and can not be examined in the scope of this work. A good and comprehensive overview is given in the light-sterile-neutrino whitepaper [57].

1.4.1 Vacuum Oscillation

The introduction of additional sterile flavors is trivial in vacuum oscillation framework. (1.9) and (1.12) are still the valid wave equation and plane-wave solution, respectively. However, the dimension of the system increases to accommodate the additional particles. Scenarios with n additional sterile flavors are usually named $3 + n$, but sometimes the scheme is modified to indicate the mass hierarchy of the new flavors, such as $1 + 3$ for a sterile flavor lighter than the conventional neutrinos, or $1 + 3 + 1$, etc. This work focuses on the $3 + 1$ scenario as explained in section 4.2.

When adding additional flavors, the PMNS matrix U gets extended to $(3+n) \times (3+n)$, and the Hamiltonian gains n additional squared-mass differences. As mentioned before, a PMNS matrix of rank \tilde{n} has $\tilde{n}(\tilde{n}-1)/2$ mixing angles and the same number of CP-violating phases, of which $(\tilde{n}-1)(\tilde{n}-2)/2$ are of Dirac type. Because of this, U gains $\frac{1}{2}n^2 + \frac{5}{2}n$ mixing angles, n Majorana CP phases and $\frac{1}{2}n^2 + \frac{3}{2}n$ Dirac-type CP phases. Putting in numbers, in the $3 + 1$ scenario these are three mixing angles θ_{14} , θ_{24} , and θ_{34} , the squared-mass difference Δm_{41}^2 , the Majorana phase α_4 , and two Dirac-type phases, e.g., δ_{24} and δ_{34} . In the $3 + 2$ scenario, there are yet four more mixing angles θ_{i5} , Δm_{51}^2 , α_5 , and three more Dirac-type phases δ_{i5} .

The selection of where to put the additional Dirac-type phases is somewhat arbitrary, but it is convenient to chose a parameterization that simplifies to the usual three-flavor convention (1.4) when the additional mixing angles vanish (and analog for $3 + 2 \rightarrow 3 + 1$), so the new phases are placed into the additional matrix elements. Still, CP-violating phases are set to zero for this work as detailed in appendix A.

Again, the ordering of the rotation matrices that parameterize U is important as those matrices do not commute in general. Choosing a different order can strongly change the effect of the individual mixing angles and phases and thereby lead to misunderstandings. The parameterization used in this work is

$$U = R_{34}(\theta_{34}, \delta_{34})R_{24}(\theta_{24}, \delta_{24})R_{14}(\theta_{14})R_{23}(\theta_{23})R_{13}(\theta_{13}, \delta_{13})R_{12}(\theta_{12})J,$$

which is identical to the lexicographically ordered

$$U = R_{34}(\theta_{34}, \delta_{34})R_{24}(\theta_{24}, \delta_{24})R_{23}(\theta_{23})R_{14}(\theta_{14})R_{13}(\theta_{13}, \delta_{13})R_{12}(\theta_{12})J$$

because R_{14} and R_{23} commute. Effectively, the PMNS matrix used later on simplifies to

$$U = R_{24}(\theta_{24})R_{23}(\theta_{23})R_{13}(\theta_{13})R_{12}(\theta_{12})$$

as the other parameters have been dropped for this work as explained in chapters 4 and 5.

For $3+2$ models, the lexicographical ordering is a similarly good and common choice. All these parameterizations simplify to (1.4) when the sterile mixing angles vanish.

As the solution to the higher-dimensional Schrödinger equation in vacuum is still the plane-wave equation (1.12), the expected signature of a $3+1$ scenario is still a superposition of harmonic oscillations as in the three-flavor case. However, the vacuum mixing with the newly introduced flavors has to be small because it is constrained by the 3×3 PMNS matrix unitarity limits: As a mixing matrix, U has to be unitary by construction, because otherwise the sum of the probabilities does not add up to one. However, the known 3×3 matrix might be a submatrix of a unitary $(3+n) \times (3+n)$ U , and as such does not have to be unitary by itself [58]. Experimental verification of the 3×3 unitary cannot constrain the dimension of U , but it can impose strong limits on the amount of mixing between the three known flavors and potential additional ones. According to global fits [58],

$$|UU^\dagger| = \begin{pmatrix} 0.951 \dots 1.000 & < 0.017 & < 0.074 \\ < 0.017 & 0.962 \dots 1.000 & < 0.064 \\ < 0.074 & < 0.064 & 0.762 \dots 1.000 \end{pmatrix} \quad (1.22)$$

at 90% confidence level⁹ (C.L.), where U is the conventional 3×3 PMNS matrix. The small deviation of the product $|UU^\dagger|$ from identity does not leave much room for mixing of the known conventional flavors with additional ones.

Note that more limiting constraints such as the following can often be found, but those rely on additional assumptions such as the Minimal Unitarity Violation scheme (MUV), which means that they assume that all additional flavors are much heavier than the electroweak energy scale ($\gg 246$ GeV) [59]:

$$|UU^\dagger| \stackrel{\text{MUV}}{=} \begin{pmatrix} 0.9979 \dots 0.9998 & < 10^{-5} & < 0.0021 \\ < 10^{-5} & 0.9996 \dots 1.0000 & < 0.0008 \\ < 0.0021 & < 0.0008 & 0.9947 \dots 1.0000 \end{pmatrix}$$

This assumption is invalid for this analysis, so only the limits given in (1.22) apply.

Figure 1.21 displays the vacuum solution for the muon neutrino survival probability for a sterile-neutrino benchmark model allowed by (1.22) (for $\theta_{14} = \theta_{34} = 0$, $\theta_{24} < 14.9^\circ$ at 90% C.L.). It shows that the deviations from the no-sterile case (figure 1.14) are relatively small. The next section will demonstrate that matter effects strongly apply to sterile neutrinos and that they can increase oscillation and thereby detection probabilities significantly.

⁹Based on figures 2 and 3 of [58], verified by private conversation.

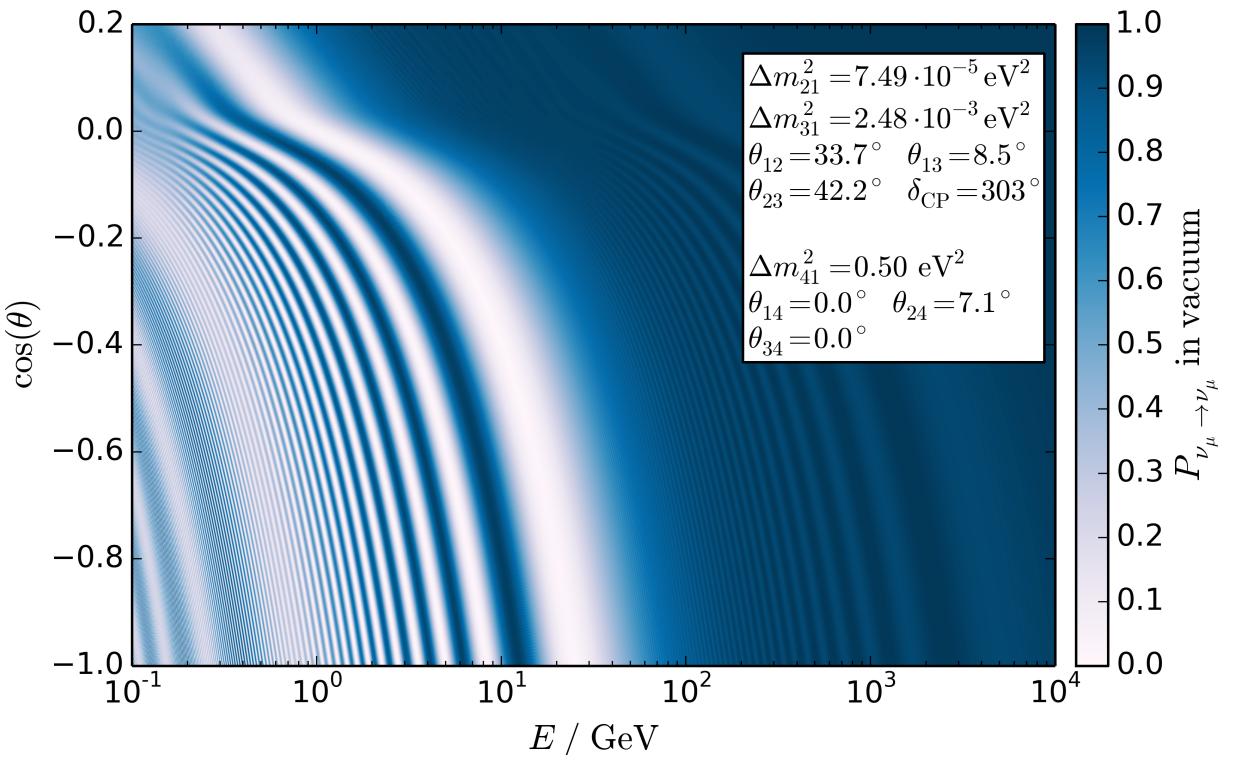


Figure 1.21: Survival probability of atmospheric muon neutrinos with an additional flavor of sterile neutrinos as a benchmark, without matter effects. The oscillation parameters are displayed in the plot.

1.4.2 Matter Effects

The reason for the existence of matter effects for sterile neutrinos is an induced matter potential as introduced in section 1.3.2. The main difference is that not only charged-current but also neutral-current interactions contribute to the potential. The conventional flavors can all interact by the neutral-current exchange of Z^0 bosons with nucleons and electrons as illustrated in figure 1.15, while sterile neutrinos cannot. Because of this, the induced matter potential A has to be modified to

$$\begin{aligned} A &= \text{diag}(A_{\text{CC}} + A_{\text{NC}}, A_{\text{NC}}, A_{\text{NC}}, 0) \\ &\hat{=} \text{diag}(A_{\text{CC}}, 0, 0, -A_{\text{NC}}), \end{aligned} \quad (1.23)$$

yet again using that terms proportional to identity do not matter for oscillations. For additional sterile flavors, the last element in the diagonal simply has to be repeated.

A_{CC} remains unmodified and relies on the electron density of the matter. For A_{NC} , nucleons have to be taken into account. However, the Earth is made up of electrically neutral matter and both the charge and the third component of the weak isospin of electrons and protons are opposite. Because of this, the potentials induced by protons and electrons are identical except for their signs and cancel out each other, and only the neutron density in the traversed matter needs to be taken into account [60]:

$$A_{\text{NC}} = \mp \sqrt{2} G_F \frac{\rho}{m_N} Y_n E_\nu = \mp 7.5669 \cdot 10^{-5} \text{ eV}^2 (1 - Y_e) \frac{\rho}{\text{kg/dm}^3} \frac{E_\nu}{\text{GeV}}, \quad (1.24)$$

analog to (1.16), with $Y_n = 1 - Y_e$. Again, the upper sign applies to neutrinos, the lower sign to antineutrinos. The remaining derivation is analog to section 1.3.3, where the non-constant A causes both the amplitude and the frequency of the oscillation to change along the neutrino trajectory. Because of the inverted sign of A_{NC} in relation to A_{CC} , the muon-sterile resonance occurs in the antineutrino channel if $m_4 > m_2$.

In the sterile case, the importance of the MSW resonance is very high since the vacuum angles are small. If the additional mass eigenvalue m_4 is large enough such that $\Delta m_{42}^2 \gg \Delta m_{32}^2$, the muon-sterile resonance defined by $A_{\text{NC}} - \cos(2\theta_{24})\Delta m_{42}^2 = 0$ will lie in a neutrino energy region where muon neutrinos are not affected by other oscillations, so the resonance can lead to full muon neutrino disappearance as illustrated in figure 1.22. The plot shows a distinct oscillation minimum around 2 TeV that differentiates this signature clearly from figure 1.18. The analysis presented here will exploit this resonance behavior to achieve high sensitivity in a large parameter region. With the theoretical background explored, the specific effects of the oscillation parameters will be analyzed in more detail in chapter 4.

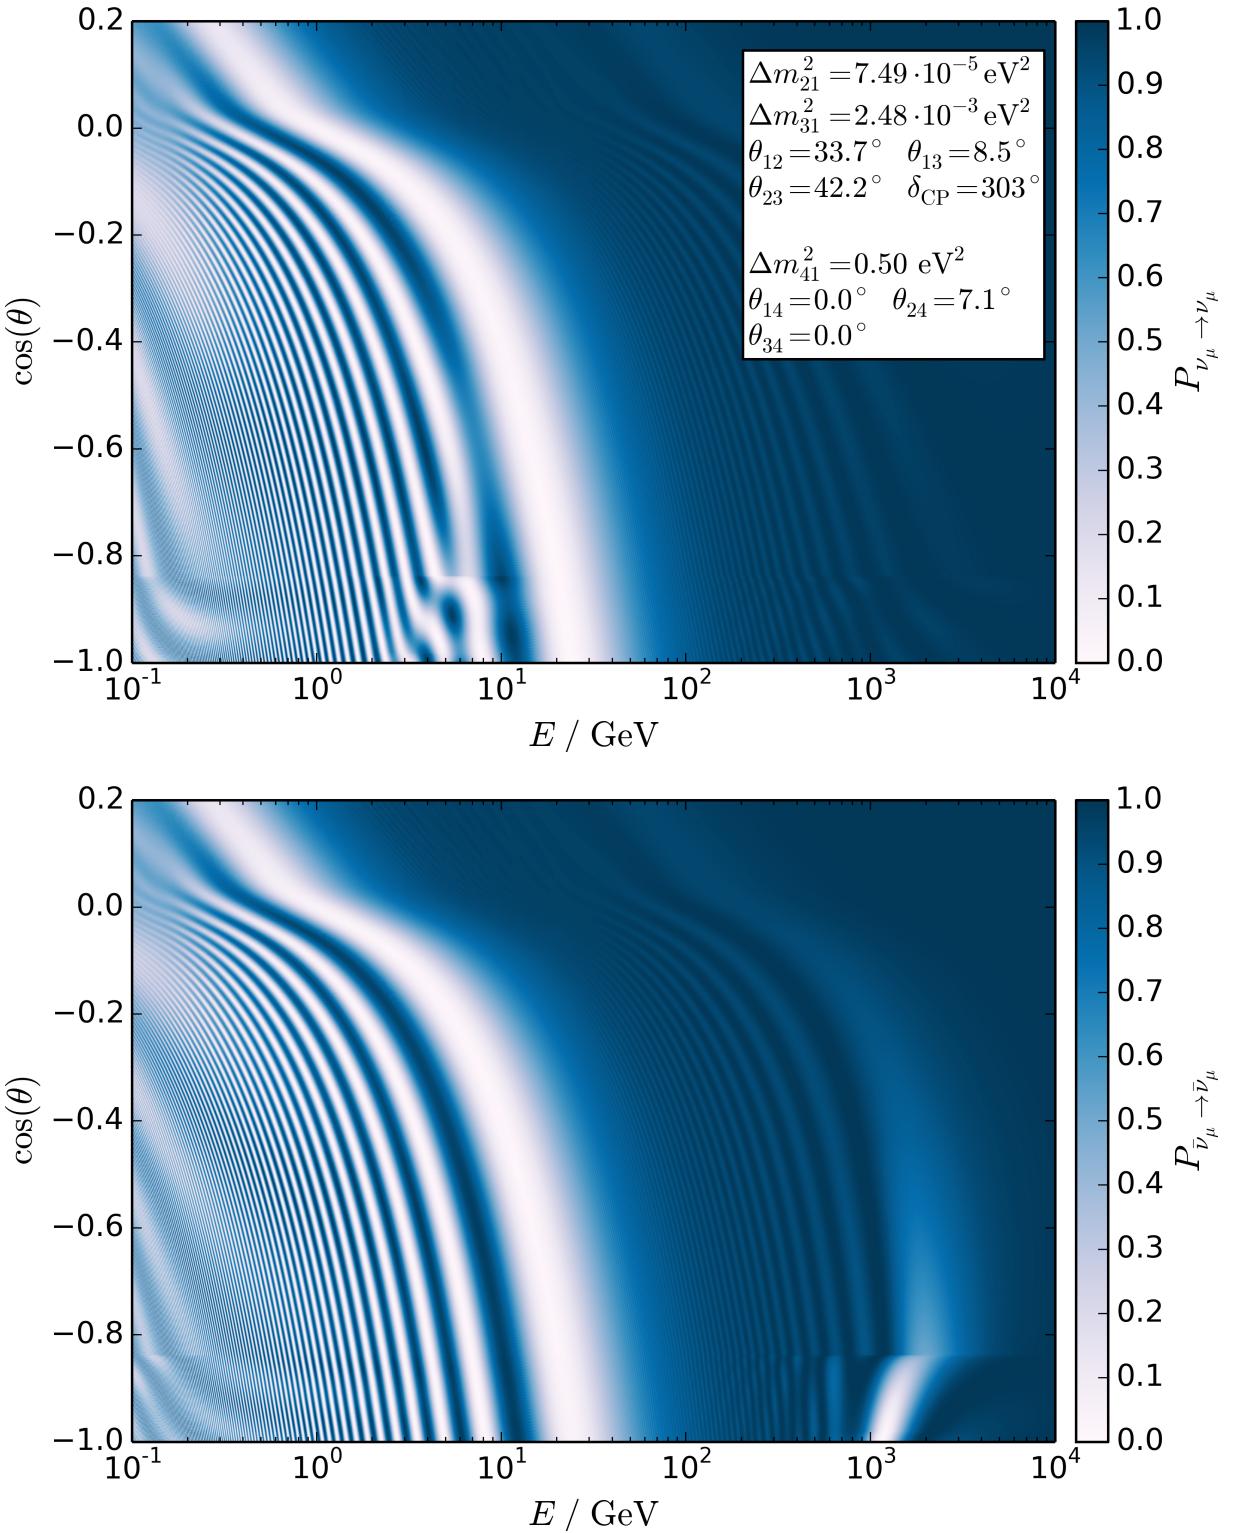


Figure 1.22: Survival probability of atmospheric muon neutrinos (**top**) and muon antineutrinos (**bottom**) with an additional flavor of sterile neutrinos as a benchmark, taking into account matter effects. The oscillation parameters are identical to figure 1.21.

1.5 Current State of Research

The state of research regarding sterile neutrinos is complicated because possible evidence for sterile neutrinos is often contradictory, and established limits can often be circumvented by additional theoretical assumptions. This section will give a small overview, but providing a comprehensive picture is out of scope for this work.

Experimental data hinting at the existence of sterile neutrinos originally comes from LSND and MiniBooNE [61][62] ($\bar{\nu}_e$ and $(\bar{\nu})_e$ appearance, respectively), Gallex and SAGE [63] (ν_e disappearance), and reactor neutrino experiments [64] ($\bar{\nu}_e$ disappearance). Individually, the findings could be explained by the inclusion of one or more light sterile neutrinos, but they are at best marginally compatible with each other as shown by global fits [65][66].

More recent results reinforce the anomaly that is seen by MiniBooNE [68]: New ν_e appearance data has a 4.7σ preference for a two-flavor oscillation model involving a sterile neutrino compared to a no-oscillation model, and adding $\bar{\nu}_e$ appearance data raises this to 4.9σ . The paper claims compatibility to the old LSND results, and states a combined significance of 6.0σ , with a best fit at $\sin^2(2\theta_{e\mu}) := 4|U_{e4}||U_{\mu 4}| = 4\sin^2(\theta_{14})\cos^2(\theta_{14})\sin^2(\theta_{24}) = 0.96$ and $\Delta m_{42}^2 = 0.041\text{ eV}^2$.

The reactor neutrino anomaly was claimed to be mostly explainable by neglected fuel evolution and wrongly predicted yields for inverse beta decay of ^{235}U [69], but other publications object that this is only true for the Daya Bay experiment and that the sterile-neutrino solution is still favored globally, especially including very-short-baseline data from the new NEOS and DANSS experiments [70]. Also, it was pointed out that the remaining anomaly in Daya Bay is still compatible with the existence of a sterile neutrino [71]. Other explanations such as non-linear energy scale miscalibrations were similarly proposed and rejected [72].

Cosmological limits, for example on the number (N_{eff}) or on the total mass ($\sum m_\nu$) of neutrinos, can be circumvented by various means, among which are dark-matter particles that interact more strongly with electrons or photons than with neutrinos (for example through an electric or magnetic dipole moment) [73], or so called “secret interactions” that allow sterile neutrinos to decay quickly [74].

In the following years the sterile parameter space will be probed by many additional experiments, among which are independent reproductions of the existing experiments (e.g., JSNS2 for LSND [75]), very-short-baseline reactor experiments such as NEOS, DANSS, NEUTRINO-4, STEREO, SoLid, and PROSPECT [65], and non-oscillation experiments such as KATRIN [76]. However, until the ostensible contradictions have been solved, an approach that is as model-independent as possible seems preferable and thus is pursued in this work.

Chapter 2

The IceCube Neutrino Observatory

The IceCube Neutrino Observatory is a 1 km^3 particle detector located at the Amundsen-Scott South Pole Station at the Geographic South Pole [77]. Like its predecessor, the much smaller Antarctic Muon And Neutrino Detector Array (AMANDA) [78], IceCube uses the clear and deep Antarctic ice as detector medium to look for Čerenkov radiation of charged particles caused by neutrino interactions inside or close to the detector.

The first section of this chapter explains the detection principle of Čerenkov telescopes in general, the next sections outline the hard- and software of IceCube in particular, and the last section briefly describes how neutrino and muon data is simulated for comparison with measured experimental data.

2.1 Detection Principle

Similar to water Čerenkov detectors such as the Japanese KamiokaNDE series [79] or the Mediterranean ANTARES telescope [80], ice Čerenkov detectors rely on the detection of Čerenkov radiation of neutrino interaction products in a transparent medium in dark surroundings.

2.1.1 Čerenkov Radiation

Čerenkov radiation emerges when a charged particle travels through a medium faster than the phase speed of light in that medium.

Within the scope of our current knowledge, the speed of light in vacuum is constant and universal and its value is described by well-measured constant $c = 299\,792\,458 \text{ m/s}$; since 1983, it is used to define the length of a meter, so the quoted value is exact by definition

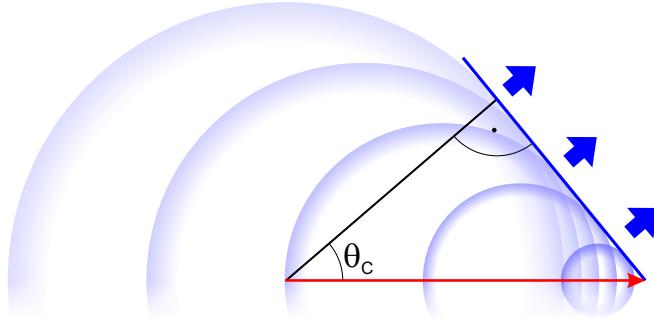


Figure 2.1: Sketch of a charged particle (red arrow) traveling faster than the phase speed of light in the medium. Blue circles indicate electromagnetic fields emitted by excited electrons. The blue line indicates the bow shock that is the Čerenkov radiation, emitted at the Čerenkov angle θ_c [83].

[7]. However, since light is electromagnetic radiation, if it travels through a medium, it exerts small forces on the electrons and nuclei of that matter. The electrons then emit electromagnetic fields at the same frequency, but with shifted phases. These fields get superimposed on the original field and result in an electromagnetic field that is shifted in phase (and differs in amplitude if the phase difference is not 90°) [81]. The phase shift increases continuously and results in a field that effectively propagates slower¹⁰ than the original field.

Macroscopically, this effect defines the refractive index $n_r := \frac{c}{v_p}$, where v_p is the phase speed of light in the medium. The refractive index plays a role in many optical phenomena such as reflection, refraction, or the occurrence of Čerenkov radiation:

If a charged particle travels through a medium, it exerts small forces on the electrons similar to an electromagnetic field as described above. The electrons again emit electromagnetic fields which get superimposed onto each other. When the particle is slow enough, the electromagnetic fields cancel out each other because on average their interference is destructive. If however the particle propagates faster than the speed of light in that medium, i.e., $v \geq \frac{c}{n_r}$, there is a certain angle at which the individual electromagnetic fields are in phase and can only add up, comparable to a bow shock of a ship or supersonic booms (see figure 2.1) [82]. This angle is called the Čerenkov angle θ_c and is given by

$$\cos(\theta_c) = \frac{c}{n_r v}. \quad (2.1)$$

The frequency spectrum of the emitted Čerenkov radiation is continuous and gets

¹⁰Taking into account resonances and dispersion, refractive indices smaller than 1 are possible, but unusual especially in the visible régime and not relevant to this work.

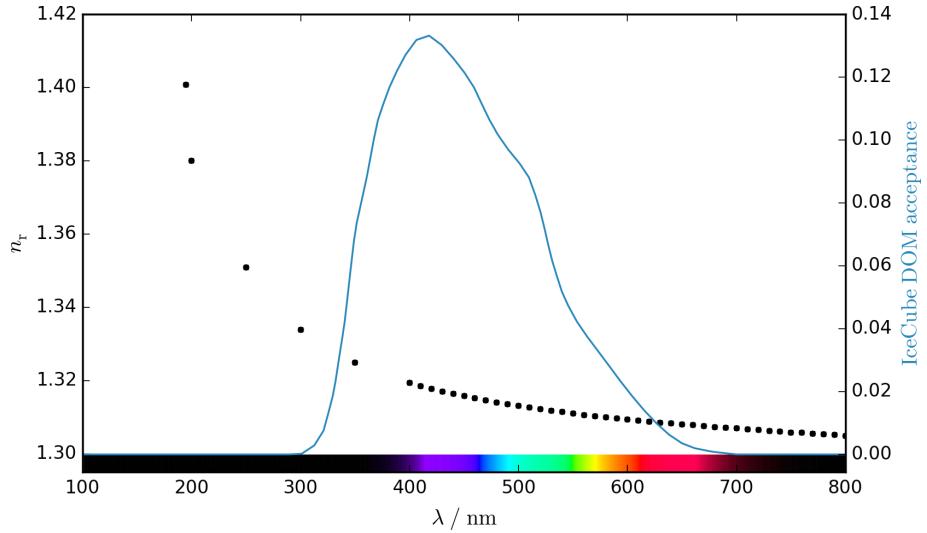


Figure 2.2: The refractive index n_r of water ice at -7°C according to ref. [85]. Uncertainties are smaller than the data points, and usage of the data is recommended between -60°C and 0°C . The blue line shows the relative spectral acceptance of IceCube’s Digital Optical Modules for vertically arriving photons according to ref. [86].

parameterized well by the Frank-Tamm formula [84]:

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi\alpha z^2\mu(\lambda)}{\lambda^2} \sin^2(\theta_c(\lambda)), \quad (2.2)$$

where λ denotes the wavelength, z is the particle’s electric charge, $\alpha \approx \frac{1}{137}$ is the fine-structure constant, and μ the magnetic permeability as a function of the wavelength (which is usually omitted as $|1 - \mu| \leq 10^{-3}$ for common transparent media). Integrating (2.2) over the wavelength shows that approximately $\frac{dN}{dx} \propto \frac{1}{\lambda}$, so the intensity increases linear with the frequency. For this reason, Čerenkov radiation occurs to be blue to the human eye, and increases towards the ultraviolet spectrum until the absorption in the medium becomes too strong.

The refractive index n_r of water ice lies between 1.32 and 1.33 in the blue and near ultraviolet regions (see figure 2.2), so the corresponding Čerenkov angle for ultrarelativistic particles is $\theta_c = 41^\circ$. The refractive index rises strongly in the middle UV range (300 nm to 200 nm), but the plot also shows the spectral acceptance of the optical sensors used by IceCube (including glass and optical gel, see section 2.2.1) to illustrate that this is not relevant to this work. Taking into account the spectral acceptance, (2.2) can be used to compute that an ultrarelativistic particle causes about 2450 theoretically detectable Čerenkov photons per meter without taking into account secondary interactions [86].

The usage of water ice as Čerenkov medium for a large neutrino detector is motivated by its high availability, its cheapness, its lack of bioluminescent lifeforms, and especially for the old Antarctic ice its small abundance of radioactive isotopes, together with relatively

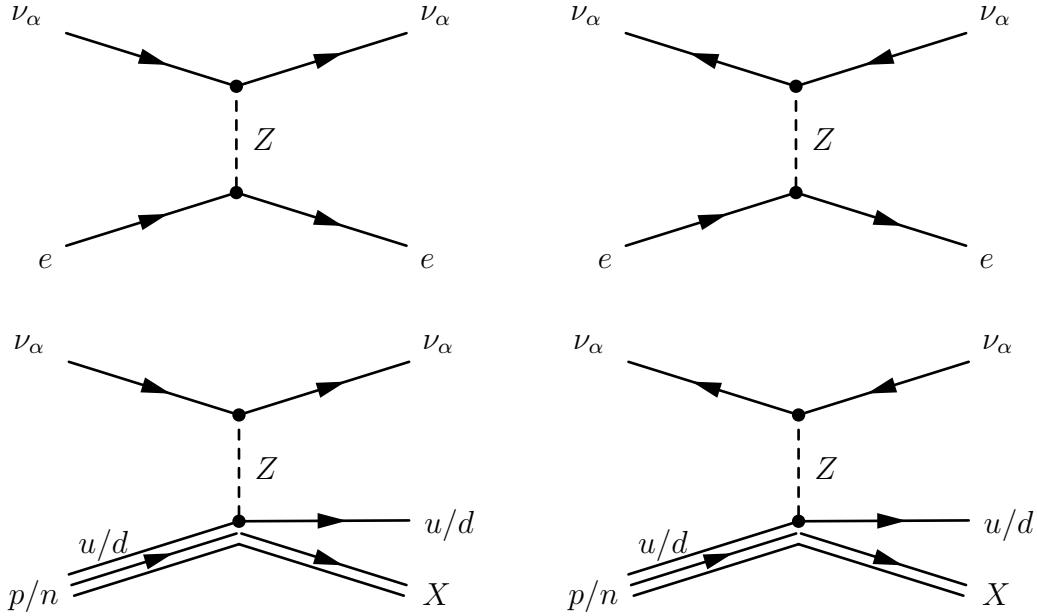


Figure 2.3: Feynman graph of neutral-current non-elastic interactions of conventional neutrinos and antineutrinos on matter. Equivalent interactions with sea quarks are omitted.

good optical quantities (see section 2.2.2). Deployment of photosensors is more difficult than in liquid water, but once frozen in there is little corrosive stress and no issues with leakage. Compared to gasses, the much higher refractive index means that the energy threshold implied by $v \geq \frac{c}{n_r}$ is low enough to usually not be of importance, and the higher density means that neutrinos are more likely to interact inside the Čerenkov medium.

2.1.2 Neutrino Interaction Products

As neutrinos are electrically neutral, they can not be detected directly with Čerenkov detectors. However, they can interact with matter to produce detectable secondary particles. Relevant interactions are shown as Feynman graphs in figure 2.3 (NC) and figure 2.4 (CC); some of these look identical to the Feynman graphs of elastic scattering in section 1.3.2, but were repeated here because the processes can be inelastic at high neutrino energies and high momentum transfers. Figure 2.5 shows that outside of one resonance region around 6.3 PeV, the cross-sections for neutrino-nucleon interactions are much higher than those for neutrino-electron interactions, causing the latter to be negligible. The resonance is caused by the on-shell creation of a W^- boson in the process $\bar{\nu}_e + e^- \rightarrow W^-$, which can then decay into a neutrino-lepton pair (approximately 10.9% probability per flavor) or hadronically (67.4%) [7]; it is named after Sheldon L. Glashow, who predicted it in 1959

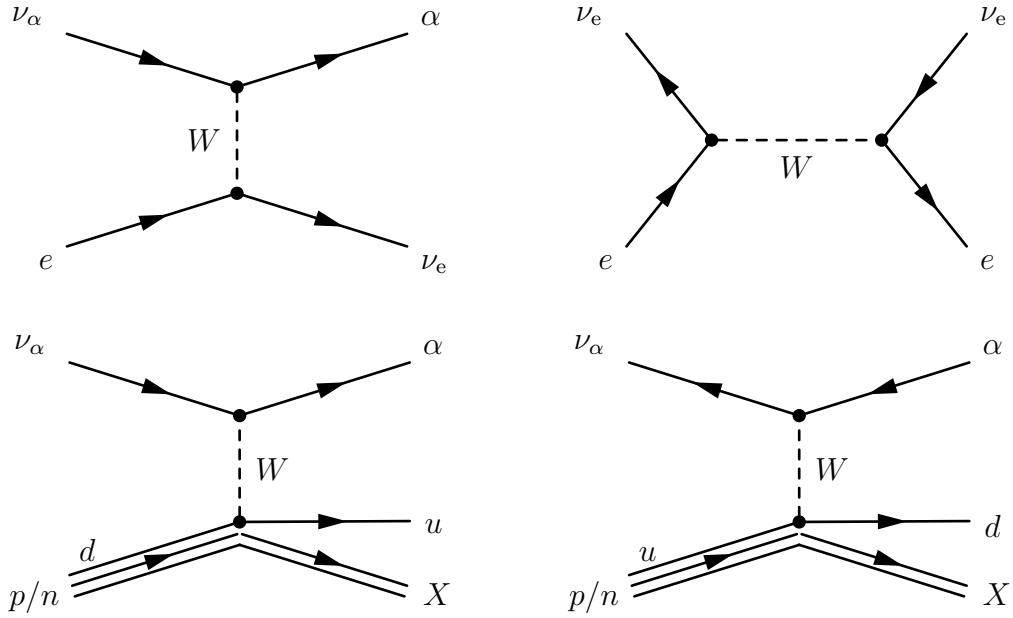


Figure 2.4: Feynman graph of charged-current non-elastic interactions of conventional neutrinos and antineutrinos on matter. Equivalent interactions with sea quarks are omitted.

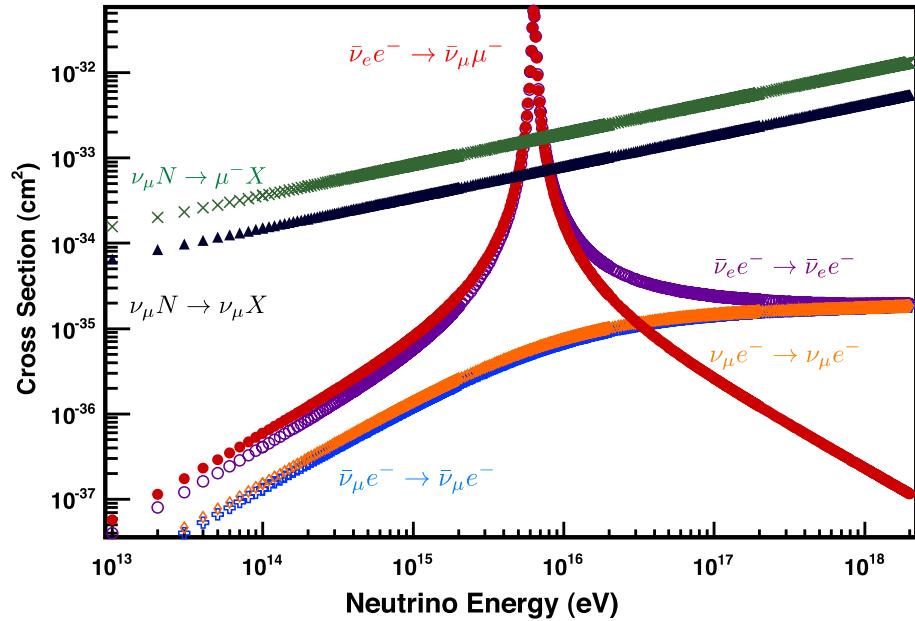


Figure 2.5: Calculated total neutrino and antineutrino cross-sections for various interactions with electrons or nucleons at high energies [87].

[88]. The Glashow resonance lies well beyond the energy range relevant for this analysis, so this work will focus on interactions with nucleons.

Figure 2.6 shows charged-current cross-sections of neutrinos with nucleons at low energies, separated into three components: quasi-elastic interactions (QE), resonance production (RES), and deep inelastic scattering (DIS); see ref. [87] for details:

In quasi-elastic interactions (and elastic interactions, which are their NC counterpart), neutrinos cannot resolve individual quarks and instead transfer momentum to one or multiple nucleons, e.g.:

$$\nu_\alpha + n \xrightarrow{\text{CC}} \alpha^- + p, \quad \bar{\nu}_\alpha + p \xrightarrow{\text{CC}} \alpha^+ + n, \quad (\bar{\nu})_\alpha + h \xrightarrow{\text{NC}} (\bar{\nu})_\alpha + h,$$

where h (hadron) stands for protons or neutrons. Their cross-sections become constant above 10 GeV (even above 2 GeV for most interactions), so they quickly become irrelevant compared to other interactions at high energies.

Resonance production is the name for processes of the types

$$(\bar{\nu})_\alpha + h \xrightarrow{\text{CC}} \alpha^\mp + h^* \quad \text{and} \quad (\bar{\nu})_\alpha + h \xrightarrow{\text{NC}} (\bar{\nu})_\alpha + h^*,$$

where h^* denotes an excited hadronic state that will quickly decay, usually into one or more pions and kaons. Their cross-sections become constant above 10 GeV as well, quickly losing importance.

Deep inelastic scattering is the dominant process at high energies, and the only one represented in the Feynman graphs. It occurs when the neutrino resolves and interacts with an individual quark (q) instead of a whole nucleon, e.g.:

$$\nu_\alpha + d \xrightarrow{\text{CC}} \alpha^- + u, \quad \bar{\nu}_\alpha + u \xrightarrow{\text{CC}} \alpha^+ + d, \quad (\bar{\nu})_\alpha + q \xrightarrow{\text{NC}} (\bar{\nu})_\alpha + q$$

The target quark, which can also be a sea quark or sea antiquark, is expelled from its nucleon, and QCD confinement causes both it and the remainder to produce a hadronic cascade, which is often indicated as $\dots \rightarrow \dots + X$. Below center-of-mass energies corresponding to the W^\pm and Z^0 boson masses, the cross-sections approximately increase linearly with the neutrino energy; above, they get damped to $\sigma_{(\bar{\nu})q} \propto E_\nu^{0.363}$.

In summary, at high energies ($\gtrsim 100$ GeV), almost all neutrino interactions produce hadronic cascades, while charged-current interactions also produce a charged lepton corresponding to the neutrino flavor. The hadronic cascade inherits on average just below 50% of the neutrino energy or about 35% of the antineutrino energy as illustrated in figure 2.7, although these distributions are wide. The angle between the incident neutrino and the outgoing lepton (either a charged one or a neutrino in the neutral-current case) can be parameterized as [90]

$$\psi = 0.7^\circ \left(\frac{\text{TeV}}{E_\nu} \right)^{0.7}. \quad (2.3)$$

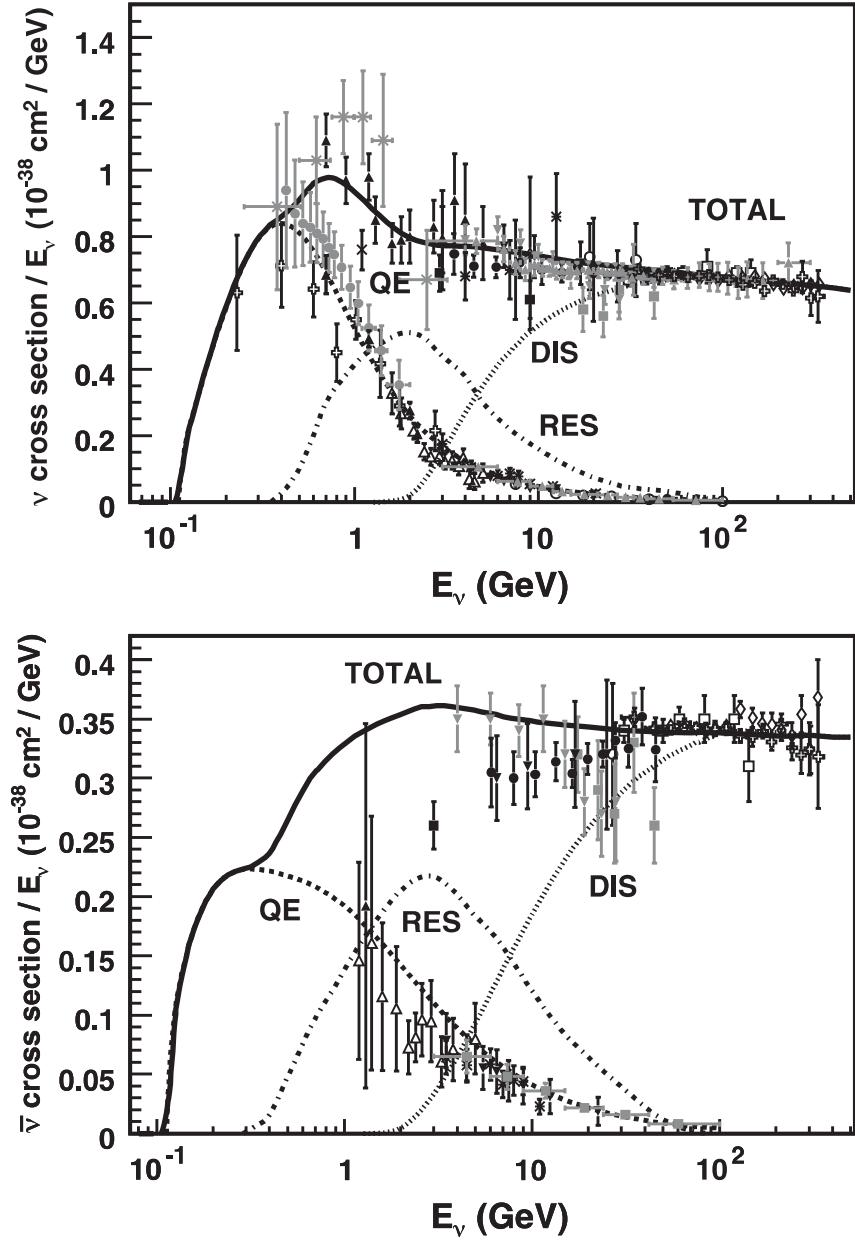


Figure 2.6: Predicted and measured total neutrino and antineutrino cross-sections for charged-current interactions with isoscalar nucleons, multiplied with the neutrino energy E_ν [87].

The isoscalar nucleon is a hypothetical hadronic singlet state with isospin $I_3 = 0$ whose particle distribution function (PDF) is the arithmetic mean between those of a proton ($I_3 = +\frac{1}{2}$) and a neutron ($I_3 = -\frac{1}{2}$). It is used as an approximation of average nucleons in light elements; the true isospin singlet is the deuteron with twice its PDF.

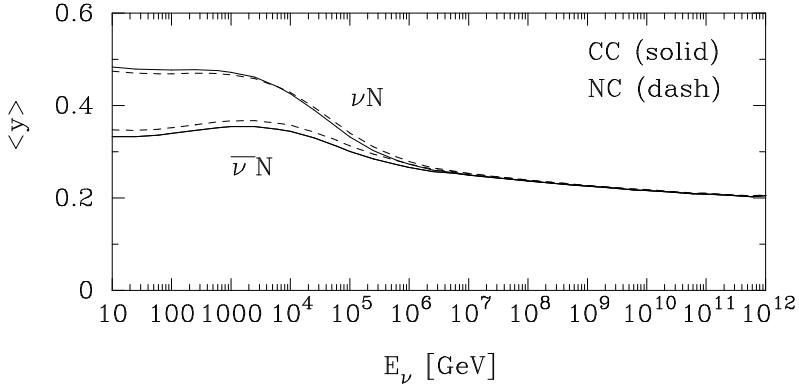


Figure 2.7: Predicted average of the kinematic inelasticity parameter Bjorken-Y, defined as $y = E_{\text{hadronic}} / E_\nu$, for CC and NC neutrino and antineutrino interactions as a function of the incident neutrino energy E_ν [89].

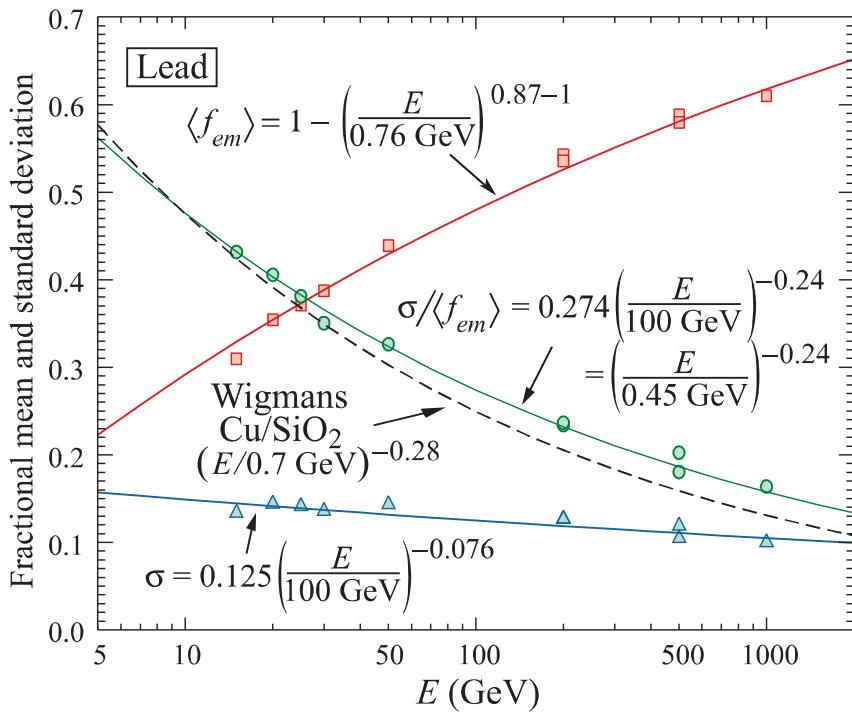


Figure 2.8: Mean energy fraction $\langle f_{em} \rangle$ of the electromagnetic component of a hadronic cascade in lead calorimeters as simulated by FLUKA90 (red), with absolute (σ , blue) and relative ($\sigma/\langle f_{em} \rangle$, green) standard deviations and parameterizations, as function of hadron energy [91].

Hadronic cascades can be seen clearly in Čerenkov detectors such as IceCube: Many of their constituent particles (mostly pions and kaons, but also protons and neutrons) are charged and therefore produce Čerenkov light. More importantly however, a large fraction of the uncharged particles are π^0 that decay into photons and cause electromagnetic cascades (see section 2.1.3). Figure 2.8 shows the fraction f_{em} of energy that goes into this electromagnetic component of hadronic cascades, simulated for a lead calorimeter. Although not directly applicable to this work, it shows the qualitative energy development of f_{em} and demonstrates that the fraction becomes substantial for high-energy events, which is also true for cascades in ice. Of the energy in the hadronic component, roughly 30% is invisible to Čerenkov detectors or causes photons that are highly delayed in time ($\gtrsim 1 \mu\text{s}$) because of elastic and non-elastic nuclear processes, neutron and neutrino production. Overall, the light yield of hadronic cascades in water ice lies between about 50% at 10 GeV and up to 90% over 10 TeV of the light yield of electromagnetic cascades (see [92]).

2.1.3 Lepton Propagation

The charged leptons that are produced by charged-current neutrino interactions produce the most important signatures for Čerenkov detectors as they carry most of the energy and either allow calorimetric energy reconstruction or much more precise angular reconstruction of the primary neutrino than the hadronic cascade at the primary interaction vertex. These signatures are distinct for the flavors.

Electrons

Because of their relatively low mass, electrons quickly lose energy through bremsstrahlung when traveling through dense matter such as ice. These high-energy photons then interact with the matter to produce new electron-positron pairs. These new particles again produce bremsstrahlung, which causes the number of particles to grow exponentially, until the particle energies lie below the energy required for another pair production [7]. Naïvely, this energy would be the minimum photon energy required for pair production, i.e., $2m_e = 1.022 \text{ MeV}$; however, the cascade is described better by defining a critical electron energy E_c^e at which the ionization energy loss over one radiation length X_0 is sufficient to stop the electron, because at that energy it becomes unlikely that the electron will emit another bremsstrahlung photon. These critical energies are 78.60 MeV for electrons and 76.50 MeV for positrons in water ice [7]. At $X_0 \rho_{\text{ice}} = 39.3 \text{ cm}$, the radiation length in water ice is small compared to the size of IceCube.¹¹ It is defined as the distance over which an electron loses $1 - \frac{1}{e}$ of its energy due to bremsstrahlung, but it also is $\frac{7}{9}$ of the photon mean-free

¹¹Water ice is basically incompressible, so $\rho = \text{const.}$ can be assumed safely for bubble-free ice.

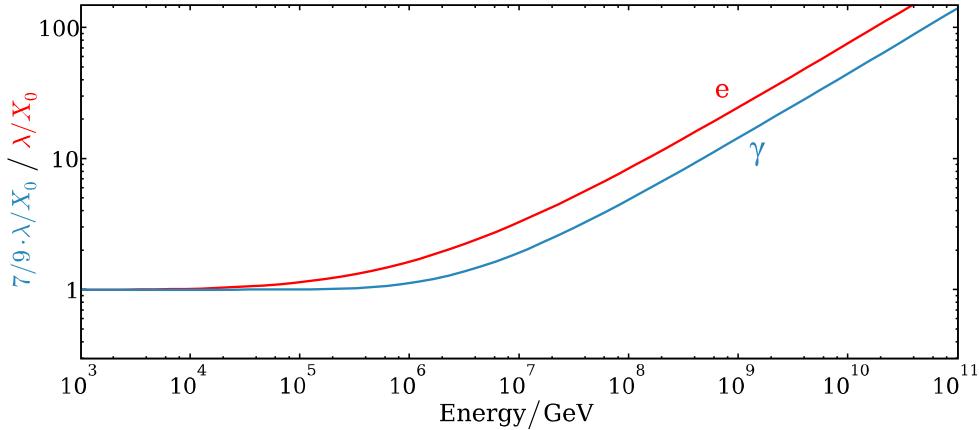


Figure 2.9: Radiation length (red) and photon mean free path (blue) in water ice modified by the LPM effect in units of the unmodified radiation length X_0 (based on [93]).

path for pair production. Therefore, the total length of an electromagnetic cascade can be approximated by

$$L = X_0 \rho_{\text{ice}} \log_2 \left(\frac{E_0}{E_c^e} \right),$$

which results in about 2.8 m for 10 GeV, 5.4 m for 1 TeV, and 8.0 m for 100 TeV.

This parameterization holds up to a few PeV, where the LPM effect starts to become relevant. Named after Landau, Pomeranchuk, and Migdal, this effect describes the increase of the radiation length at high energies: Both pair production and bremsstrahlung emission do not occur instantly, but over a short formation length. The formation length becomes substantial for ultra-relativistic particles, and additional interactions during the effect formation can suppress the initial interaction [94]. Figure 2.9 shows the relative increase of the radiation length due to the LPM effect. Note that at around 100 EeV, nuclear interactions start to dominate for both photons and electrons, but this is not relevant for IceCube or this work.

Muons

Muons are much more massive than electrons ($m_\mu = 206.8 m_e$), so their bremsstrahlung losses are smaller, resulting in a higher critical energy for muons in water ice: $E_c^\mu = 1031 \text{ GeV}$ [7]. Below this energy, ionization losses dominate and cause a steady loss of particle energy; above, electron pair production, nuclear interactions, and bremsstrahlung prevail, summarized as radiative losses. Especially bremsstrahlung often leads to large (so-called hard) losses in which over 10% of the muon energy is transferred into a single electromagnetic cascade.

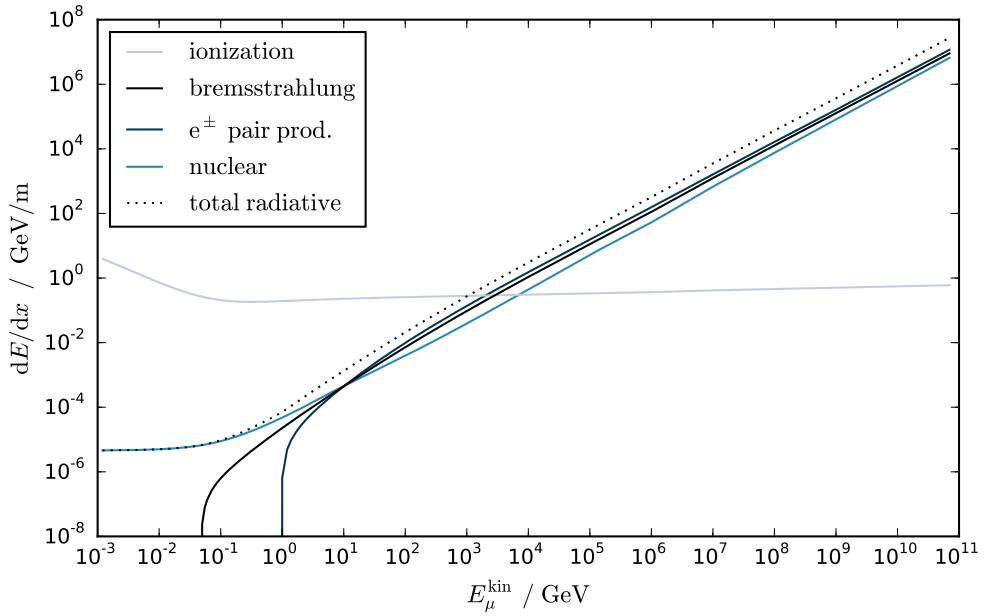


Figure 2.10: Parameterization of the energy loss of a muon in water ice due to ionization and three radiative processes, based on data from [7] and [95].

The energy loss of muons in ice can be parameterized as

$$\frac{dE}{dx} = -a - bE, \quad (2.4)$$

with the ionization losses given by $a(E)$ and radiative losses $b(E)E$ [7]. Both parameters are functions of the muon energy, but are fairly constant above 100 GeV, see figure 2.10. Therefore, the energy loss (evaluated at 1 TeV) can be written as

$$\frac{dE}{dx} = -0.28 \text{ GeV/m} - 0.272 \cdot 10^{-3} E/\text{m}.$$

The expected range due to the muon energy loss can be approximated by solving (2.4):

$$\begin{aligned} E(x) &= (E_0 + E_c^\mu) e^{-bx} - E_c^\mu \\ \Rightarrow L(E) &= \frac{1}{b} \ln\left(\frac{E}{E_c^\mu} + 1\right) = 3676 \text{ m} \ln\left(\frac{E}{1031 \text{ GeV}} + 1\right) \end{aligned} \quad (2.5)$$

This range is significantly smaller than the muon range due to its limited life time,

$$L^{\text{decay}}(E) = \tau_0 c \frac{E}{m_\mu} = 6233 \text{ m} \frac{E}{\text{GeV}} > 1000 L(E),$$

so the muon decay can be neglected as it only takes place when the muon has been stopped already, at energies too low for IceCube to detect.

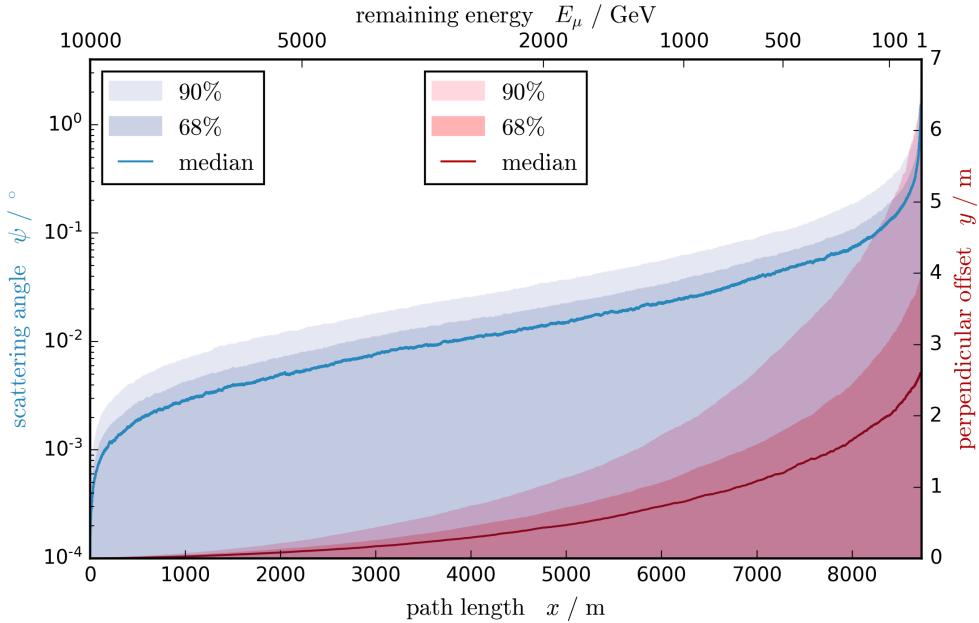


Figure 2.11: Distributions of the planar scattering angle ψ and of the offset y due to scattering of 10 TeV muons down to 1 GeV. The upper horizontal axis indicates the expected muon energy after traveling the path length x indicated on the lower axis.

The change of the direction of the muon due to multiple scattering can be neglected as well. Because of the many repeated scattering processes, the central-limit theorem applies, so the distribution of the planar scattering angle ψ can be approximated by a Gaussian. According to ref. [96], the width of this distribution is given by

$$\sigma_\psi(E, x) = \frac{0.0136 \text{ GeV}}{E} \sqrt{\frac{x}{X_0}} \left(1 + 0.088 \log_{10} \left(\frac{x}{X_0} \right) \right), \quad (2.6)$$

where X_0 is the electromagnetic radiation length (see section 2.1.3). This parameterization works well for 98% of the particles and only fails at describing the tails, which are better described by a Molière distribution. The larger problem however is the muon energy loss that is not taken into account, but becomes important for large x . Figure 2.11 shows the approximate distribution of the (absolute) planar scattering angle ψ for a 10 TeV muon in ice in the Gaussian approximation. Even at 10 GeV, the median is 0.44° and 90% of the particles scattered less than 1.12° .

Note that the angle describes the direction of flight of the muon at that position; the more relevant quantity is the offset perpendicular to the unscattered track, which is also shown in the figure. This offset increases the most at the end of the track, but even if it were picked up by the track reconstruction, it would still produce effective scattering angles significantly smaller than ψ .

Taus

Tau leptons are even more massive than muons ($m_\tau = 16.82m_\mu = 3477m_e$), so their energy losses are smaller. Equation (2.4) also applies to taus, with $a = 0.18 \text{ GeV/m}$ and $b \approx 0.55 \cdot 10^{-4} \text{ m}^{-1}$ [97]. At energies above $1 \cdot 10^8 \text{ GeV}$, the energy dependence of b has to be taken into account, but this is far above the relevant energy range of this work; the critical energy of taus in ice is approximately $E_c^\tau = 3300 \text{ GeV}$.

With a lifetime of $\tau_\tau = 290 \text{ fs}$ [7], the limiting factor of tau track length is its decay length of $L^{\text{decay}} = 49 \text{ m} \frac{E}{\text{PeV}}$ up to an energy of about 5 EeV, where the energy loss begins to become relevant. Not counting the outgoing neutrino, most taus decay hadronically into one or more charged and possibly additional neutral mesons; only about 35.2% decay purely leptonically. Of these leptonical decays, 17.8% produce an electron and the remaining 17.4% produce a muon.

Muonic decays are difficult to tell apart from muon neutrino events; on average, they would feature seemingly large Bjorken-Y values (section 2.1.2) because two neutrinos carry away energy that would have otherwise gone to the muon.

In the other cases, both hadrons and electrons cause either hadronic or electromagnetic cascades in ice, so the main signature of a tau neutrino event in IceCube is a hadronic cascade at the neutrino interaction followed by a cascade caused by the tau decay, connected by a relatively faint Čerenkov track. Because of the tau's short range at energies below a few PeV, those cascades usually overlap and in that case cannot be separated by usual cascade reconstruction methods. Methods were developed to discriminate overlapping double cascades from single cascades even at very small distances ($\gtrsim 5 \text{ m} \doteq 100 \text{ TeV}$) [98], but so far no tau neutrino interaction has been unambiguously identified with IceCube, mostly because of uncertainties in the ice properties.

2.2 Design and Hardware

The IceCube Neutrino Observatory consist of an in-ice array used as Čerenkov detector for neutrino detection, and an on-ice cosmic-ray array called IceTop, both connected to a central building called IceCube Laboratory (ICL) [77]. While IceTop has produced many important results on its own and can also be used as a veto to discriminate downgoing astrophysical neutrino events from atmospheric muons and atmospheric neutrinos, it is not relevant for this work and will not discussed in the following. Instead, IceCube refers to the in-ice array unless explicitly specified otherwise.

The in-ice array is built of 86 vertical cables called strings, which are instrumented with 60 Digital Optical Modules (DOMs) in a depth of 1450 m to 2450 m each. 78 of these strings are located on a triangular grid with an approximately hexagonal outline and 125 m horizontal spacing. On these strings, the vertical spacing between the DOMs

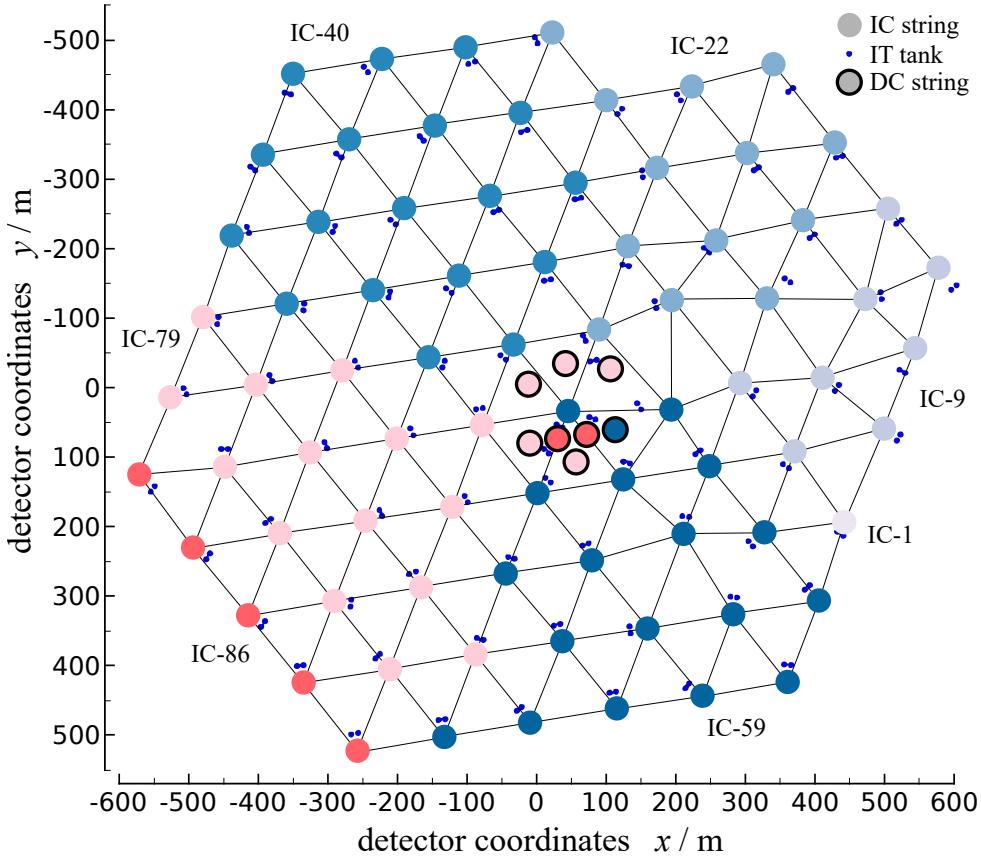


Figure 2.12: Footprint of standard IceCube strings, DeepCore strings, and IceTop tanks. Colors indicate the season of deployment. This work is based on data taken with the blue parts of the detector (IC-59).

is constant at 17 m [77].

This is different for the remaining 8 strings, which together with the 7 surrounding regular strings form the low-energy extension DeepCore. DeepCore is located close to the center of IceCube and features a much denser instrumentation, lowering the detection energy threshold by about one order of magnitude to below 10 GeV. The extension strings are instrumented evenly with 50 DOMs between depths of 2100 m to 2450 m (i.e., 7 m spacing), and the 10 remaining DOMs are located between 2000 m and 1900 m with 10 m spacing to form a veto against downgoing muons [99].

Located at the Geographical South Pole, the deployment of IceCube was a technical and logistical challenge. During Southern Hemispheric winter, average temperatures lie between -55°C and -63°C and can reach down to -80°C [100]. Because of this, deployment could only take place in the Antarctic summer months. The first string was installed in 2004-2005, and in the following seasons, the number of strings increased to 9, 22, 40, 59, 79, and finally 86 when the construction was finished in 2010-2011. Because of the modular design of IceCube, data taking was already possible with non-final detector

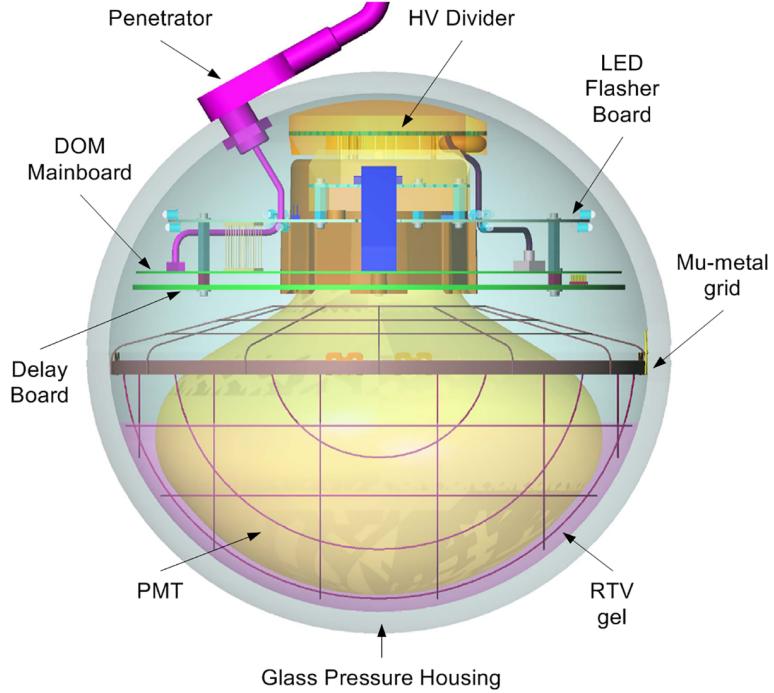


Figure 2.13: Layout of an IceCube Digital Optical Module (DOM) [101].

configurations; this work is based on data taken with the 59-string configuration (IC-59) between May 2009 and May 2010, as shown in figure 2.12. This configuration consists of 58 standard strings and the first DeepCore string, arranged in a concave shape, somewhat similar to that of the popular video game character Pac-Man.

2.2.1 Digital Optical Modules

The digital optical modules consist of a photomultiplier tube (PMT), an electronics board for triggering and digitization, a board for voltage transformation, and a board with LEDs for calibrations (the so-called flasher board), all housed in a waterproof and pressure resistant glass sphere with an outer diameter of 13" (33 cm) as displayed in figure 2.13 [77]. The PMT is 10" model (25.4 cm) from Hamamatsu (R7081-02 or R7081MOD for DeepCore DOMs), which is shielded from the terrestrial magnetic field by a mu-metal grid and embedded into gel for stability and optical coupling. The overall spectral acceptance of the DOM including glass and optical gel is illustrated in figure 2.2.

The DOMs are capable of autonomously triggering (*launching*) the recording of a digitized waveform when the voltage at the PMT surpasses a threshold set at the equivalent of 0.25 photo electrons. To not lose the beginning of the waveform, the PMT signal is divided and routed through the trigger logic while the other route leads to a delay board.

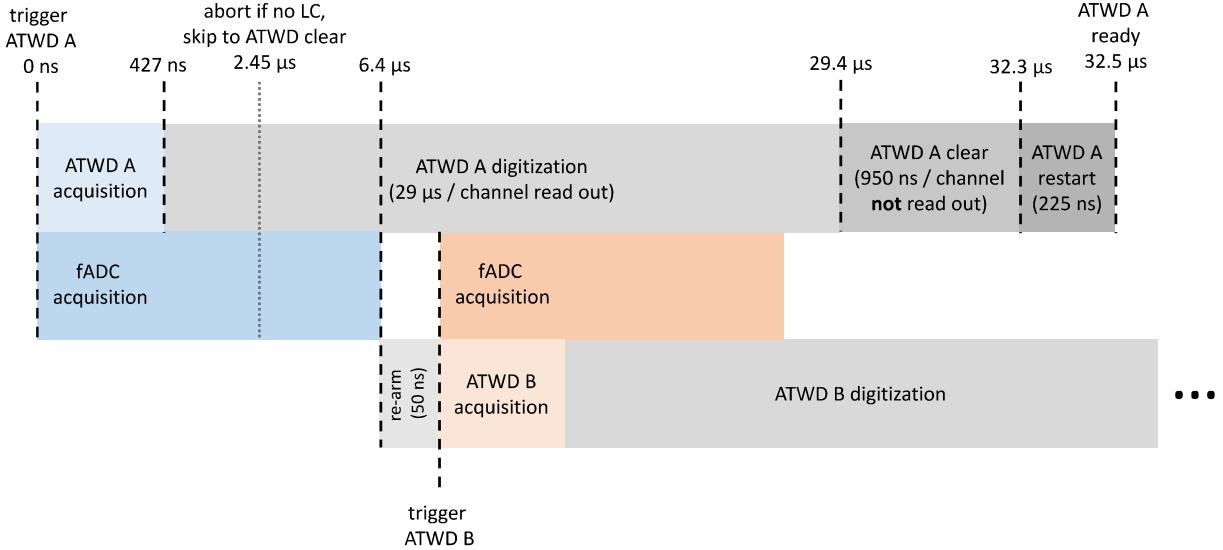


Figure 2.14: Diagram of the DOM digitizer launch logic, with two adjacent HLC launches as described in the text [77]. The time axis is not to scale.

When the DOM launches, the digitizers can access the data out of the delay board and start recording before the triggering part of the signal, with the delay being 75 ns [77]. DOMs are also capable of directly communicating with DOMs on the same string. A launch only gets completed when one of the neighboring and next-to-neighboring DOMs also launches within 1 μ s, which is called *hard local coincidence* (HLC)¹². Otherwise, the digitization is stopped and only very condensed information is transferred to the surface, referred to as *soft local coincidence* (SLC) charge stamp.

The digitization is performed by three digitizers. Two ATWD chips are available, and one fADC. The fADC samples continuously and saves 256-bin waveforms with a bin width of 25 ns, therefore spanning 6.4 μ s. Its input signal is smeared out intentionally to spread it over multiple bins to allow for sub-bin precision after reconstruction. Its effective downtime is 50 ns after one waveform has been saved.

The two ATWD chips are identical and take turns because their downtime after digitization is substantial at 32.5 μ s after launch. Both chips have three different gain channels, which when combined have a much higher dynamic range than the fADC, with relative gains of 0.21 to 15.7. They only launch together with the fADC, so there is at least 6.45 μ s between any two ATWD waveforms. The launch logic is summarized in figure 2.14, more information can be found in ref. [77] and [83].

The flasher board features twelve violet LEDs ($\lambda = 399$ nm), which are arranged uniformly in a horizontal ring in pairs of two. In every pair, one LED is positioned such that after refraction, it emits light horizontally, and the other one such that it

¹²Hardware-wise there is the possibility for a second condition called self-local coincidence based on the recorded charge [102], but it is not implemented at this point because the current scheme works well.

emits upwards at an angle of 48°. All LEDs can be flashed individually or in groups at customizable intensities and pulse widths, and their resulting photon outputs correspond to neutrino interactions between 7 GeV and 1 PeV [77].

The light signals produced by those boards can be received by the same or other DOMs and have been used for precise measurements of the relative DOM positions (seeded with the positions estimated during deployment) and the properties of the ice. To be able to measure wavelength-dependent ice properties, eight DOMs in one of the DeepCore strings and eight DOMs on one of the edge strings have had their LEDs replaced by LEDs of four different colors, but data taken with those *color-DOMs* has not been incorporated into ice models yet.

2.2.2 Ice Properties and Models

The Čerenkov medium of the IceCube Neutrino Observatory is Antarctic ice that has accumulated at the South Pole naturally over a time scale of roughly 165000 years, with IceCube extending to layers approximately 100000 years old [104]. During this time, climate and air properties changed often, and most importantly the concentrations of dust and volcanic ash varied strongly, leading to significant variations in the optical ice properties on a scale of less than 50 m.

The ice properties relevant to IceCube can be parameterized well by the combination of two phenomena: scattering and absorption [103]. Scattering can be described by a mean free path length λ_s and a mean scattering angle θ_s . The mean scattering angle – calculated using the Mie model of scattering – is quite small with $\langle \cos(\theta_s) \rangle = 0.94$. Because of this, the more useful and better accessible quantity is the effective scattering length:

$$\lambda_e = \lim_{n \rightarrow \infty} \left(\lambda_s \sum_{i=0}^n \langle \cos(\theta_s) \rangle^i \right) = \frac{\lambda_s}{1 - \langle \cos(\theta_s) \rangle}$$

It describes the average distance after which the direction of the light has become isotropic. Absorption can similarly be modeled by a single parameter, namely the absorption length λ_a . It is defined as the distance after which $1 - 1/e \approx 63.2\%$ of the photons have been absorbed [103].

Figure 2.15 shows the inverse effective scattering length and the inverse absorption length as functions of depth for multiple wavelengths as measured by AMANDA. The effective scattering length rises dramatically from about 1400 m upwards because of bubbles in the ice. At higher pressures, those bubbles have been compressed into clathrates that do not influence light propagation significantly. In these depths, the main reason for scattering are dust particles that cause several non-trivial features. Four layers of increased dust concentration have been identified. The deepest one – named D or simply *the dust layer* – at a depth of about 2050 m has the strongest impact by far; it is the

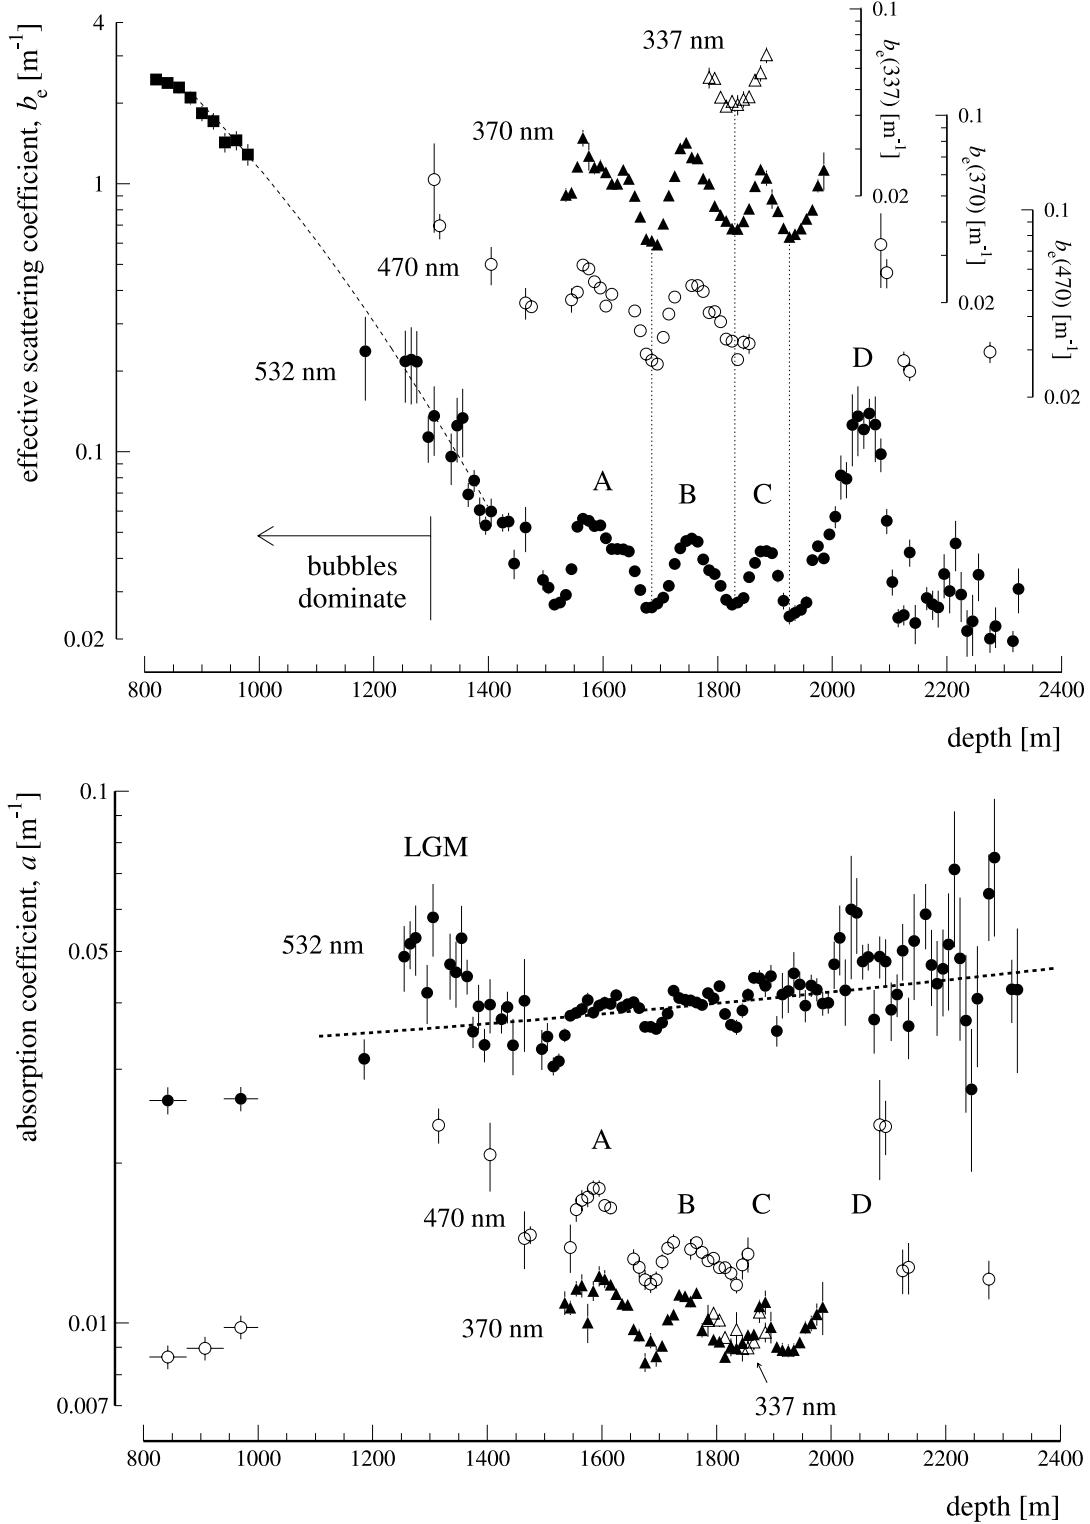


Figure 2.15: Inverse effective scattering length $b_e = \lambda_e^{-1}$ (**top**) and inverse absorption length $a = \lambda_a^{-1}$ (**bottom**) for ice at the South Pole vs. depth for different wavelengths as measured by AMANDA. Scattering becomes dominated by air bubbles at low depths and is largely wavelength independent at 1000 m and above [103].

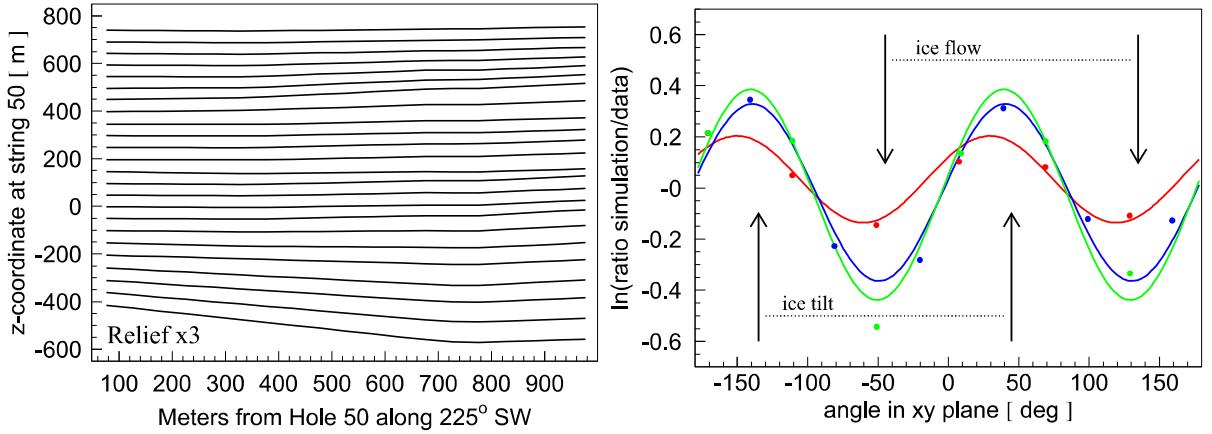


Figure 2.16: Left: Depths of ice layers along the average direction of the ice tilt gradient. For clarity the changes have been exaggerated by a factor of three, so the y-axis only applies to the left-most points [86].
Right: Average charge ratio between experimental data and simulated data without taking into account anisotropy (using Spice Mie). The amplitudes increase with distance (red $\doteq 125$ m; blue $\doteq 217$ m; green $\doteq 250$ m) [105].

reason why DeepCore strings have their main instrumentation below 2100 m and the top veto above 2000 m, omitting the dust layer. The wavelength dependence of scattering is relatively weak, with large wavelengths being scattered less.

In ultra-violet regions, the absorption length shows the same dust-induced structure as the scattering length, but the absorption by pure ice starts to dominate at about 480 nm, reducing the importance of the dust concentration at longer wavelengths. Overall, the absorption length at the DOMs' peak sensitivity near 400 nm is close to 100 m and therefore larger than the effective scattering length by a factor of four to five.

These wavelength dependencies are described by six global parameters that have been fitted with AMANDA data: one parameter for scattering, one parameter for absorption of dust, two for absorption on pure ice, and two parameters that provide a relation between scattering and absorption on dust. However, all but the last two of these parameters are universal for all contemporary ice models, so they are of little interest in the following; for more detail, see ref. [86].

Older IceCube ice models such as AHA (2007 update of the model described in ref. /citeAMANDAIce), Spice 1 [106] (2009, adds IceCube flasher data), and WHAM! (2011 update of AHA) described the ice with depth as the only spatial dimension, mostly because of limitations in simulation software and available computational power. With the advent of direct-photon-propagation (i.e., ray-tracing) software that can run performantly on GPUs (such as PPC [107] or clsim [108]), it became possible to take into account additional phenomena and degrees of freedom such as a tilt of the ice layers or anisotropy of the optical properties. The tilt of the ice has been measured to 2 mm accuracy during deployment by dust loggers on eight different strings [106] (figure 2.16), and

has been compared to Antarctic ice cores to model regions below the detector [86]. It was included in Spice 2 with the caveat that at that point no tilted high-energy datasets (i.e., datasets with a generated neutrino spectrum of E_ν^{-1}) could be produced with the available hardware. Large-scale high-energy production was only possible using tabulated ice properties with the software *photronics*, which could not be extended by another dimension because of memory constraints. This limitation dropped over time when higher numbers of increasingly stronger GPUs became available, and additional improvements to the ice models were made. Many of those changes were related to software and methodological improvements of the fitting algorithm, but there were also two major changes in the underlying physics model:

Spice Mie introduced a new parameterization for individual scattering processes. Instead of using exclusively the Henyey-Greenstein approximation to the Mie model that itself is too computationally intensive to be used on production scales, it uses a linear combination of Henyey-Greenstein with another approximation called Simplified Liu to improve the parameterization especially for small angles up to 35° [86].

The other change came with Spice Lea (stylized as λ_e -a), which introduced anisotropy of the optical parameters [105]. Figure 2.16 shows the impact of this anisotropy on the amount of charge (or light) captured by IceCube’s DOMs as a function of azimuth angle towards the flashing string at three distances: Simulations that do not take into account the anisotropy (Spice Mie) are off by up to 50%, overpredicting light parallel to the ice tilt and underpredicting it perpendicular to the tilt. One possible explanation for this phenomenon is the flow of the ice that happens perpendicular to the ice tilt at a speed of roughly $10 \frac{\text{m}}{\text{a}}$; it might influence the preferred alignment of either ice crystals or dust particles [105]. The baseline ice model for this analysis is Spice Lea (see section 6.1.3).

One more detail of the ice model is hole ice. When IceCube strings were deployed, holes were molten into the ice and the water was left to freeze again. This process removes the layered structure, might introduce contaminants, and could be responsible for the formation of bubbles. The deployment of cameras together with DOMs brought some information about the properties of this so-called hole ice, and there are efforts underway to improve the understanding of hole ice [109]. The model used in this work assumes the hole ice to be homogenous, 60 cm in diameter, and to differ from the remaining ice by strongly increased scattering; the default value is $\lambda_e^{\text{hole}} = 50 \text{ cm}$.

2.3 Processing and Event Reconstruction

The data processing and event reconstruction chains in IceCube are fairly complex and have been explained in various other works. For this reason, only a short overview will be given with a focus on the tools and algorithms relevant for this work. For more details, see ref. [77], [110], and [111].

2.3.1 Triggering, Low-Level Reconstruction and Filtering

DOMs that capture enough charge to get launched as described in section 2.2.1 either send HLC waveforms or SLC charge stamps to the computer cluster located inside the ICL at the surface. The data acquisition system (DAQ) checks the HLC launches for various predefined temporal and spatial patterns. If at least one of these so-called trigger conditions is fulfilled, all HLC and SLC launches in a time window (usually $-10\ldots 10\ \mu\text{s}$) around the triggering HLC launches are collected into an event.

To decide whether an event could be useful for one of the many physics analyses that are conducted with IceCube, events are passed to filtering. Filters are designed to cover a broad number of physics cases to keep the number of required filters (and with it CPU time requirements) as low as possible while still being specific enough to allow the necessary reduction of unwanted background events because of limited satellite bandwidth. Examples for IC-59 filters are the muon filter that specializes on muon tracks and is the only one relevant to this work, the cascade filter, or the extremely-high-energy (EHE) filter. Note that in the following, the procedures are described for IC-59; for later seasons, various aspects have been improved slightly.

Every filter is responsible for running the reconstructions that are necessary to make a decision, although they can share reconstructions with each other. Firstly, the triggered events are passed to low-level reconstruction, where waveforms and charge stamps are calibrated to compensate for individual DOM gain settings, electrical droop and similar effects. The resulting calibrated waveforms and charge stamps are then feature extracted to obtain so called pulses that contain the arrival times and numbers of photons on each DOM (measured as electrical charge at the PMT), using the algorithms of FeatureExtractor described in ref. [83].

Many of these extracted pulses are not associated with muon or neutrino events but are caused by PMT noise or radioactive decays (mostly in the DOM glass spheres). Additionally, there is no guarantee that triggered events contain only one particle signature; instead, about 10% of the events are expected to contain coincident atmospheric muons, irrespective of whether the other particle is a neutrino or an atmospheric muon. Often, those coincident muons do not have enough energy to produce clearly visible tracks by themselves and in this regard produce pulses that are comparable to noise.

To allow or improve the reconstruction of muon tracks, potential noise pulses are removed by various cleaning algorithms; these algorithms include general time window cuts, time-space causality cuts (e.g., SeededRTCleaning [112]) and more advanced cluster algorithms (e.g., TopologicalTrigger [113]). During filtering, only a dynamic time window cut is used that selects pulses in a $10\ \mu\text{s}$ window such that the sum of their charges is maximized.

After the cleaning, events are passed to the initial-guess track reconstruction algorithm LineFit [110] that is based on a plane-wave approach. A more advanced iterative likelihood reconstruction (SPE [110], see next section) is then run for a single iteration based on the LineFit output, and the directions of both fits as well as the number of launched

DOMs, the total PMT charge, and the goodness of fit of the reconstructions are used as filter criteria to decide whether the event passes the filter and gets sent to the Northern Hemisphere for further processing.

2.3.2 High-Level Reconstruction

During the IC-59 season, CPU power at the ICL at the South Pole was much more limited than it is now, so only easy reconstructions were done at the Pole. After filtering, low-level reconstructions were repeated in the North with more advanced cleaning and more specialized settings for different filter streams. Most importantly for the muon stream, TopologicalTrigger (TT) was used to identify possible coincident events prior to track reconstruction. The general idea behind this algorithm is to start with the earliest pulse and to add other pulses that are causally connected by a space-time requirement $|\Delta t - \Delta \vec{x}/c| < \alpha$ [113], with $\alpha = 450\text{ ns}$ for this work. Individual pulses and small clusters (less than five different DOMs) are discarded as noise, while larger clusters are returned as coincident subevents, for which individual reconstructions can be run. After analyzing the algorithm’s performance, it was decided to discard all events for which TopologicalTrigger identified multiple clusters in favor of fewer systematical uncertainties, because this happens for less than 5% of the events at that stage [114], and at that time no simulated neutrino data with coincident cosmic rays was available. The remaining events are categorized in two streams: *split* means that TopologicalTrigger found and discarded pulses that do not belong to the main cluster, while for *unsplit* all pulses were found to be causally connected. This distinction plays a small role in the event selection (chapter 3), and the two subsamples have been thoroughly tested for systematic differences on analysis level.

Based on the TopologicalTrigger output, the LineFit algorithm is run again and its output is again used as the initial guess of a single iteration of the likelihood-based track reconstruction SPE. SPE stands for Single Photo-Electron and means that the employed probability density function (PDF) describes the probability for a single photon to arrive from the track hypothesis to a given DOM at a given time [110]. Unlike LineFit, it uses tabulated probabilities that take into account non-homogenous ice models, as well as the Čerenkov cone geometry. The reconstruction uses only the first pulse per DOM, because unscattered (*direct*) photons provide the best information about the track geometry. Seeded with this single-iteration fit, the SPE fit is run for eight more iterations to increase the likelihood of the fit finding the true maximum-likelihood solution, and with it a good reconstruction of the actual muon track. The result of this fit is used for a single iteration of a more advanced likelihood fit called MPE, which stands for Multi-Photo-Electron. Its PDF takes into account that the most likely arrival time of the earliest photon depends on the number of photons that arrive at a given DOM, which is a more realistic hypothesis, but narrows the time PDF for high photon numbers and thereby complicates convergence.

With the geometrical properties of the muon track reconstructed, various energy reconstruction algorithms can be employed. Algorithms that do not rely on track information can only give very rough energy estimates because of the inhomogeneous ice and very diverse track geometries (most importantly the track length inside the detector). With a given track geometry, most algorithms employ tabulated photon arrival probabilities binned in zenith angle and depth, obtained by Monte-Carlo (MC) simulations. The algorithm used for this analysis is called muEx and instead only uses the input track and a one-dimensional ice model to do a likelihood fit of the most probable muon energy. It is an improved version of the analytical algorithm described in ref. [111] that takes into account ice properties and hard (stochastic) energy losses as described in section 2.1.3. Its resolution is discussed in section 4.3.1. For in-depth explanations of IceCube energy reconstructions, see ref. [111].

Other algorithms (most importantly TruncatedEnergy [115]) have been tested but have been found to perform on average at best as good as muEx between 100 GeV and a few hundred TeV.

Finally, quality parameters are calculated for every event to help assessing the reconstruction quality. Those quantities can then be used for the event selection that is described in the next section.

2.4 Neutrino and Muon Simulations

Most of IceCube’s analyses rely on the comparison between experimentally measured and simulated data (i.e., Monte-Carlo data, in short MC). The simulation chain consists of many tools; this section will only give an overview and point out relevant features of simulated data.

The simulation of a MC event starts with the generation of the primary particle. For muon datasets this is done by the cosmic-ray simulation code CORSIKA [116]; for neutrinos of all flavors, the proprietary code neutrino-generator is used, which is originally based on ANIS [117]. Both tools handle all relevant physics processes except for oscillations up to the point where a charged lepton enters the vicinity of the detector.

In order to create neutrino datasets without wasting huge amounts of CPU time, neutrino generation exclusively employs weighted events: Instead of generating huge numbers of neutrinos and discarding all those that would not interact, every neutrino is forced to interact close to the detector and a weight is assigned to the event to take into account the (usually very small) interaction probability. Additionally, the energy spectrum at which events are generated is set to a simple power law (usually E^{-2} or E^{-1}) to accumulate high statistics in desired energy regions and to allow easy reweighting to different physical fluxes (e.g., atmospheric models, astrophysical sources, etc.).

CORSIKA can be run in weighted or unweighted modes, and datasets from both modes are used for this analysis.

The charged leptons are passed to the lepton propagation code to simulate stochastic energy losses, absorption, decay, and secondary particles. This is usually done by MMC (*Muon Monte Carlo* [118]), but can also be done by Geant4 [119] when using the `clsim` photon propagation code. The resulting tree of particles and losses is then passed to the photon propagation code as described in section 2.2.2. The output of these tools are the numbers and timings of photons hitting IceCube DOMs.

Additional photons are added to this output by the noise simulation module Vuvuzela (version 2). It adds both uncorrelated Poissonian noise caused by thermal effects and radioactive decay, as well as so-called correlated noise that is following other hits and might be caused by scintillation in the PMTs and glass pressure spheres [112].

The resulting lists are passed to proprietary code to turn the numbers and timings of photons into voltage distributions like those recorded by the actual PMTs, and afterwards both the DOM triggering and subsequently the IceCube event triggers are simulated to match the triggering described in section 2.3.1. From that point onward, MC events contain all information present in experimentally measured events, so the remaining reconstruction chain is identical.

Note that neutrino interactions in the Earth including energy losses and absorption (which does not eliminate the particle, but instead reduces its weight) are handled without taking into account oscillation effects, which are multiplied to the weights at a later stage. For a discussion of this approximation, see appendix C.1.

Another approximation often encountered in IceCube MC data is DOM oversizing. It refers to the upscaling of DOMs during GPU photon propagation, perpendicular to the current direction of the photon (resulting in an oblate) and usually by a factor of five. Photons that hit this oversized volume are registered as hits and get assigned a lower weight and an additional time offset to compensate for the oversizing. They are not terminated unless they hit the non-oversized DOM to not artificially obstruct DOMs that are further away. This oversizing reduces the number of propagated photons that is required to get a precise sampling, but can also distort hit times for DOMs where five DOM radii are not negligible compared to the distance between DOM and photon source. Because of the higher-than-usual susceptibility of this analysis to systematic effects, all datasets produced for this analysis had oversizing turned off.

Chapter 3

Event Selection

As indicated in section 3.1, the baseline event selection has been developed in close cooperation with Anne Schukraft and has previously been published in ref. [114]. Besides the common development and benchmarking of the direct cuts, the author of this work especially contributed the alternative event selection based on machine learning, as well as work on the employed tools.

Event selection in the context of this work is the task of identifying well-reconstructed muon neutrino events to discriminate them from badly reconstructed ones and most importantly from atmospheric muon events, which form the largest background for this analysis. The main idea for analyses that require a high sample purity – defined as fraction of neutrino events over the total number of events – is to select only upgoing events, or more precisely only events with enough overburden to shield the detector from atmospheric muons. This cut on the zenith angle alone reduces the rate of atmospheric muons by many orders of magnitude.

However, a small fraction of atmospheric muons gets wrongly reconstructed as upgoing, usually because of pathologically scattered light, noise pulses, or coincident atmospheric muon events with timings that are compatible with the causality requirements of TopologicalTrigger (see section 2.3.2) and track hypotheses. Even though the relative fraction of these events is small, the low cross-section of neutrinos means that they are still more abundant than actual muon neutrino events by a factor of about 10^3 . By carefully analyzing and cutting on various quantities, atmospheric muon events can largely be removed from an event sample while still retaining a large fraction of muon neutrino events.

3.1 Baseline Event Sample

The baseline event sample used for this analysis was developed in collaboration with Anne Schukraft, and is described in great detail in her dissertation [114]. It is optimized for high signal efficiency at high energies to search for astrophysical neutrino events (with an expected energy spectrum of E^{-2}) with a special emphasis at high purity and good data/MC agreement. However, it also features a good passing rate for atmospheric neutrinos because firstly it was designed also with this search for sterile neutrinos in mind, and secondly good statistics for atmospheric neutrinos over a large energy range help to constrain systematic uncertainties even for high-energy analyses.

The selection consists of twelve one-dimensional cuts that have been carefully tested for redundancy and have been tuned independently for the split and unsplit streams mentioned in section 2.3.2 [114]. Earlier cuts that are part of the muon filter or of the subsequent common processing of this data stream (*muon Level 3*) are strictly weaker and therefore irrelevant for this work. The cuts are summarized in table 3.1 and explained in the following.

The first cut discards downgoing events based on the MPE reconstruction. This cut was already performed for the Level 3 selection and therefore only affects the split stream where the MPE reconstruction has changed due to the TopologicalTrigger hit cleaning.

Cut 2 does the same but uses LineFit instead of MPE. While MPE usually gives better reconstructions, LineFit often performs more robustly for coincident atmospheric muon events, partly because it is allowed to fit the particle speed, which is locked to c for SPE and MPE.

Cut 3 rejects events with bad MPE quality of fit. It uses the reduced negative log-likelihood (rlogl), which is the (neg.) log-likelihood value at the best fit divided by the degrees of freedom, i.e., the number of hit DOMs minus the five independent fit parameters.

Cut 4 discards events for which the angular uncertainty of MPE estimated by the paraboloid code is larger than 5° . For high-quality neutrino events, this requirement is quite weak (compare figure 4.12 for the zenith component), so this cut again mostly discards misreconstructed atmospheric muon events.

Cut 5 rejects events for which LineFit and MPE disagree by more than 15° in their angular reconstructions, for the same reason as cut 2.

Cut 6 performs a likelihood ratio test (see section 5.3) between the regular SPE fit and an SPE fit that was forced to perform a downgoing reconstruction (*Bayesian*). The SPE fit was selected for computational reasons, as it is faster than the MPE fit, which furthermore requires SPE as a seed.

Table 3.1: Cuts performed to get the baseline event sample, starting from the more generic *muon Level 3* sample [114]

n_{cut}	cut	passing rate atm. $(\bar{\nu})_{\mu}$		passing rate exp. data	
		indiv.	cumul.	indiv.	cumul.
1	$\cos(\theta_{\text{MPE}}) < 0$	100.00 %	100.00 %	100.00 %	100.00 %
2	$\cos(\theta_{\text{LineFit}}) < 0$	90.78 %	90.78 %	67.34 %	67.34 %
3	$r\log l_{\text{MPE}} < 11$	74.34 %	67.49 %	35.51 %	23.91 %
4	$\sigma_{\text{paraboloid}} < 5^{\circ}$	67.28 %	45.41 %	14.74 %	3.525 %
5	$\Delta\psi(\text{LineFit}, \text{MPE}) < 15^{\circ}$	91.05 %	41.35 %	35.33 %	1.246 %
6	$\Lambda(\text{SPE, Bayesian}) > 29$	67.35 %	27.85 %	22.91 %	0.2853 %
7	$\min(\theta_{\text{geo1}}, \theta_{\text{geo2}}, \theta_{\text{t1}}, \theta_{\text{t2}}) > 80^{\circ}$	65.40 %	18.21 %	14.24 %	0.0406 %
8	$n_{\text{dir}} > 6$	94.88 %	17.28 %	83.65 %	0.0340 %
9	$l_{\text{dir}} > 250 \text{ m}$	97.36 %	16.82 %	96.64 %	0.0329 %
10	$ S_{\text{dir}} < 0.45$	93.83 %	15.79 %	92.92 %	0.0305 %
11	$\text{cog}_z < 400 \text{ m}$	98.40 %	15.53 %	98.89 %	0.0302 %
12	$\text{cog}_z > -450 \text{ m}$	98.30 %	15.27 %	97.90 %	0.0296 %
(split)					
1	$\cos(\theta_{\text{MPE}}) < 0$	97.12 %	97.12 %	95.14 %	95.14 %
2	$\cos(\theta_{\text{LineFit}}) < 0$	91.53 %	88.90 %	56.13 %	53.40 %
3	$r\log l_{\text{MPE}} < 11$	75.50 %	67.12 %	36.97 %	19.74 %
4	$\sigma_{\text{paraboloid}} < 5^{\circ}$	67.58 %	45.36 %	16.41 %	3.240 %
5	$\Delta\psi(\text{LineFit}, \text{MPE}) < 15^{\circ}$	93.80 %	42.55 %	42.04 %	1.362 %
6	$\Lambda(\text{SPE, Bayesian}) > 29$	64.99 %	27.65 %	38.94 %	0.5305 %
7	$\min(\theta_{\text{geo1}}, \theta_{\text{geo2}}, \theta_{\text{t1}}, \theta_{\text{t2}}) > 80^{\circ}$	71.24 %	19.70 %	13.54 %	0.0718 %
8	$n_{\text{dir}} > 6$	92.11 %	18.14 %	81.64 %	0.0586 %
9	$l_{\text{dir}} > 250 \text{ m}$	97.56 %	17.70 %	97.73 %	0.0573 %
10	$ S_{\text{dir}} < 0.6$	99.42 %	17.60 %	98.23 %	0.0563 %
11	$\text{cog}_z < 400 \text{ m}$	98.66 %	17.36 %	98.31 %	0.0553 %
12	$\text{cog}_z > -450 \text{ m}$	97.59 %	16.95 %	97.02 %	0.0537 %

For cut 7, the hit DOMs are split into two equally sized sets in two different ways: the geometrical split is performed in a plane perpendicular to the MPE fit, while the time split evenly splits the hit DOMs in an early and a late half. This cut then discards all events for which at least one of the four SPE fits made on these four samples points further downwards than 10° from the horizon, i.e., has a zenith angle smaller than 80° . A cut at the horizon would have seem more natural, but would have cost too many neutrino events with little gain in muon suppression.

Cut 8 discards events with small numbers of direct hits n_{dir} . A direct hit is a DOM hit in a time window of $[-15, 45]$ ns around the expected arrival time of unscattered photons according to the MPE reconstruction. These hits are especially valuable for the reconstruction algorithms, and their number is a good measure of quality for the resulting fits.

Cut 9 removes events for which the distance between the direct hits (perpendicularly projected onto the MPE reconstruction) is too small. Small distances correspond to small leverage for the reconstruction algorithms and therefore larger angular uncertainties.

Cut 10 rejects events with a bad absolute *smoothness* of direct hits, i.e., events where the direct hits cluster at the beginning or end of the reconstructed track instead of being distributed evenly.

Finally, cuts 11 and 12 (“11 top” and “11 bottom” in ref. [114]) are containment cuts in the z-direction. They discard all events for which the center of gravity (cog) of the hit DOMs is too close to the top or bottom of the detector which often happens when the particle did not or only shortly enter the detector volume.

The resulting sample consists of 21943 events. In ref. [114] the purity has been estimated to be 99.85%, with a statistical error of 0.06% based on the statistics of the MC datasets, and a systematical error of 0.04% based on a 25% uncertainty in the cosmic-ray flux (at 1 TeV) according to the HKKMS2006 model for atmospheric neutrinos [32]. 13440 of those are in the TopologicalTrigger unsplit stream, and 8503 in the split stream.

An alternative event selection based on machine learning (i.e., boosted decision trees) was developed in parallel to the cuts, providing useful insights for the optimization of the cuts. In the end it was discarded because while it still reached about 10% higher passing rates of atmospheric neutrinos than the final cuts at the same purity, the cuts allowed for easier checks regarding unsimulated background (e.g., by removing each cut while keeping the others, so called $n - 1$ plots [114]), and smaller systematic uncertainties are more important for our analyses than the increased statistics.

3.2 Adaption to this Analysis

A few additional cuts were applied to the baseline event sample. Firstly, a cut was introduced to remove all events for which TruncatedEnergy could not be computed successfully in the AllDOMs method. The reason for this to happen is less than eight hit DOMs within a 60 m radius cylinder around the track¹³, which can happen for low-energy events, badly reconstructed events, or for tracks that were reconstructed to have barely entered the instrumented volume. It was an implicit cut in the search for astrophysical neutrino events that was not mentioned explicitly in ref. [114], but is also useful here as an additional quality cut. It reduces the number of events by 59 to 21884.

To reduce the influence of the astrophysical neutrino flux that dominates at very high energies, all events above $10^{4.7}$ GeV have been removed from the sample, resulting in the loss of 5 high-energy events. Similarly, events below $10^{2.1}$ GeV have been removed because in this region, the statistics of the event sample is lacking (22 events), and the sub-par performance of reconstruction algorithms that are not specialized on low energies could introduce unwanted systematic effects in this region with little gain. At this point, the sample contains 21857 events.

Finally, the event rate of coincident atmospheric-muon events that had been misreconstructed as vertically upgoing low-energy tracks had been underestimated by the CORSIKA simulation. Table 3.2 lists all CORSIKA background events that remain in the baseline sample, based on two unweighted datasets of 11.5 days lifetime each, and three weighted datasets with roughly comparable statistics (see section 2.4). Many variables have been examined, and the two quantities n_{dir} and l_{empty} are the most promising ones to further suppress background events: n_{dir} is the number of direct hits explained above, and l_{empty} is the largest distance along the track between two successive hit DOMs within a cylinder with 150 m radius around the track.

Figures 3.1 and 3.2 show these events for the split stream and for the unsplit stream, together with experimental data and Monte-Carlo predictions for neutrino-induced muon events. Aside from the total normalization that is too low because the TopologicalTrigger split decision (see section 2.3.2) is underpredicted by MC at all energies (figure 3.3), the data-MC agreement in the unsplit stream is good, with no noticeable excesses of experimental data that could be attributed to unsimulated background. In contrast, the plots for the split stream show excesses for vertically upgoing tracks and large values of l_{empty} . Those regions feature above-average numbers of coincident MC background events, but not enough to explain the excess. However, examining experimental events in these regions by eye revealed more coincident atmospheric muon events than expected from MC predictions. Consequently, events with $l_{\text{empty}} \geq 350$ m have been cut from the split stream, removing 233 events, for a new total of 21624 events and an estimated purity of over 99.9%.

¹³More precisely, the 60 m cylinder starts and ends with the first and last hit DOMs inside a hollow cylinder 10...80 m from the reconstructed track, i.e., leading and trailing hits less than 10 m from the track are not counted.

Table 3.2: Simulated atmospheric muon neutrino background events before final cuts

<i>n</i>	TT stream	<i>n</i> _{dir}	<i>l</i> _{empty} /m	<i>E</i> _{muEx} /GeV	cos (θ)	weight	MC type	survives
1	split	8	472.1	138.4	0.918	2.81	unw. coinc.	no
2	split	8	757.0	276.3	0.899	2.81	unw. coinc.	no
3	split	9	385.5	672.7	0.513	2.81	unw. coinc.	no
4	unsplit	7	94.5	2943.1	0.091	0.04	w. single	yes
5	unsplit	8	94.9	6508.4	0.003	5.39	w. single	yes
6	unsplit	7	123.1	544.9	0.041	0.20	w. single	yes
7	unsplit	8	121.7	4148.6	0.019	0.00	w. single	yes
8	unsplit	9	112.0	1391.8	0.008	0.16	w. single	yes
9	unsplit	9	164.6	2335.6	0.027	0.91	w. coinc.	no*
10	split	7	238.7	8459.4	0.485	1.23	w. coinc.	yes
11	unsplit	8	200.5	1892.6	0.591	0.04	w. coinc.	yes
12	split	7	438.4	2302.8	0.957	0.01	w. coinc.	no
13	unsplit	9	91.2	10127.0	0.232	0.03	w. single	yes
14	split	7	770.5	194.8	0.974	1.81	w. coinc.	no
15	split	7	502.7	305.5	0.927	0.05	w. coinc.	no
16	split	8	101.4	1071.1	0.153	0.32	w. single	yes

* cut because TruncatedEnergy reconstruction failed

For comparison, final plots of the reconstructed energy and zenith angle after these cuts are shown in figure 7.7.

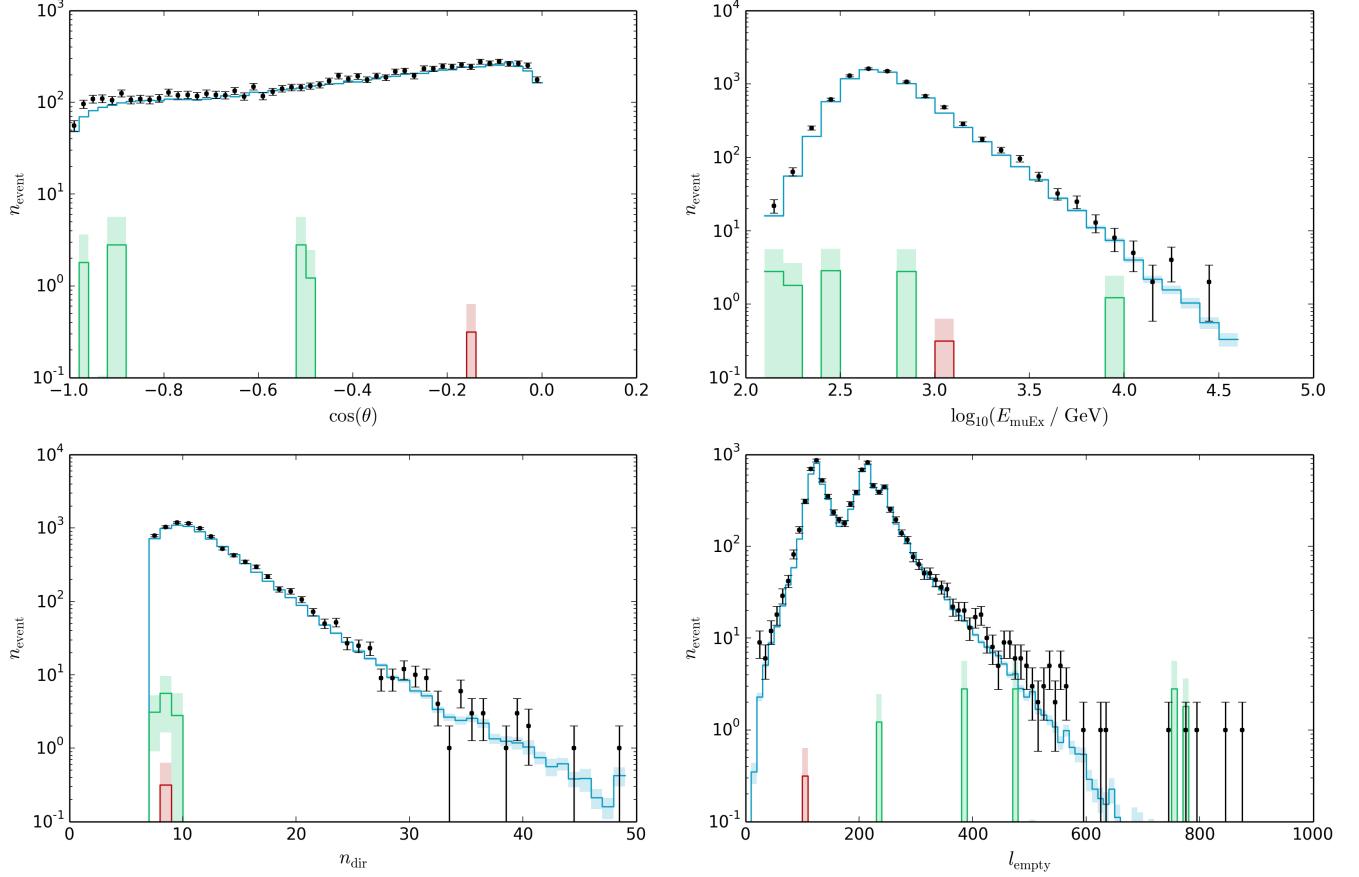


Figure 3.1: Distributions of the atmospheric muon background in the split stream before the final cuts; shown are zenith angle (**top left**), reconstructed energy (**top right**), number of direct hits (**bottom left**) and l_{empty} (**bottom right**), see text. Experimental data in black, MC neutrino predictions in blue, single-muon background in red and coincident-muon background in green. Note that the neutrino predictions have been fitted to the whole sample, not to the split stream alone.

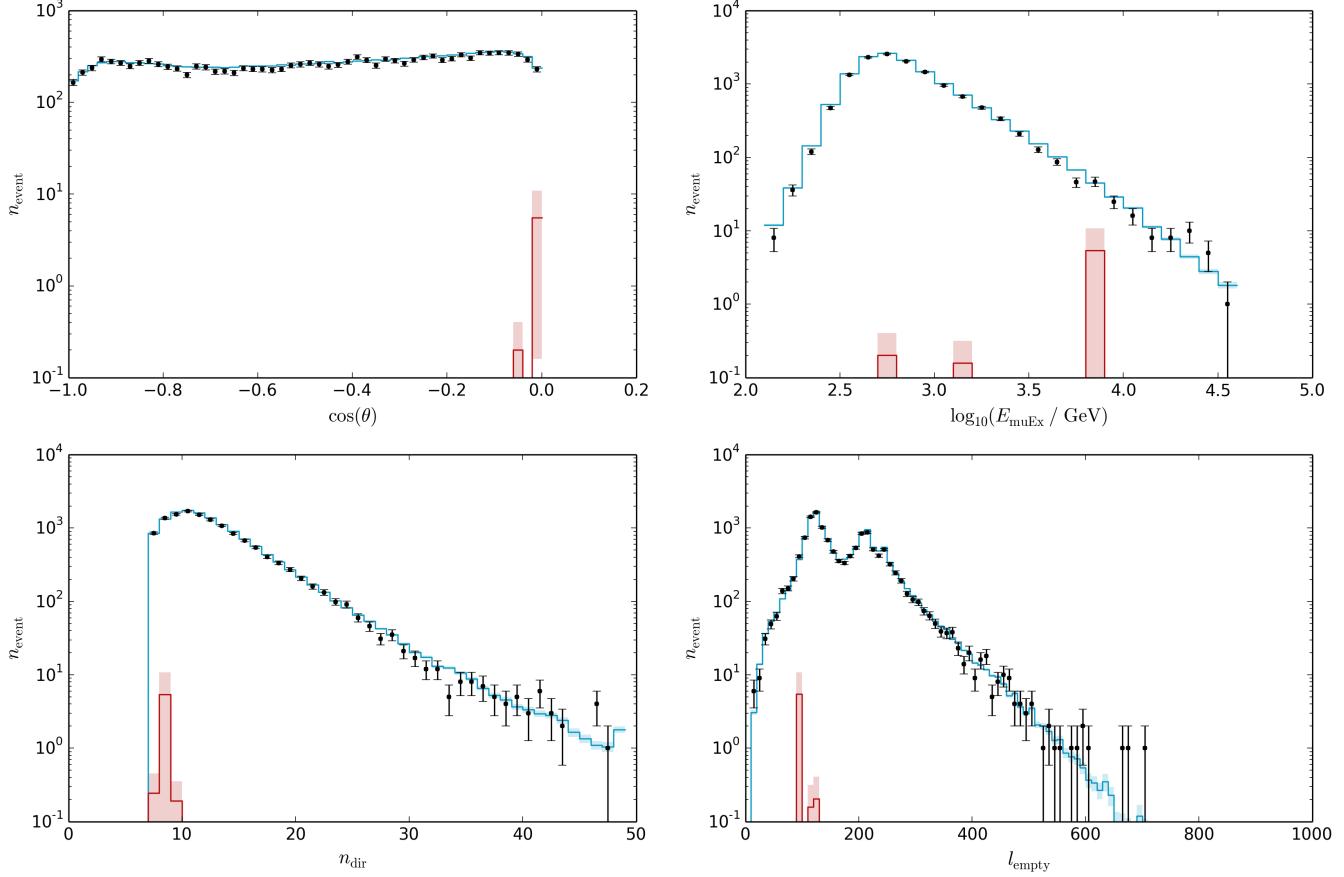


Figure 3.2: Distributions of the atmospheric muon background in the unsplit stream before the final cuts, analog to figure 3.1.

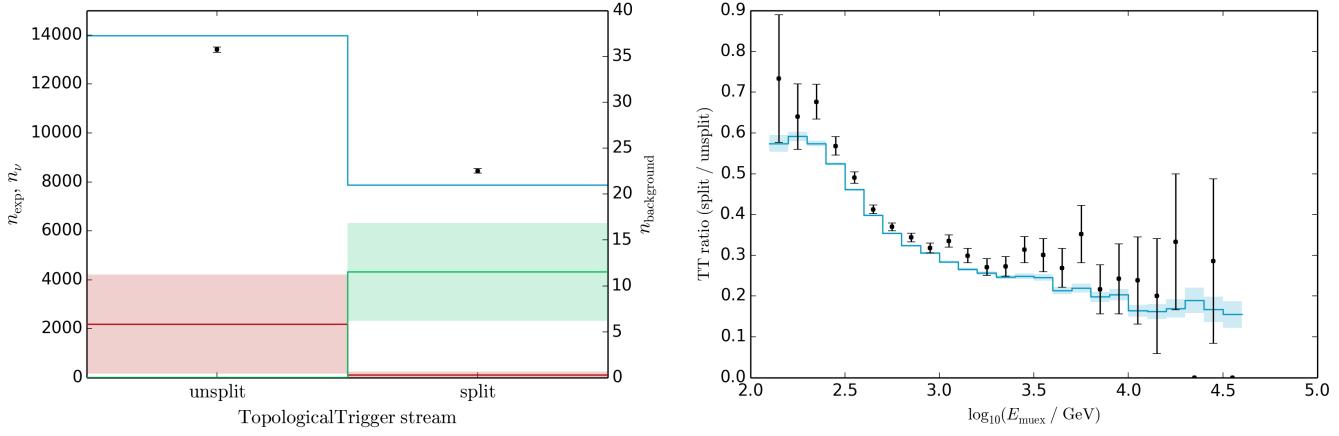


Figure 3.3: Distributions of the TopologicalTrigger decision (section 2.3.2) before the final cuts, analog to figure 3.1 (left; note the second y-axis for background events, shown in red and green), and ratio versus energy (right, without background).

Chapter 4

Signatures of Sterile Neutrinos in IceCube

A less detailed discussion of signatures has previously been presented at conferences and has accordingly been published in proceedings by the IceCube Collaboration (ref. [120] and [121], also published in ref. [122] and [123], respectively). The publications have been written as corresponding author. Everything here is work of the author unless specifically noted.

A similar study of signatures and parameter dependencies has been conducted in an associated bachelor's thesis (ref. [124]), and is appropriately cited.

This chapter is intended to give a short overview over the signatures that sterile neutrinos can produce in IceCube to motivate the analysis strategy that is outlined in the next chapter. The first two sections describe theoretical signatures, the third section accounts for the experimental resolution.

4.1 Four-Flavor Models

4.1.1 $\theta_{14} = \theta_{34} = 0$

This is the model that is investigated in this thesis. Only one of the mixing angles that get introduced with an additional flavor is different from zero, which is θ_{24} . As a consequence, in the last column of the mixing matrix, U_{14} and U_{34} are still zero, and only muon neutrinos can oscillate to the sterile flavor.

The two free parameters involved in this model compared to a standard three-flavor model are Δm_{42}^2 and θ_{24} . As pointed out in section 1.4, the sine-pattern vacuum oscillations of sterile neutrinos are expected to be very weak according to existing limits, so the most important part of the signature is the resonant appearance caused by matter effects.

The (squared-)mass difference mostly controls at which neutrino energy the resonance takes place. Individual rows in figure 4.1 show how a signature changes for different values of Δm_{42}^2 . Despite the entanglement between Δm_{42}^2 and θ_{24} due to strong matter effects as explained in section 1.3.2, the impact of different values of Δm_{42}^2 can easily be estimated as it is just a shift of the signature along the logarithmic energy axis; the oscillation could be described well in a two-flavor scheme. This only changes when the disappearance takes place at energies low enough for it to merge with the standard muon-tau disappearance [124], but this happens below the energy range relevant for this analysis.

The effect of the mixing angle θ_{24} is slightly more complicated, as it impacts both the amplitude of the oscillation and the zenith angle at which the resonance occurs. The columns of figure 4.1 show how the signature changes when θ_{24} is varied between 2.5° and 11.7° . With increasing mixing angle, the signature becomes more pronounced outside the resonance region, and the resonance becomes stronger up to total disappearance and then wanders towards the horizon (because of the lower matter density along more shallow trajectories). Figure 4.2 displays this behavior for vertically upgoing events; note that after the disappearance has become maximal around 7.25° , it becomes weaker again. At large mixing angles, the signatures can become very strong. Figure 4.3 shows signatures for 22.8° , 33.9° , and 45° . The resonance is still visible and heavily distorts the first oscillation minimum, but the disappearance pattern is strong for all energies below this first minimum.

4.1.2 $\theta_{34} = \theta_{24}$, $\theta_{14} = 0$

Similar to $\theta_{14} = \theta_{34} = 0$, $\theta_{14} = 0$ and $\theta_{34} = \theta_{24}$ is another relatively simple model often encountered in theoretical publications. It features the same number of free parameters, but also involves tau neutrinos, with $U_{34} = \cos(\theta_{24})U_{24}$ (i.e., $U_{34} \approx U_{24}$ for small θ_{24}).

The main difference to the model discussed in the previous section is that $\bar{\nu}_\mu$ now oscillate mostly into $\bar{\nu}_\tau$ at the high-energy resonance (figure 4.4); this transition does not depend on the value of θ_{23} [124]. Outside the resonance, $\langle\bar{\nu}\rangle_\mu$ still oscillate mostly into $\langle\bar{\nu}\rangle_s$ with little tau contribution. Figure 4.5 shows the differences between those two models for the muon (anti-)neutrino disappearance: the shape of the resonance changes slightly, and the additional mixing causes the conventional low-energy muon-tau oscillations to shift apart for neutrinos and antineutrinos. Overall, the pattern remains similar: The differences in the low-energy regime are diametrical between particles and antiparticles and partially cancel out if the detector can not discriminate between those; they also

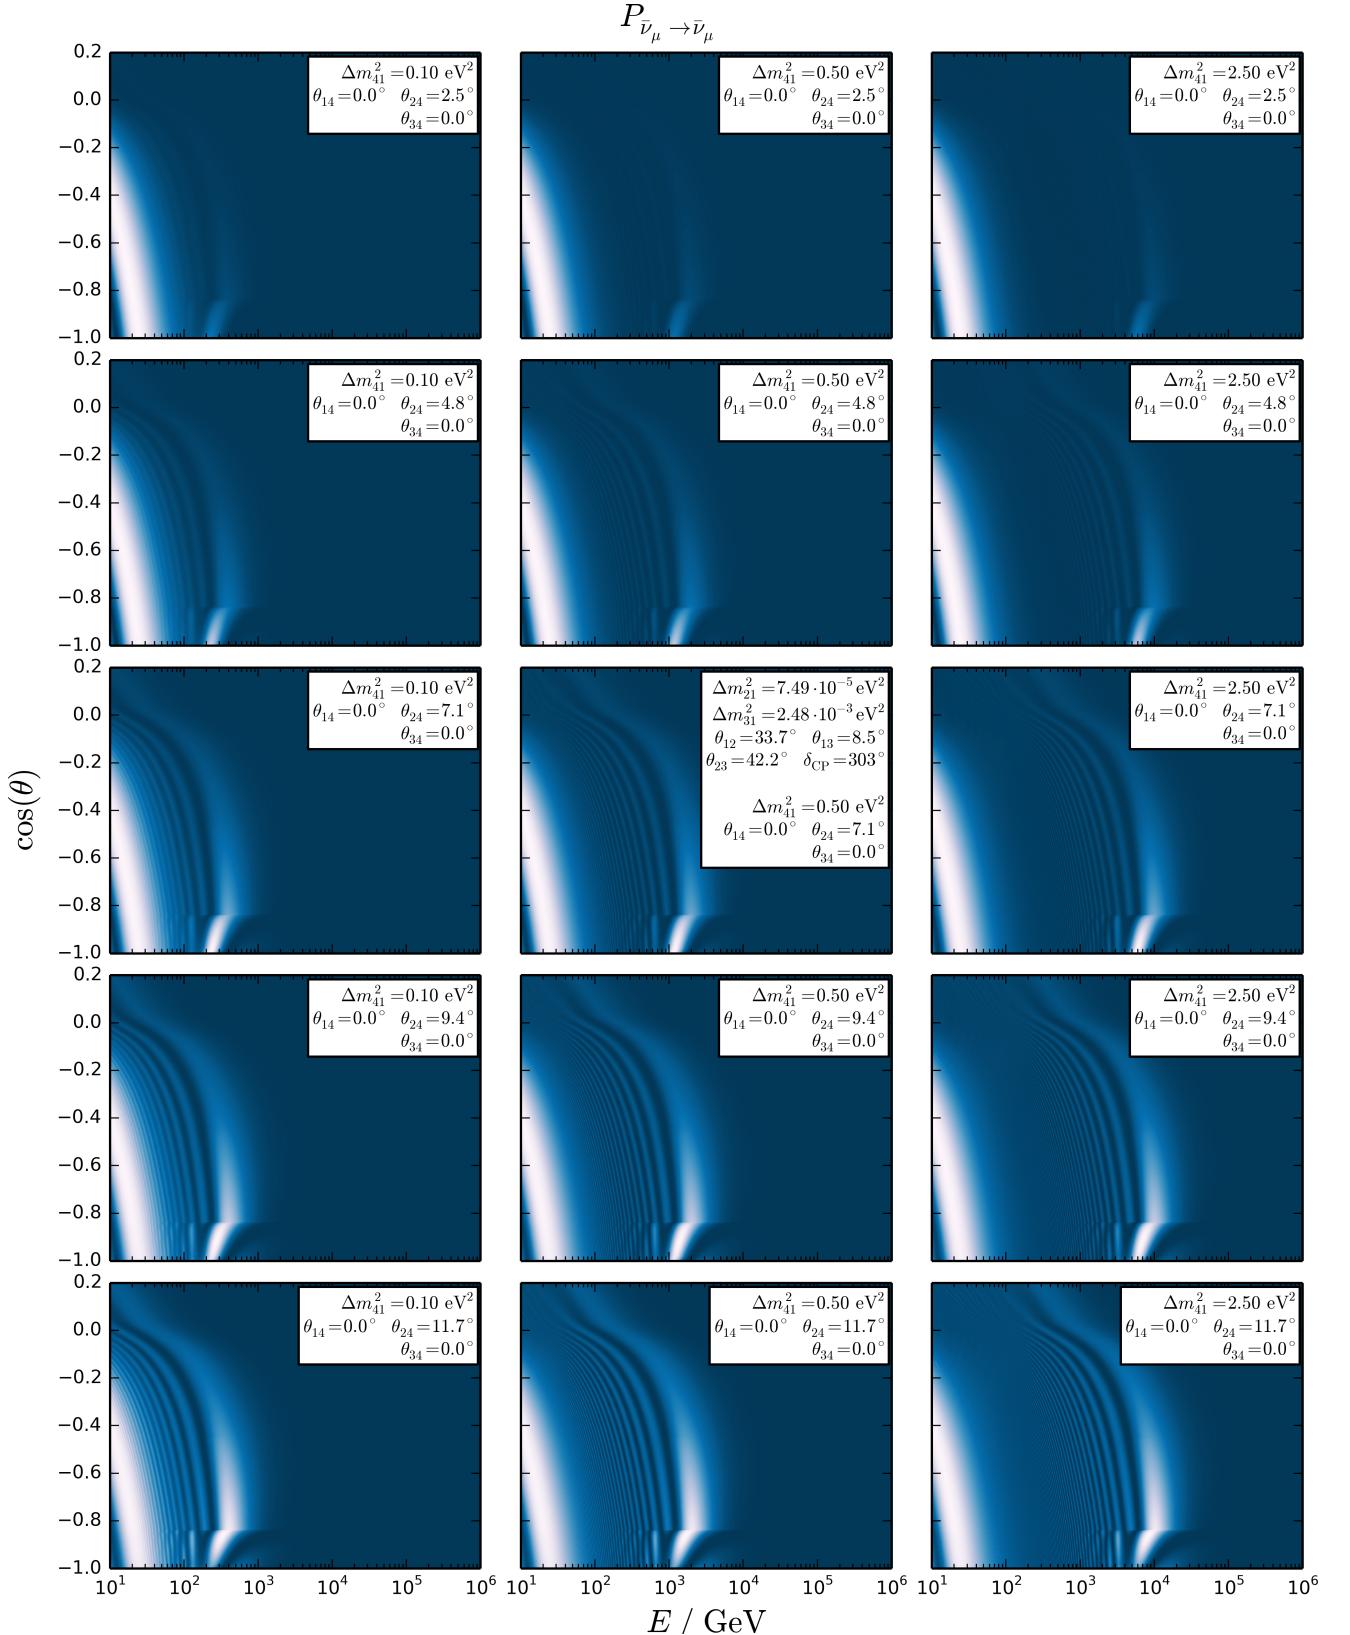


Figure 4.1: Change of muon antineutrino disappearance signatures for varying values of Δm_{42}^2 and θ_{24} . The center plot was shown already in figure 1.22. Rows feature varied Δm_{42}^2 ($\approx \Delta m_{41}^2$), columns have varied θ_{24} , the remaining parameters are fixed.

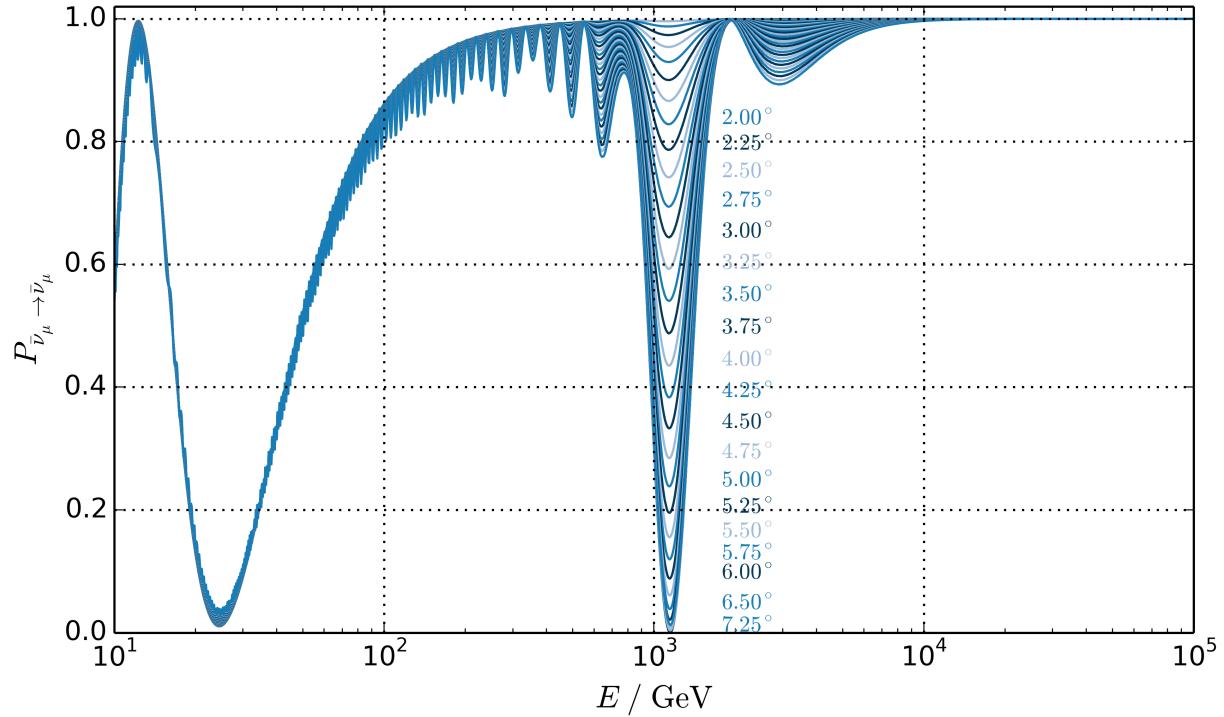


Figure 4.2: Muon antineutrino survival probabilities for vertically upgoing particles for $\theta_{24} = 0.25\ldots 8^\circ$ in steps of 0.25° at $\Delta m_{42}^2 = 0.5 \text{ eV}^2$; all other parameters as in figure 4.1. Labels are printed on the level of the local minimum; not all lines are labeled.

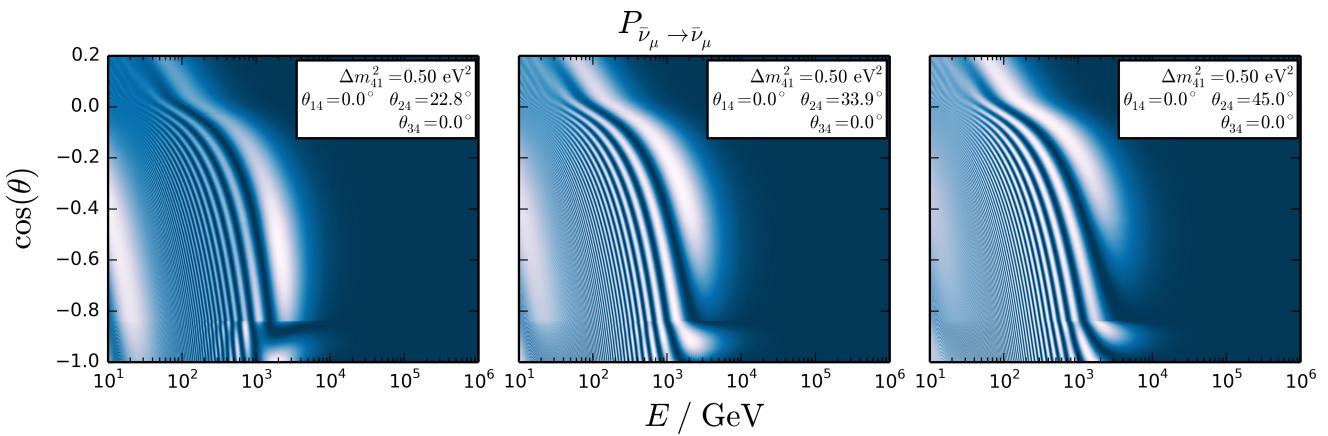


Figure 4.3: Muon antineutrino disappearance signatures for large values of θ_{24} , analog to figure 4.1.

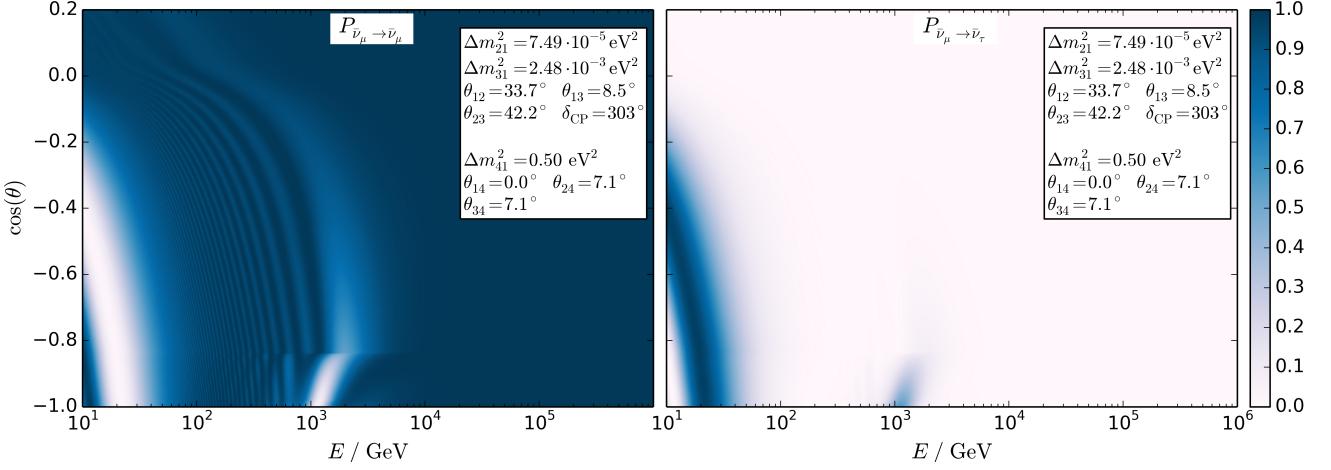


Figure 4.4: **Left:** Muon antineutrino disappearance signature for the $\theta_{34} = \theta_{24}$ scheme, analog to the center plot of figure 4.1.

Right: Corresponding tau antineutrino appearance probability.

mostly lie below 100 GeV, so they are not relevant for this analysis, but can be used as a complimentary way to explore sterile neutrinos with IceCube [124][125]. The change at the resonance seems noticeable for vertical tracks, but strongly washes out when mediocre energy resolution is taken into account (section 4.3).

Because of the similarity of the models in the context of high-energy muon disappearance analyses, it was decided not to spend CPU time into testing this model. Limits on θ_{24} for $\theta_{34} = \theta_{24}$ could be derived from those of $\theta_{34} = 0$, and are expected to be fairly similar.

4.1.3 $\theta_{34} \neq \theta_{24}$, $\theta_{14} = 0$

Non-zero values for θ_{34} different from those for θ_{24} are a generalization of the model from section 4.1.2 and introduce an additional parameter. For small values of θ_{34} , the qualitative description of the previous section holds: there is a strong $\bar{\nu}_\tau$ appearance at the high-energy resonance, and with increasing mixing angle, the low-energy oscillations shift apart between ν_μ and $\bar{\nu}_\mu$. While the extent of the $\bar{\nu}_\tau$ appearance increases with θ_{24} and θ_{34} , the $\bar{\nu}_\mu$ disappearance at the resonance gets smaller for increasing values of θ_{34} for a given value of θ_{24} , see figure 4.6. However, the figure also shows that the formerly sharp resonance becomes wider towards lower energies. Assuming a basic $E^{-3.7}$ spectrum for the atmospheric neutrino flux, figure 4.7 shows the relative development of the total vertical muon (anti-)neutrino flux above 100 GeV: the $\bar{\nu}_\mu$ disappearance decreases, but the ν_μ flux is reduced significantly for most values of θ_{34} , which is a non-resonant phenomenon.

Overall, the additional free parameter θ_{34} allows much variation of the oscillation behavior both at low and high energies if large angles are permitted. The effects are

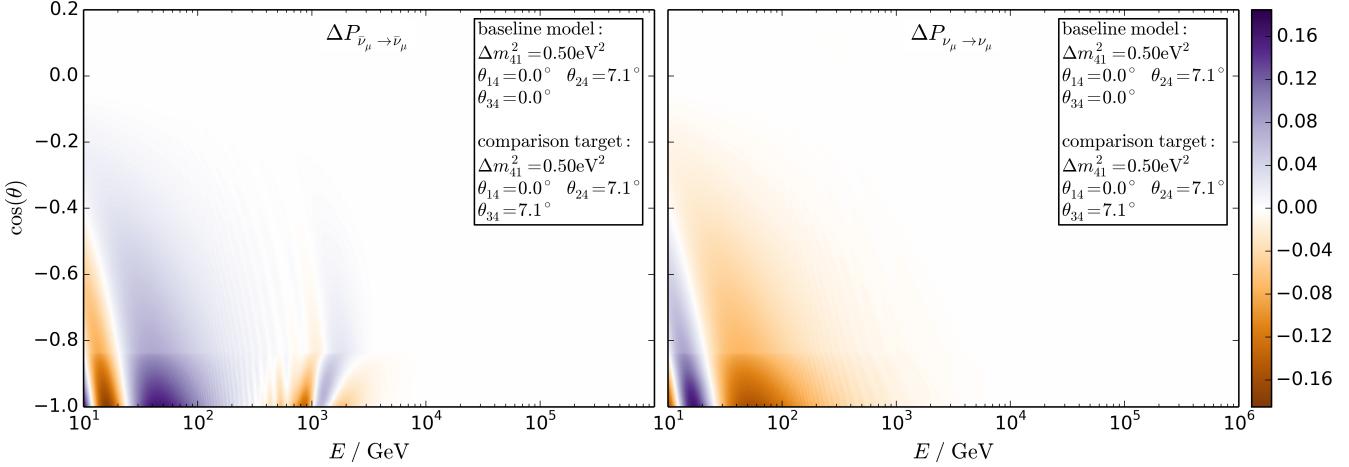


Figure 4.5: Differences in the muon antineutrino (**left**) and muon neutrino (**right**) disappearance probabilities for the model of figure 4.4 compared to the central model of figure 4.1, i.e., $P_{\text{baseline model}} - P_{\text{comparison target}}$.

non-linear and difficult to parameterize; for example, using the effective mixing angles and squared-mass differences introduced in section 1.3.3, one obtains [51]

$$\begin{aligned}
 P_{(\bar{\nu})\mu \rightarrow (\bar{\nu})\tau} &= \sin^2(2\theta_{24}^{\text{mat}}) \sin^2(\theta_{34}^{\text{mat}}) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right) \\
 P_{(\bar{\nu})\mu \rightarrow (\bar{\nu})s} &= \sin^2(2\theta_{24}^{\text{mat}}) \cos^2(\theta_{34}^{\text{mat}}) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right) \\
 \Rightarrow P_{(\bar{\nu})\mu \rightarrow (\bar{\nu})\mu} &= 1 - \sin^2(2\theta_{24}^{\text{mat}}) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right),
 \end{aligned} \tag{4.1}$$

so according to this model, the $(\bar{\nu})_\mu$ survival probability does neither explicitly nor implicitly depend on θ_{34} , as the effective parameters are defined using two-flavor approximations. This obviously contradicts the more accurate numerical results shown in figures 4.5 to 4.7.

Very large values of θ_{34} can be excluded because they would induce strong effects at low energies even for $\theta_{24} = 0$. According to the global-fit constraints on the unitarity of U (1.22) [58], $\theta_{34} \leq 18^\circ$ at 90% C.L. for $\theta_{24} = 0$. This limit deteriorates for larger values of θ_{24} , e.g., down to $\theta_{34} \leq 24^\circ$ for $\theta_{24} = 7.1^\circ$, $\theta_{34} \leq 25.5^\circ$ for $\theta_{24} = 11.7^\circ$, and $\theta_{34} \leq 27^\circ$ for $\theta_{24} = 15.2^\circ$, but beyond that, θ_{24} can not be increased further because the now limiting matrix elements of $|UU^\dagger|$ approximately scale with $\sin(2\theta_{24}) \sin(\theta_{34})$ [51].

Because of the highly increased complexity and CPU requirements, it was decided to not use θ_{34} as an additional free parameter in this initial search for sterile neutrinos with IceCube. Deriving limits on these models should still be possible with results for $\theta_{34} = 0$ and $\theta_{34} = \theta_{24}$, but will not be a topic covered in this thesis.

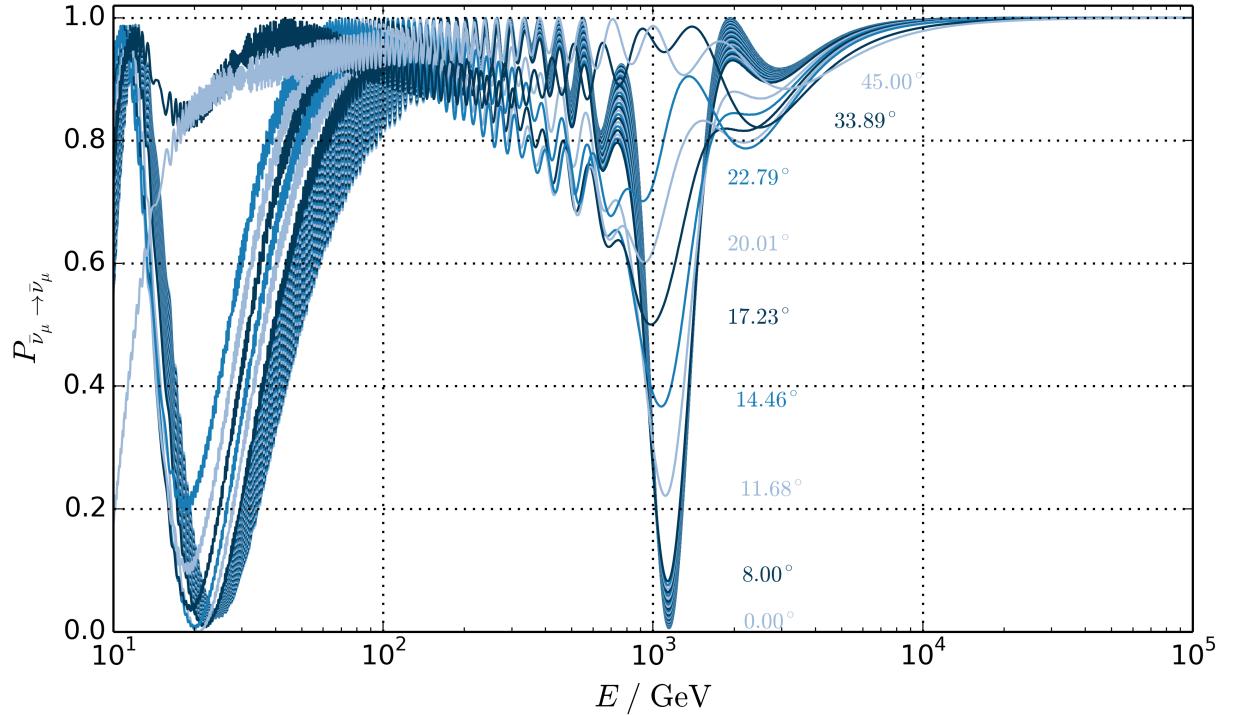


Figure 4.6: Muon antineutrino survival probabilities for vertically upgoing particles for various values of θ_{34} ; all other parameters as in figure 4.1 (esp. $\theta_{24} = 7.1^\circ$). Labels are printed close to the corresponding line; not all lines are labeled. The family of curves that are close together corresponds to the variation of θ_{24} in figure 4.2 (0.0° ... 8.0°).

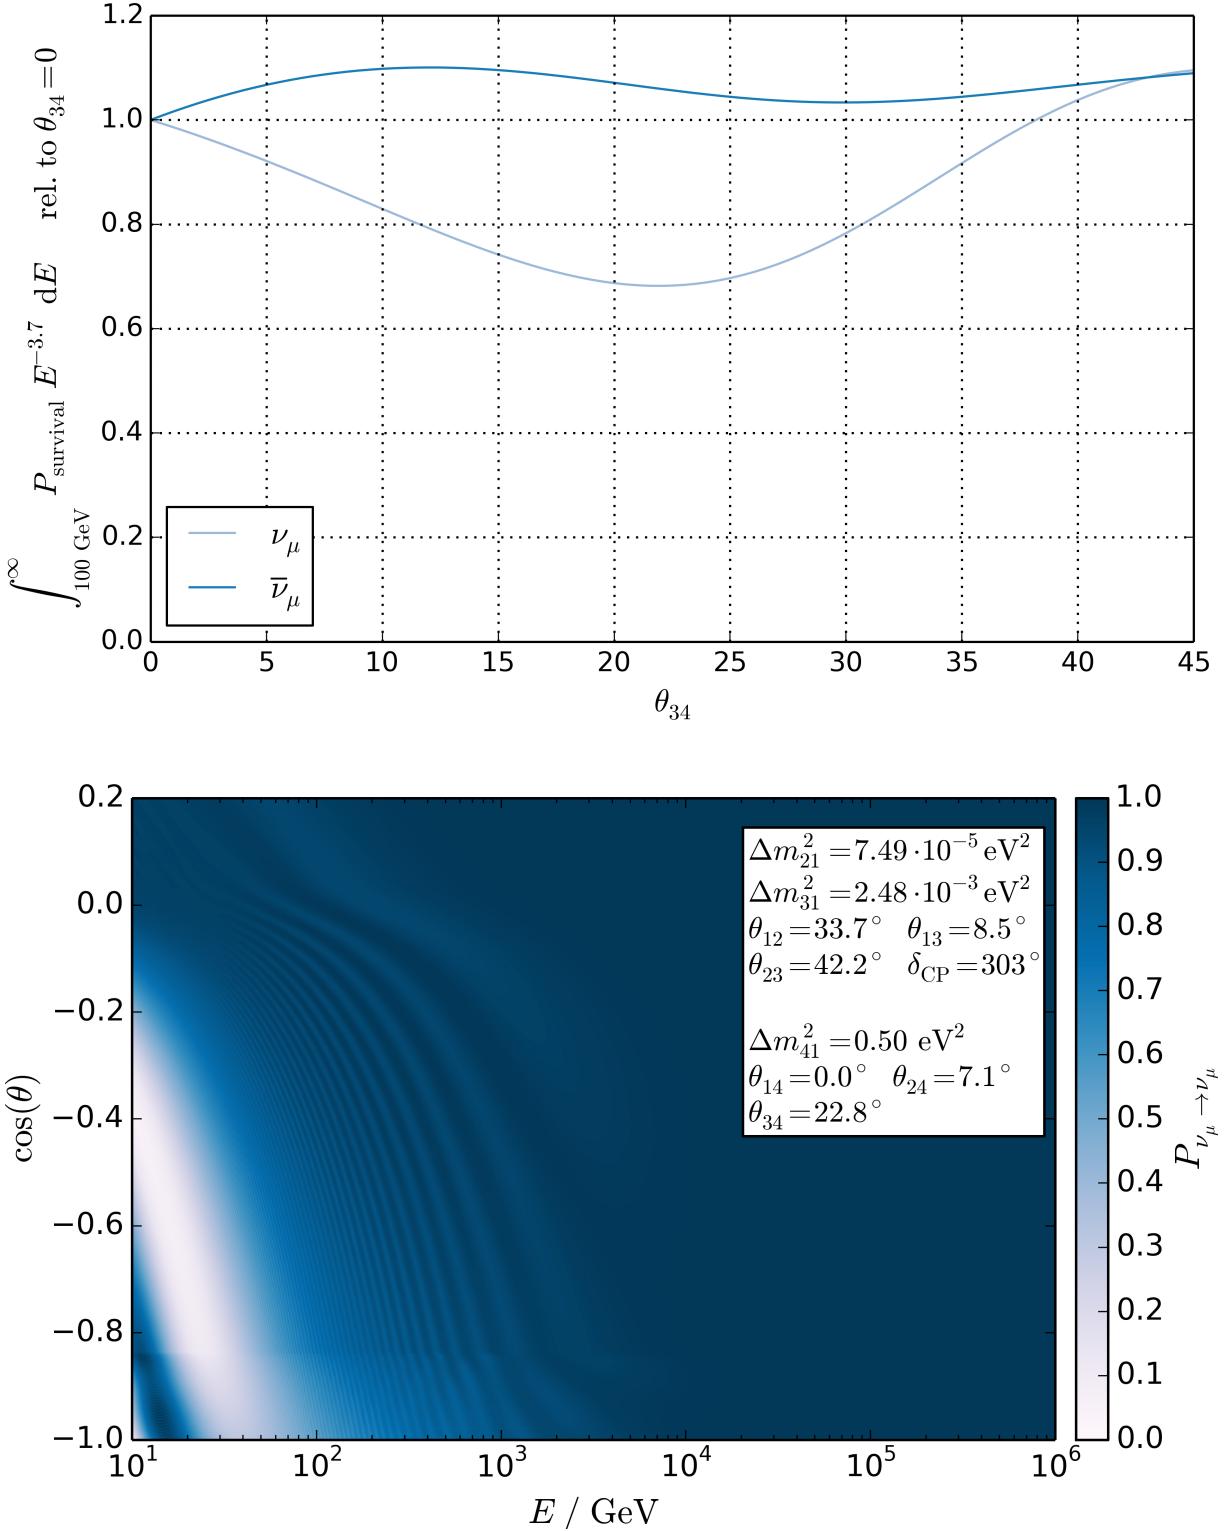


Figure 4.7: **Top:** Integral muon (anti-)neutrino survival probabilities for vertically upgoing particles for various values of θ_{34} above 100 GeV, assuming an $E^{-3.7}$ spectrum; all other parameters as in figures 4.1 and 4.3 (esp. $\theta_{24} = 7.1^\circ$).

Bottom: Corresponding muon neutrino disappearance signature at $\theta_{34} = 22.8^\circ$.

4.1.4 $\theta_{14} \neq 0$

Allowing non-zero values for θ_{14} introduces yet another degree of freedom to the $3 + 1$ model compared to the previous sections. Luckily, θ_{14} has relatively weak influence on the $(\bar{\nu})_\mu$ disappearance probabilities accessible to IceCube. The PMNS matrix elements U_{24} , U_{34} , and U_{44} all scale with $\cos(\theta_{14})$. Therefore, analog to (4.1) [51],

$$\begin{aligned} P_{(\bar{\nu})_\mu \rightarrow (\bar{\nu})_e} &= \sin^2(2\theta_{14}^{\text{mat}}) \sin^2(\theta_{24}^{\text{mat}}) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right) \\ P_{(\bar{\nu})_\mu \rightarrow (\bar{\nu})_\tau} &= \cos^2(\theta_{14}^{\text{mat}}) \sin^2(2\theta_{24}^{\text{mat}}) \sin^2(\theta_{34}^{\text{mat}}) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right) \\ P_{(\bar{\nu})_\mu \rightarrow (\bar{\nu})_s} &= \cos^2(\theta_{14}^{\text{mat}}) \sin^2(2\theta_{24}^{\text{mat}}) \cos^2(\theta_{34}^{\text{mat}}) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right) \\ \Rightarrow P_{(\bar{\nu})_\mu \rightarrow (\bar{\nu})_\mu} &= 1 - \left(\sin^2(2\theta_{14}^{\text{mat}}) \sin^2(\theta_{24}^{\text{mat}}) + \cos^2(\theta_{14}^{\text{mat}}) \sin^2(2\theta_{24}^{\text{mat}}) \right) \sin^2\left(\frac{\Delta m_{41}^{2, \text{mat}} L}{4E_\nu}\right). \end{aligned}$$

The model still suffers from the oversimplified handling of matter effects as mentioned above, but it shows why small values of θ_{14} only have a small influence on the $(\bar{\nu})_\mu$ survival.

For θ_{14} , proper handling of matter effects does not falsify this statement: Figure 4.8 shows the change in the probabilities if the value of θ_{14} is set to 7.1° . By far the largest effect is a skewing of the resonance towards slightly lower energies. The effect is very small with an amplitude of only 2.5%, and would be further washed out by IceCube's limited energy resolution.

Figure 4.9 shows the same for $\theta_{14} = 22.8^\circ$. This value is already slightly larger than what is allowed for θ_{14} by the unitarity conditions of (1.22) for all possible combinations of sterile mixing angles, but the effects are still relatively small and could be described well by $\theta_{14} = \theta_{34} = 0$ with a slightly shifted Δm_{42}^2 .

Finally, figure 4.10 shows how the $\theta_{34} = \theta_{24}$ model is affected by a non-zero θ_{14} . The differences are even smaller than in the $\theta_{34} = 0$ case, namely at the 1.7% scale.

Overall, θ_{14} is by far the least important of the sterile mixing angles. It is very slightly correlated with Δm_{42}^2 , but it has relatively small effects even at the largest allowed value. To reduce complexity and CPU requirements, it is set to zero for this analysis, and unless future global limits severely constrain θ_{34} , θ_{14} should receive the lowest priority for future IceCube analyses.

4.2 Models with More Flavors

The theoretical background for more than one additional flavor of light sterile neutrinos has been explained already in section 1.4 and does not differ significantly from the situation

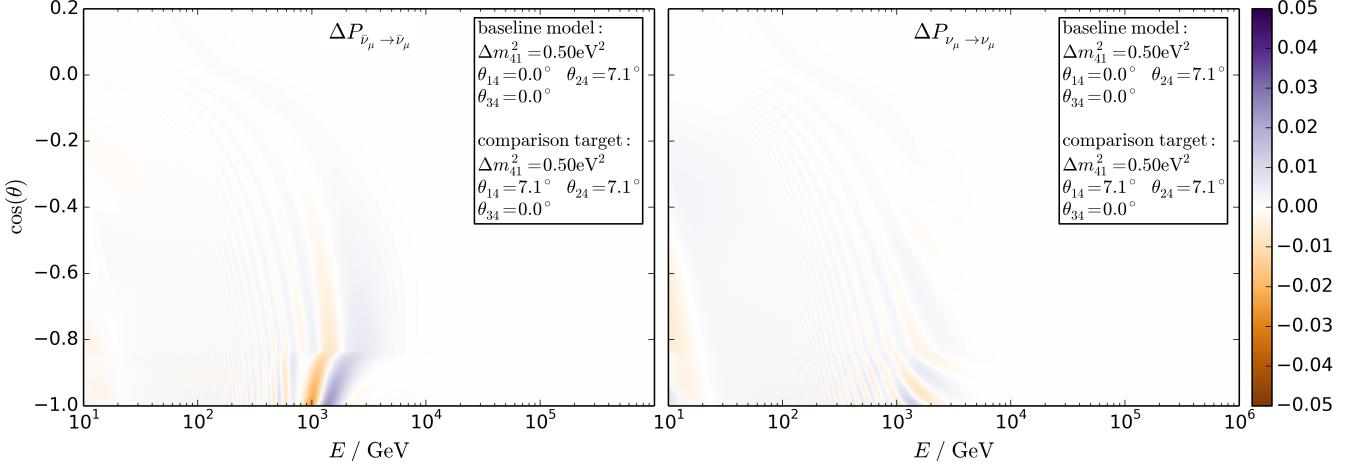


Figure 4.8: Differences in the muon antineutrino (**left**) and muon neutrino (**right**) disappearance probabilities for $\theta_{14} = \theta_{24} = 7.1^\circ$ compared to the central model of figure 4.1, i.e., $\theta_{14} = 0$ ($P_{\text{baseline model}} - P_{\text{comparison target}}$).

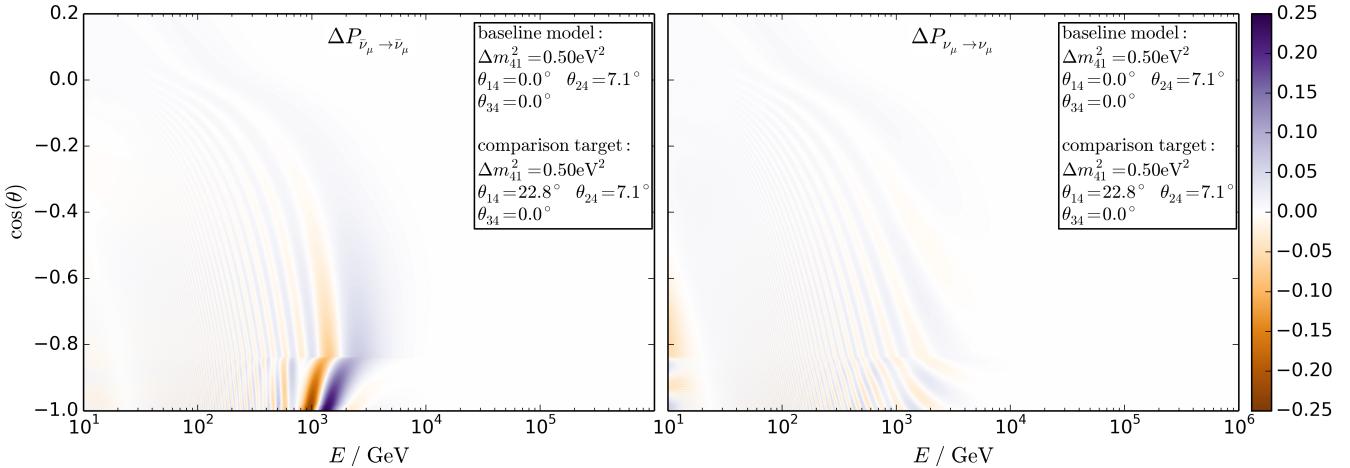


Figure 4.9: Differences in the muon antineutrino (**left**) and muon neutrino (**right**) disappearance probabilities analog to figure 4.8, but for a large value of θ_{14} .

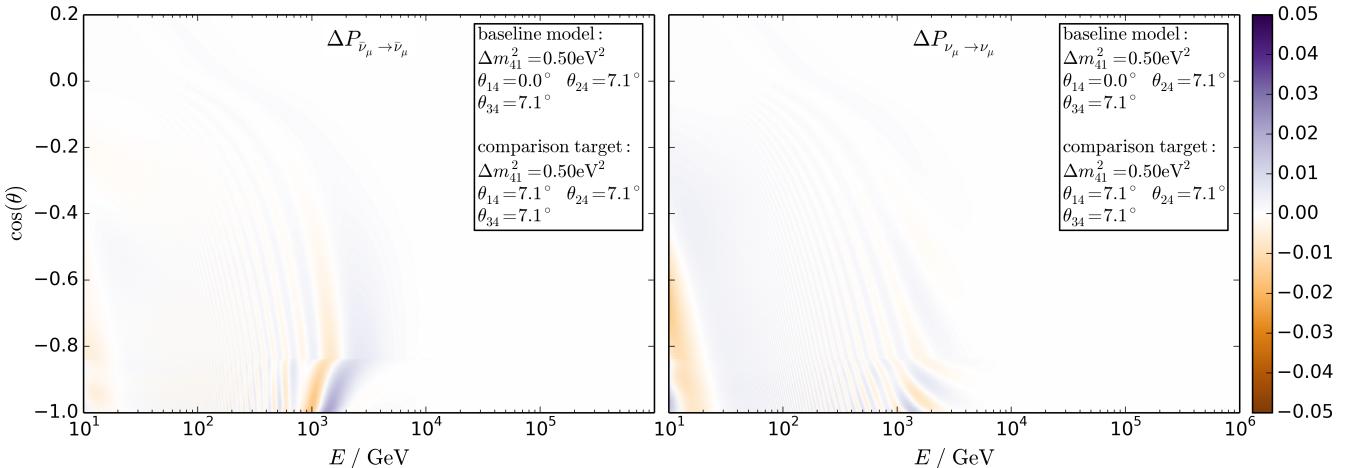


Figure 4.10: Differences in the muon antineutrino (**left**) and muon neutrino (**right**) disappearance probabilities analog to figure 4.8, but comparing zero and non-zero θ_{14} for the $\theta_{34} = \theta_{24} = 7.1^\circ$ baseline model.

described above. Most importantly, every flavor of light sterile neutrinos can have its own matter resonance depending on mass and mixing angles. Depending on the mass hierarchy of the additional flavors, these multiple resonances can be distributed between the neutrino and antineutrino channels. Resonances in different channels (e.g., in the $1 + 3 + 1$ case) barely influence each other and can be approximated well by the superposition of models with fewer flavors, while resonances in the same channel show stronger non-linear effects the closer the resonances are (figure 4.11). If for two mass states in the same channel $\Delta m_{j2}^2 / \Delta m_{i2}^2 \lesssim \frac{3}{2}$, their resonances fall together and cannot be resolved [124].

Models with more than one flavor of light sterile neutrinos are not covered explicitly in this analysis, even though they are favored compared to $3 + 1$ because they could more easily explain the measurements of LSND and MiniBooNE [57]. The main reasoning for this is that this analysis (as described in the next chapter) aims to exclude regions in the sterile-neutrino parameter space, and is able to do so when fitting $3 + 1$ templates to $3 + n$ ($n > 1$) data: Even with the non-linear effects mentioned above, $3 + 1$ templates with their resonances at the positions of the $3 + n$ resonances (or for very close masses, even in between) match the data much better than the no-sterile template as illustrated in figure 4.11. The data-MC agreement would be relatively bad in these cases, but it would still be better than the agreement for the no-sterile template. Therefore, the analysis has sensitivity to $3 + n$ models without having to use $3 + n$ templates, and is most importantly not at risk to produce wrong limits.

4.3 Impact of Detector Resolution

4.3.1 Angular and Energy Resolutions

The detector resolution plays an important role in this analysis. IceCube's median zenith angle resolution for muon neutrino tracks lies below 1° for all energies in this sample (figure 4.12). Even a relatively small deviation between reconstructed and true zenith angles can lead to large deviations in the calculated path lengths of atmospheric neutrinos near the horizon (above 10% for $|\cos(\theta) < 0.2|$). The impact of the angular resolution on the sensitivity of the analysis is nonetheless small because the sharp features that need to be resolved are the matter resonances, which lie far away from the horizon for interesting models (i.e., for small mixing angles). Also, the impact of the finite angular resolution gets overshadowed by the energy resolution.

Figure 4.13 shows the reconstructed energy values (using the algorithm muEx) versus the true neutrino energy values, as well as the spread of the reconstructed values for every range of true values. The reconstructed energy correlates decently with the true energy (with a correlation coefficient of $\rho = 0.452$). The absolute values can be far off; especially at high energies, reconstructions systematically underestimate the neutrino

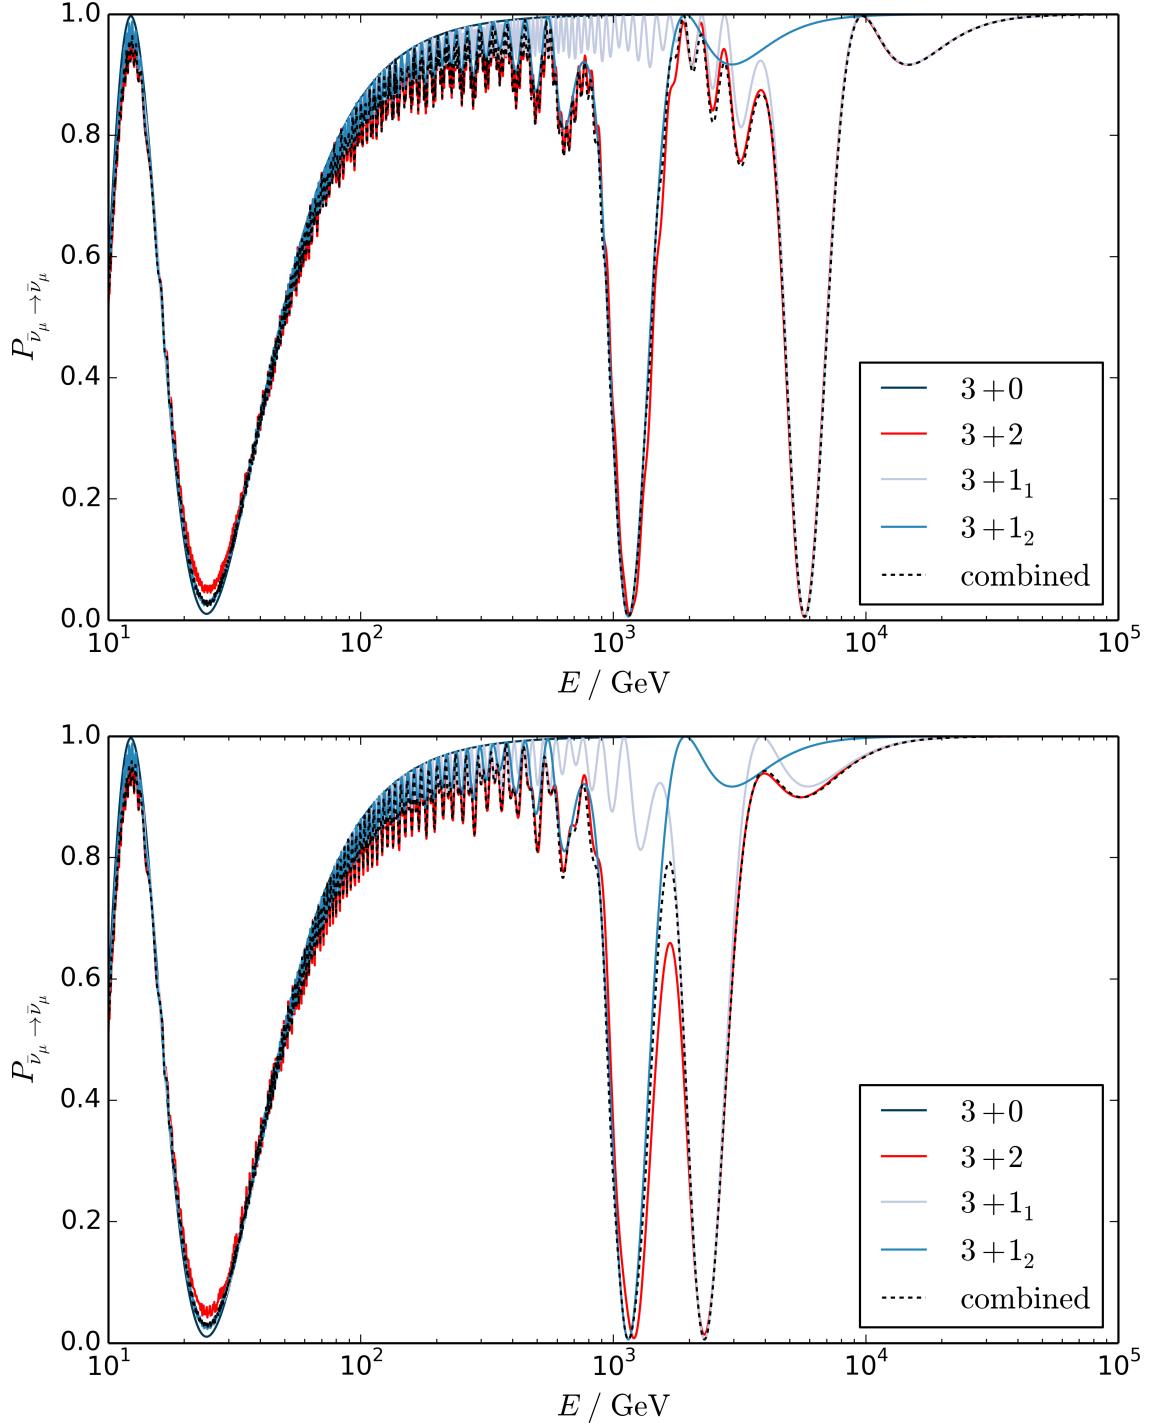


Figure 4.11: Muon antineutrino survival probabilities for vertically upgoing tracks for a $3+2$ model compared to those for the corresponding $3+1$ models and to the combination of those $3+1$ models:

$$P_{\text{combination}} = P_{3+0} - ((P_{3+0} - P_{3+1_1}) + P_{3+1_1}(P_{3+0} - P_{3+1_2}))$$

$$\Delta m_{42}^2 = 0.5 \text{ eV}^2, \quad \theta_{24} = \theta_{25} = 7.1^\circ$$

Top: $\Delta m_{52}^2 = 2.5 \text{ eV}^2$; the combination reproduces the $3+2$ probabilities well
Bottom: $\Delta m_{52}^2 = 1.0 \text{ eV}^2$; non-linear behavior \Rightarrow more deviation

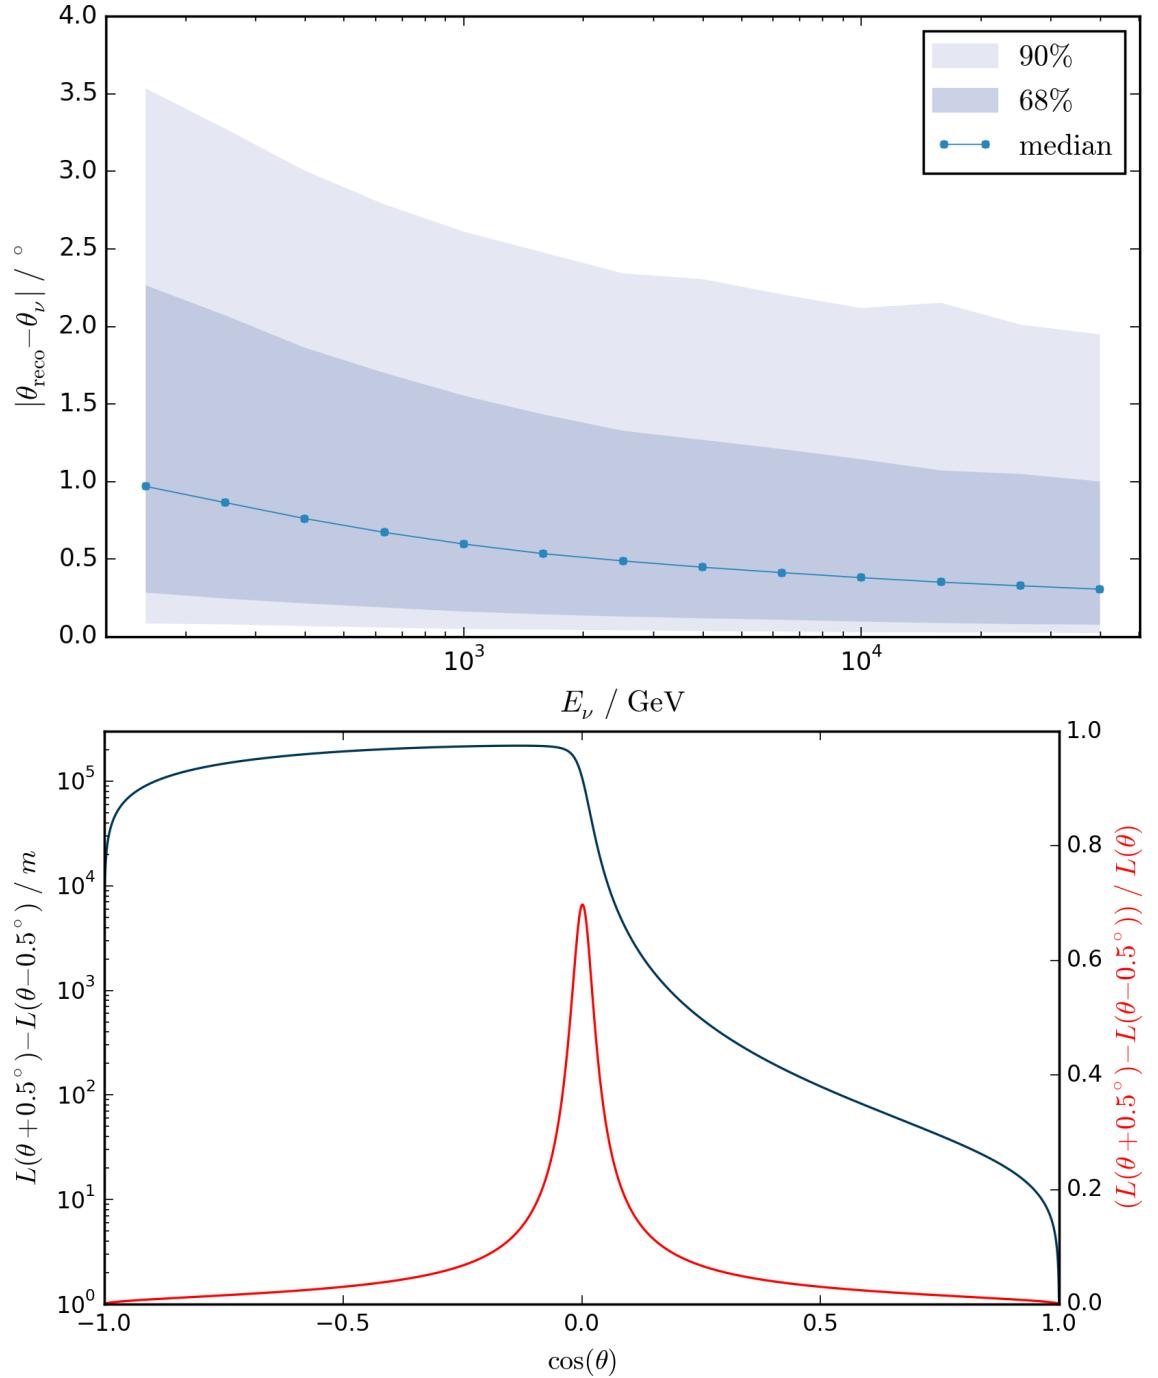


Figure 4.12: IceCube angular resolution for muon neutrino track reconstruction as a function of energy (**top**), and absolute and relative changes of the baseline L of atmospheric neutrinos due to a 1° zenith angle misreconstruction (**bottom**).

energies of events that start outside of the detector, because the induced muons have lost an unknown part of their energy already when they reach the detector. It would have been possible to calibrate the reconstruction results with a linear function to produce values that lie in a more accurate energy range, but this would not increase the amount of information yielded by the algorithm. As this analysis employs forward-folding (see chapter 5), it does not rely on the absolute energy values, so no energy calibration was done to maintain easy comparability to other samples.

The large spread of reconstructed values per true value that is visible in figure 4.13 is similarly caused by neutrinos depositing an unknown ratio of their initial energy outside of the detector. The strongest effect again comes from interactions far outside of the detector volume, but also the energy carried away by muons leaving the volume as well as the randomness of hard energy losses (section 2.1.3) complicate calorimetric reconstructions. Various algorithmic improvements have been made to compensate for these effects by not using purely calorimetric methods, but the spread remains large and much of it is intrinsic.

4.3.2 Effect on the Signatures

Since the analysis is conducted on two-dimensional histograms in reconstructed energy and zenith angle, the two finite resolutions shown above act as smoothing in their respective dimensions. Figure 4.14 shows histograms of the ν_μ and $\bar{\nu}_\mu$ survival probabilities of all MC events in the baseline dataset versus true neutrino energy and true zenith angle to be compared to the example signature from figure 1.22 (and figure 4.1 center). The histograms in figure 4.15 show the same entries, but rebinned by reconstructed energy (muEx) and zenith angle (MPE). The previously sharp signatures (especially the resonance for $\bar{\nu}_\mu$) are washed out severely, and features extend far towards lower reconstructed energies.

Note that the z-scales of these plots are dynamic to retain optimal contrast, and note that IceCube can not distinguish between ν_μ and $\bar{\nu}_\mu$ for individual events.

The limited energy resolution of IceCube is detrimental to the sensitivity of this analysis and turns highly specific oscillation signatures into fairly unspecific signatures that could more easily be confused with effects caused by other phenomena, such as those introduced in chapter 6. However, the signatures retain their characteristically bended shape and a relatively sharp edge towards energies above the resonance. Therefore, even the washed out signatures are clear and unambiguous enough over a large parameter range and allow a meaningful analysis with good sensitivity as demonstrated in the next chapters. Only at very large Δm_{42}^2 (around 10 eV^2), the high-energy edge is too high in energy to be resolved anymore, and the signature turns into global constant relative event rate deficit, which is much harder to distinguish from other phenomena.

A way to circumvent the limitation by selecting only events starting inside the detector because of their better energy resolution is discussed in section 8.2.1.

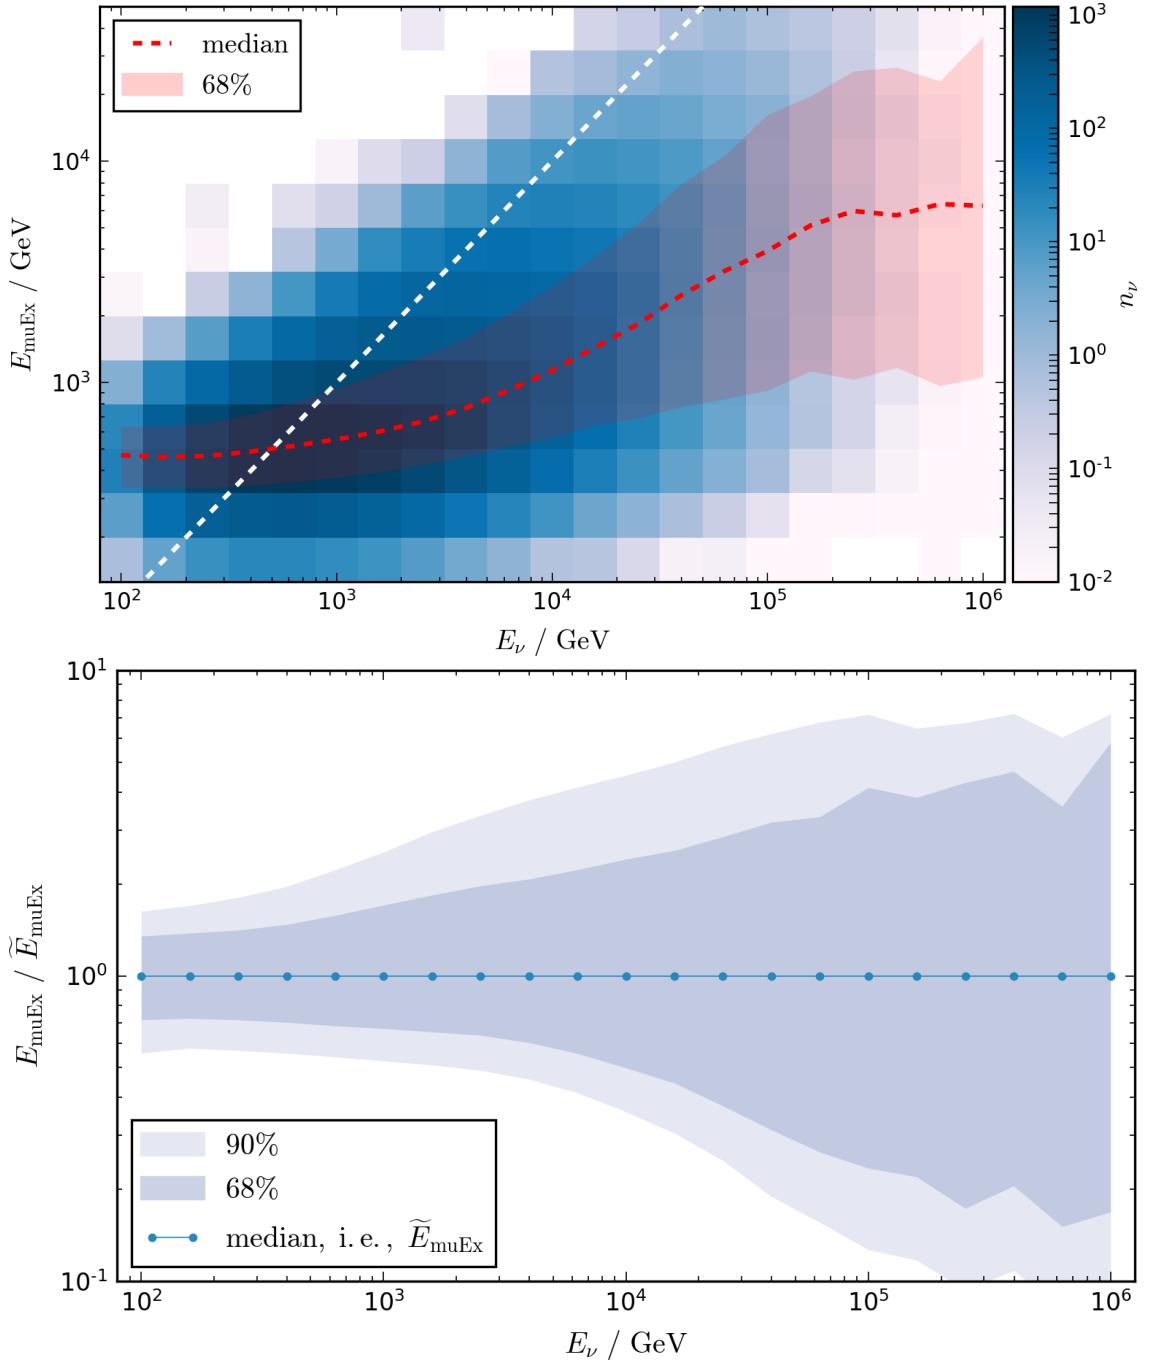


Figure 4.13: Reconstructed muon neutrino energy according to the muEx algorithm as a function of true neutrino energy (**top**), and relative spread of the reconstructed values as function of true neutrino energy (**bottom**).
 The white dashed line marks the axis diagonal, i.e., the behavior of a perfect reconstruction.

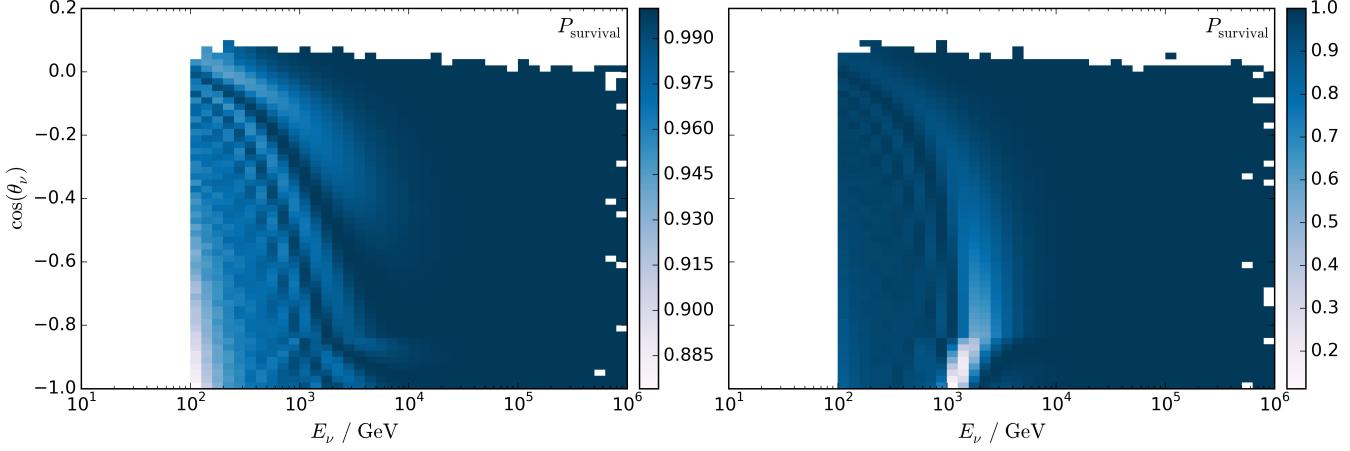


Figure 4.14: Muon neutrino (**left**) and antineutrino (**right**) survival probabilities sampled over all final-level MC events versus true MC zenith angle and neutrino energy, equivalent to figure 1.22 (i.e., $\Delta m_{42}^2 = 0.492 \text{ eV}^2$, $\theta_{24} = 7^\circ$).

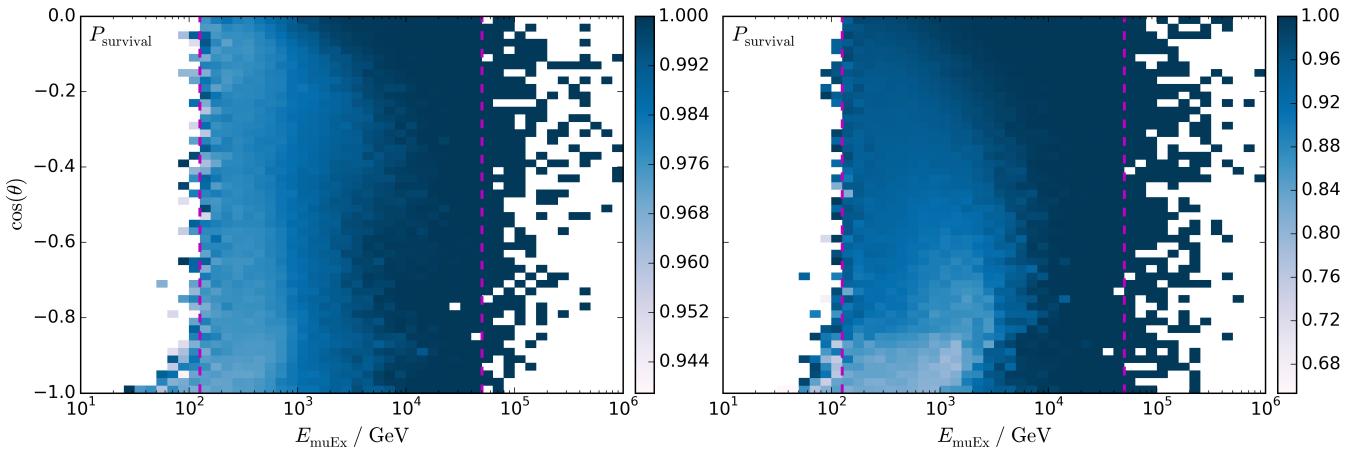


Figure 4.15: Muon neutrino (**left**) and antineutrino (**right**) survival probabilities sampled over all final-level MC events versus reconstructed zenith angle and neutrino energy, analog to figure 4.14.

Dashed magenta lines indicate the energy range used for this analysis.

Chapter 5

Analysis Strategy

The analysis method has previously been presented at conferences and has accordingly been published in proceedings by the IceCube Collaboration (ref. [120] and [121], also published in ref. [122] and [123], respectively). The publications have been written as corresponding author. The descriptions in this work go into far more detail, and everything is work of the author unless specifically noted.

The analysis presented here is a muon (anti-)neutrino disappearance analysis for atmospheric neutrinos at high energies (90% interval of 236 GeV to 10.6 TeV). It is implemented as a likelihood ratio test in two physics parameters (θ_{24} and Δm_{42}^2), conducted on two-dimensional histograms of neutrino events versus reconstructed energy and zenith angle. The null hypothesis of the likelihood ratio test is the non-existence of light sterile neutrino flavors.

5.1 Binning and Probability Calculation

The binning of the histograms employed in the likelihood ratio test is 25 bins linearly in $\cos(\theta)$ between -1 and 0 , and 13 bins linearly in $\log(E_{\text{muEx}}/\text{GeV})$ between $10^{2.1}$ GeV and $10^{4.5}$ GeV as illustrated in figure 5.1. Using a much smaller binning reduces the statistics per bin to the point where the resulting contours become more unstable, whereas much larger binnings significantly reduce the sensitivity of the analysis because features of the model signatures become washed out.

To illustrate the statistics of the baseline MC dataset (consisting of weighted events, see section 2.4), figure 5.2 shows its numbers of effective MC events per bin, defined as

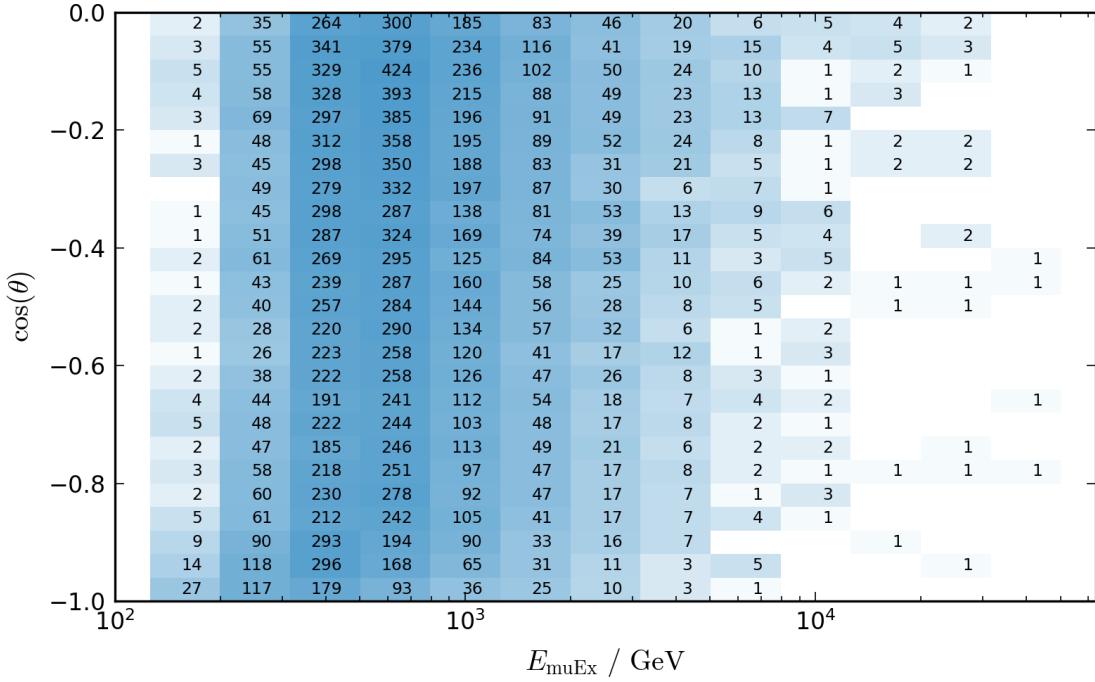


Figure 5.1: Event numbers of the final-level event sample versus reconstructed zenith angle and energy in the analysis binning.

$n_{\text{eff}} = \frac{(\sum w_i)^2}{\sum w_i^2}$, equaling the number of unweighted events one would need for the same statistical power [126]. These numbers are at least 12 times higher than the experimental event counts per bin almost everywhere and usually significantly better (see figure 5.3).

The calculation of the oscillation probabilities is done per MC event, using the tool nuCraft (see appendix C). The calculation has to be done using the true MC values for energy and zenith angle as illustrated in figure 4.14, but the rebinning to the experimentally observable reconstructed values can introduce relatively large fluctuations in the resulting histograms. The reason for this is that the migration into distant bins can be large because of the mediocre energy resolution, and the weights of migrating events can be large because of the nature of weighted simulation. Figure 5.4 shows the fraction $\max(w_i) / \sum(w_i)$ in the binning of figure 4.14 (four times finer than the analysis binning): At high or very low energies, single events can have a large influence on MC distributions. This can become problematic at energies far below the first oscillation minimum for given parameters, where oscillations are so fast that they should appear smeared out. The oscillation probabilities of individual events are chaotic, i.e., deterministic, but highly dependent on the oscillation parameters in a way that makes them seem random, which can cause large unphysical fluctuations for small parameter changes. As a solution to this problem, the individual event probabilities are filled into a histogram in true neutrino energy and zenith angle in the fine binning of figure 5.4, gently smoothed using a locally weight-preserving Gaussian blur algorithm explained in appendix D, and then probabili-

5.1. BINNING AND PROBABILITY CALCULATION

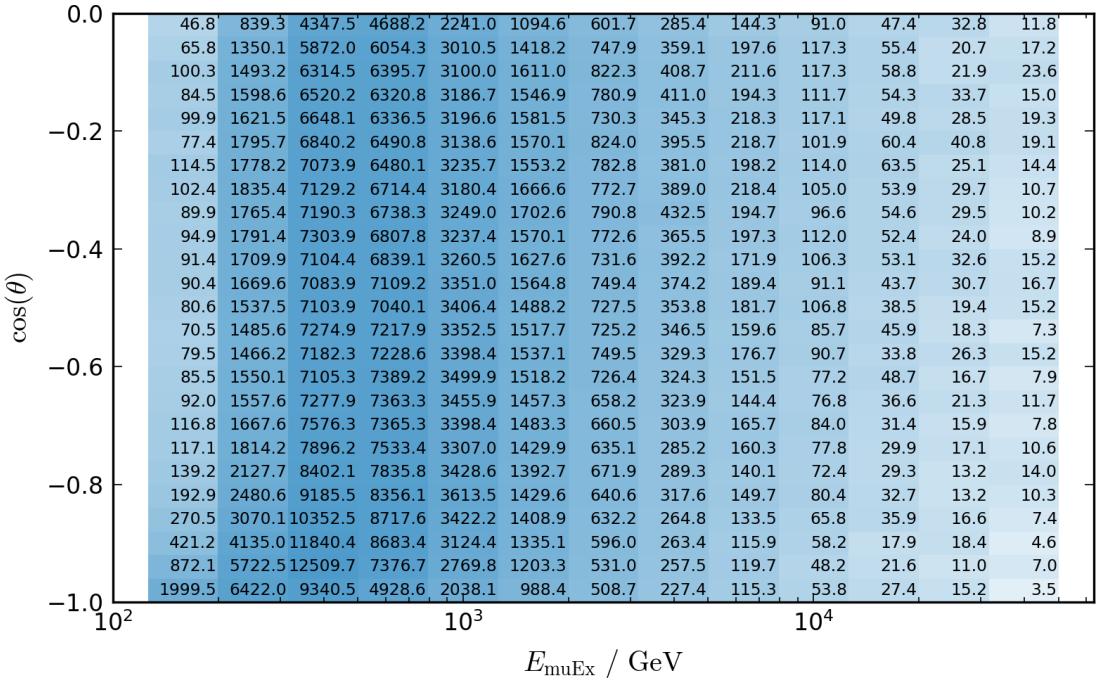


Figure 5.2: Effective number of events n_{eff} per bin for the baseline MC dataset at final level (with fitted nuisance parameters).

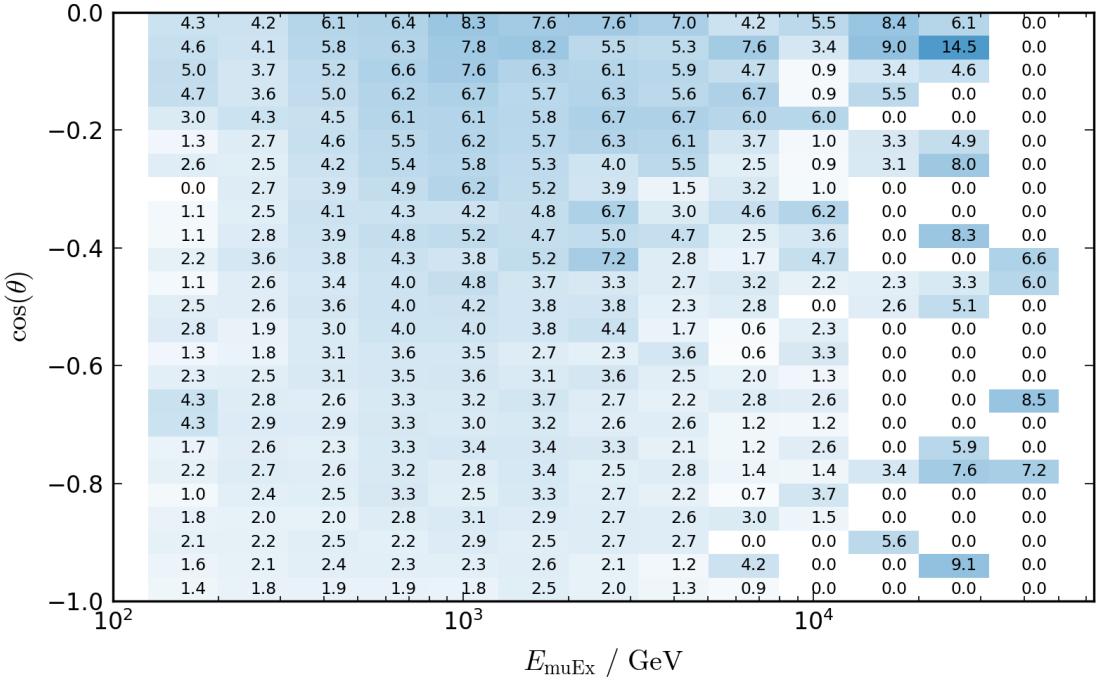


Figure 5.3: Percentual ratio $100 \cdot n_{\text{exp}} / n_{\text{eff}}$ of experimental event numbers to effective event numbers for the baseline MC dataset at final level, i.e., ratio of figure 5.1 to figure 5.2 in percent.

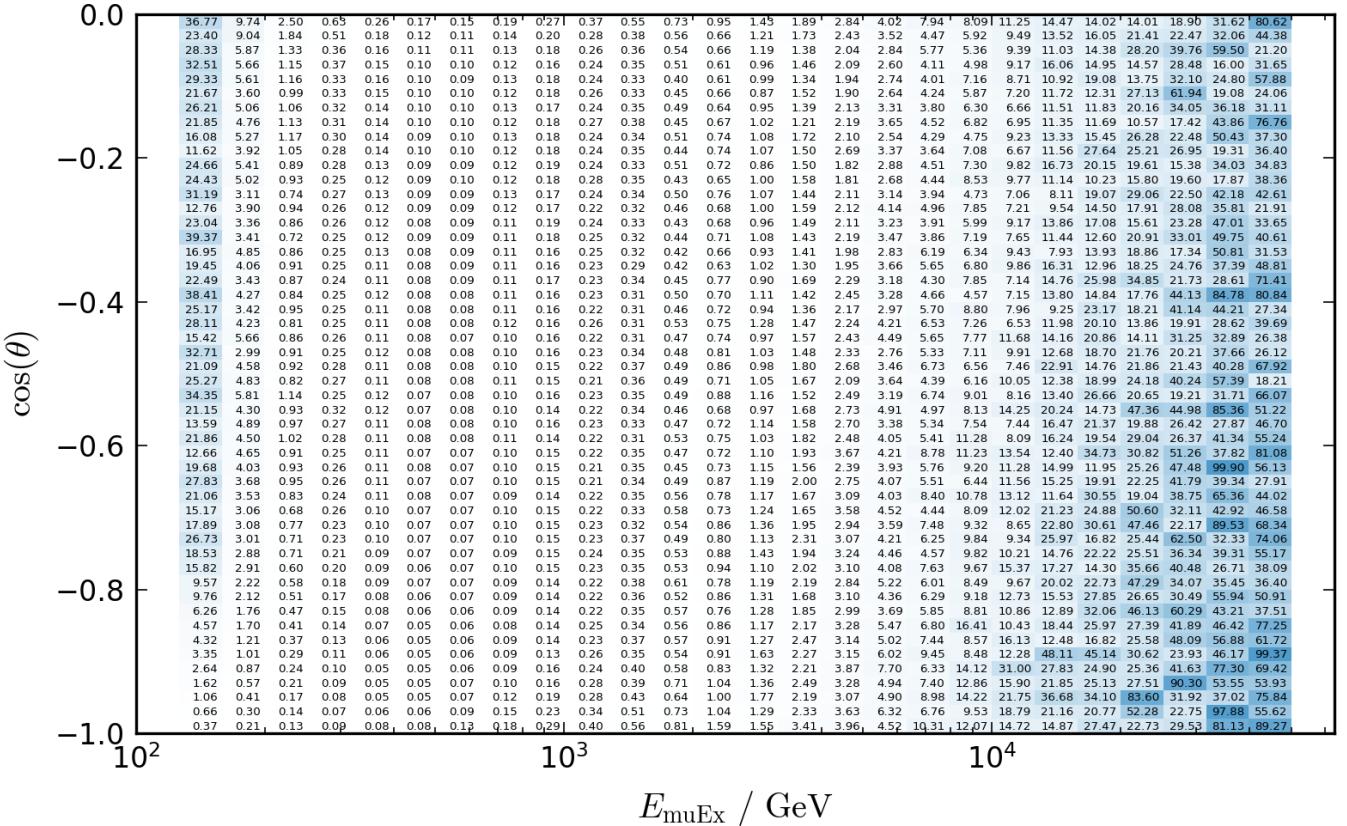


Figure 5.4: Fraction of the maximum event weight among the summed event weights per bin in percent for the baseline MC dataset at final level at a four times finer binning, i.e., $100 \cdot \max(w_i) / \sum(w_i)$.

ties are drawn from this histogram for rebinning into reconstructed quantities.

This procedure has very little effect for small mixing angles because the amplitude of fast oscillations is small in that case, but it is important for larger angles. Figure 5.5 shows its effects for an intermediate angle of $\theta_{24} = 19^\circ$. The smoothing is visible in the MC truth histograms, but does not wash out features in the observable histograms because the smoothing kernel is fine compared to the reconstruction uncertainties. It does however reduce the fluctuations at very low energies as intended.

The effect can be seen more clearly in the standard deviations of the MC event probabilities per bin as shown in figure 5.6: the standard deviations are basically unchanged in the region of the first oscillation minimum, to which events with various different oscillation weights migrate, but they are reduced strongly at lower energies, where oscillations are expected to be washed out.

5.1. BINNING AND PROBABILITY CALCULATION

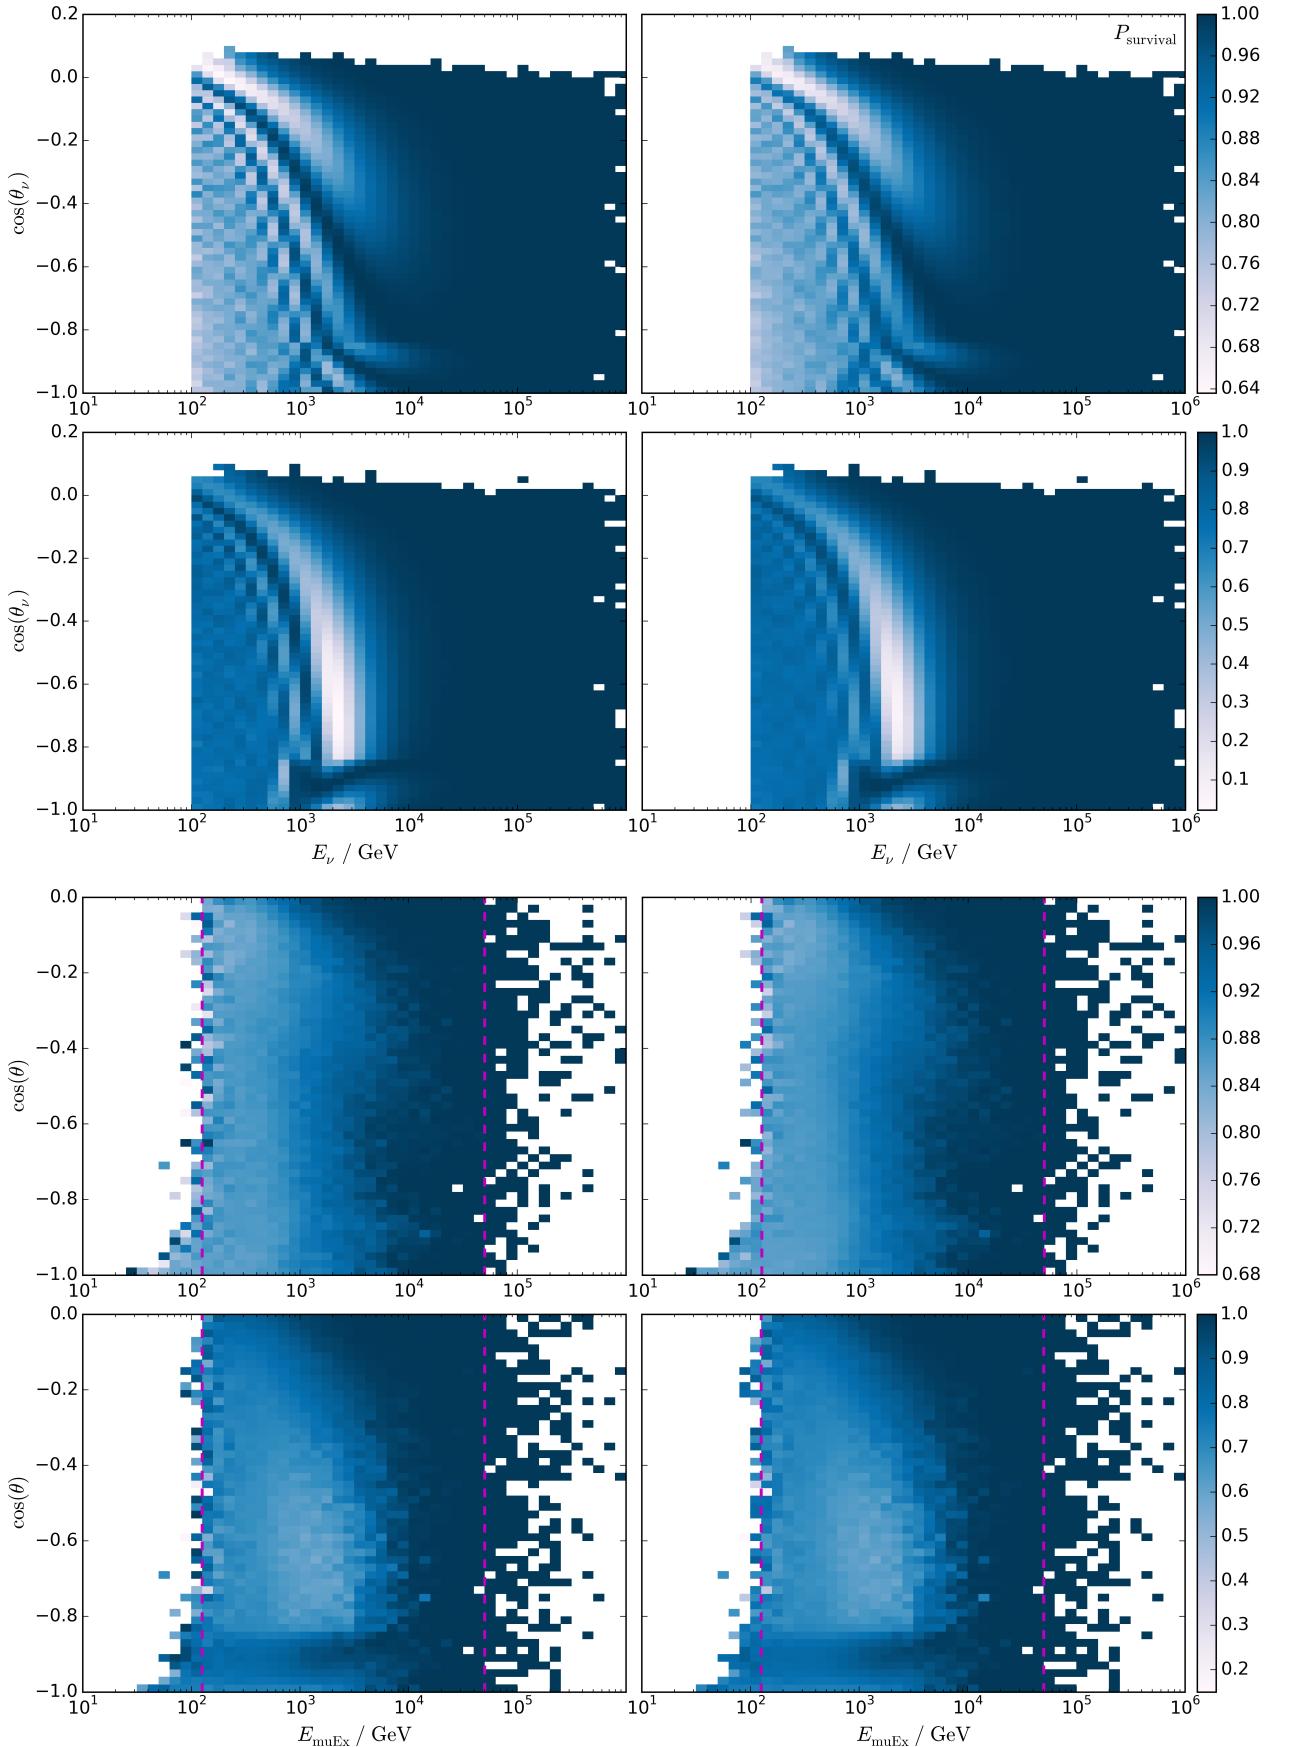


Figure 5.5: ν_μ (upper) and $\bar{\nu}_\mu$ (lower) survival prob. vs. true (top) and reconstr. (bottom) energy and zenith angle for $\Delta m_{42}^2 = 0.492 \text{ eV}^2$ and $\theta_{24} = 19^\circ$. Left side is before, right side after the smoothing and histogramming procedure described in the text.

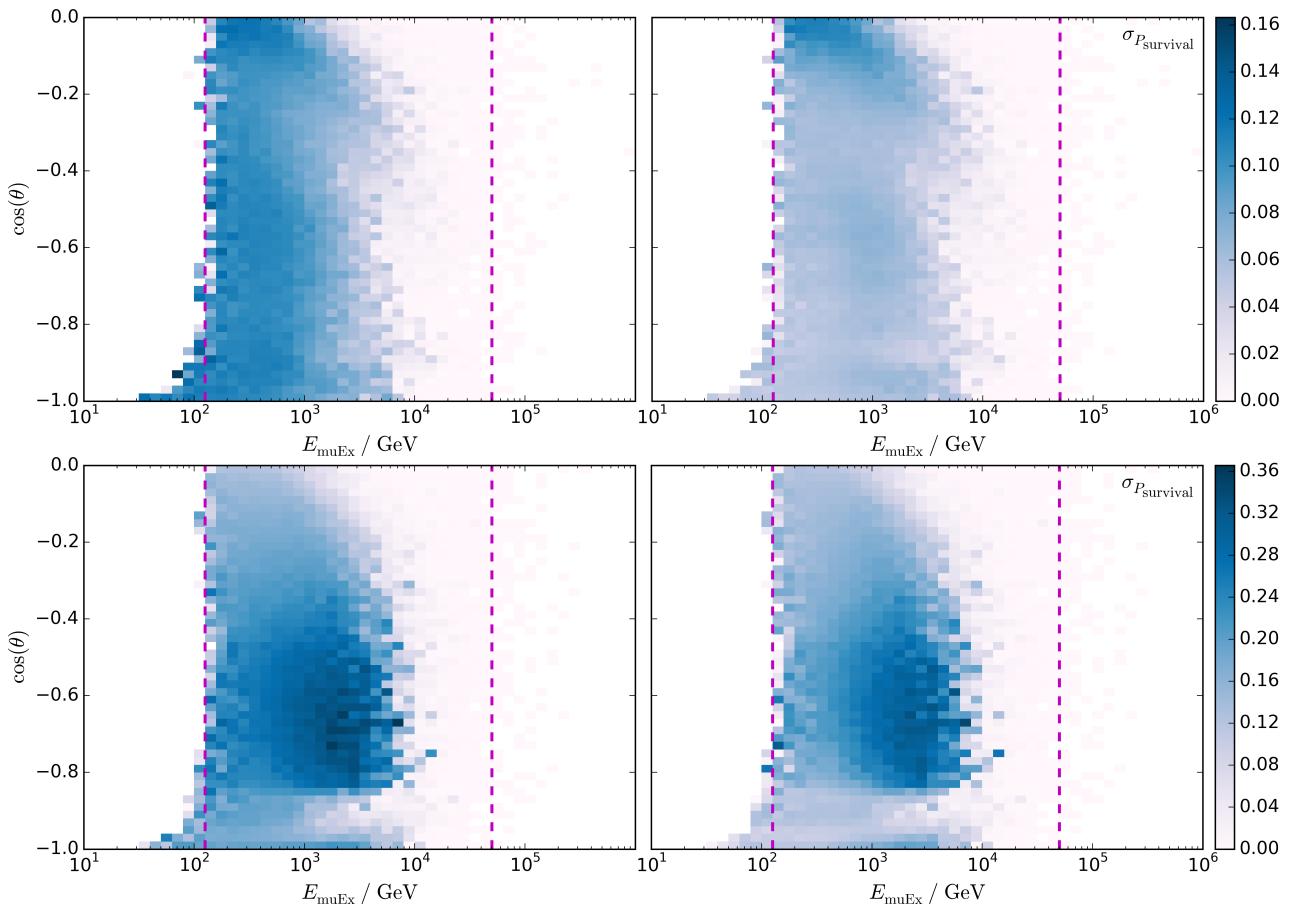


Figure 5.6: Standard deviations of the individual $(\bar{\nu})_\mu$ survival probabilities of all MC events per bin for the lower four plots of figure 5.5 (i.e., ν_μ **top**, $\bar{\nu}_\mu$ **bottom**).

5.2 Likelihood Function

A likelihood function \mathcal{L} can be used to compute the goodness of fit between two histograms. The underlying statistical question is how likely it is that the histogram of experimental data (i.e., the observation x) is a representation of the distribution s described by the second histogram (the simulated data which is defined by the parameters $\xi \in \Xi$). The likelihood function answers this by providing the probability of the outcome x as function of the parameters ξ [127]:

$$\mathcal{L} : \Xi \rightarrow \mathbb{R} : \xi \rightarrow \mathcal{L}(\xi|x) = P(x|\xi) \quad (5.1)$$

The per-bin-likelihood function used in this analysis is the Poisson likelihood

$$\tilde{\mathcal{L}}_{i,j}(\xi|x_{i,j}) = \frac{s_{i,j}(\xi)^{x_{i,j}} e^{-s_{i,j}(\xi)}}{x_{i,j}!}, \quad (5.2)$$

which correctly takes into account the statistical fluctuations of the experimental data per bin (i, j) by assuming a Poisson distribution [127] (i.e., assuming a constant event rate). For numerical reasons, the function is transformed to a negative log-likelihood function (LLH) before a sum is taken over all bins:

$$\begin{aligned} \mathcal{L}(\xi|x) &= \sum_{i,j}^{s_{i,j}(\xi) \neq 0} -\log(\tilde{\mathcal{L}}_{i,j}(\xi|x_{i,j})) \\ &= -\sum_{i,j}^{s_{i,j}(\xi) \neq 0} (x_{i,j} \log(s_{i,j}(\xi)) - s_{i,j}(\xi) - \log \Gamma(x_{i,j} + 1)) \end{aligned} \quad (5.3)$$

This log-likelihood function (5.3) works very well for this analysis even though it is flawed in the regard that it does not take into account statistical uncertainties on the simulated data $s(\xi)$. During the development of this analysis, older MC datasets with significantly lower statistics were used, at times with smaller effective event numbers n_{eff} than in experimental data. For these datasets, the likelihood was switched to LSSL [128], which is a modified Poisson likelihood that also takes into account the statistics of $s(\xi)$. The drawbacks of LSSL are that it needs a more complicated data structure as input (namely the list $W_{i,j} := \{w(\xi)\}_{i,j}$ of all MC event weights in bin (i, j) instead of $s_{i,j}(\xi) = \sum W_{i,j}$), and that per evaluation, a non-linear function has to be solved, which was implemented using Brent's method [129][130]. Because of this, the CPU time requirements increase by about 40%. Since LSSL converges to the Poisson likelihood in the case of infinite MC statistics, and no significant difference between using LSSL and (5.2) had been found using the final MC datasets, LSSL has been dropped.

5.3 Likelihood Ratio Test

With the log-likelihood function (5.3) defined, a likelihood ratio test can be performed to compare the goodness of fit of various hypotheses against each other quantitatively and derive exclusion regions in the parameter space Ξ . The likelihood ratio is defined as the ratio between the likelihood value of a best-matching hypothesis in Ξ and the likelihood value of an hypothesis $\xi^0 \in \Xi$ (or of the best hypothesis in $\Xi_0 \subset \Xi$). For negative log-likelihood functions, this translates to a subtraction of these two functions:

$$\Lambda = \inf_{\xi \in \Xi} (\mathcal{L}(\xi|x)) - \mathcal{L}(\xi^0|x)$$

The infimum (inf) describes the best-matching hypothesis as it lies at the same position as the supremum of the original (non-negative) likelihood function $\tilde{\mathcal{L}}$ (because the logarithm is a strictly monotone function).

According to Wilks' theorem [131], the distribution of 2Λ follows a χ^2 distribution with its degrees of freedom n_{DOF} equal to the dimensionality of Ξ that gets lost when transitioning from a general hypothesis to the null hypothesis ξ_0 (or null hypotheses Ξ_0 , respectively). In the context of this analysis with its two physics parameters Δm_{42}^2 and θ_{24} , this means $n_{\text{DOF}} = 2$.

The caveat is that Wilks' theorem is only applicable if the null hypothesis is identifiable in Ξ , i.e., if the true values of all parameters could be learned given an infinite number of observations. This is not the case in this analysis:

The null hypothesis is the non-existence of light sterile neutrinos that mix with muon neutrinos. Parameter-wise, this could be implemented either as $\theta_{24} = 0$, in which case Δm_{42}^2 can not be determined, or as $\Delta m_{42}^2 = 0$ with arbitrary θ_{24} .

Nonetheless, 2Λ can be used a test statistic to calculate significances and exclusion regions, but the distribution has to be inferred from ensemble tests instead of relying on Wilks' theorem. Results of these tests and data challenges are shown in section 7.1.2 and show that Wilks' theorem holds in good approximation, but final limits are based on the more robust ensemble tests.

5.4 Nuisance Parameters

Nuisance parameters are a well-established method to correctly take into account systematic uncertainties [7]. They are parameters that are added to the model s , but unlike the physics parameters Δm_{42}^2 and θ_{24} they are not of immediate interest for the respective analysis and only have to be added because they influence s and might be correlated to the physics parameters. With the addition of these parameters, the null

hypothesis in this analysis is not described by a single ξ^0 , but instead by a subset $\Xi_0 = \{\xi_0 \in \Xi \mid \Delta m_{42}^2 = 0 \vee \theta_{24} = 0\}$. This means that other parameters of s , such as the total flux normalization of the atmospheric neutrinos or hole ice models, can be optimized also for the null hypothesis, and the likelihood ratio therefore becomes

$$\Lambda = \inf_{\xi \in \Xi} (\mathcal{L}(\xi|x)) - \inf_{\xi^0 \in \Xi_0} (\mathcal{L}(\xi^0|x)). \quad (5.4)$$

The individual nuisance parameters are explained in detail in chapter 6.

It is important to understand the nature of nuisance parameters. The phenomena described by those parameters are not intrinsically uninteresting. They can however be uninteresting for a given analysis, because their measurements are not the declared goal of this analysis. The scope of an analysis can be widened by turning nuisance parameters into physics parameters, but this might cost sensitivity to the original parameters and requires careful examination of all related parameters:

- The sensitivity can be reduced because the distribution of the test statistic becomes wider as n_{DOF} rises. The effect of the original physics parameters compared to the null hypothesis therefore has less impact on the derived p-value unless it is strongly correlated to the additional physics parameter.
- Nuisance parameters are allowed to absorb effects of various phenomena even when they are named after only one specific phenomenon. It is only important that all systematic effects that could influence the measurement of physics parameters get absorbed by the ensemble of nuisance parameters. Because of this, the measured value of a nuisance parameter may not be interpreted as a measurement of the phenomenon after which it was named, and it may only be turned into a physics parameter when all correlated phenomena have been modeled individually.

Special care has to be taken when nuisance parameters are constrained by priors. Priors can be introduced to improve the sensitivity of an analysis and to facilitate the convergence of the numerical minimization towards the infimum of the negative log-likelihood functions. In this analysis, Gaussian priors are implemented as additional addends in (5.4) as in

$$\Lambda = \inf_{\xi \in \Xi} \left(\mathcal{L}(\xi|x) + \sum_{i=1}^{n_{\text{prior}}} \frac{1}{2\sigma_i^2} (\xi_i - \mu_i)^2 \right) - \inf_{\xi^0 \in \Xi_0} \left(\mathcal{L}(\xi^0|x) + \sum_{i=1}^{n_{\text{prior}}} \frac{1}{2\sigma_i^2} (\xi_i^0 - \mu_i)^2 \right) \quad (5.5)$$

with parameter expectation values of μ_i and prior widths of σ_i . Some nuisance parameters were left unconstrained (i.e., without prior¹⁴) because convergence was unproblematic and no a-priori knowledge should be imposed on some critically important parameters

¹⁴This also often called a flat prior, because adding a constant of a uniform distribution is irrelevant to the likelihood ratio method.

to ensure an unbiased result. In general, weak priors were chosen (i.e., large σ_i ; see chapter 6 and especially section 6.3) to allow the nuisance parameters to properly absorb systematic deviations between standard MC simulations and experimental data without allowing them to run to completely unphysical values.

The best-fit values of the nuisance parameters were checked on the null hypothesis and later on also on the best-fit hypothesis while keeping the physics parameters blind to ensure that no parameters are strongly constrained by priors, which could be a hint for an overlooked systematic uncertainty or for wrong priors. These checks are shown in chapter 7.

5.5 Software Implementation

The analysis has been implemented as a parameter scan that applies the likelihood ratio test to a 18×24 grid of possible physics parameter values. The grid is rectilinear and has nodes in θ_{24} at all odd numbers from 1° to 29° inclusive as well as 32° , 37° , and 42° . In Δm_{42}^2 , the nodes are located at the linear middle points between eight evenly logarithmically distributed values per decade, starting at 10^{-2} eV^2 and ending at 10^1 eV^2 ; this means the scanned values lie between 0.0167 and 8.75 eV^2 . In addition, the null hypothesis is included ($\Delta m_{42}^2 = 0$ and $\theta_{24} = 0$).

The range in Δm_{42}^2 was determined by estimates for the analysis' sensitivity from Asimov studies (see section 7.1 and especially figure 7.1) and ensemble tests without taking into account theoretical predictions to remain model independent. Around $4 \cdot 10^{-2} \text{ eV}^2$, the first oscillation minimum lies below the effective energy range of this dataset. Simple studies with an IC-79 sample that was optimized for low energies and standard oscillation analyses indicate that this analysis would gain sensitivity down to about 10^{-2} eV^2 [132] (see figure 8.1). However, event selection techniques and therefore systematic uncertainties are very different for such energies, so an extension of the energy range was not attempted for this initial analysis.

Above the optimum sensitivity at around 0.2 eV^2 , the sensitivity decreases steadily but more slowly, because the first oscillation minimum moves into regions with smaller statistics. For very high Δm_{42}^2 , only horizontally incoming high-energy events are unaffected by oscillation effects, with other events already lying in the smeared-out fast-oscillation region. The high- Δm_{42}^2 cut at 8.75 eV^2 was selected because of the low remaining sensitivity, but also because it corresponds to a resonance energy of about 33 TeV for vertically upgoing events, where the approximation that oscillation effects and absorption in the Earth are decoupled slowly begins to fail (see appendix C.1).

For testing purposes the Δm_{42}^2 range was extended by seven values up to 65.6 eV^2 (middle point towards $10^{1.875} \text{ eV}^2$ as described above) as shown in figure 7.1, but the additional parameter space is not considered for final exclusion results.

For every physics parameter pair in the scan, a minimizer has to be used to find the infima of all nuisance parameters at that position. These values are obtained simultaneously by the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm for bound-constrained optimization (L-BFGS-B) as implemented in SciPy [133][130]. Being a quasi-Newtonian method, it relies on an estimate of the inverse Hessian matrix to iteratively find the minimum. It was tested to reliably identify the global minimum and is faster than other algorithms such as MIGRAD of the Minuit(2) software package [134].

With the nuisance parameters optimized for every scan point (including the null hypothesis), the best physics parameters can be determined as the one with the smallest LLH contribution, and Λ can be approximated as difference between the best point (again including null) and the null hypothesis. To get more precise results, the LLH landscape is interpolated fivefold between the grid points using third-degree splines as provided by SciPy's multi-dimensional image processing package (`scipy.ndimage.interpolation.zoom`) [130]. This approach is valid because the LLH landscape is smooth as can be seen in the result plots such as figure 7.1, which display the original LLH landscapes as well as exclusion contours that are based on the interpolations.

Chapter 6

Handling of Systematic Uncertainties

The handling of systematic uncertainties can be split into two categories. Making up the first category, parameterizable uncertainties for the purpose of this analysis are those that are describable by a parameterization with one or more real-valued, continuous parameters. They are handled as nuisance parameters as described in section 5.4, so a minimizer is used to simultaneously optimize all their parameters, taking into account correlations and potentially complex, non-linear behaviors. Each parameterization yields one number per MC event by which that event weight is modified multiplicatively, so the minimization is implemented as unbinned likelihood method.

The second category are unparameterizable uncertainties. It only contains phenomena for which continuous parameterizations could not be found or are not necessary. For those phenomena that have multiple discrete options (e.g., atmospheric models), all combinations are run through individual likelihood scans with all nuisance parameters minimized as usual, and the best combination of unparameterizable uncertainties (i.e, the one with the smallest LLH value) is selected per bin of the scan before the interpolation is conducted.

Because of this nesting and because the atmospheric models are the basis for many other uncertainties, the unparameterizable ones are described first.

6.1 Unparameterizable Uncertainties

6.1.1 Atmospheric Models

In the context of this work, “atmospheric model” means flux predictions for atmospheric (muon) neutrinos and antineutrinos. Since this is a disappearance study on atmospheric

neutrinos, their distribution in zenith angle and energy has to be described as accurate as possible. Various parameterizable uncertainties (namely the total flux normalization, the spectral index, the kaon-pion ratio and the antineutrino-neutrino ratio) are designed to adjust the atmospheric model to match the measurements, but nonetheless different atmospheric models are tested to find the best matching as they offer an additional degree of freedom.

The first model that was implemented and the one whose implementation differs slightly from the others is HKKMS2006¹⁵ [32] by Honda, Kajita, Kasahara, Midorikawa, and Sanuki. The model is not used as provided in the source tables [33], but instead was fitted to the Gaisser formula (1.3) above 700 GeV and to a fourth-degree polynomial in the two variables $\cos(\theta)$ and E_ν below this energy, to account for higher-order corrections [135]. It was then modified to be based on Gaisser’s H3a cosmic-ray model [136] instead of a simple unbroken power law by applying a modified Elbert equation [137][138] to H3a that parameterizes the muon neutrino yield of cosmic-ray showers. Details of this modification can be found in ref. [114] and ref. [139], but the impact of this change to the original HKKMS2006 model is small at neutrino energies below the knee and therefore largely negligible for this work. The model is known to describe the data reasonably well as it was the preferred model of the earlier search for an astrophysical flux on this data [114].

The other models were generated for another analysis by Gabriel Collins, using Anatoli Fedynitch’s tool **MCEq** [140]. They are freely available at ref. [141]. Similar to the modification of HKKMS2006, **MCEq** combines existing hadronic-interaction models with existing cosmic-ray models. The three atmospheric models taken into account are an older version of Honda’s model (from a 2002 review paper [142]) fitted to balloon data and modified according to H3a (called “combined Gaisser-Honda and Hillas-Gaisser (H3a)” and described in ref. [143]), Hörandel’s poly-gonato [144], and a model by Zatsepin and Sokolskaya [145] refitted to PAMELA data as described in ref. [146]. The three hadronic-interaction models that were used are QGSJet II.04, release candidate 1 of Sibyll 2.3, and a modification of this Sibyll release candidate that assumes charm quark production to be point-like as described in ref. [140]. This leads to nine possible combinations, all of which were implemented in this analysis. **MCEq** offers a few other options for both kinds of models, but those are either outdated versions of the other models or known to describe existing data worse in the relevant energy region [141][143].

¹⁵The model is also known as HKKMS2007 (as the paper was published in that year), HKKM2007 (omitting Sanuki because he was not part of the team that made the progenitor model HKKM2004), or Honda 2006 after the first author.

6.1.2 Seasonal Variations

As described in section 1.2, the production of atmospheric neutrinos depends on the air density and therefore temperature of the upper atmosphere, most importantly around a height of 15 km [27]. This analysis uses the tool `temp_flux_reweight` [147] together with weather data from the AIRS instrument of the Aqua satellite Earth Observation System [148] to take into account this effect. The tool works by computing the correlation coefficient α between the relative temperature change and the relative neutrino flux change:

$$\alpha = \frac{T}{\phi} \frac{\partial \phi}{\partial T}$$

The flux is approximated by the Gaisser equation (1.3), and an effective temperature T_{eff} is computed that describes the temperature of the atmosphere weighted with the production height and effective area of IceCube as functions of neutrino energy and zenith angle (details are given in ref. [147]). The correction factor is then computed relative to the US Standard Atmosphere model [149] that is assumed during the production of the MC datasets, so it can easily be used to reweight those datasets:

$$w_i = 1 + \frac{\phi - \phi_{\text{US}}}{\phi_{\text{US}}} = 1 + \alpha \frac{\langle T_{\text{eff}} \rangle - T_{\text{eff,US}}}{T_{\text{eff,US}}}$$

The temperatures and therefore effects are averaged, taking into account the weather during the exact lifetime of that year of IceCube data taking [147].

An upper bound for the time-averaged effect of seasonal variations can be estimated by assuming maximum correlation: with a maximum temperature difference of about 10 K in a 220 K environment the effects are on the order of five percent. Figure 6.1 shows the actual effects obtained by the method described above. The neutrino rate change can reach up to four percent at high energies and varies non-trivially as function of the zenith angle. The pattern is not compatible with a sterile signature because it does not have a minimum for vertical tracks such as all moderate sterile hypotheses while its amplitude is much too small for an extreme sterile hypothesis. It could nonetheless influence the fit results, so seasonal variations are important to take into account.

The uncertainties in the calculation of the effects of seasonal variations are expected to be negligible. The statistical fluctuation of Aqua’s measured temperatures is 1 K to 2 K, but these fluctuations get canceled out by averaging over heights and azimuth angles. The temporal coverage of only two temperatures per location and day is a larger issue, but this is compensated by averaging over the whole year of data taking, and by the fact that temporal fluctuations are small in the two extreme seasons (winter and summer) [147][150].

For the atmospheric models generated by `MCEq` (see section 6.1.1), fluxes were provided with Aqua-based temperature corrections for the first year of data taking with

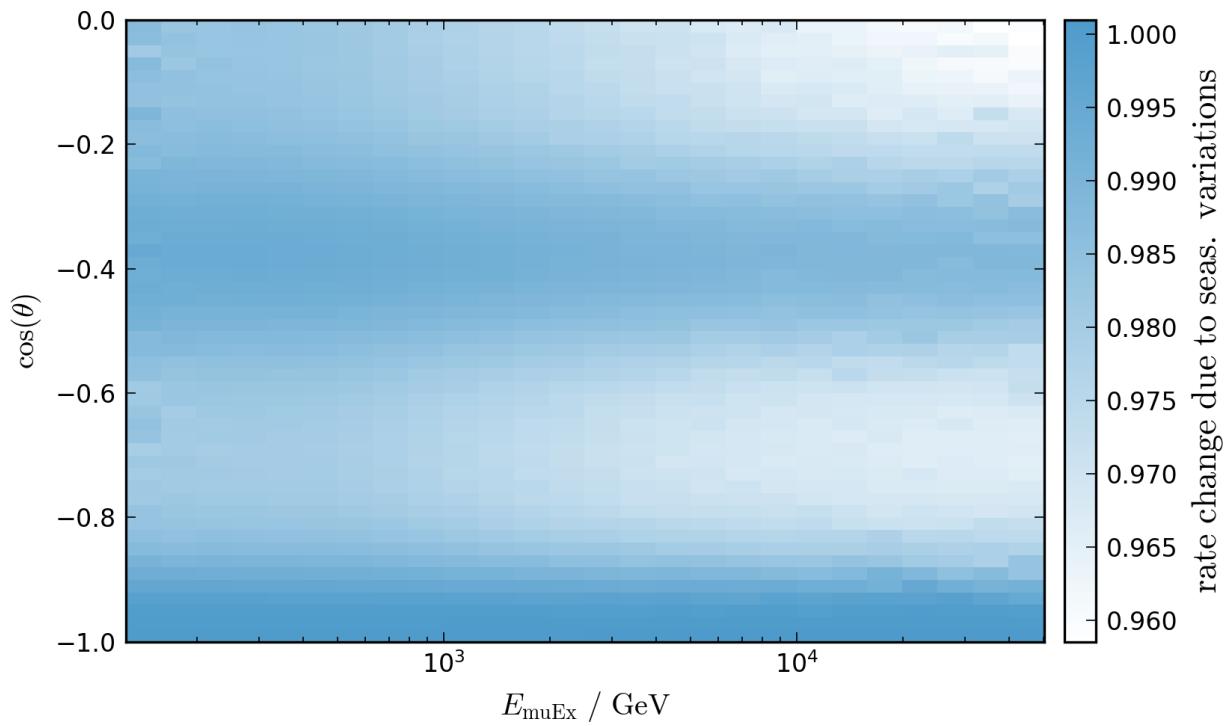


Figure 6.1: Effect of the time-averaged seasonal variations on the neutrino event rate as a function of reconstructed energy and zenith angle, including all systematic uncertainties but atmospheric models.

IC-86, which went from May 2011 to May 2012. These corrections were reverted using `temp_flux_reweight`, and the corrections for the IC-59 season were applied.

6.1.3 Ice Models

The concept of and the differences between the most important ice models were already introduced in section 2.2.2. As explained there, this analysis uses the most advanced ice model available at this time, named Spice Lea, which models the effective scattering and absorption lengths as tilted layers with anisotropic properties. Tests with older ice models revealed that the tilt is instrumental in correctly modeling ice on the level of accuracy that is needed for the search for sterile neutrinos, as its absence can induce a fake signature as shown in figure 7.16 (section 7.3).

Spice Lea was selected as baseline because it describes the ice better than its predecessors. A measure for this is the *model error*, defined as the width of the distribution of the differences between the predicted and measured DOM charges larger than 10 PE in the flasher data the Spice models were fit to [86]; its value is 0.20 for Spice Lea, 0.28 for Spice Mie, 0.29 for Spice 1, and 0.42 for WHAM! [106]. Other ice models were tested during the development of the analysis, but are not used as discrete nuisance parameters because differences between the most recent models are smaller than the allowed variations of the ice absorption and scattering lengths as parameterized uncertainties, and were found to not significantly influence results of pseudo experiments. Especially the anisotropy that was introduced in Spice Lea is of little importance for this analysis, as was tested in cross-checks between Spice Mie and Spice Lea.

6.2 Parameterized Uncertainties

6.2.1 Total Flux Normalization

The total (flux) normalization is the simplest uncertainty with respect to its implementation, being just a constant factor that gets applied uniformly to all MC events. It is correlated with various other uncertainties and also signal hypotheses because those similarly influence the number of expected events. However, all these other phenomena have unique patterns in energy and zenith angle that reduce the correlation, so the total normalization is not fully absorbed by the other uncertainties.

The physical justification for this nuisance parameter is the uncertainty in the incoming cosmic-ray flux normalization. The prior has been set to $\sigma = 40\%$ around the flux predicted by the atmospheric model. This is much larger than the uncertainties of those models (e.g., a conservative estimate for the HKKMS2006 model is 25% at 1 TeV [32]),

but it has to encompass the correlated uncertainties mentioned above. The influence of the prior only becomes important for very high Δm_{42}^2 ($\gtrsim 30 \text{ eV}^2$) well outside the standard range, where the only observable effect of sterile neutrinos is the very fast and therefore washed out disappearance across all baselines and energies; in these parameter regions, all of the analysis' sensitivity is due to this prior.

6.2.2 Cosmic-Ray Spectral Index

The spectral index γ of the incoming cosmic rays that interact in the atmosphere to produce neutrinos has been introduced in section 1.2.1. It controls the steepness of the cosmic-ray and neutrino power spectrum and therefore directly influences the energy distribution. A nuisance parameter was implemented to allow γ to change within the uncertainty quoted in ref. [20] (i.e., $\gamma = 2.65 \pm 0.05$ below the knee, which is the relevant energy region for this analysis). To uncouple γ from the total flux normalization, the point of reference was chosen to be the median neutrino energy of the final MC sample \tilde{E}_ν , so the factor by which the MC weight of event i gets modified is

$$w_i = \left(\frac{E_i}{\tilde{E}_\nu} \right)^{-\Delta\gamma}.$$

The total number of events is only strictly conserved for infinitesimally small $\Delta\gamma$ because of the non-linearity of the power function; however, the variations in γ are small enough to make this limitation unproblematic. The alternative method of computing

$$w_i^* = E_i^{-\Delta\gamma} / \sum_i E_i^{-\Delta\gamma}$$

was tested to give identical results, but consumes slightly more memory and CPU time.

Note that this correction factor gets applied on top of a complex model of atmospheric neutrinos, so the analysis does not assume a simple power law for the neutrinos. A simple power law is a model assumption that ignores the complexity of the cosmic-ray spectrum, which with its non-trivial chemical composition is a superposition of various broken power laws with different cut-off energies. The described implementation only assumes that the sources and propagation effects affect the components of the cosmic rays similarly, which is a reasonable approach to account for first-order effects.

6.2.3 Kaon-Pion Ratio

The kaon-pion ratio, too, has been introduced in section 1.2.1: the different kinematic properties and decay processes of pions and kaons mean that the relative abundance of

these mesons in air showers has a significant influence on the properties of atmospheric neutrinos. In contrast to most other nuisance parameters, the kaon-pion ratio is not a multiplicative correction to the weights, but instead directly changes the calculation of those weights. For the atmospheric models generated by `MCEq`, contributions of neutrinos from pions and kaons are available separately and can be scaled arbitrarily. For the model HKKMS2006, the calculation is described below. In both cases the nuisance parameter scales the kaon contribution by a factor of $r \approx 1$. The Gaussian prior has been set to 10% in accordance with the uncertainty given by HKKMS2006 [114].

Above 700 GeV, HKKMS2006 is described by the Gaisser formula (1.3) with empiric values for $\mathcal{A}_{\pi\nu}$, $\mathcal{A}_{K\nu}$, $\mathcal{A}_{\pi\bar{\nu}}$, $\mathcal{A}_{K\bar{\nu}}$, \mathcal{B}_π , \mathcal{B}_K , ϵ_π , ϵ_K , $\cos(\theta^*)$, and one minor zenith-angle-dependent normalization correction for ν_μ and $\bar{\nu}_\mu$ each. In this energy region, the two addends in (1.3) can therefore simply be split to compute weights for pion and kaon distributions separately, and the kaon term is scaled by the kaon-pion ratio to obtain the modified parameterization.

Below 700 GeV, the model does not use (1.3) (which still applies in principle), but instead fits a fourth-degree polynomial in $\cos(\theta)$ and E_ν to account for higher-order corrections as described in section 6.1.1. To be able to apply this nuisance parameter to events below 700 GeV without discontinuities, the high-energy parameterization is extrapolated to these lower energies, and the ratio of it with and without kaon rescaling is multiplied to the original low-energy parameterization.

6.2.4 Antineutrino-Neutrino Ratio

The antineutrino-neutrino ratio has been introduced as source of systematic uncertainties in section 1.2.4. The baseline antineutrino-neutrino ratio for every fit is given by the Monte-Carlo data at final level. It takes into account all cross-sections, the selected atmospheric model, and all nuisance parameter settings, which can differ from the defaults for pseudo experiments (see section 7.1.2). Even without neutrino interactions and detector effects, it can have a complicated shape as shown for HKKMS2006 in figure 1.11.

For the uncertainties on this ratio, this analysis uses the predictions of Barr et al. [38], shown in figure 1.10. They take into account the cancellation of uncertainties due to neutrinos and antineutrinos being produced by the same parent particles. However, they are overestimated in the context of this analysis, because correlations between the various hadron production and flux quantities have not been taken into account, and some – such as the kaon-pion ratio – are being fitted directly in this analysis, with full propagation of the effects on the antineutrino-neutrino ratio. Therefore and because of the unphysically large values obtained by naïve extrapolation that is required beyond 1 TeV, only the shape of those predictions is used, while the amplitude has been reduced.

To implement the uncertainties, two nuisance parameters are used. The first one is a constant $w_1 \approx 100\%$ (MC prediction, taking into account all other nuisance parameter de-

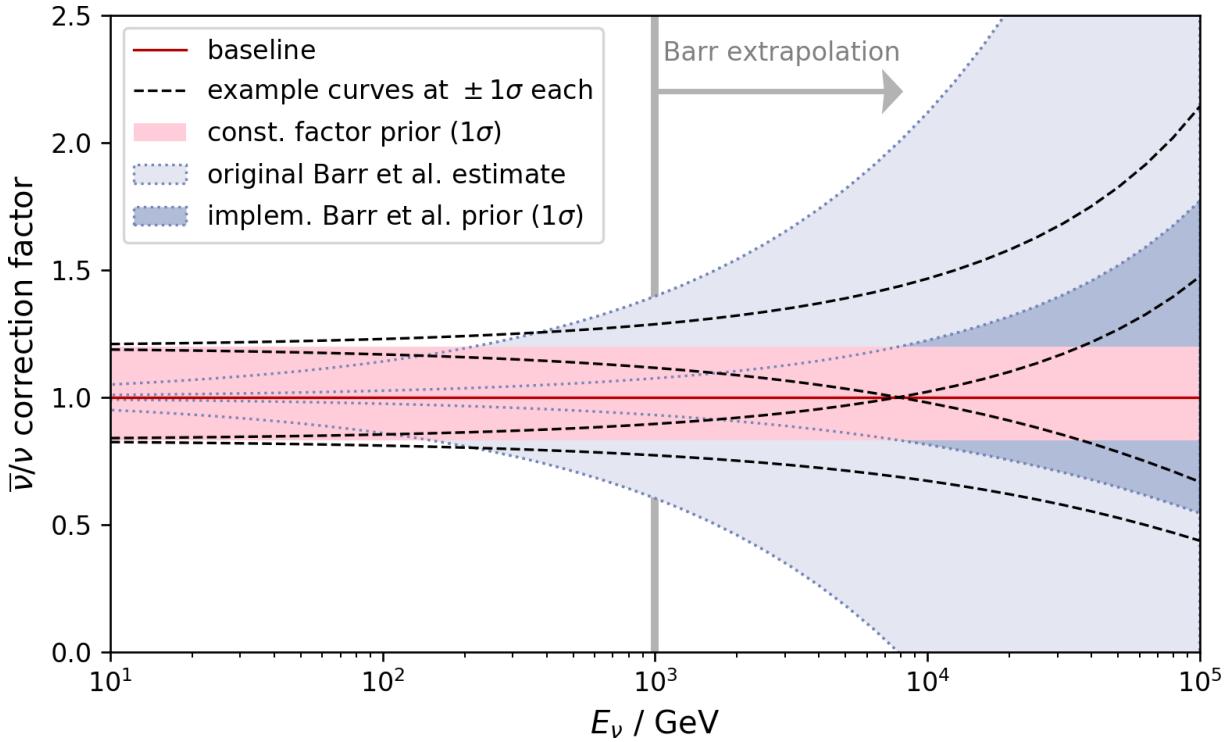


Figure 6.2: Antineutrino-neutrino ratio uncertainty parameterization. Shown is the extrapolation of the Barr et al. predictions to energies above 1 TeV, the 1σ regions of both nuisance parameters, and four example curves for this ratio, given by the combinations of $\pm 1\sigma$ deviations in both parameters.

fault settings) with a Gaussian prior of 10%. The second one ($w_2 \approx 0$) is centered around zero with the same prior and scales the extrapolation of the energy dependence of the maximum uncertainty given by Barr et al., i.e., $w_B(E) = 1.77\% \cdot (E/\text{GeV})^{0.45}$. The two weights are added up to $w(E) = w_1 + w_2 w_B(E)$, which is the factor by which all antineutrino events get scaled. The (non-anti-)neutrino events get scaled down appropriately to conserve the total flux normalization, which leads to the asymmetry visible in figure 6.2.

Combining two nuisance parameters in this manner gives the fit the freedom to, e.g., increase the number of antineutrinos at lower energies while decreasing them at higher energies, shown in the example lines in figure 6.2.

6.2.5 DOM Efficiency

The DOM efficiency is the nuisance parameter that describes the uncertainty in the DOMs' capabilities to detect photons that hit them. It is also known as optical efficiency because it absorbs effects that are not directly related to the DOMs.

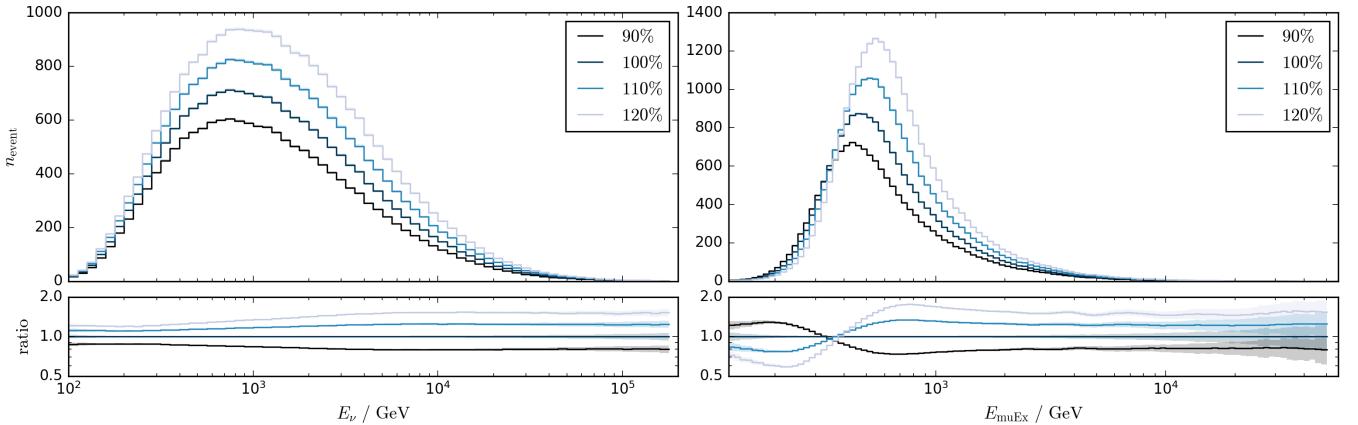


Figure 6.3: Effects of the DOM efficiency on the distribution of true neutrino energy (**left**) and reconstructed energy (**right**).

The optical efficiency of the PMTs has been measured in various laboratory setups as functions of wavelength, angle, impact parameter, or temperature, both with and without glass sphere and optical gel [77]. Its wavelength dependence is shown in figure 2.2. However, systematical uncertainties and DOM-to-DOM variations are at the 10% level in these measurements [151], and in the detector the DOM efficiency is highly correlated with hole ice effects (section 6.2.6) and cable shadowing (i.e., photons hitting the strings cables and getting absorbed).

Dedicated IceCube analyses using abundant and well-understood minimally-ionizing (stopping) atmospheric muons as light source suggested that the overall DOM efficiency is about 10% larger than what was used in Monte-Carlo settings based on the laboratory measurements, and the baseline value was adjusted appropriately for newer simulations (such as the baseline set used in this analysis).

The effects of DOM efficiency variations on the dataset are non-linear. Increased efficiency means that more light is detected, which causes sub-threshold events to now pass trigger and selection cuts to get included into the sample, increasing the number of low-energy events. The higher light yield also influences reconstruction algorithms to assign higher reconstructed energies to all events. Together, these two effects reduce the number of events below the most common reconstructed energy and strongly increase the event number above when the DOM efficiency is increased, and vice versa (figure 6.3). The most common reconstructed energy itself shifts towards higher energies, and there is a weak effect on the zenith angle distribution as well.

Because of the large uncertainties and the relatively strong and non-trivial effects, the DOM efficiency is an important systematic uncertainty for this analysis. For parameterization, an improved version of what was done in the IC-59 search for extragalactic neutrinos [114] was developed, which requires datasets with various values of DOM efficiency. To decrease statistical fluctuations between these datasets (as well as computational effort), all datasets have been produced with strongly increased DOM efficiency (120% of the

nominal value) until after photon propagation as described in section 2.4. At that point, photon hits are discarded randomly until the desired DOM efficiency is reached, before the simulation chain is continued.

Using this method, final-level datasets were created for 120%, 110%, 100%, 90%, and 80% of the nominal DOM efficiency. These datasets were used to produce Kernel Density Estimations (KDEs, with bootstrapping for error estimates, see appendix E) in reconstructed energy and zenith angle, which were evaluated at the positions of all MC events in the baseline dataset (100% DOM efficiency). The resulting five weights (including errors) are normalized to the baseline weight and fitted by least squares for every event, using the log-quadratic function

$$w(x) = 10^{ax^2+bx+c}$$

with DOM efficiency x and free parameters a , b , and c , to get an unbinned per-event parameterization. The log-quadratic function was selected because it describes the weights well according to χ^2/n_{DoF} ; a log-linear parameterization fails around the modal energy value (which was not obvious in ref. [114] because of insufficient statistics at low energies), while a spline of first order can lead to unphysical edges in the parameterization (which is acceptable for low-energy analyses because the effect of DOM efficiency is weaker and more linear below the modal energy).

Examples for these per-event parameterizations are shown in figure 6.4, and χ^2/n_{DoF} values and the two most problematic bins are shown in figure 6.5. The reason for the bad goodness of fit in these bins is an overfluctuation in the baseline dataset and an underfluctuation in the 110% dataset combined with small errors, which gets smeared out over various bins by the KDEs. Even in these bins, the parameterizations look reasonable, and the zenith/energy region where this happens is not critical for sterile neutrinos.

6.2.6 Ice Model Parameters

Absorption length and scattering length have been introduced as the relevant ice parameters in section 2.2.2; this section will focus on their implementations in this analysis. The scattering length is parameterized and fitted separately for the bulk ice, i.e., the ice that constitutes the large volume between the IceCube strings, and for the 60 cm hole ice columns.

The three parameterizations are similar to that of the DOM efficiency in that they are per-event fits made with very similar techniques. One major difference is that the two bulk ice parameterizations are entangled by sharing three common MC datasets: one with $+10\%$ absorption length, one with $+10\%$ scattering length, and one with both $-7.1\% = -10\% \cdot \sin(45^\circ)$ absorption and scattering lengths, each with 1/5 of the baseline MC statistics. Unlike the DOM efficiency datasets, ice model datasets are independent and therefore suffer larger statistical uncertainties.

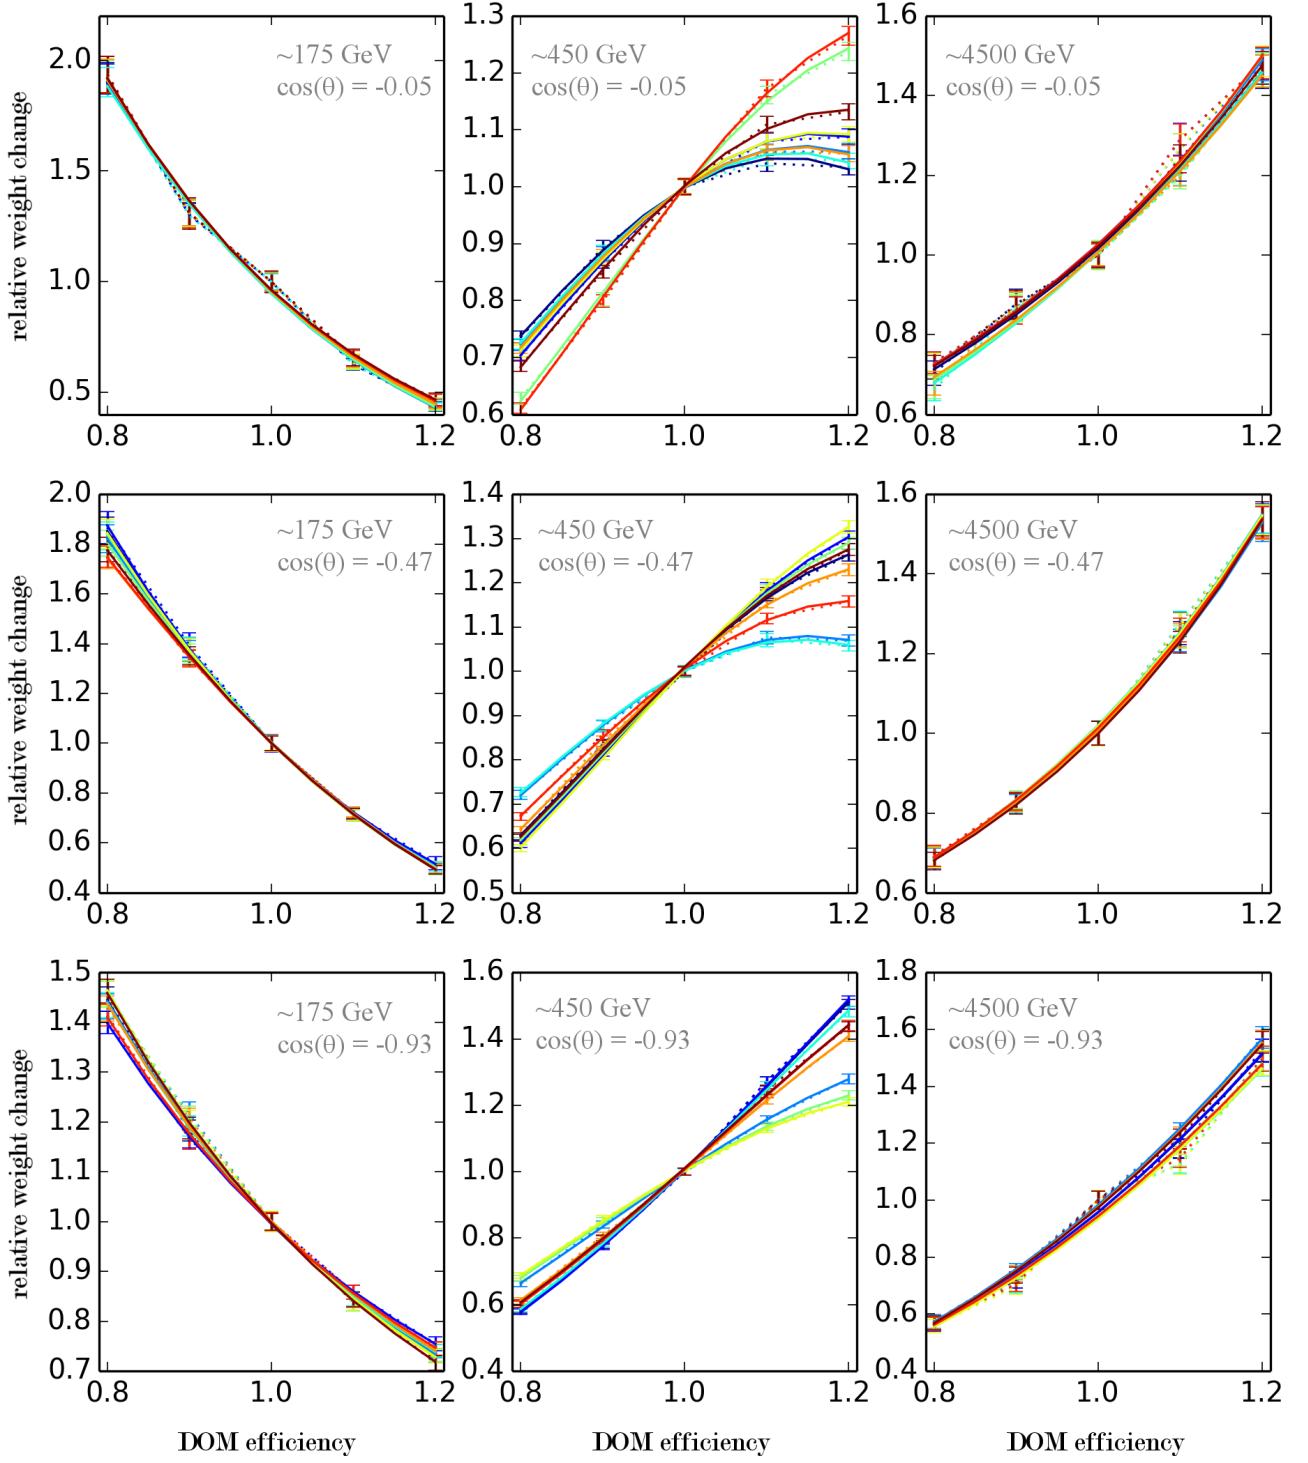


Figure 6.4: Examples for the per-event DOM efficiency parameterizations in nine different regions of energy and zenith angle. Every plot shows ten randomly selected events of a bin of $\Delta \cos(\theta) = 0.2$ and $\Delta \log_{10}(E_{\text{muEx}} / \text{GeV}) = 0.1$. Points are obtained from the KDEs of MC datasets at that DOM efficiency, error bars come from bootstrapping, dashed lines connect points that belong to one event to guide the eye, and solid lines show the log-quadratic parameterizations.

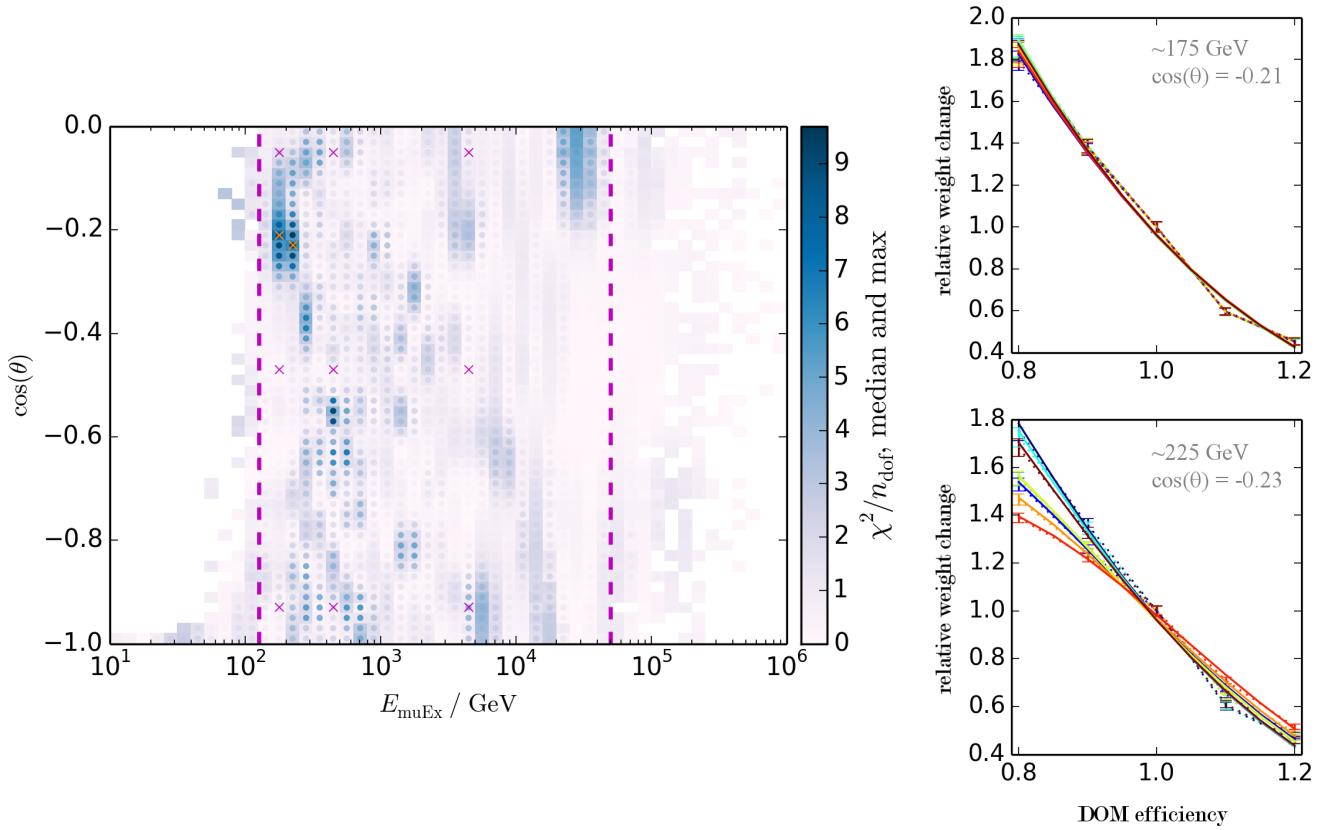


Figure 6.5: **Left:** χ^2/n_{DoF} distribution of the DOM efficiency parameterization, showing the median (field color) and the maximum (dot) of all event parameterizations per bin. Magenta lines indicate the used energy region, magenta crosses indicate the coordinates of figure 6.4, orange crosses indicate the coordinates of the right plots. **Right:** Ten randomly picked events each from the two bins with the highest median and maximum χ^2/n_{DoF} values, as in figure 6.4.

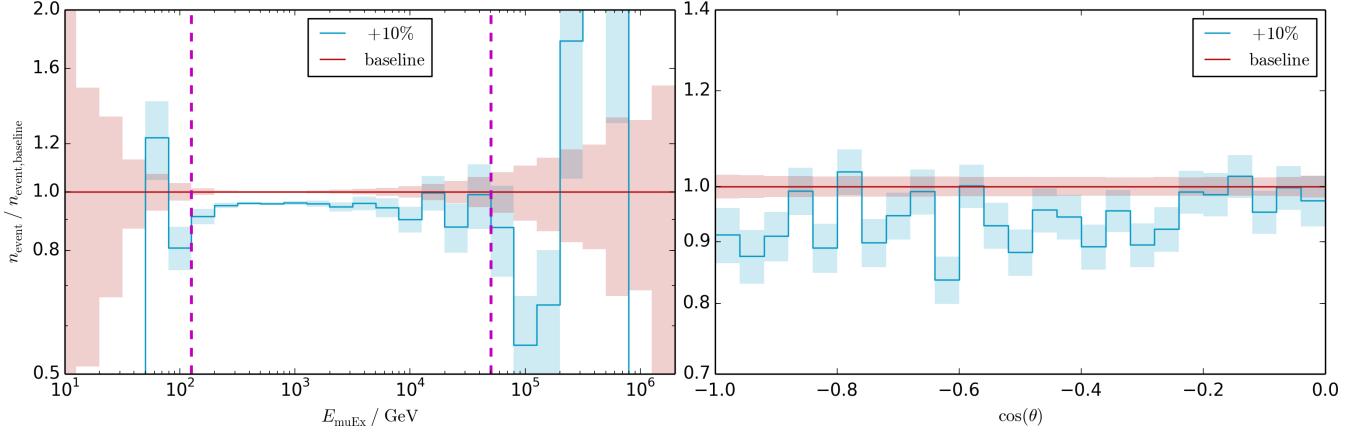


Figure 6.6: Ratio plot of reconstructed energy (**left**) and zenith angle (**right**) for the two datasets used for the bulk ice scattering parameterization. Magenta lines indicate the energy range of the analysis.

Bulk Ice Scattering Length

For violet light around 400 nm, the effective scattering length is about five times shorter than the absorption length, and the assumed values vary much more strongly due to different dust concentrations (see figure 2.15). Nonetheless, the scattering length is much less important as systematic quantity for this analysis compared to the absorption length because it is highly correlated to the total flux and does not have a clearly distinguishing structure either in energy or zenith angle (figure 6.6). Therefore, the effective scattering length has been parameterized first, using only the baseline dataset and the +10% scattering one, such that the mixed dataset could then be corrected for scattering to be used for the absorption length parameterization.

As described for the DOM efficiency in section 6.2.5, the two datasets were used to create two-dimensional KDEs (with 64-fold bootstrapping) in reconstructed energy and zenith angle. The KDEs were normalized to the baseline dataset and evaluated at the baseline event coordinates to obtain two weights per event, and these were interpolated linearly (and extrapolated) to generate per-event correction factors for arbitrary scattering lengths.

Bulk Ice Absorption Length

To get a third data point for the parameterization of the bulk ice absorption length, the scattering length parameterization as described in section 6.2.6 was also done for the events of the mixed absorption and scattering dataset, and was used to apply a +7.6% scattering length correction to that dataset to get back to the baseline value for scattering. The baseline dataset, the +10% absorption dataset, and the resulting -7.1%

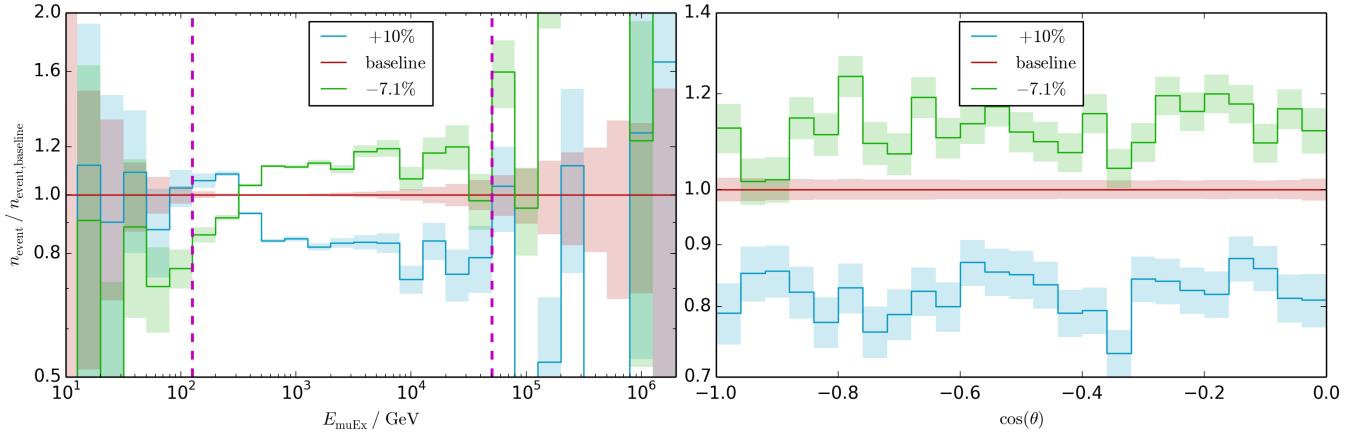


Figure 6.7: Ratio plots of reconstructed energy (**left**) and zenith angle (**right**) for the three datasets used for the bulk ice absorption parameterization. Magenta lines indicate the energy range of the analysis.

absorption dataset (shown in figure 6.7) were again used for two-dimensional KDEs (with 128-fold bootstrapping), and the normalized weights obtained for the baseline dataset event coordinates were fitted linearly to obtain the per-event parameterization.

Hole Ice Scattering Length

For this parameterization, two dedicated datasets were available at effective scattering lengths of 30 cm and 100 cm in addition to the 50 cm baseline dataset, again at 1/5 of the baseline MC statistics. The KDEs were generated using 64-fold bootstrapping, and the resulting normalized weights for the baseline dataset event coordinates were interpolated by a first-order spline, because the effect of hole ice is non-linear as can be seen in figure 6.8.

Overall, changing the scattering length in hole ice columns has little effect on this analysis as shown in figure 6.8. A shorter scattering length seems to favor reconstructed energies above $316 = 10^{2.5}$ GeV, but does not have a significant systematic impact on the zenith angle distribution at the level of the MC statistics. A better parameterization of this weak effect would require higher MC statistics, but since the effect is small and correlated with other nuisance parameters such as bulk ice absorption and DOM efficiency, the available parameterization is sufficient.

6.3 Prior Values

Table 6.1 summarizes the values of the Gaussian priors of all parameterized nuisance parameters used in the analysis. In general, the priors are large as explained in section 5.4

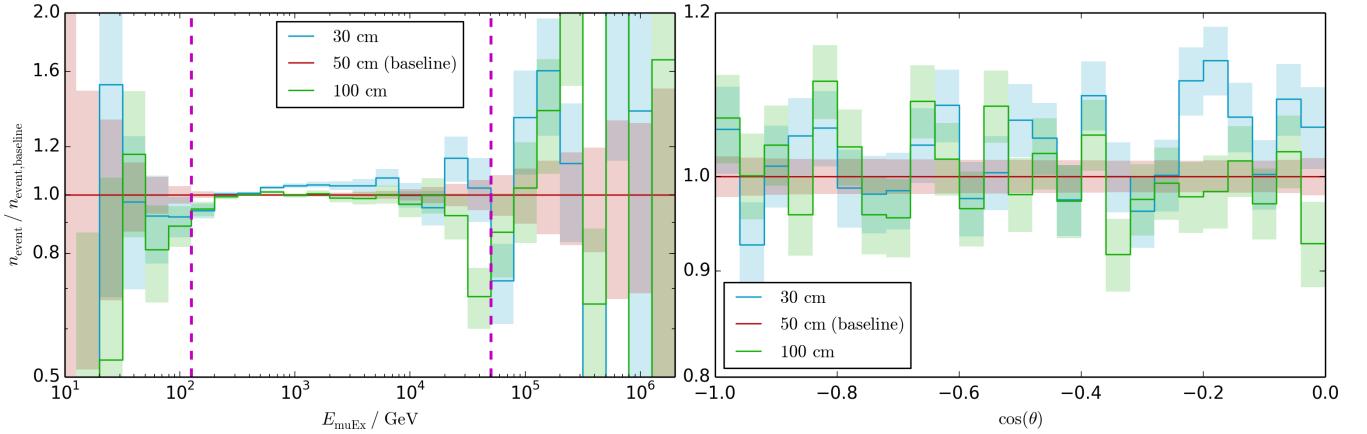


Figure 6.8: Ratio plots of reconstructed energy (**left**) and zenith angle (**right**) for the three datasets used for the hole ice scattering parameterization. Magenta lines indicate the energy range of the analysis.

Table 6.1: Values of nuisance parameter priors, pseudo experiment randomization parameters, and allowed ranges

parameter	μ_{prior}	σ_{prior}	μ_{pseudo}	σ_{pseudo}	allowed range		
total flux normalization	1.0	0.4	1.0	0.1	0.1	...	10
cosmic-ray spectral index	2.65	0.05	2.65	0.03	2.31	...	2.91
kaon-pion ratio	1.0	0.1	1.0	0.1	0.1	...	10
antineutrino-neutrino ratio w_1	~ 1.0	0.1	1.00	0.05	10^{-5}	...	6.66
antineutrino-neutrino ratio w_2	0.0	0.1	0.0	0.1	-0.3	...	0.3
DOM efficiency	1.0	∞	1.0	0.1	0.8	...	1.3
bulk ice scattering length	1.00	0.03	1.00	0.03	0.8	...	1.2
bulk ice absorption length	1.00	0.03	1.00	0.03	0.8	...	1.2
hole ice scattering length / m	0.5	0.1	0.5	0.1	0.1	...	1.3

to account for all systematic deviations between data and MC while ensuring convergence of the minimizer and non-unphysical values.

Also given are the ranges from which values are randomly sampled for the generation of pseudo experiments (which are often smaller than the very conservative priors), and the allowed ranges for the values. The latter limits are verified to be too large to be reached by the minimizer and to not affect the final minimization results, but realistic ranges can speed up the minimization because they can reduce the number of iterations far away from the true minimum (and affect the initial approximation of the inverse Hessian matrix in non-L-BFGS-B methods).

Chapter 7

Sensitivity and Results

Sensitivity plots of older versions of this analysis have previously been presented at conferences and have accordingly been published in proceedings by the IceCube Collaboration (ref. [120] and [121], also published in ref. [122] and [123], respectively). The publications have been written as corresponding author. Everything here is work of the author unless specifically noted.

Very conservative exclusion limits have been published in ref. [152], citing this work. See section 7.4 for details.

7.1 Sensitivity and Full-Circle Tests

7.1.1 Asimov Sensitivity Estimates

The Asimov estimate for this analysis’ sensitivity is shown in figure 7.1 and was mentioned in section 5.5 to justify the choice of binning for the parameter scan. Asimov estimates are named after the biochemist and science fiction author Isaac Asimov as the method was inspired by his short story *Franchise* [153]. The idea behind it is to replace the ensemble of pseudo experiments that is normally used to compute the analysis sensitivity by a singular (*Asimov*) dataset that is representative to the whole ensemble. It can be shown that the dataset that reproduces the true input parameters when being plugged into the analysis is the one that also reproduces the median LLH ratio for every parameter [154]. In practice, that means that simply running the analysis on the full (weighted) MC histograms without an interposed Poisson process – essentially comparing them against themselves – produces the median sensitivity as result of the LLH scan.

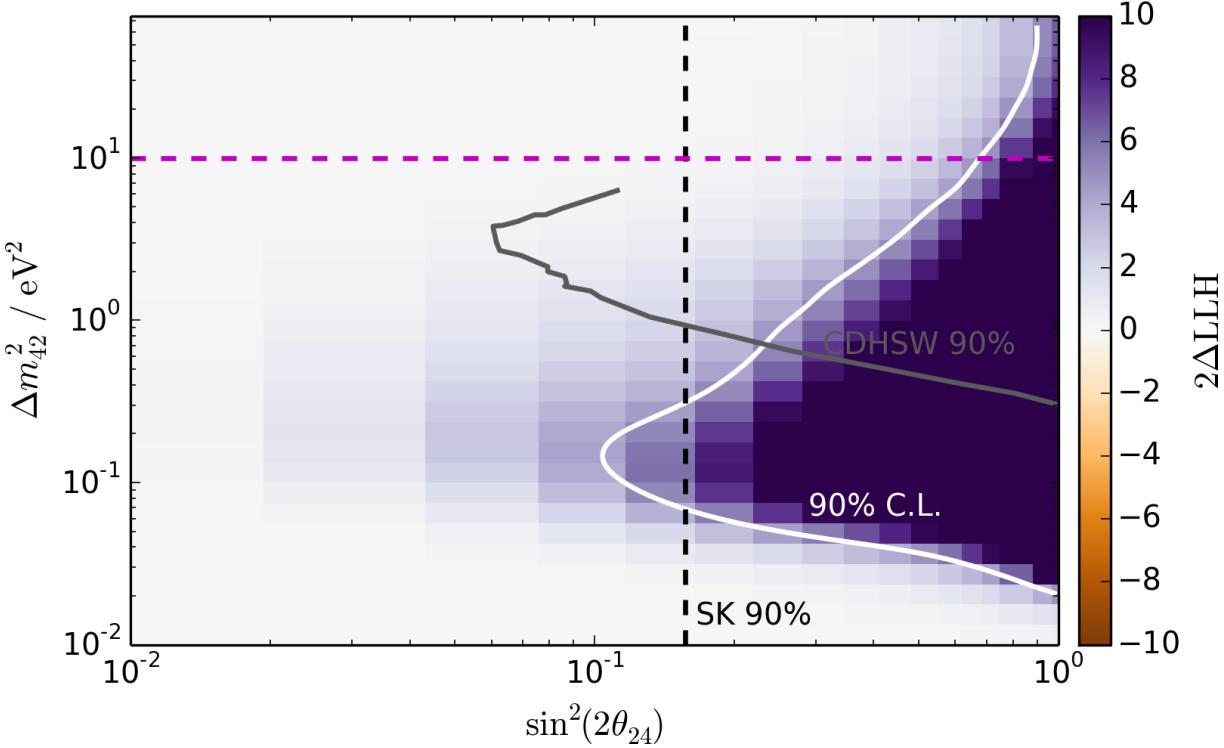


Figure 7.1: Asimov estimate for the analysis' sensitivity, i.e., for the log-likelihood landscape of a null result, for the extended parameter space (see section 5.5). Lines are exclusion limits at 90% confidence levels (C.L.) for this analysis (white, using Wilks' theorem), CDHSW [57], and Super-Kamiokande¹⁶ (black) [155].

The method greatly reduces the computational effort of computing median analysis sensitivities and uses only insignificant approximations; generally it produces a better estimate than, e.g., the Fisher information matrix [154]. A caveat is that the method usually has to be combined with using Wilks' theorem to calculate sensitivities with the limitations mentioned in section 5.3. Therefore, Asimov estimates are very useful for relatively quick tests during the development of the analysis, but final estimates are based on ensembles of pseudo experiments.

7.1.2 Ensembles of Pseudo Experiments

Null Hypothesis

For the final analysis, 300 pseudo experiments were generated for the null hypothesis to get a reasonably good conversion curve between ΔLLH and analysis sensitivity. These pseudo

¹⁶The Super-Kamiokande limit was derived from a limit on $|U_{\mu 4}|$ under the optimistic assumption of $\theta_{14} = 0^\circ$. It would deteriorate down to $\sin^2(2\theta_{24}) = 0.30$ for maximal θ_{14} .

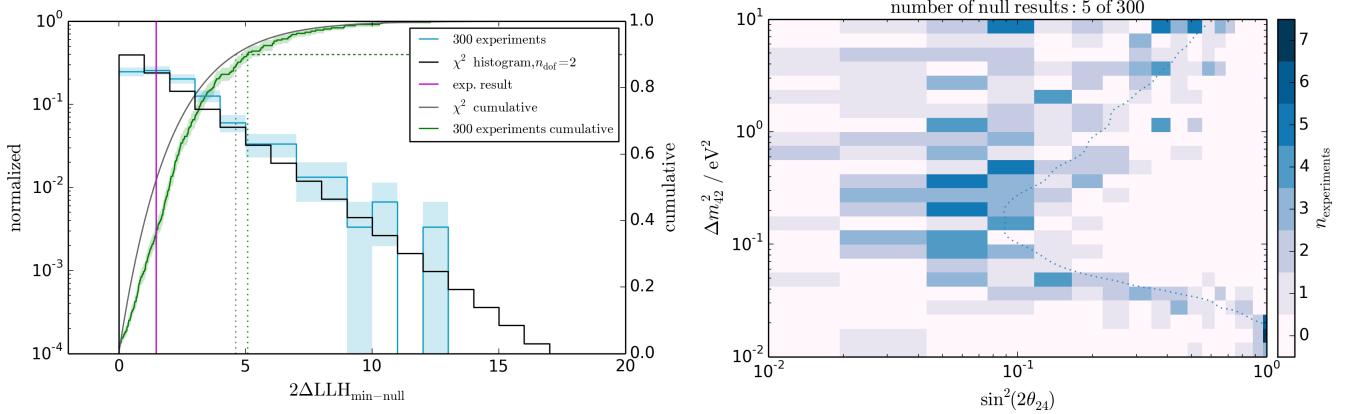


Figure 7.2: Distributions of the test statistic values (2 Λ , **left**) and best-fit positions (**right**) of 300 pseudo experiments for the null hypothesis.

experiments were generated with randomized true values for the nuisance parameters according to the values given in table 6.1. Figure 7.2 shows their distribution of 2Λ (as in (5.5)) where Ξ_0 is the parameter space of the null hypothesis, and compares it to the approximation of Wilks' theorem. The cumulative distribution lies slightly to the right side of the Wilks curve almost everywhere, which was observed also in all other iterations of this analysis' development, which means that assuming Wilks' theorem slightly overestimates sensitivities and significances. Errors were estimated by bootstrapping (1000-fold, see appendix E).

Figure 7.2 also shows the best-fit positions for the null hypothesis, which are clustered around the one-sigma significance level. This behavior is unproblematic and can be explained by the larger parameter space of this region: The hypotheses to the left of this region differ only weakly from each other or the null hypothesis. All pseudo experiments are expected to have both upward and downward fluctuations in some bins, and almost always a small cluster of fluctuations on the 1σ level can be found that coincides weakly with a signal hypothesis, given the very varied signatures of sterile neutrinos. As the fluctuations do not follow the pattern of sterile signal, the goodness of fit and therefore significance are low, but the hypothesis still matches slightly better than the null hypothesis.

The analysis sensitivity estimate obtained from this ensemble is shown together with experimental results in figure 7.8.

Signal Hypotheses

100 pseudo experiments each were generated for three signal hypotheses for full-circle tests and to get conversion curves at other points of the sterile-neutrino parameter space. The hypotheses are

- weak: $\Delta m_{42}^2 = 0.656 \text{ eV}$, $\theta_{24} = 7^\circ$, $\Lambda_{\text{Asimov}} = 0.509 \hat{=} 0.317 \sigma$,
- intermediate: $\Delta m_{42}^2 = 0.066 \text{ eV}$, $\theta_{24} = 19^\circ$, $\Lambda_{\text{Asimov}} = 5.789 \hat{=} 2.935 \sigma$, and
- strong: $\Delta m_{42}^2 = 0.207 \text{ eV}$, $\theta_{24} = 23^\circ$, $\Lambda_{\text{Asimov}} = 13.396 \hat{=} 4.808 \sigma$,

classified by the sensitivity estimates obtained from the Asimov sample. Note that the largest significance had to be obtained from Wilks' theorem because of insufficient statistics in the null hypothesis sample. Figures 7.3, 7.4, and 7.5 show the distributions of $2\Lambda_{\text{sig}}$, i.e., 2Λ against the true signal hypothesis instead of null, and the best-fit positions of the signal hypothesis ensembles.

The curves of $2\Lambda_{\text{sig}}$ match the Wilks curve much better than 2Λ did for the null hypothesis. This is not only true for these three signal hypotheses, but also for all other non-null hypotheses tested in previous iterations of the analysis. The reason for this is that as pointed out in section 5.3, Wilks' theorem is not strictly applicable for the null hypothesis, but it is for signal hypotheses as these are well identifiable. Since a full Feldman-Cousins approach [156] that gets correct coverage for all hypotheses is computationally prohibitive for this analysis, the conversion curve obtained for the null hypothesis is used for all hypotheses as it gives the most conservative results.

The positions of the best fits cluster around the injected true parameters as expected, with sharper clustering for stronger hypotheses. This means that the full-circle test is passed successfully: the analysis correctly identifies sterile signal hypotheses, and more importantly does not erroneously exclude true signal hypotheses. The dashed orange lines indicate the median contours of the pseudo experiments. They are closed for the two hypotheses inside the 90% C.L. sensitivity estimate, which shows that the null hypothesis is usually excluded at 90% C.L. in these cases; the numbers of experiments that exclude null are 35/100 (weak), 96/100 (intermediate), and 100/100 (strong). Accordingly, for the weak hypothesis, the contour is merely pushed away from the true parameters.

It is noteworthy that the median contour of the weak hypothesis becomes stronger at very low and very high Δm_{42}^2 compared to the estimated sensitivity. This is natural behavior and happens because the injected signal creates a valley in the LLH landscape that is deeper than the null hypothesis, so unrelated (distant) hypotheses are excluded more strongly; in other words, they are more similar to the null hypothesis than to a very different signal hypothesis.

7.2 Best Fit, Resulting Limits, and Checks

The likelihood contour obtained from applying the analysis on the experimental data is shown in figure 7.6. The position of the best fit is the scan point of $\theta_{24} = 13^\circ$ and $\Delta m_{42}^2 = 0.0369 \text{ eV}^2$ with $\Lambda = 0.733$, which corresponds to a p-value of $64.3\% \hat{=} 0.463 \sigma$ ($23.1\% \hat{=} 0.706 \sigma$ using Wilks' theorem). Using interpolation, the best fit shifts slightly to $\theta_{24} = 12.08^\circ$ and $\Delta m_{42}^2 = 0.0401 \text{ eV}^2$. The result is compatible with the null hypothesis and lies outside the 90% C.L. of the expected sensitivity.

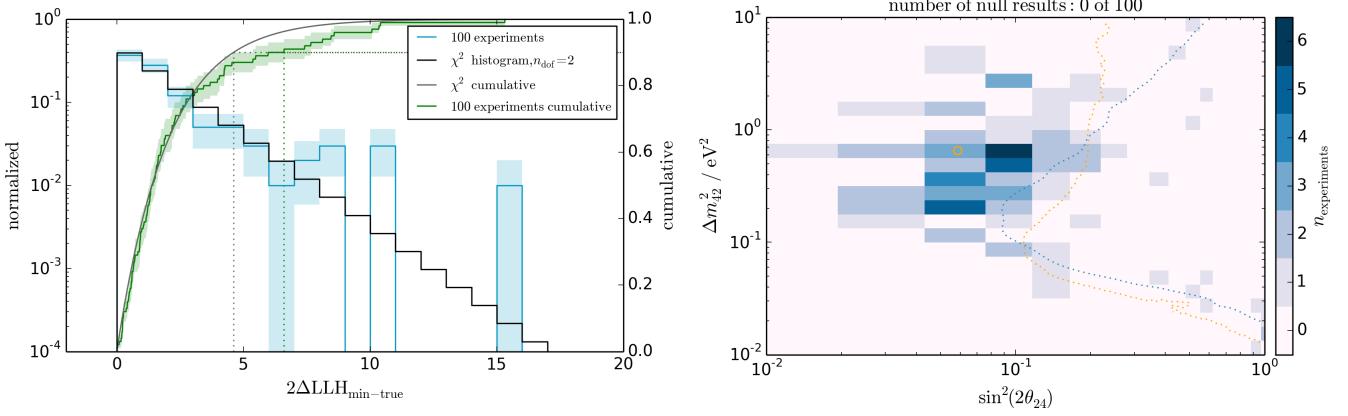


Figure 7.3: Distributions of $2\Delta\text{LLH}_{\text{min-true}}$ (left) and best-fit positions (right) of 100 pseudo experiments for a weak signal hypothesis (marked by the orange circle).

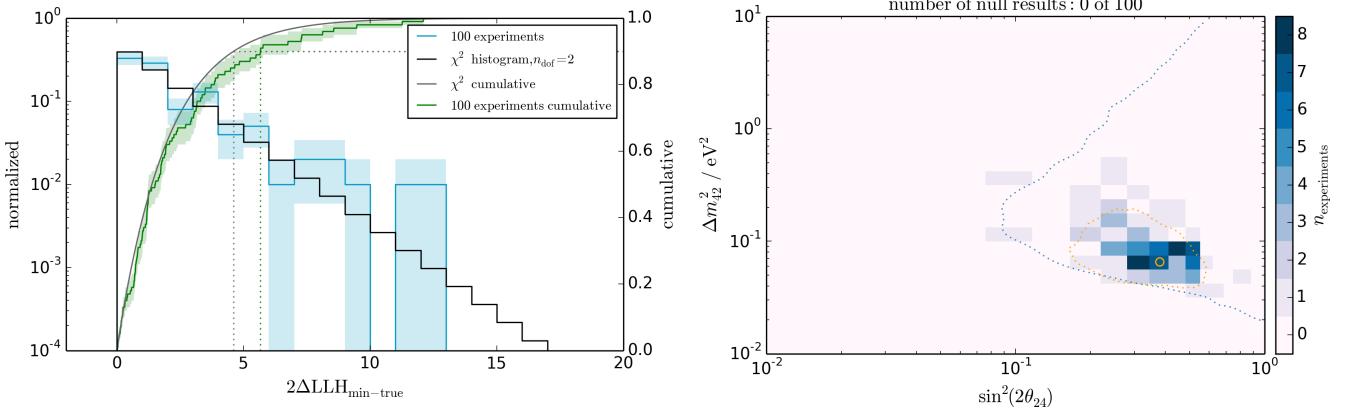


Figure 7.4: Distributions of $2\Delta\text{LLH}_{\text{min-true}}$ (left) and best-fit positions (right) of 100 pseudo experiments for an intermediate signal hypothesis (marked by the orange circle).

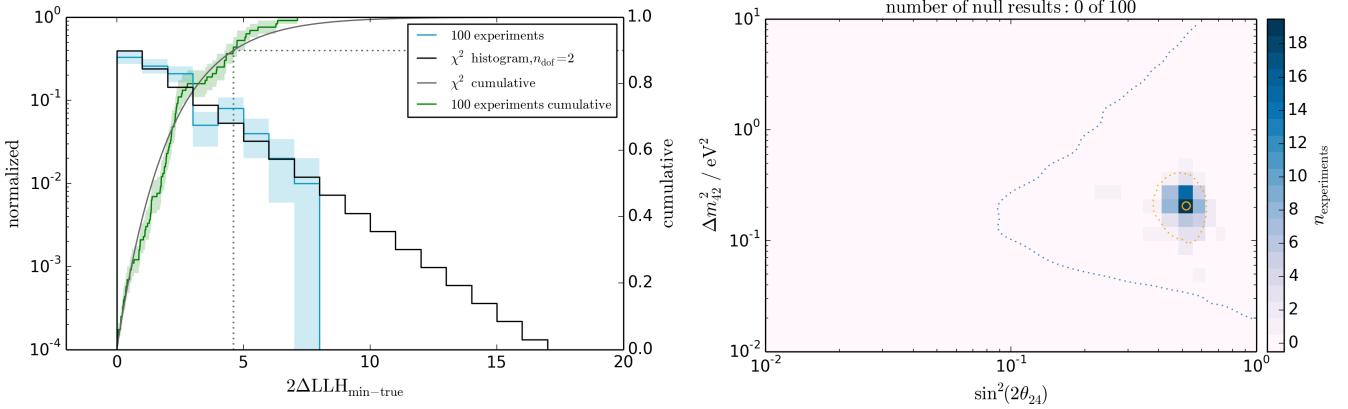


Figure 7.5: Distributions of $2\Delta\text{LLH}_{\text{min-true}}$ (left) and best-fit positions (right) of 100 pseudo experiments for a strong signal hypothesis (marked by the orange circle).

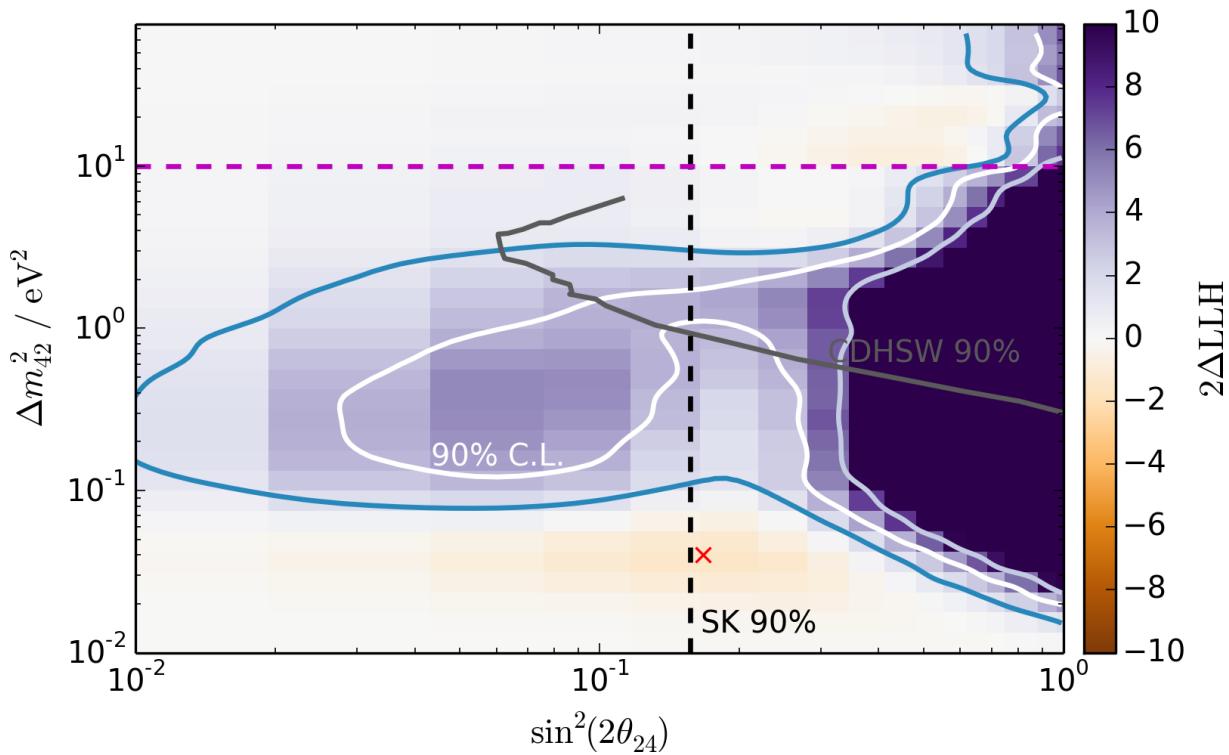


Figure 7.6: Likelihood landscape and 90% C.L. exclusion contour for the experimental data with the extended parameter space, analog to figure 7.1. Additionally shown are 68.3% (1 σ , blue) and 99% (light blue) C.L. exclusion contours.

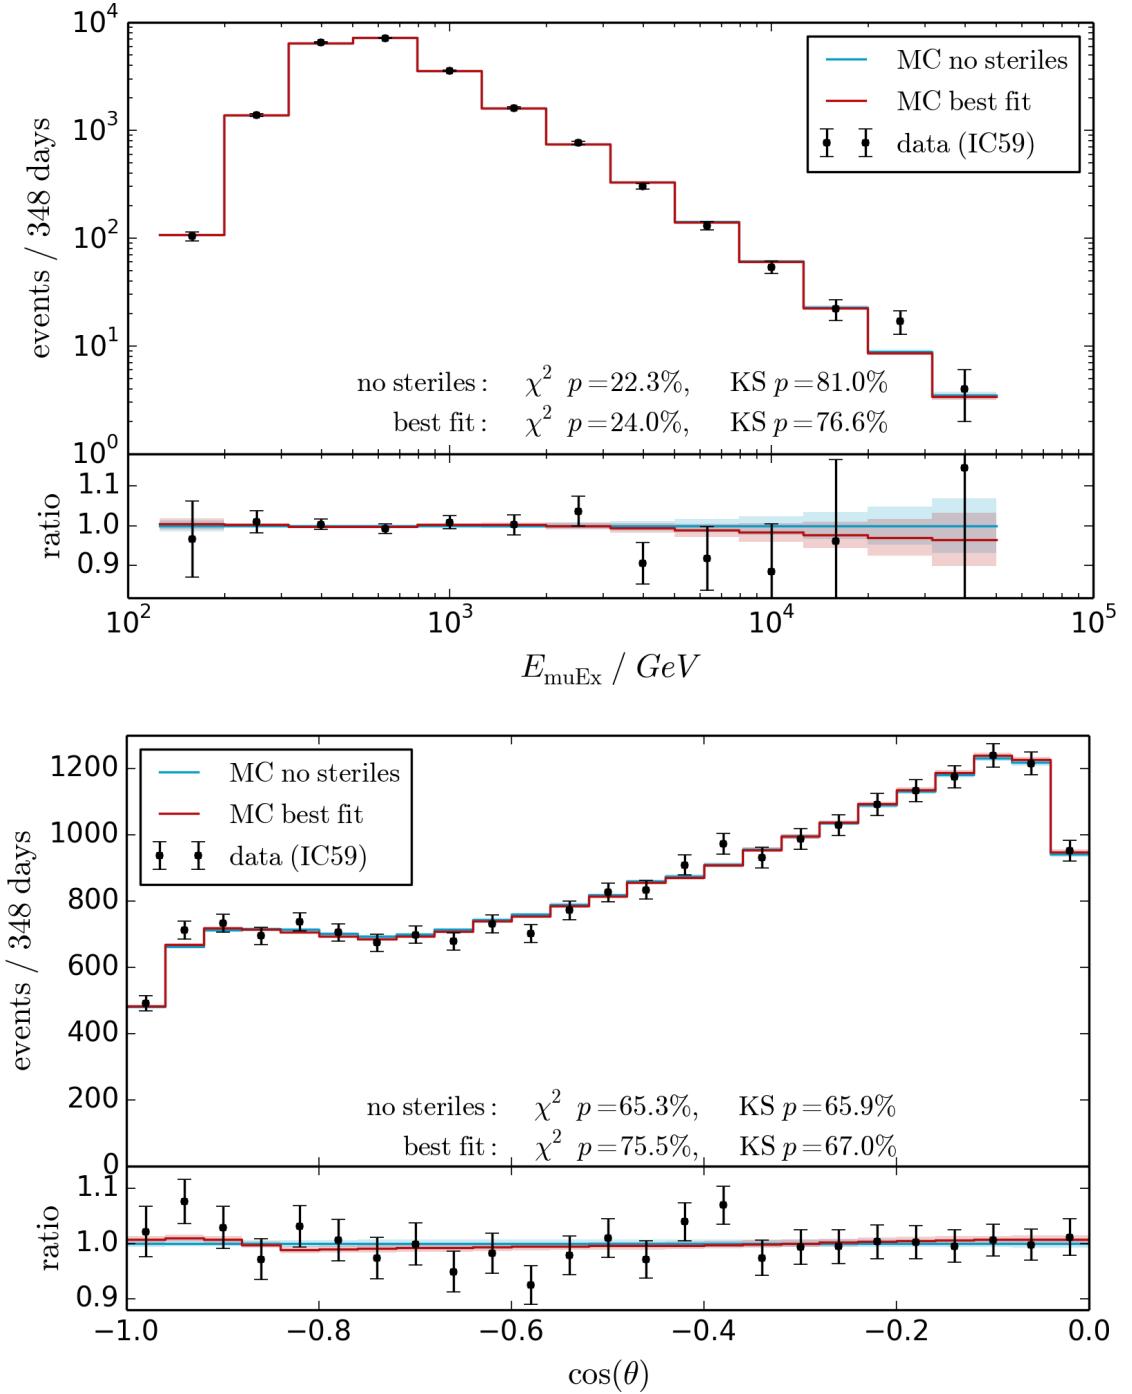


Figure 7.7: Energy (top) and zenith angle (bottom) distributions of the experimental data compared to the best-fit MC predictions for the null hypothesis (“no steriles”) and the best fit, including numbers for goodness of fit. The ratio plots on the bottom are relative to the null hypothesis.

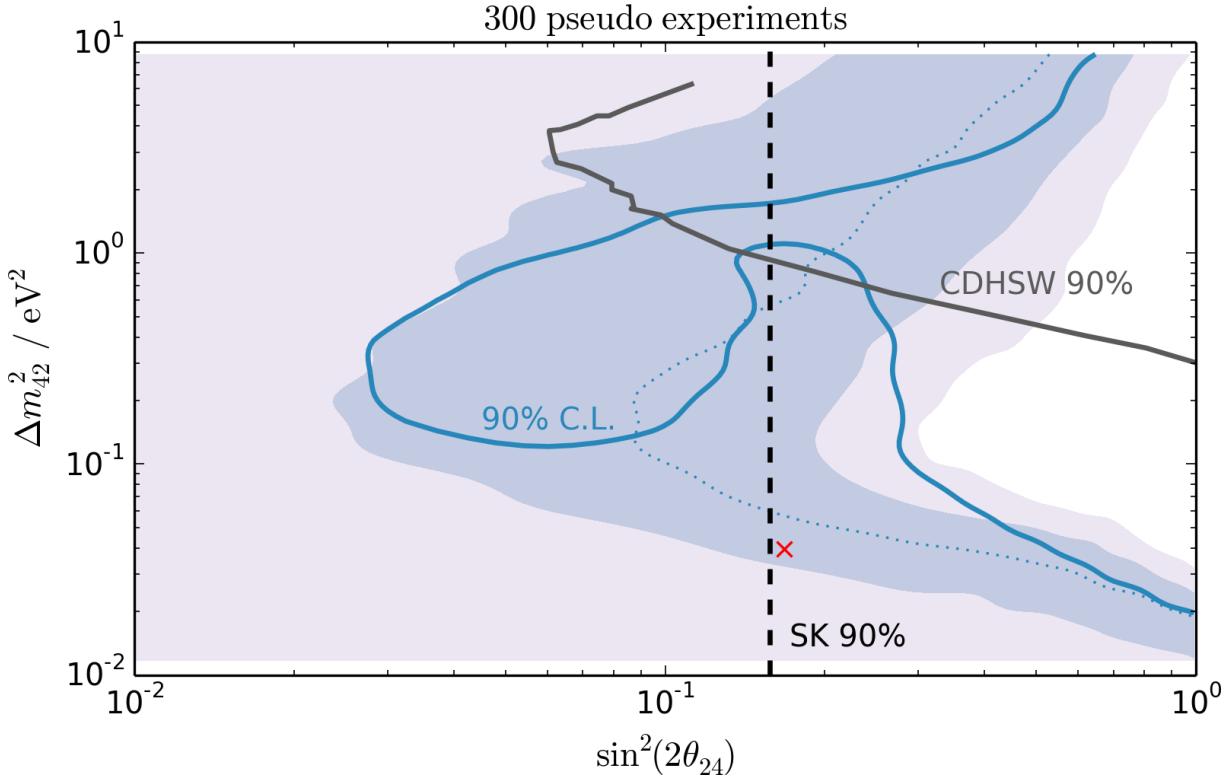


Figure 7.8: Experimental contour and distribution of contours from null hypothesis pseudo experiments. Light blue indicates the central 95.5% (2σ) percentile, dark blue the central 68.3% (1σ) percentile, and the dashed line marks the median sensitivity from pseudo experiments. See also appendix G.

The expected significance of a best fit at this position from the Asimov estimate is $\Lambda = 0.666$, which indicates that there is not much tension in the fit. This is confirmed by the projected energy and zenith angle distributions, which both feature a relatively good goodness of fit as shown in figure 7.7. As measures for the goodness of fit, both χ^2/n_{dof} (χ^2) and a Kolmogorov-Smirnov (KS) test were used. The χ^2 test was only applied to bins with at least 5 predicted events because otherwise its assumption that the numbers of entries follow a normal distribution is not fulfilled well [157]; the KS test is an unbinned test of the underlying cumulative distributions and therefore complements the χ^2 test by being sensitive to deviations on a multi-bin scale [158]. The difference between the numbers for the null hypothesis and the best fit are small because the hypotheses are fairly similar as indicated by the low significance, and also because the nuisance parameters are fitted independently for both hypotheses, so they further absorb the differences.

Figure 7.8 shows the experimental contour with the ensemble of contours of the null hypothesis pseudo experiments. The algorithm used to create the plot is explained in appendix G. This plot shows that the experimental result is fairly strong between $\Delta m_{42}^2 = 0.1 \text{ eV}^2$ and $\Delta m_{42}^2 = 1 \text{ eV}^2$ for small mixing angles ($\theta_{24} \lesssim 10^\circ$) compared to the analysis

sensitivity, but it has large indent around $\theta_{24} \approx 13^\circ$ towards low values of Δm_{42}^2 , because that is where the best fit lies. Indents are not uncommon for this analysis; roughly 20% of the pseudo experiments feature indents, not counting experiments that exclude the null hypothesis, and about 6% have indents of a similar size or sharpness, although this is a subjective assessment. A sample of randomly selected pseudo experiment contours can be seen in figure G.1.

Even in this strongly excluded region, the experimental contour is only negligibly stronger than a 1σ upward fluctuation of the pseudo experiment ensemble.

7.2.1 Pull Distribution

Pulls are deviations of data from reference data in units of the (Poissonian) standard deviation of that reference data, i.e., $(x - x_{\text{ref}})/\sqrt{x_{\text{ref}}}$. They are often more useful than residuals (i.e., $x - x_{\text{ref}}$) to visualize deviations in histograms when the statistical power changes strongly between bins. The main goal of these plots is to check for an uneven distribution of pulls that could indicate unhandled systematic uncertainties.

Figure 7.9 shows the pulls of the experimental data compared to the null hypothesis and to the best fit. Again the differences between both plots are very small for the same reasons given above. Pulls were computed without taking into account the statistical uncertainties on the MC data, so their values are slightly overestimated. Some bins are noteworthy because of values over 3σ (and sometimes over 4σ). All but one of these bins lie at reconstructed energies above 20 TeV, where MC expectation values are very low and singular experimental events can cause such excesses (compare figure 5.1). The large pull in the middle of the plots is caused by 53 experimental events at expectation values around 31 events; while this is significant, no systematic reason for this excess (such as previously unidentified background) could be found. The result is not significantly influenced by this bin, as it does not correspond to a signature of sterile neutrinos.

A more elaborate approach of analyzing the pulls is to search for clusters of excesses or deficits. The method was originally devised by David Jason Koskinen during the discussion of the analysis within the IceCube collaboration, and later improved by him, Sebastian Böser, Ken Clark, Benjamin Jones, and the author of this work. Events and MC predictions for the null hypothesis of all possible rectangular clusters of bins are added up and compared with each other, assuming Poisson statistics, and the resulting pulls are histogrammed and compared to pseudo experiments. An advantage of this method is that clusters can be visualized to search for hotspots of bad data-MC agreement, although it has to be noted that these clusters are highly correlated and do not necessarily indicate problems. Only clusters with a MC expectation of larger than 2.1 events are taken into account, because at lower expectation values the distribution of Poissonian pulls becomes strongly non-Gaussian [157].

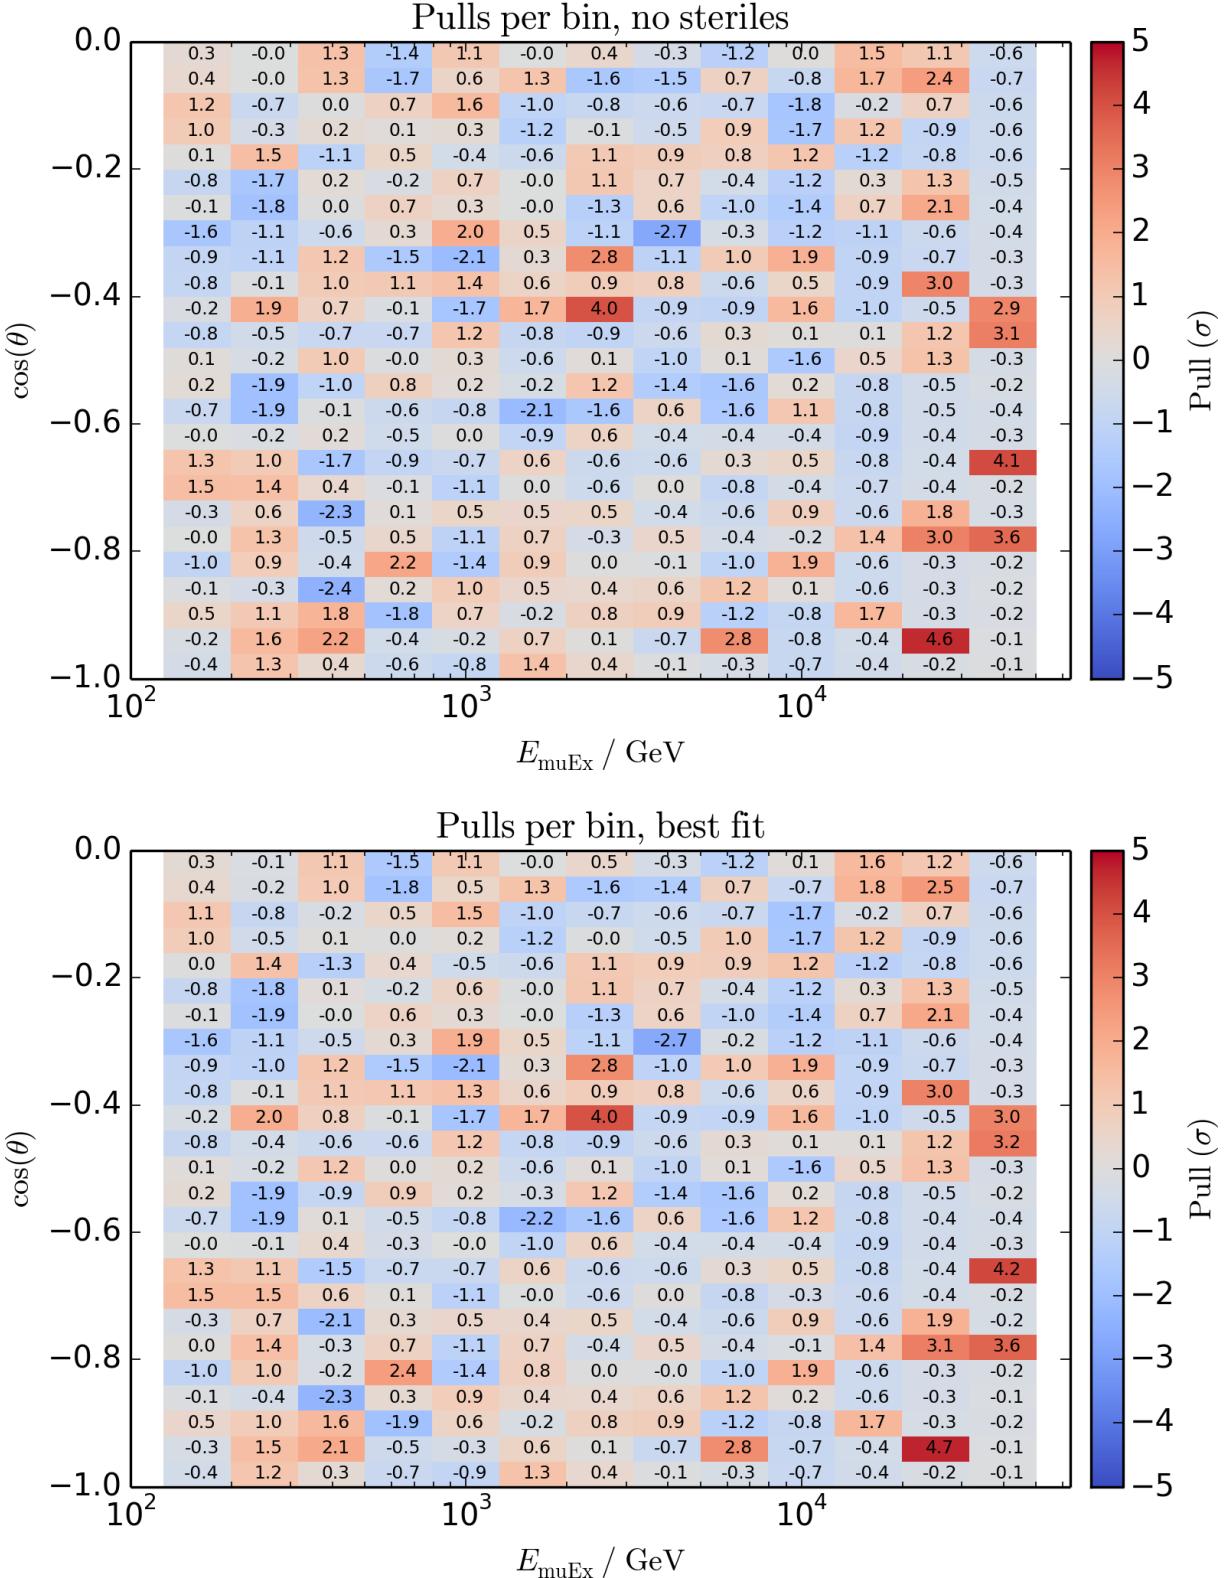


Figure 7.9: Pulls of the experimental data (i.e., deviations in units of Poissonian standard deviations) compared to the null hypothesis (top) and to the best fit (bottom). MC statistics are not taken into account.

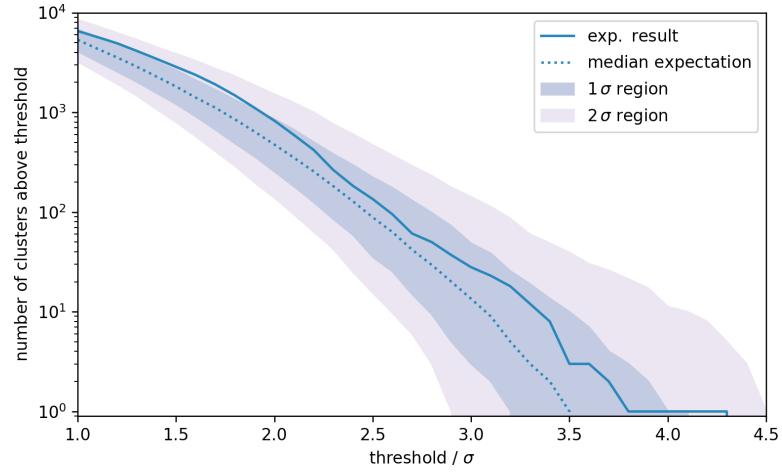


Figure 7.10: Distribution of the number of rectangular bin clusters exceeding a given Poissonian pull value vs. pull thresholds, showing central percentiles of the 300 pseudo experiments.

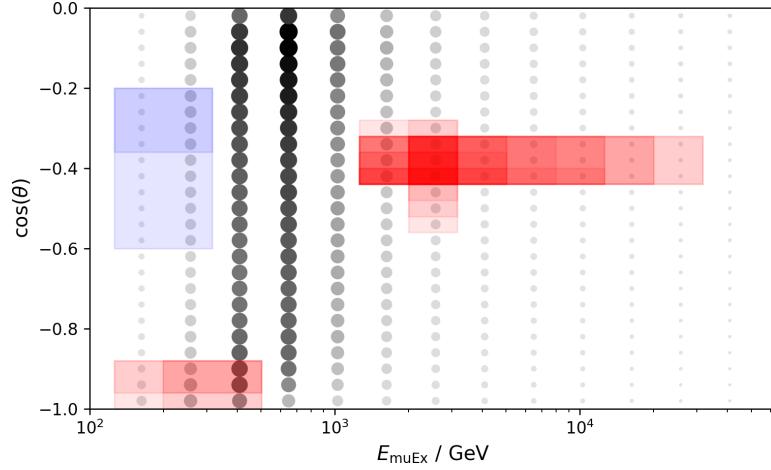


Figure 7.11: Positions of the clusters above the 3.0σ pull threshold in experimental data; red indicates excesses of experimental data compared to the expectations, blue indicates deficits, gray dots show the MC expectations for orientation.

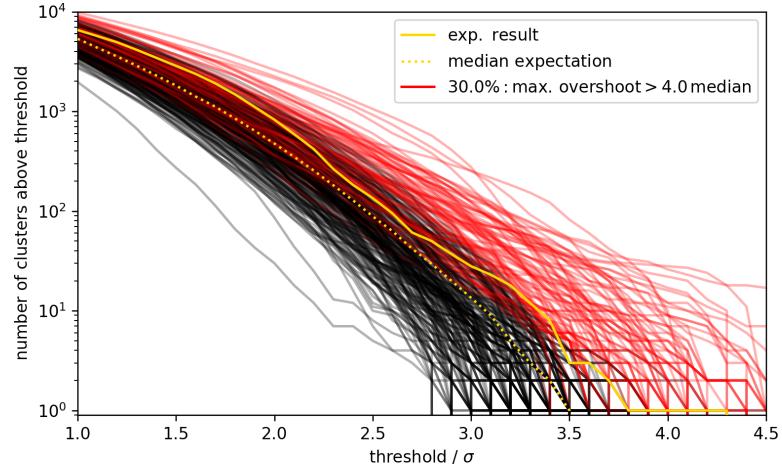


Figure 7.12: Same as figure 7.10, but with individual lines for all pseudo experiments; lines exceeding the maximum relative experimental deviation from the median of 4.0 are drawn in red.

Figures 7.10 to 7.12 show the results of this approach. The same method was applied to the 300 fitted pseudo experiments for comparability. The line of the experimental data lies above the median of the pseudo experiments, but only barely exceeds the $+1\sigma$ region around 1.7σ . Because of the high correlation between the clusters, the line is not expected to lie equally below and above the median.

The largest relative deviation from the median is a factor of 4.0 at 3.3σ and 3.4σ , but as shown in figure 7.12, 30% of the pseudo experiments have larger relative deviations than the experimental data at some point (lines drawn in red), so this is deemed unproblematic. Finally, figure 7.11 shows the distribution of the clusters that exceed $\pm 3.0\sigma$. The largest group of clusters is centered around the large excess that was mentioned above. While large, the prevalence of even larger deviations in 30% of the pseudo experiments asserts the assessment that this excess is unproblematic. A smaller group of excesses can be found in the vertically upgoing low-energy region. In previous iterations of this analysis (see section 7.3), this excess was much stronger and led to the discovery and removal of the unsimulated background discussed in section 3.2. The remaining excess in this region was suspicious, but the cuts and the remaining events were carefully checked in various variables and by eye using an event viewer, and no anomalies were found, leading to the conclusion that the remaining excess is statistical or physical. Lastly, a small cluster of deficits can be found at non-vertical low-energy events. Again, no systematic problems were found in this region.

7.2.2 Nuisance Parameters and $n - 1$ Plots

Figure 7.13 shows the values of the nuisance parameters for all scanned sterile-neutrino hypotheses, i.e., scan points. Overall, the plots feature smooth landscapes that do not raise concerns regarding failed fits. The only discontinuities visible in some parameters lie at the transitions between discrete atmospheric models (and are therefore physically motivated), or are located at very high mixing angles and are caused by jumps in the bulk ice absorption parameter that reacts to chaotic oscillation patterns at those extreme mixing parameters.

The prevailing discrete atmospheric model (described in section 6.1.1) is HKKMS2006, which encompasses both the null hypothesis and the best fit (469 scan points plus null hypothesis). The only other relevant model is poly-gonato with QGSJet II.04, which is preferred over much of the low-mass edge of the final contour (85 scan points). A third model, poly-gonato with Sibyll 2.3rc1, is preferred in four scan points at very large mixing angles and has no influence on the resulting contour.

The continuous nuisance parameters show no surprises. The flux normalization is increased at very high mixing angles and masses because these models flatly reduce the event rate as discussed at the end of section 4.3. The cosmic-ray spectral index is increased towards the HKKMS2006 default value, and is somewhat correlated to the signal strength.

7.2. BEST FIT, RESULTING LIMITS, AND CHECKS

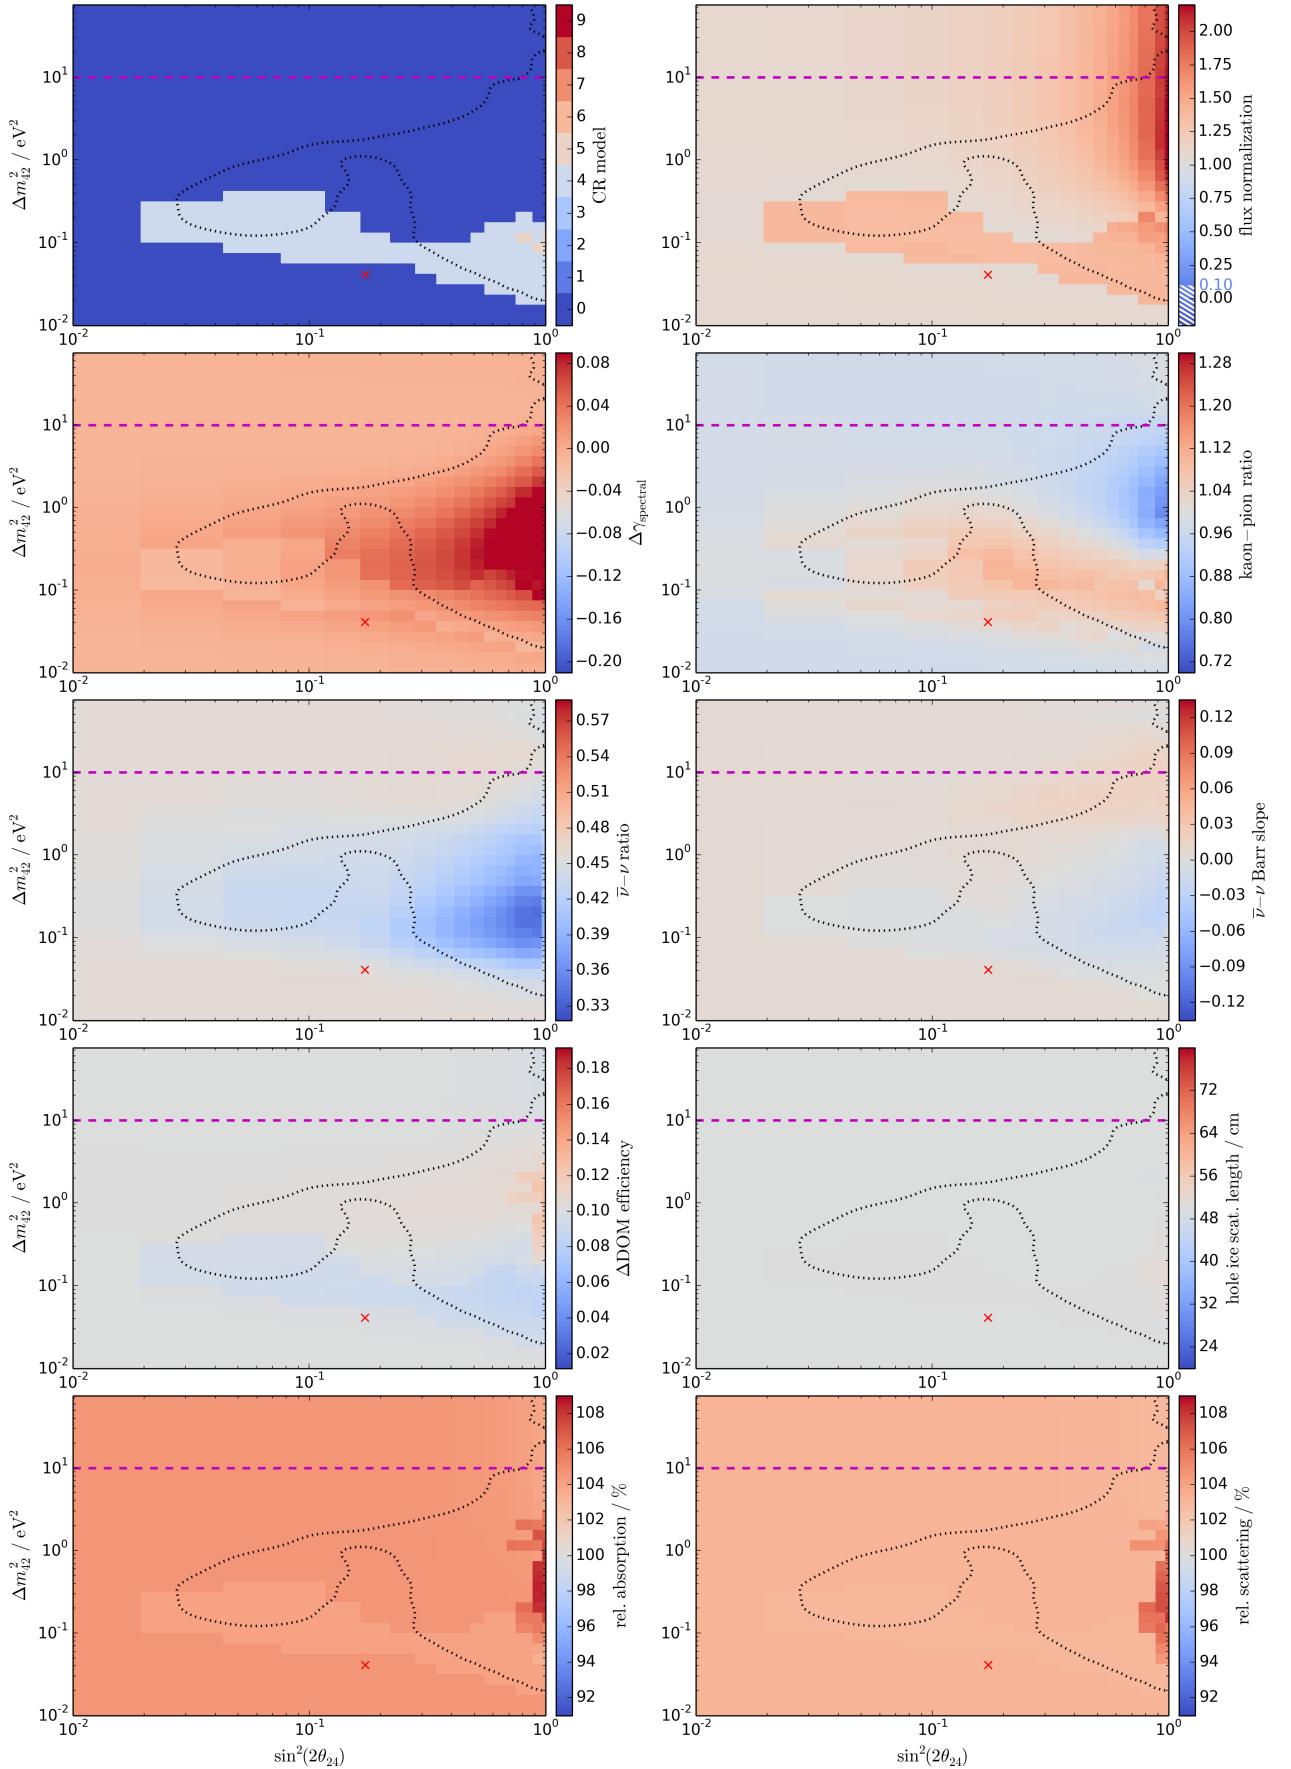


Figure 7.13: Values of all nuisance parameters over the full scanned parameter space.

The z-axes correspond to $-3\sigma \dots 3\sigma$ priors where applicable (see table 6.1). 125
 The dotted line indicates the experimental 90% C.L. contour for orientation.

The antineutrino ratio (with both parameters) is slightly reduced in regions of high signal strength because this weakens the resonance part of the signature, which is not present in experimental data. The kaon-pion ratio and the DOM efficiency are anti-correlated and differ between the low-mass edge and the high-mass/high-mixing-angle region.

Finally, the ice parameters are flat almost everywhere, with their values only changing at very high signal strengths in a relatively unsteady manner. This indicates that their parameterizations might not be smooth enough for the minimizer to freely leave local LLH minima because of the relatively low statistics of the ice parameter MC datasets as discussed in section 6.2.6. However, the fact that there are variations at high signal strengths or different atmospheric models shows that the parameterizations are good enough to allow variations if those were reasonably beneficial for the fit, so it can be concluded that the ice parameters are described adequately and do not influence the fit greatly.

Figure 7.14 shows how the contours change when individual nuisance parameters are removed, i.e., fixed to their default values; these plots are internally dubbed $n - 1$ plots. The results are consistent with the plots and discussions above. The noteworthy plots are those for the total normalization, for the spectral index $\Delta\gamma$, and for the DOM efficiency, as they show strong effects. For the latter two, one reason for this is that their best-fit values are relatively far away from their respective default values (see figure 7.13; note that the default for Δ DOM efficiency lies at 0.0), so a large change can be expected when removing these parameters. The large deviations are unproblematic as the best-fit values are well compatible with theoretical predictions (sections 6.2.2 and 6.2.5) and are not strongly influenced by the priors (1.37 σ for $\Delta\gamma$, no prior for the DOM efficiency).

A fixed spectral index moreover strongly influences the shape of the likelihood landscape; the best fit is shifted to $\Delta m_{42}^2 = 0.024 \text{ eV}^2$ and $\theta_{24} = 1.9^\circ$ (which lies outside the plotting range that starts at $2.9^\circ \doteq \sin^2(2\theta_{24}) = 0.01$). This completely removes the indent and leads to a very strong exclusion, but is strongly disfavored compared to the solutions with $\Delta\gamma$ enabled ($\Delta\text{LLH} > 4.1$).

A fixed total normalization unsurprisingly leads to a strong exclusion of high-mass hypotheses, becoming significant for $\Delta m_{42}^2 \gtrsim 0.3 \text{ eV}^2$, because these hypotheses lead to a washed out disappearance effect over large energy regions as explained in chapter 4.

7.3 Development

The development of this analysis was long and not entirely straight-forward. As usual within the IceCube collaboration, this analysis was developed blindly, i.e., without applying it to experimental data other than a predefined low-statistics sample (about 10% of the data in this case) that can in the late stages of development be used to check for problems and data-MC mismatches, but which is not large enough to produce significant results.

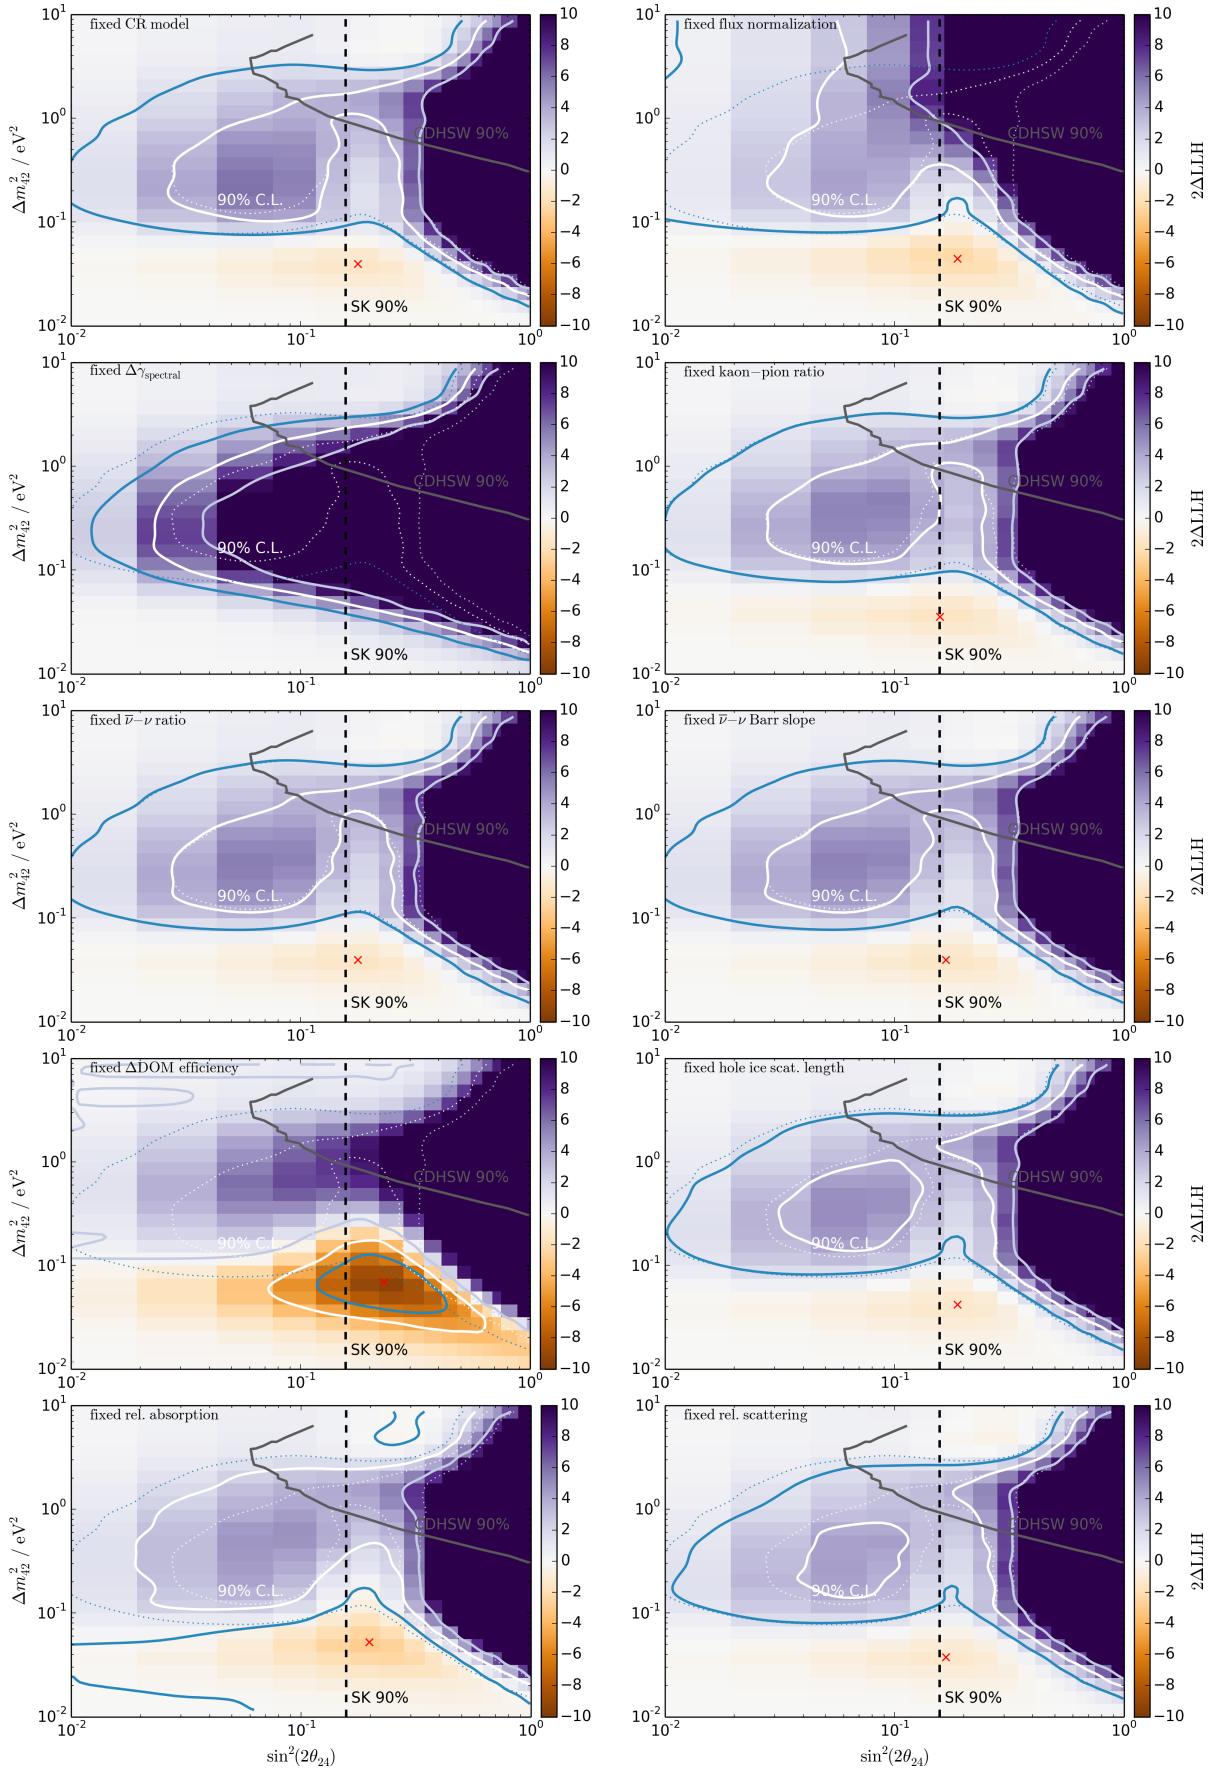


Figure 7.14: Likelihood landscapes with one nuisance parameter each fixed to the default value (called $n - 1$ plots).

The faint dotted lines indicate the experimental contours for orientation.

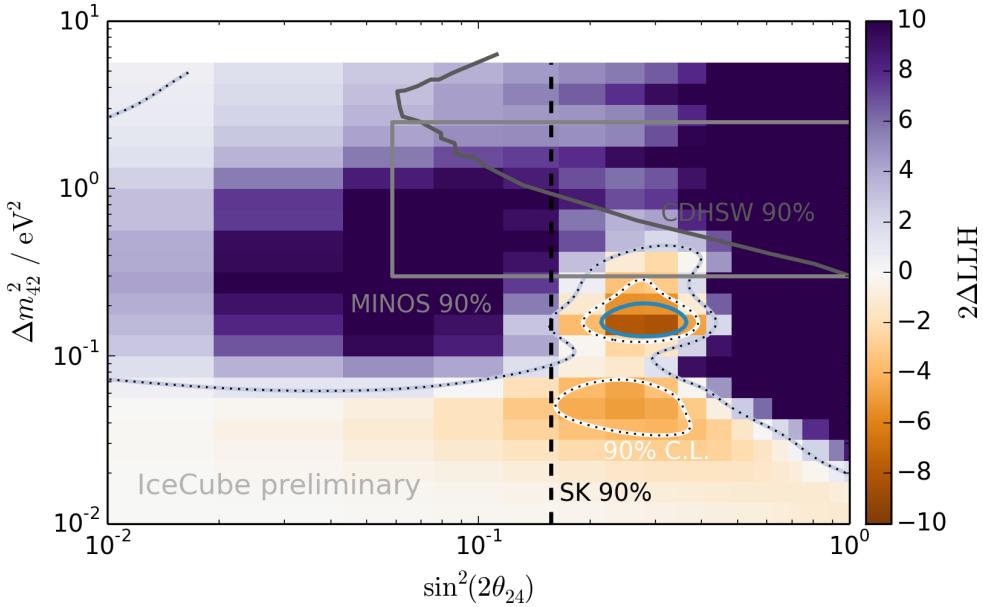


Figure 7.15: Likelihood landscape of the initial analysis result, see text. Dotted black lines have been added to increase contrast. The MINOS line is from ref. [159].

The analysis had originally been worked out up to the point of unblinding in March 2014. The result back then (figure 7.15) was $\Delta m_{42}^2 = 0.156 \text{ eV}^2$ and $\theta_{24} = 17^\circ$ at a significance of roughly 2.75σ according to 300 pseudo trials (2.42σ according to Wilks' theorem), but the data-MC agreement was so bad that neither the null hypothesis nor the best fit described the experimental data well. This was surprising because the experimental data, the MC data, and the handling of the systematic uncertainties of the original analysis were very similar to those of the search for a diffuse astrophysical neutrino flux [114], which did not suffer from these problems. However, the measure of the goodness of fit (g.o.f.) in that analysis was a two-dimensional saturated Poissonian test, which was not sensitive enough to capture the full extent of the mismatch due to small MC statistics (as already suspected in ref. [114], section 11.4).¹⁷ Note that this does in no way invalidate the cited work, because the searched signature therein lies at the high-energy tail, which is nonetheless sufficiently well modeled by the MC.

To not bias the further development of this analysis, it was reblinded, running all tests only on MC if possible (e.g., comparing different sets), and not looking at the best-fit results if experimental data was needed.

Various tests and checks were conducted in the following months to identify the problem. The DOM efficiency parameterization was found to be very important to this work, and was updated from an energy-dependent parameterization to the energy- and zenith-

¹⁷One general problem of g.o.f. tests is that an insensitive test can not easily be distinguished from one that asserts good matching. This work relies on one-dimensional χ^2 and KS tests to not be limited by MC statistics.

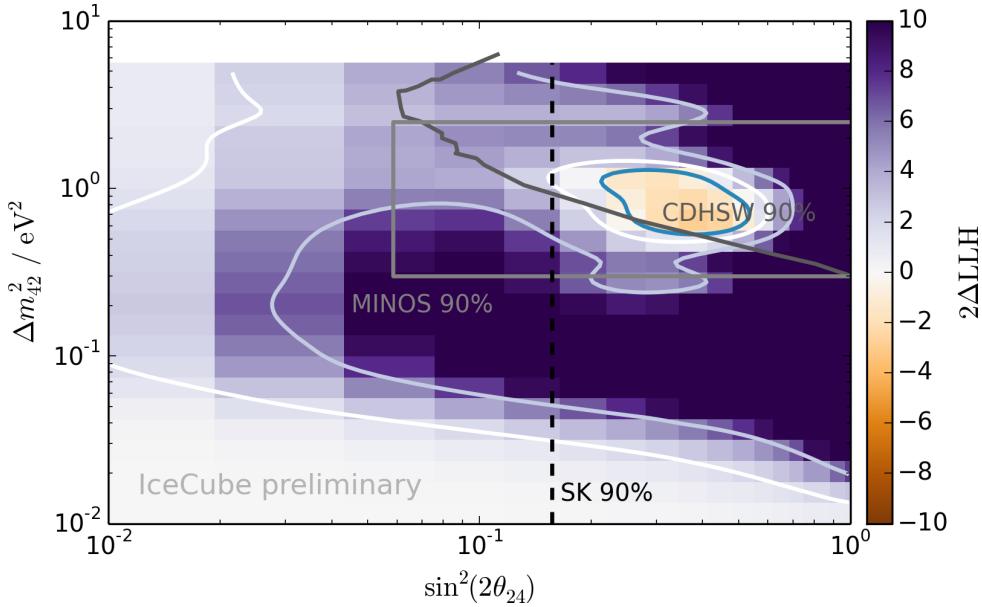


Figure 7.16: Likelihood landscape of a MC sample with ice tilt fitted with a similar sample without (with an older version of this analysis, see text). The MINOS line is from ref. [159].

angle-dependent parameterization described in section 6.2.5, with a much stronger emphasis on low-energy accuracy. The two antineutrino-ratio parameters were introduced as additional nuisance parameters, as well as the three continuous ice parameters that only became available with new MC.

The most important discovery was that the g.o.f. was good for subsets of the data that were defined with regards to the ice properties: The first subsample contained all events for which the center of gravity of hit DOMs (cog_z) was above the dust layer ($cog_z > 30$ m); events with their cog_z inside the dust layer were discarded (about 3600). Events with $cog_z < -180$ m were further split into a three samples according to their reconstructed azimuth angle ϕ , namely aligned with the ice tilt ($|\phi - 36^\circ| < 30^\circ$ or $|\phi - 216^\circ| < 30^\circ$), aligned with the ice flow ($|\phi - 126^\circ| < 30^\circ$ or $|\phi - 306^\circ| < 30^\circ$), and all remaining events lying in between (compare figure 2.16, which also illustrates that the tilt is very weak for the upper part of the detector). Split like this, the existing MC could describe the experimental data well as the nuisance parameters were able to absorb unmodeled ice properties without leaving physical parameter regions; only the tensions between those samples did not allow a good fit of the unsplit sample.

To test the hypothesis that insufficient ice modeling and especially the unsimulated ice tilt was the reason for the bad g.o.f., a small MC sample of Spice Mie ice generated with PPC (section 2.2.2, 64160 events) and thus featuring ice tilt was fitted with a Spice Mie sample generated with the older photonics software that is incapable of simulating tilt (310772 events). The result was a relatively strong fake signal: a best fit at $\theta_{24} = 19^\circ$

and $\Delta m_{42}^2 = 0.7 \text{ eV}^2$ at a significance of 1.2σ (according to Wilks' theorem) compared to an Asimov prediction of 3.2σ (figure 7.16), indicating a bad g.o.f.. Note that this fit was performed with an early version of this analysis, before the nuisance parameter improvements listed above had been implemented.

The production of sufficient amounts of updated MC proved to be difficult. When the deficiency of the older MC was discovered, GPU resources were relatively sparse and production of MC for newer detector configurations was prioritized. Additionally, modern MC contained many improvements besides the ice model (such as, e.g., better noise and DOM simulations) that made it incompatible with the processing software that was used with the IC-59 data. Since reprocessing the complete IC-59 experimental data with newer software was completely unfeasible, the new MC had to be backported to become compatible to old processing software, which took additional time. The resulting state-of-the-art datasets were carefully vetted and then used to produce the final results.

7.4 Previous Publication of Results

99% C.L. exclusion limits based on this work have been published together with results from a similar analysis run on newer IC-86 detector data in ref. [152]. The results differ from those reported in this work, because they are based on an older and overly conservative version of this analysis. The change responsible for the differences is an additional cut in this previous publication that discards events below $E_{\text{muEx}} = 10^{2.5} \text{ GeV}$ (1493 events) because of the relatively strong underfluctuation discussed in section 7.2.1 that pushes the exclusion region beyond the median sensitivity of the analysis. It was not physically well motivated because no indication for remaining background or systematic data-MC disagreement could be found after various tests had been conducted. Moreover, the cut was designed a posteriori to have a large effect at reducing the 90% C.L. exclusion contours, although it was later decided to publish the 99% C.L. contours, which were largely unaffected (figure 7.17). The additional cut has been removed for this work.

For reference, figure 7.18 shows the exclusion contours and sensitivities from pseudo experiments for the IC-86 analysis at 90% and 99% C.L. Its sensitivity is very close to that of IC-59 (compare figure 7.8) because both methodology and the event sample are very similar (20145 events compared to 21624 for IC-59 [152]), but the experimental exclusion contour is larger in the interesting region between 0.1 and 2 eV² because the best fit ($\Lambda = 0.75 \hat{=} 47\% \text{ p-value}$) lies at $\Delta m_{42}^2 = 10 \text{ eV}^2$ and $\theta_{24} = 22.5^\circ$ and therefore pushes the exclusion region towards lower values of Δm_{42}^2 (see appendix F).

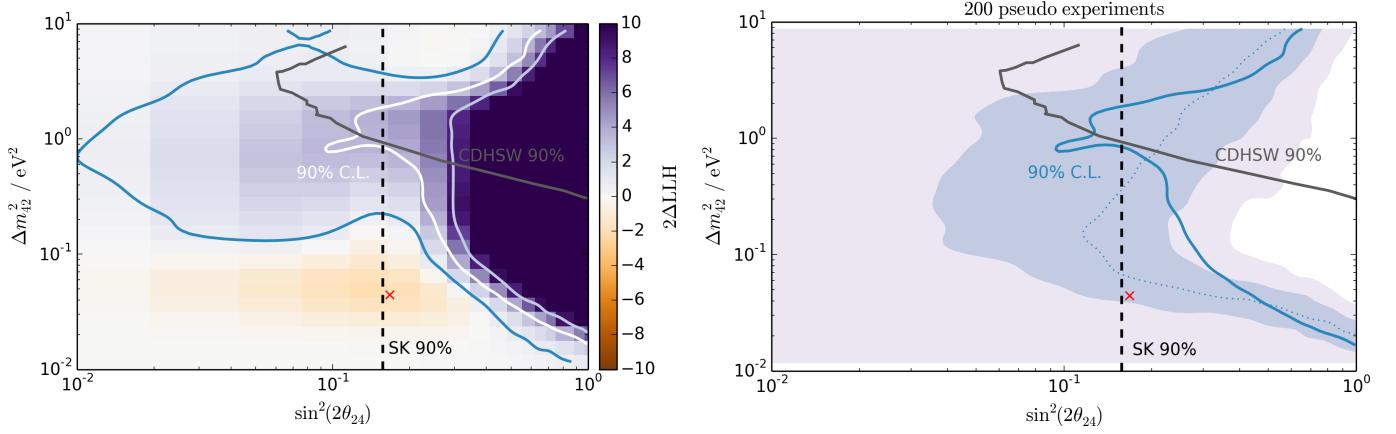


Figure 7.17: Likelihood landscape (left) and distribution of pseudo experiment contours (right) of an overly conservative result suggested by parts of the IceCube collaboration, see text.

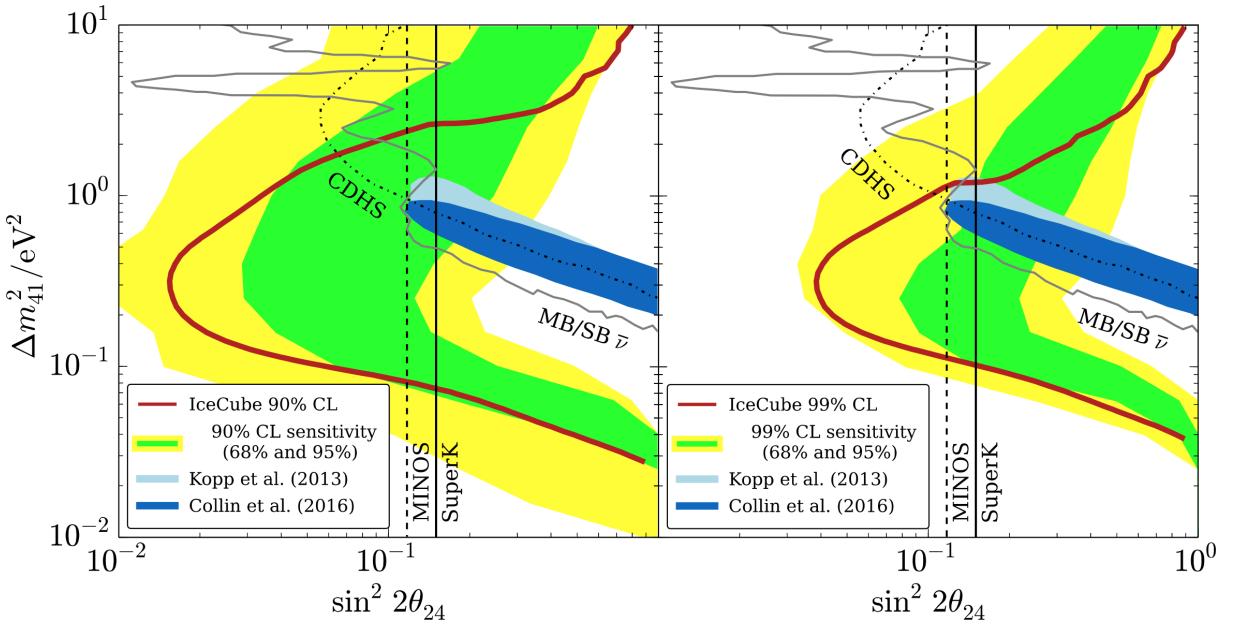


Figure 7.18: Contours and sensitivities of the IC-86 analysis [152], for 90% C.L. (left) and 99% C.L. (right).

Chapter 8

Conclusions and Outlook

8.1 Conclusions

The IceCube Neutrino Observatory is exceptionally well suited to search for sterile neutrinos in the Δm_{42}^2 range of roughly 0.05 eV^2 to 2 eV^2 if they mix with muon neutrinos, because it detects large amounts of atmospheric neutrinos over a variety of energies and baselines. This sensitivity mostly comes from matter effect resonances that can strongly increase the oscillation probability of muon (anti-)neutrinos to sterile (anti-)neutrinos in sharply defined baseline/energy regions for a given set of neutrino parameters.

The existing data regarding the existence of sterile neutrinos is inconclusive, so this analysis is designed to search for and exclude sterile neutrinos in the 3+1 framework over a large parameter region in a model-independent way. The resulting signatures for IceCube are diverse in their strength and position in the parameter space of observables, but luckily they retain a specific shape which allows the distinction between sterile-neutrino signatures and other phenomena, even after they have been washed out by IceCube's limited energy resolution for through-going neutrinos. The analysis remains highly sensitive to systematic effects, which necessitated the generation of much newer and strongly improved Monte-Carlo data sets after the bad data-MC agreement of the older data sets became apparent. Additional systematic uncertainties were added and existing parameterizations were reevaluated and improved.

The final analysis result provides strong exclusion limits and cuts far into the allowed region of the LSND sterile neutrino hypothesis, although it depends on the mixing angle θ_{14} how well that hypothesis can be excluded [160], on which IceCube is not sensitive. The best fit ($\theta_{24} = 12.08^\circ$, $\Delta m_{42}^2 = 0.0401\text{ eV}^2$) is compatible with the null hypothesis at a significance of 0.463σ and lies outside the median expected 90% C.L. sensitivity at the lower boundary of the analysis' energy range. Accordingly, the exclusion limits are relatively strong for high Δm_{42}^2 (corresponding to high energies) and relatively weak at

low Δm_{42}^2 compared to the median expected sensitivity (figure 7.8); for the same reason they feature an indent around $\theta_{24} \approx 13^\circ$ towards low Δm_{42}^2 , which in this size is unusual, but compatible with MC predictions, with about 6% of all pseudo experiments featuring similarly sized indents.

The analysis results were carefully vetted after the final unblinding. Besides various low-level checks, the nuisance parameters introduced to take care of systematic uncertainties are well-behaved; the goodness of fit was assured by applying the χ^2 and Kolmogorov-Smirnov methods to one-dimensional distributions to retain high statistics and therefore a good sensitivity for mismatches, and getting good results (all p-values above 65%); and the two-dimensional distribution of the observables was checked for data-MC mismatches by a statistical analysis of rectangular clusters of pulls, which showed that the analysis outcome is well compatible with MC expectations.

8.2 Outlook

This analysis was the first to explore searching for sterile neutrinos with IceCube (besides an early study by Warren Huelsnitz), and despite the good results, several improvements could be explored for future analyses.

Firstly, the quality of IceCube Monte-Carlo data is continuously improved by various people who work on software and low-level studies to learn more about and limit systematic uncertainties. This progress will greatly help future analyses as it will allow them to tighten the priors on nuisance parameters, increasing the analysis sensitivity. Together with improved event selection techniques, this might also allow to safely include events from slightly above the horizon. Currently, downgoing events are prone to slight data-MC mismatch because it is difficult to produce enough atmospheric-muon background. While sufficiently good for almost all IceCube analyses, such mismatch could severely bias searches for sterile neutrinos (compare appendix F).

Another obvious improvement to this analysis is to include more data. When the development of this analysis started, only data for incomplete detector configurations was available, and simulation and event selection improved so quickly that these one-year samples differed greatly from each other. Today, IceCube’s construction is long finished and data taking is very stable, making it easy to use data from multiple years without complicating systematic uncertainties. The increased statistics not only directly increase the sensitivity, but also help to limit systematic uncertainties by allowing better fits of the nuisance parameters.

8.2.1 Starting Events

An approach that could be promising especially in combination with the improved statistics of a multi-year analysis is to limit the analysis to starting events, i.e., to events where the neutrino interacts inside the instrumented volume of IceCube. This allows a reconstruction of the energy that goes into the initial cascade, which eliminates the largest source of energy reconstruction uncertainty (compare figure 2.7). The highly increased energy resolution allows to more clearly resolve the resonance signature of sterile neutrinos, but it is bought by strongly reduced statistics (and thereby reduced power to fit nuisance parameters), so this becomes an optimization problem. The approach was not explored with this work because of the limited available statistics, but a dedicated analysis is under development within the IceCube collaboration.

8.2.2 Low-Energy Sample

In addition, the analysis could be extended to higher and lower reconstructed neutrino energies to increase the sensitivity to higher and lower values of Δm_{42}^2 , respectively. The extension to higher energies requires a good knowledge of the astrophysical neutrino flux, which becomes dominant above about 300 TeV [13] (figure 1.7 is accurate in this regard). Because of the power-law spectrum, not many events and therefore not much sensitivity can be gained at high energies. At low energies however, high event numbers are available, and the region is relatively well understood. Problems arise from the intrinsically higher background contamination and from very different event selection techniques required for events below about 100 GeV, which severely complicate analyses that want to combine low- and high-energy samples. Also, additional systematic uncertainties become relevant at low energies [161], and some existing ones require different parameterizations. Despite these challenges, the gain could be significant. Figure 8.1 shows results from a simple sensitivity study with an IC-79 low-energy sample (8117 events) [132], using the analysis method of this work but without most nuisance parameters, and without combining it with an high-energy sample. The obtainable contours are smaller because of lower statistics, but they complement the high-energy analysis well in the low-energy region as expected.

A different approach to use low-energy samples for IceCube’s search for sterile neutrinos is to search for changes in the conventional oscillation pattern due to sterile neutrinos, as described in section 4.1.2. This kind of analysis only weakly depends on the value of Δm_{42}^2 , because it assumes to be in the fast-oscillation regime with fully washed out signatures, but instead depends relatively strongly on θ_{34} . Figure 8.2 shows results from a three-year analysis using this approach [162]. The exclusion of θ_{24} is weaker than the exclusion of this work in its optimal region, but the exclusion extends further towards high Δm_{42}^2 and surpasses it at about $\Delta m_{42}^2 = 1.4 \text{ eV}^2$. Both analyses should be continued and developed further to complement each other and to provide additional insights for the search for light sterile neutrinos.

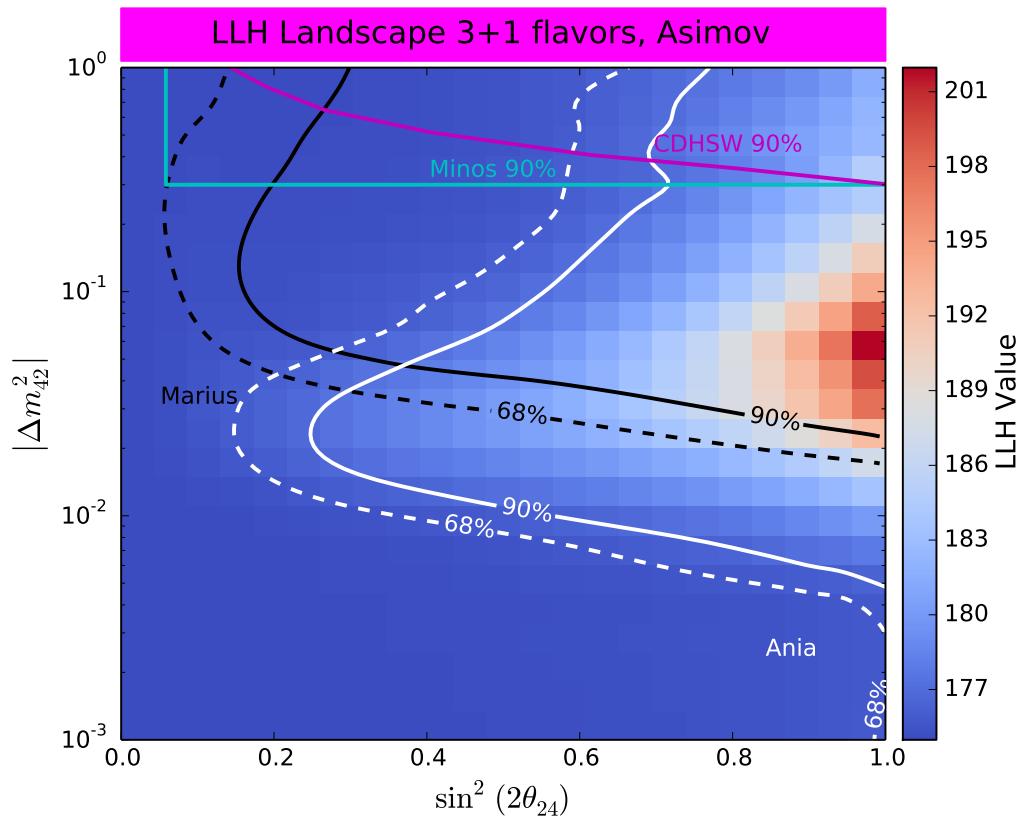


Figure 8.1: Estimated sensitivity and likelihood landscape of a simple study with an IC-79 low-energy sample by Ania Koob [132]. The black lines indicate earlier but comparable Asimov estimates of this work (figure 7.1). The MINOS line is from [159].

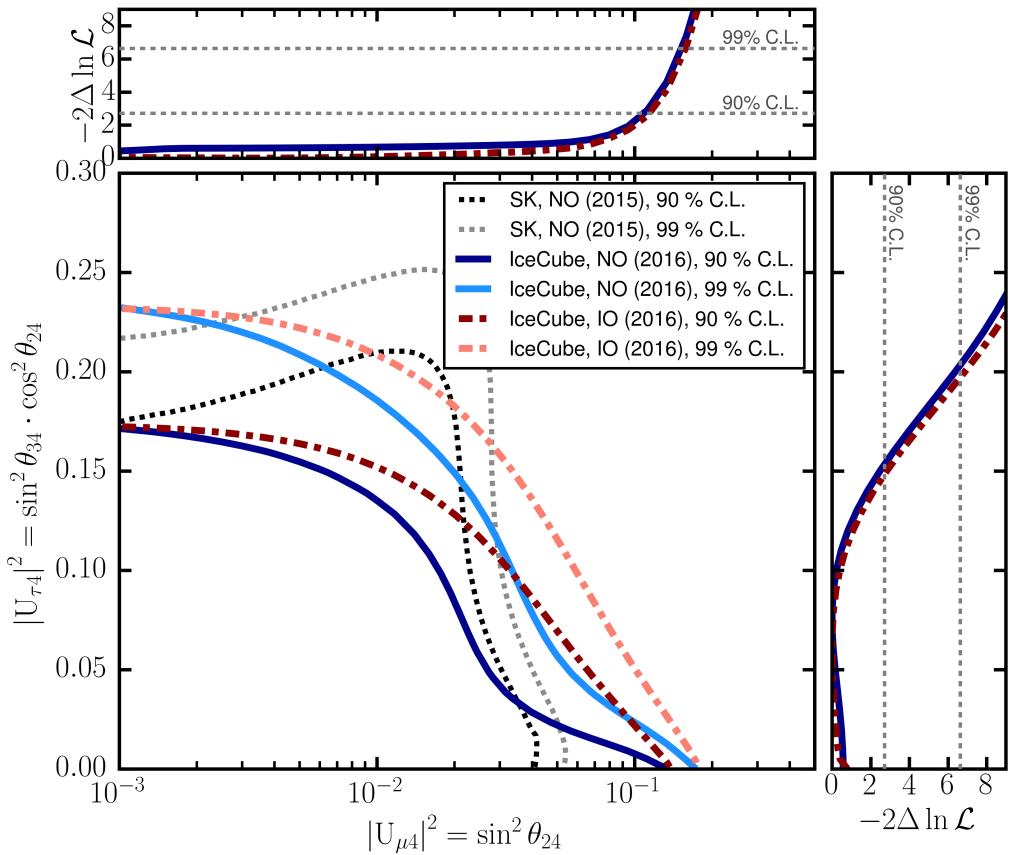


Figure 8.2: Exclusion limits of another IceCube analysis searching for sterile neutrinos [162], using a low-energy sample. The limits were produced with $\Delta m_{42}^2 = 1 \text{ eV}^2$, but are valid roughly between 0.1 eV^2 and 10 eV^2 .

Appendix A

CP Violation

When the concept of CP violation was introduced in section 1.3, it was postulated that CP-violating phases do not matter for this analysis and that they are being set to 0. For Majorana-type phases this is clear (as $|J| = 1$, J commutes with all other matrices), so this appendix justifies the omission of CP violation for Dirac-type phases δ_{ij} .

Dirac-type phases start to become relevant for three-flavor oscillations. This can be shown using the Jarlskog invariant $J_{\alpha\beta}$, which is proportional to the amount of CP violation in the oscillation between flavors α and $\beta \neq \alpha$ [163]. It can be expressed as

$$J_{\alpha\beta} \equiv \Im(U_{\alpha 1} U_{\alpha 2}^* U_{\beta 1}^* U_{\beta 2}) = \pm \cos(\theta_{12}) \sin(\theta_{12}) \cos(\theta_{23}) \sin(\theta_{23}) \cos(\theta_{13})^2 \sin(\theta_{13}) \sin(\delta_{13}),$$

with + for cyclic oscillations (e to μ , μ to τ , τ to e). All three mixing angles have to be non-zero for $J_{\alpha\beta} \neq 0$, so CP violation does not occur in two-flavor approximations.¹⁸

From this one can deduce that in cases that can be described well as two-flavor approximations, CP violation has no effect, which is the case in this work. A thorough MC study has been conducted in ref. [124], which will be summarized here to give a more complete overview.

For $\theta_{14} = \theta_{34} = 0$ (which is the case in this work), the oscillation can be described well as a two-flavor process, and accordingly CP violating phases do not matter in this case. When all mixing angles are non-zero, CP violation occurs and the effects of the different phases can be seen. Figure A.1 shows the effect of the phases δ_{24} and δ_{34} on the survival probabilities of muon neutrinos and antineutrinos. The effects of the phases are nearly identical: both phases have no effect on the matter resonance, but they fully switch the effects of the sterile neutrino on the conventional atmospheric oscillations from the neutrino to the antineutrino channel and vice versa. Effects of numerical values between

¹⁸Another approach is to evaluate the formula for the non-absorbable Dirac-type phases from section 1.3, i.e., $(n-1)(n-2)/2$, which yields 0 for $n = 2$.

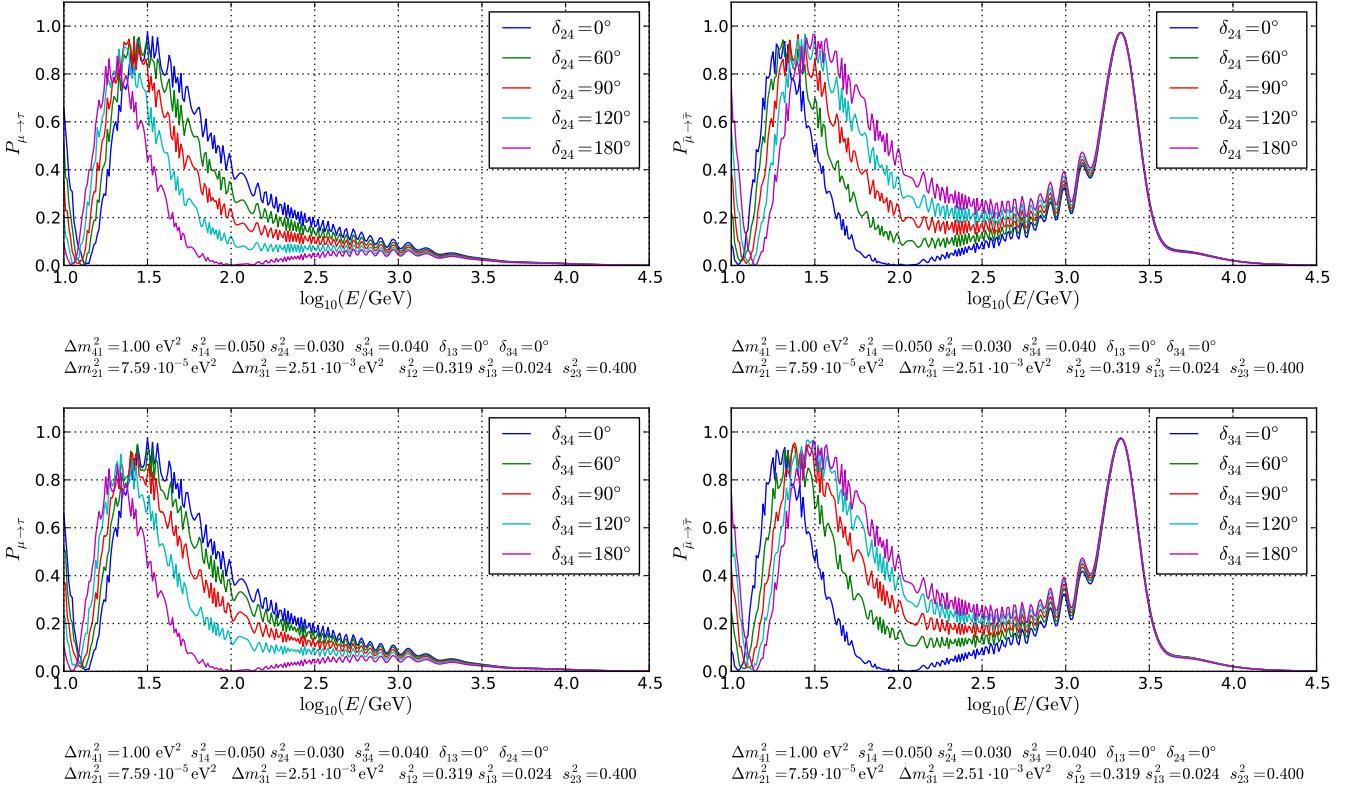


Figure A.1: Muon neutrino (**left**) and antineutrino (**right**) survival probabilities for vertically upgoing particles for different values of δ_{24} (**top**) or δ_{34} (**bottom**) vs. neutrino energy [124]. All parameters are given in the plots.

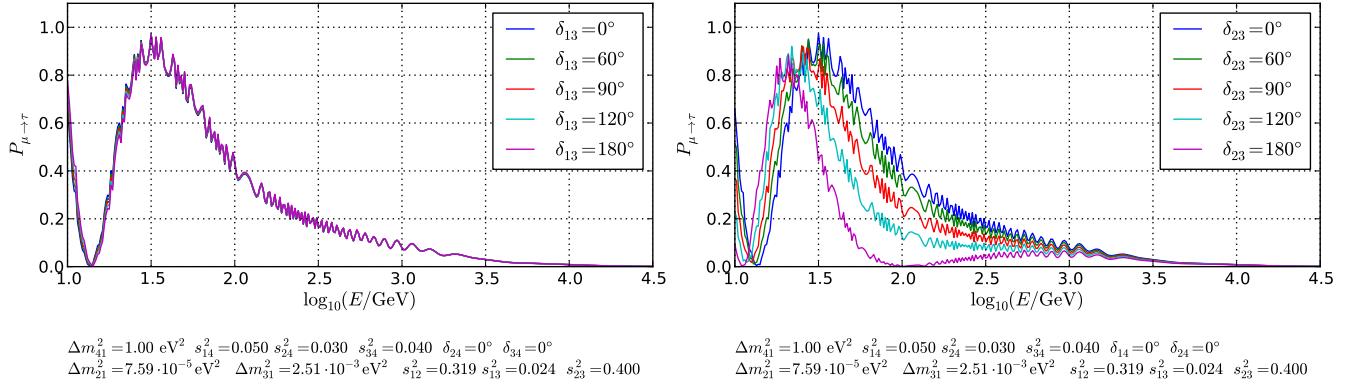


Figure A.2: Muon neutrino survival probabilities for vertically upgoing particles for different values of δ_{13} (left) or δ_{23} (right) vs. neutrino energy [124]. Note that the right plots uses an alternative CP violation parameterization scheme.

the endpoints 0° and 180° differ between the phases, but this is an effect of the ordering of the matrices that make up the PMNS matrix U (see section 1.4.1): For every value of δ_{24} , a corresponding value of δ_{34} can be found that emulates the effect, as long as the endpoints are identical; it is only a matter of scaling.

Only at very low energies ($\lesssim 15$ GeV), when oscillations to the electron flavor occur, the degeneracy of δ_{24} and δ_{34} gets lifted slightly, leading to deviations up to 2% in the 180° case [124]. This deviation disappears when $\theta_{14} = 0$, but otherwise θ_{14} does not matter regarding CP violation.

Figure A.2 shows the results of varying CP-violating phases between the conventional flavors. In the usual parameterization, this means varying δ_{13} , which can be seen to basically have no effect above about 15 GeV. If the parameterization is changed to the set $\{\delta_{14}, \delta_{23}, \delta_{24}\}$, the “conservative” phase δ_{23} becomes relevant and effectively indistinguishable from δ_{24} and δ_{34} in the standard parameterization (figure A.1).

This further confirms that even when all mixing angles are non-zero, the IceCube signature can very well be approximated as three-flavor oscillation between muon, tau, and sterile neutrinos, and because of this only has one relevant Dirac-type phase (or zero Dirac-type phases for $\theta_{34} = 0$ as in this work).

Appendix B

Neutrino Mass Hierarchy

Up to this day, the absolute neutrino masses have not been determined. The best upper limit for the mass of electron neutrinos from direct measurements is $m_{\nu_e} < 2 \text{ eV}$ from tritium decay experiments, with better results expected from KATRIN in the near future [76]. Direct measurements for other flavors are not competitive to the combination of the ν_e limit with oscillation results [7]. Indirect neutrino mass measurements based on cosmological data are often more constraining, e.g., $\sum m_\nu < 0.17 \text{ eV}$ at 95% C.L. in ref. [164], but they are model-dependent and their limits can be circumvented by various means [165], compare section 1.5.

As shown in section 1.3, neutrino oscillations are not sensitive to the signs of Δm_{ji}^2 unless there is either CP violation (which introduces terms that scale with $\sin(\Delta m_{ij}^2 L / 2E_\nu)$ instead of $\sin^2(\Delta m_{ij}^2 L / 4E_\nu)$) or matter effects (where the induced matter potential is either added to or subtracted from the term that contains Δm_{ji}^2). The observation of solar neutrino oscillations with their strong matter effects established $\Delta m_{21}^2 > 0$, but it is still unclear whether $\Delta m_{32}^2 > 0$. This problem is known as the neutrino mass ordering or neutrino mass hierarchy.¹⁹

For this analysis, the usual neutrino mass hierarchy between the known three mass eigenstates is irrelevant, as the differences between the two signatures only affects standard oscillations below the first oscillation minimum (where matter effects become relevant) [124], which is well below the energy range of this analysis.

As pointed out in section 4.2, the ordering of the sterile neutrinos determines whether matter effect resonances occur in the antineutrino ($m_4 > m_2$, 3+1) or neutrino ($m_4 < m_2$,

¹⁹Note that sometimes these terms are not used interchangeably. The term neutrino mass hierarchy is sometimes meant to include knowledge of the absolute values. In that case, there is not only the normal mass ordering (NO) in which $m_1 < m_2 < m_3$ and the inverted ordering (IO) with $m_1 < m_3 < m_2$, but also the quasi-degenerate hierarchy where $0 \ll m_1 \approx m_2 \approx m_3$ (which still needs to be differentiated into NO and IO when oscillations are relevant).

$1 + 3$) channel analog to what was shown in section 1.3.3 for conventional flavors, so it is relevant for this analysis unlike the conventional ordering. This analysis concentrates on the normal sterile ordering ($3 + 1$) because it is the conservative choice with weaker signatures, as there are fewer antineutrinos than neutrinos (section 1.2.4).

Appendix C

Calculation of Oscillation Probabilities with nuCraft

The tool used for the calculation of oscillation probabilities of atmospheric neutrinos throughout this work is nuCraft, an open-source Python tool created by the author of this work. It has originally been written specifically for this analysis, but was published independently to be available to a broader audience [166].

nuCraft works by numerically solving the flavor-basis wave equation (1.9) with the ODE solver ZVODE, using a SciPy wrapper [167][130]. It has a lightweight Python implementation that supports an arbitrary number of flavors, is highly configurable, and can easily be adapted to extend the applicability of nuCraft to problems for which it was not designed (such as neutrino beam experiments). All details can be found in the dedicated publication [166], but two of nuCraft’s features will be discussed here as the focus of the paper is the oscillation of conventional neutrinos: The modeling of the atmosphere and the modeling of the Earth’s interior.

Close to the horizon, the neutrino path length through the atmosphere can be longer than the path length through solid matter (see figure 1.13 and figure 1.6, so neglecting the atmosphere completely can lead to large errors. This can be fixed easily by letting neutrinos start about 20 km above the ground; however, the interaction height of cosmic rays in the atmosphere is random and also depends on the energy and zenith angle under which the particle hits the atmosphere. This not only leads to corrections of the path length compared to assuming a fixed interaction height, but also to smearing of the oscillation probabilities because neutrinos with the same parameters can still have different baselines.

The atmosphere is modeled in accordance to ref. [27]. Figure C.1 shows the distribution of atmospheric path lengths. The effect of the neutrino energy is much smaller than the widths of the distributions; therefore, to ease the parameterization and computational effort, the energy has been fixed to 2 GeV, i.e., the energy dependence is ignored. The

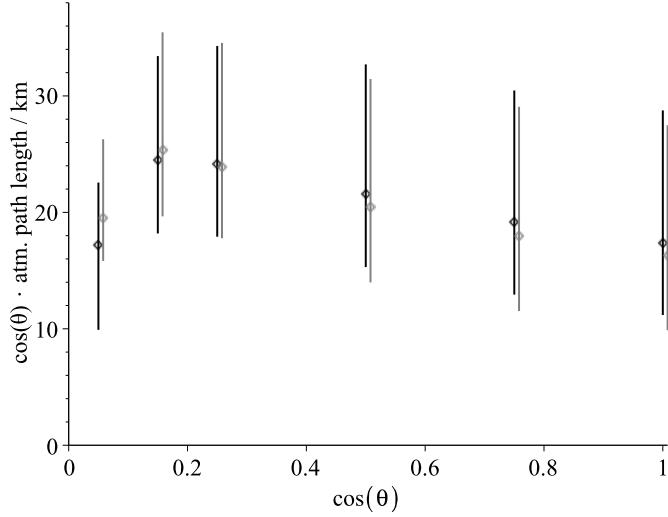


Figure C.1: Expectation value and central 68% interval of the path length distributions of neutrinos in Earth's atmosphere as a function of the zenith angle, for 200 GeV (black) and 2 GeV (gray, slightly shifted for clarity) [166].

reasoning for choosing 2 GeV is that for conventional neutrinos, the effect of the atmospheric smearing becomes strong ($\gtrsim 3\%$) only below about 4 GeV [166]. For the resonance signature of sterile neutrinos in the context of this work, 2 GeV is much too low, but firstly the difference between energies is low anyway, and secondly the atmospheric smearing is not very important in this case because for relevant mixing angles, the resonance occurs for near-vertical tracks where the impact of the atmospheric baseline is very small (figure 1.13 and figure C.2).

The interior of the Earth is modeled after the Preliminary Earth Model PREM (section 1.3.2). Unlike some other popular tools, nuCraft does not simplify the model by assuming (usually four) shells of constant density, but uses a continuous parameterization (bar intentional discontinuities) as shown in figure 1.17. This approximation gives a significant speed-up (about a factor of 1000 using, e.g., the tool Prob3++ [168]) and provides good results for non-resonant oscillations. However, the calculation of resonances can be affected strongly because they occur at precisely matching density values. Figure C.2 shows a comparison between the accurate calculation and the approximation (using four layers as in Prob3++) with all other parameters identical: The strongest differences can be observed for conventional oscillations below 10 GeV which is not relevant for this work, but a difference of up to 7% can be observed for the sterile-neutrino resonance, which corresponds to 10% of the effect (as the resonance does not reach 100% in this example [166]). This increase in accuracy warrants the drawback of the highly increased computational effort, which in turn necessitates the use of parameter scanning instead of fitting as described in section 5.5.

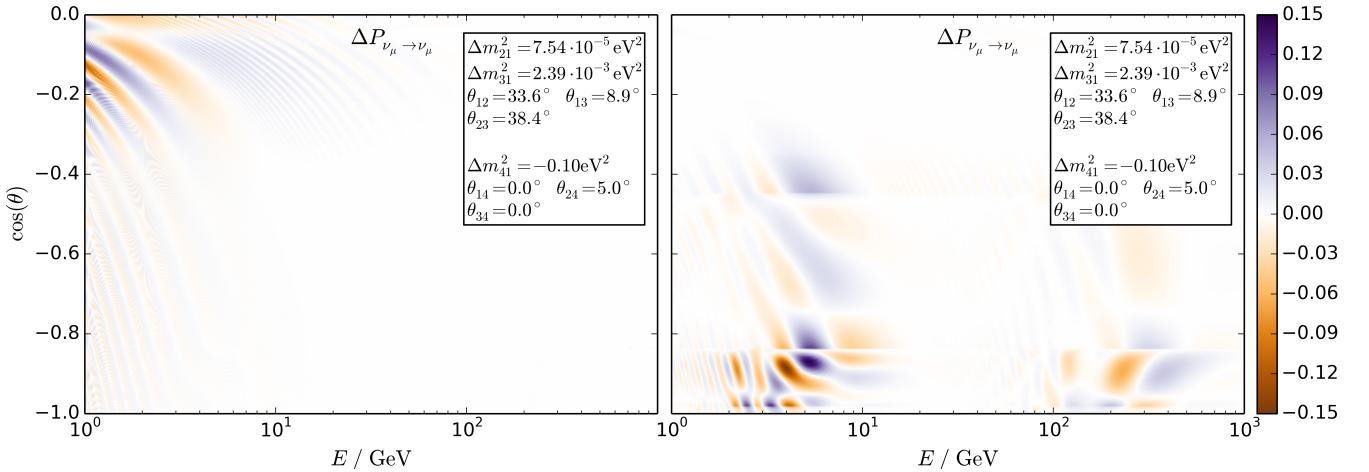


Figure C.2: Comparison of nuCraft default computations vs. common approximations [166]: assuming a fixed atmospheric height for neutrino generation (**left**), and approximating the Earth interior by four shells of constant density (**right**).

C.1 Factorization of Oscillations and Non-Coherent Interactions

One limitation of nuCraft is that it only computes the oscillation probability without taking into account non-coherent and non-elastic interactions. Instead, both small and catastrophic energy losses caused by non-elastic interactions are independently handled by the IceCube MC simulation framework, which means that using nuCraft, oscillation and non-oscillation matter effects are factorized. This approach is an approximation of reality, especially when sterile neutrinos are involved, because the probability of non-elastic interactions depends on the neutrino flavor.

The flavor that was simulated by the IceCube software for this work was muon neutrinos, because it is the flavor that is required at the time of detection. Therefore, the propagation through Earth was computed for muon neutrinos. Electron antineutrinos are more likely to interact with Earth than muon neutrinos, whereas sterile neutrinos don't interact with Earth at all (besides gravitationally, which is negligible). Moreover, tau neutrinos are able to regenerate, i.e., if they interact with matter, a resulting tau lepton might decay before much energy is lost, resulting in a new tau neutrino with similar energy compared to the original neutrino. The behavior of electron and tau neutrinos is not relevant for this work, as the respective mixing angles are set to zero, but the difference between muon and sterile neutrinos affects this work.

The cross-section of these processes highly depends on the neutrino energy. Below 1 TeV, the chance of absorption is smaller than 5% for muon neutrinos, and reaches 25% for vertically upgoing neutrinos at about 10 TeV [169]. The error caused by the approximation is only a correction to these percentages. However, an increasing number of

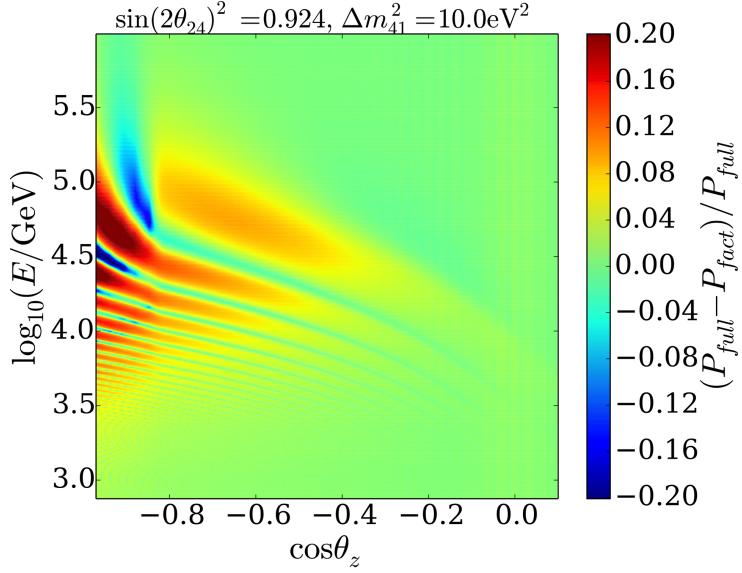


Figure C.3: Difference between the fully simulated coherent and non-coherent interactions using ν -SQuIDS [171] and the factorization approach using nuCraft and the IceCube MC framework for nearly maximal mixing angles according to ref. [170].

non-catastrophic interactions at these energies can also disrupt the coherence of the propagation and therefore influence the oscillation probabilities. Overall, the approximation holds very well up to about 19 TeV (corresponding to $\Delta m_{42}^2 = 5 \text{ eV}^2$), and locally deviates by up to 20% at 38 TeV ($\Delta m_{42}^2 = 10 \text{ eV}^2$) at very extreme (nearly maximal) mixing angles [170] (see figure C.3). The deviations alternate quickly between over- and underpredictions, which washes out the effect, and are much weaker for realistic mixing angles even at these energies, because in that case the neutrinos mostly act like muon neutrinos as which they were simulated in the factorized approximation. Therefore, the effects of the approximation are negligible for the tested mass range of this analysis ($\Delta m_{42}^2 \leq 8.75 \text{ eV}^2$).

Appendix D

Locally Weight-Preserving Gaußian Blurring

A locally weight-preserving Gaußian blurring algorithm is used to smooth oscillation weights in respect to true MC energy and zenith angle as discussed in section 5.1. The algorithm is based on the Gaußian filter of SciPy’s multi-dimensional image processing module, i.e., `scipy.ndimage.filters.gaussian_filter` [130]. Its settings are left at default besides `mode` with is set to `constant`; in combination with the default `cval=0` this means that points outside of the parameter space (i.e., outside of the array) are assumed to be zero for purposes of blurring. The kernel width (i.e., the Gaußian’s σ) is set to 0.5 bin widths in both directions (as in figure 4.14), and the filter (by unmodified default) truncates the kernel at 4σ (and renormalizes the kernel to a sum of 1).

For a given array (ν_μ or $\bar{\nu}_\mu$ disappearance probability), the filter is applied both to the array itself and to its Boolean version that is defined as having 1 as entry if the original histogram has entries in the same bin and 0 otherwise. The output of the algorithm is the filtered original array divided by the filtered Boolean array multiplied by the unfiltered Boolean array (i.e., all bins that were empty in the input array remain empty in the output array). To compensate for numerical errors, the output is clipped to only have values between 0 and 1.

The effect of this algorithm compared to simple Gaußian blurring is that the sum of the originally populated bins is preserved not only globally (which could always be achieved by renormalizing the whole array after filtering), but also locally. The filter alone would cause probability to leak out of the regions of the array that are populated by MC events. When probabilities for MC events are drawn from this smoothed histogram as described in section 5.1, this would lead to a bias because unpopulated bins are irrelevant for this drawing process, and probability that leaked into those bins would be lost. Global reweighting of the originally populated bins would still introduce a bias by underestimating probabilities around the edges and overestimating probability in the well-populated center where no probability can leak out.

Appendix E

Kernel Density Estimation

Kernel density estimation (KDE) is a technique for non-parametric estimation of probability density functions. It has been described in the 1950s by Murray Rosenblatt [172] and Emanuel Parzen [173], and is well established in, e.g., signal processing. This appendix will give a short introduction and describe the specifics of the implementation; a more detailed general mathematical description of the method can be found in ref. [174].

KDE can be seen as an alternative to the histogramming of data, in which data is not counted into bins of a fixed size, but instead functions (the so-called *kernels*) are positioned for every data point, and then added up to approximate the distribution. The kernels used in this work are Gaußian distributions; the mean μ is determined as the position of the data point, while the width σ (called *bandwidth*) is determined by the local density of data points (*adaptive KDE*). The algorithm responsible for the width determination was designed by Bernard W. Silverman [175]; it provides a bandwidth that is intentionally about 15% smaller than the optimum bandwidth for Gaußian data with Gaußian kernels. While the author notes that this still might be too wide for steeply falling or otherwise sharp distributions, it performs well for the relatively flat distributions it is used on in this work.

The result is an estimation that does not depend the selection of the widths and starting points of bins and generally also performs well in regions of small statistics, where fluctuations are smoothed out widely because the local bandwidth becomes large. Care must be taken if the estimated distribution has sharp boundaries, because in this case some probability can “flow out” of the bounded region, but nothing flows back in from outside to compensate for it. The solution to this problem for this work is to mirror the data on every boundary, such that everything that flows out equally flows back in. Effectively, this means that the estimated distribution becomes constant very close to the boundaries, which is a good approximation for the relatively flat distributions the KDEs are used on in this work.

Another challenge of using KDEs is that they are not as intuitive as histograms to most physicists in terms of error estimation, and error estimation is not as trivial as with histograms (where one usually can assume a Poissonian error on the number of entries per bin). The problem was solved by using bootstrapping: The KDEs are computed multiple times on n randomly chosen entries of the original n data points using replacement, meaning that some points are selected multiple times and others not at all. This results in an ensemble of KDEs that is centered around the original KDE and spreads by the standard error, which can therefore be computed for every point of the KDE [175].

Bootstrapping is simple and has only few requirements, the most important being that the original distribution has a finite variance (i.e., is not heavy-tailed), which is fulfilled in this case. In general, bootstrapping is well established and often recommended in combination with KDEs [175], and the number of bootstrap samples (64 or 128 as stated in sections 6.2.5 and 6.2.6) was decided both based on literature recommendations (50 to 100; e.g., [176]) and on iterative checks of the development of the bootstrapping results.

Appendix F

3×3 Bin Manipulation Study

One of the many post-unblinding studies was to check how a cluster of relatively weak systematic fluctuations can influence the result of the analysis. The size of the cluster was set to 3×3 bins of the analysis binning, inspired by the net excess in the vertically upgoing low-energy region described in section 7.2.1.

The methodology of this test is to manipulate bins of the Asimov dataset (section 7.1.1) by increasing or decreasing all bin contents of a 3×3 cluster by 1σ of their respective bin contents each. The manipulated sample is then run through the analysis code, and the resulting contour is compared to the unmodified Asimov estimate. The test was conducted for nine samples that differ in the positions of the clusters.

If all bins had the same contents, the significance of the pull cluster (not of the pseudo experiment best fit) would be $\sqrt{9}\sigma$ for every sample. Table F.1 shows how many neutrino events had to be subtracted or added for every sample (Δn), and the significance of the pull cluster, both before nuisance parameters were refitted, and after. The naming scheme for the position of the cluster is $lE/mE/hE$ for low, medium, and high energy, and $vZ/mZ/hZ$ for vertically upgoing, medium, and horizontal tracks (zenith angle). Especially at low energies and towards horizontal events, the nuisance parameters are able to absorb much of the manipulations. However, since the manipulations are not modeled according to the physical effects of the nuisance parameters, they are not fully absorbed and can still influence the outcomes of the pseudo experiments.

Figures F.1 and F.2 show the contours of the nine experiments together with the unmodified Asimov expectation. The study was conducted with the additional energy cut described in section 7.4, so the Asimov expectation differs slightly from figure 7.1. Manipulations of medium zenith angles tend to have the smallest effects on the contours, because no sterile signature lies exclusively in medium zenith angle regions. Adding vertically upgoing tracks or subtracting horizontal tracks tends to enlarge the exclusion

Table F.1: Number of added/subtracted events and resulting significance compared to the Asimov sample for all nine cluster positions before (**pre**) and after (**post**) the fits

name	Δn pre ±	sig. pre ±	Δn post +	sig. post +	Δn post -	sig. post -
lE/vZ	97.776	2.877σ	84.068	2.459σ	83.821	2.482σ
mE/vZ	29.676	2.825σ	28.706	2.721σ	28.891	2.760σ
hE/vZ	7.278	2.778σ	7.229	2.750σ	7.230	2.769σ
lE/mZ	122.734	2.968σ	107.532	2.589σ	106.952	2.598σ
mE/mZ	48.205	2.866σ	44.825	2.650σ	44.954	2.689σ
hE/mZ	14.132	2.841σ	13.755	2.744σ	13.752	2.786σ
lE/hZ	141.013	2.972σ	90.169	1.879σ	88.375	1.885σ
mE/hZ	63.586	2.893σ	52.763	2.374σ	52.798	2.430σ
hE/hZ	21.730	2.863σ	20.102	2.612σ	20.105	2.687σ

contours without introducing a strong non-null best fit, because both cases represent anti-signatures. This is obvious for the first case, as the usual signature for sterile neutrinos at small mixing angles is a deficit of vertically upgoing events (chapter 4). For the second case, one has to take into account the relatively free total normalization: For large Δm_{42}^2 , the signature is a deficit everywhere but for high-energy horizontal events, so subtracting these events (and to a lesser degree horizontal events in general) is equivalent to an excess everywhere else and therefore an anti-signature.

The two pseudo experiments lE/vZ+ and lE/hZ- are the most interesting ones; the first one is reminiscent of the result of this analysis, the latter one is reminiscent of the results of the IC-86 analysis (figure 7.18). In both cases, the clusters lie where high pull clusters were observed in the respective analyses [152], and the pseudo experiments manage to roughly reproduce the observed contours. Note that this does not mean that the analyses are systematically biased in these regions, as these clusters could also be caused by statistics or signal; the study just verifies that the analyses respond to the respective experimental data in the expected way, and shows which regions should be put under special scrutiny if one searches for possible systematic bias, as was done for IC-59.

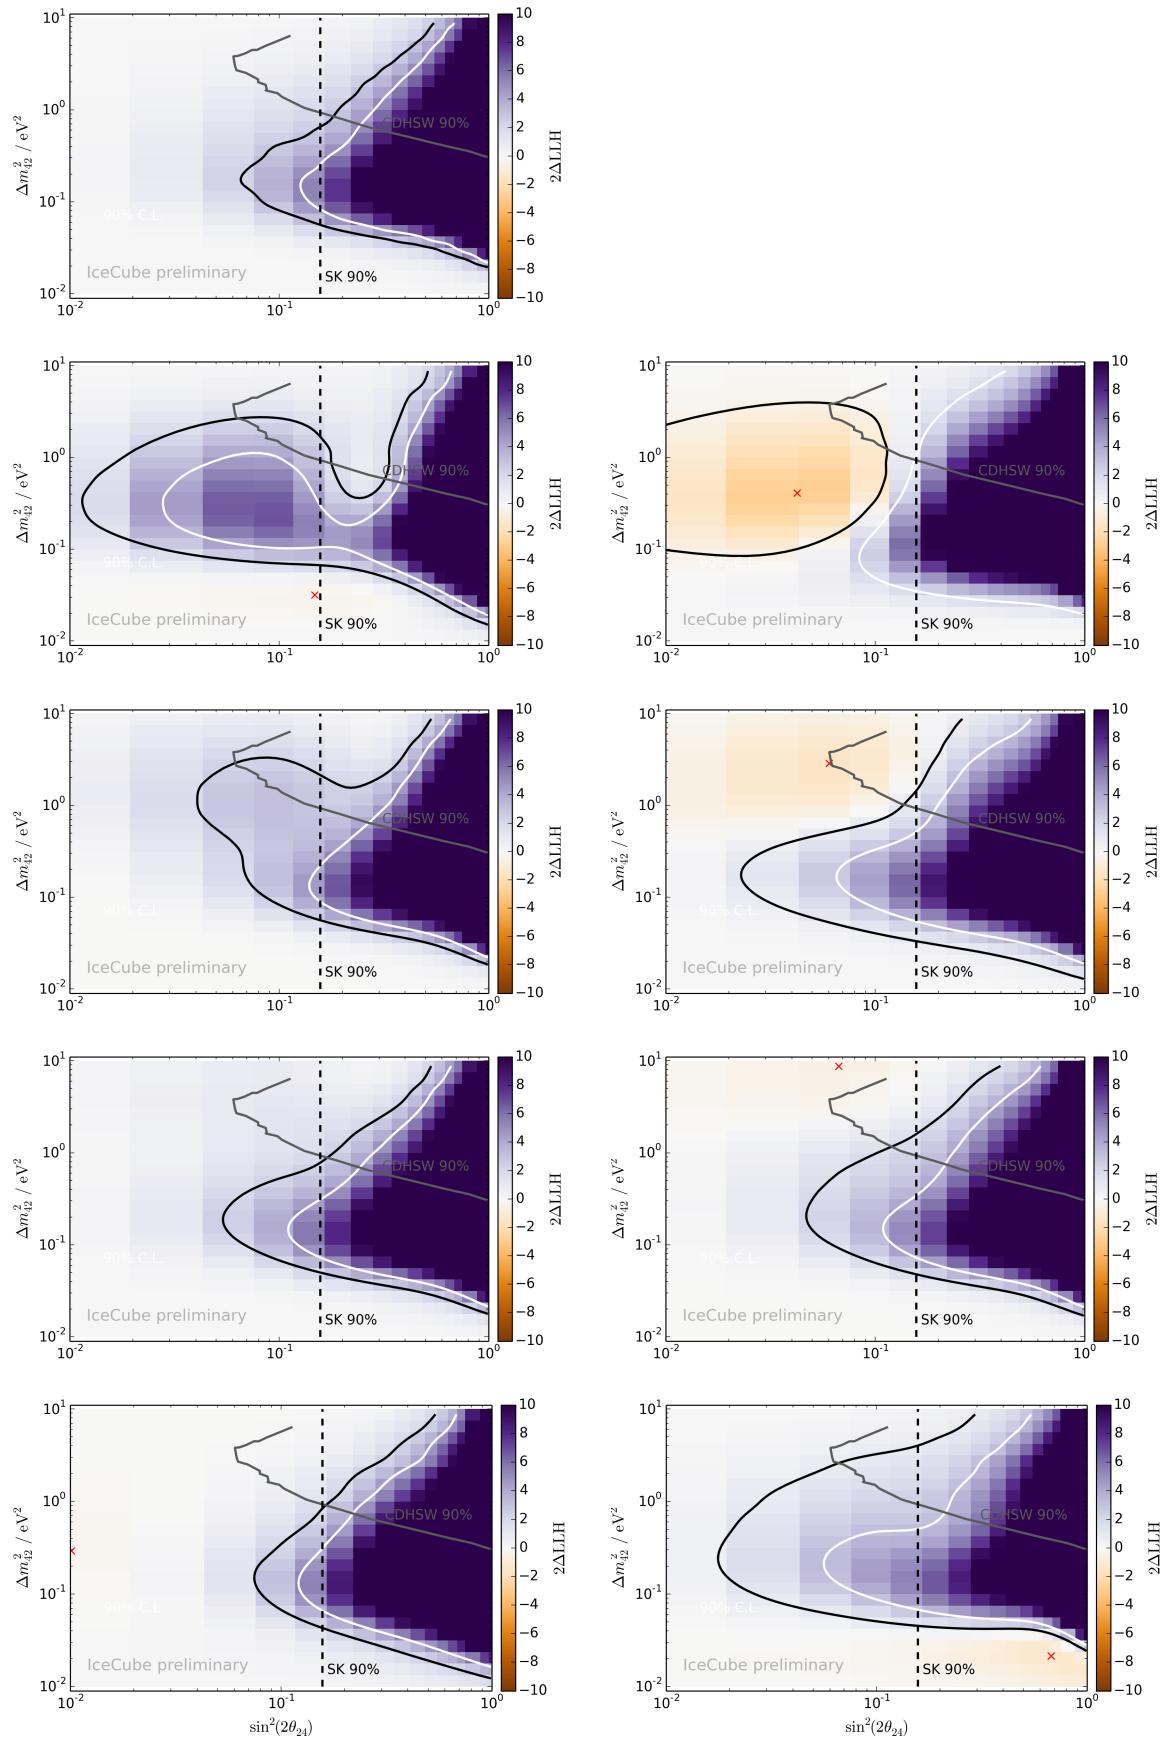


Figure F.1: Likelihood landscapes of the original Asimov expectation and manipulated ones (white: 90% C.L., black: 1σ);
left: added events; **right:** subtracted events;
top to bottom: original, lE/vZ, mE/vZ, hE/vZ, lE/mZ

APPENDIX F. 3×3 BIN MANIPULATION STUDY

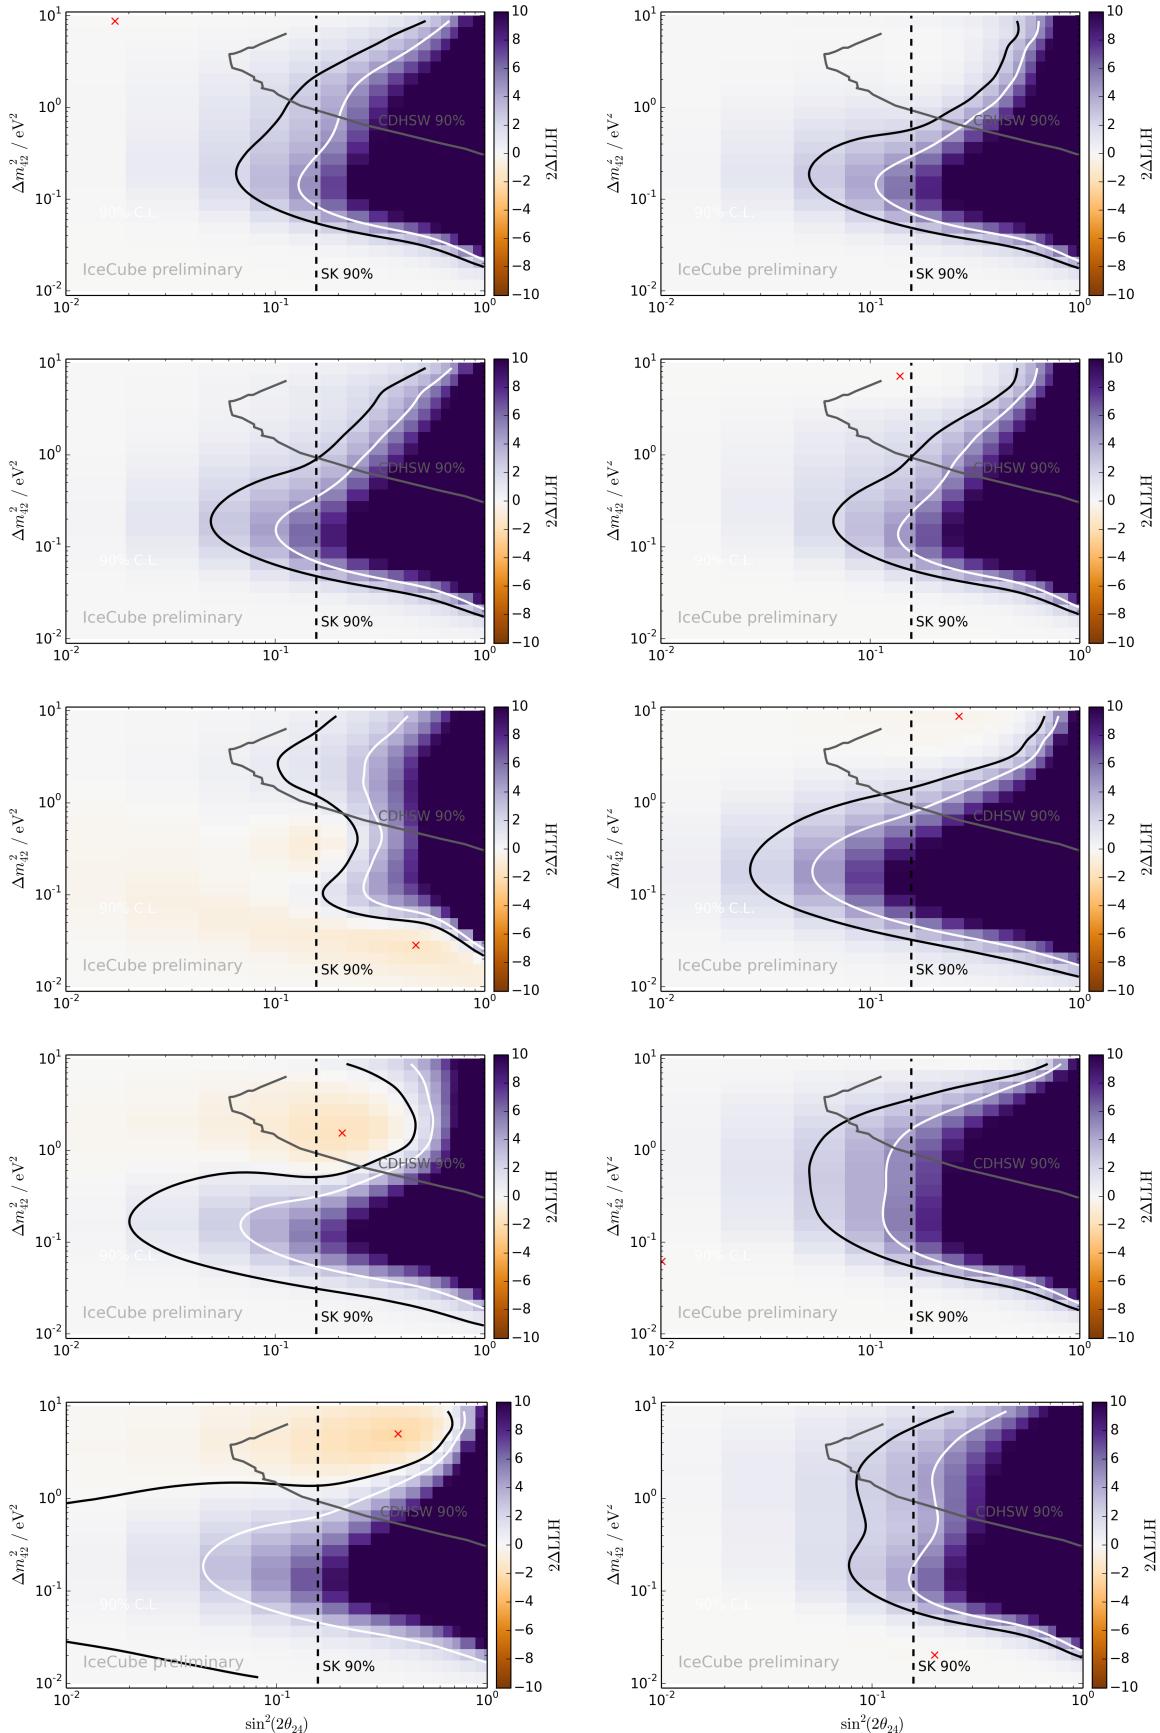


Figure F.2: Continuation of figure F.1;

156

left: added events; **right:** subtracted events;

top to bottom: mE/mZ, hE/mZ, lE/hZ, mE/hZ, hE/hZ

Appendix G

Brazilian Flag Plots

Plots such as figure 7.8 were internally dubbed Brazilian-flag plots after the color scheme used by the two big LHC detectors CMS and ATLAS for similar plots, even though the color scheme has in-between been adapted to fit the natural environment of IceCube (light blue for the central 95.5% (2σ) percentile, and dark blue for the central 68.3% (1σ) percentile).

The idea of these plots is to visualize the distribution of the (e.g., 90% C.L.) contours of pseudo experiments in a easily readable way. The natural way to think of the contours is to see them as a function of Δm_{42}^2 , because ideally (for the null hypothesis) they should map every Δm_{42}^2 to exactly one mixing angle θ_{24} that is the infimum of the mixing angles that can be excluded at this Δm_{42}^2 , as shown in the Asimov plot figure 7.1. This does not work for Δm_{42}^2 as function of θ_{24} , because the analysis sensitivity decreases for both small and large values of Δm_{42}^2 .

However, non-idealized contours often do not have this well-defined mapping even when regarded as functions of Δm_{42}^2 , as they often feature indentations such as the experimental result (figure 7.6). Also, some pseudo experiments exclude null and therefore feature closed contours for which the mapping is not well-defined. Figure G.1 shows the contours of twenty randomly picked pseudo experiments, to show that the contours often are complicated²⁰.

To produce these Brazilian-flag plots, one has to decide on an algorithm to handle contours with non-well-defined mapping. The most obvious way to produce the plots is to histogram the θ_{24} values for every value of Δm_{42}^2 , and compute median and percentiles from there. Regions with multiple contour crossings can then be handled in various ways; two extreme solutions are to histogram only the right-most (highest) or left-most (lowest) value of θ_{24} for every Δm_{42}^2 . These two algorithms produce “flags” that usually fail to visually resemble the underlying individual contours. Therefore, a more balanced algorithm was designed:

²⁰These plots were internally dubbed “Brazilian-hair plot” as reference to the similar Braizilian-flag style plots (figure 7.8) and the somewhat aligned, yet interwoven appearance of the individual contours.

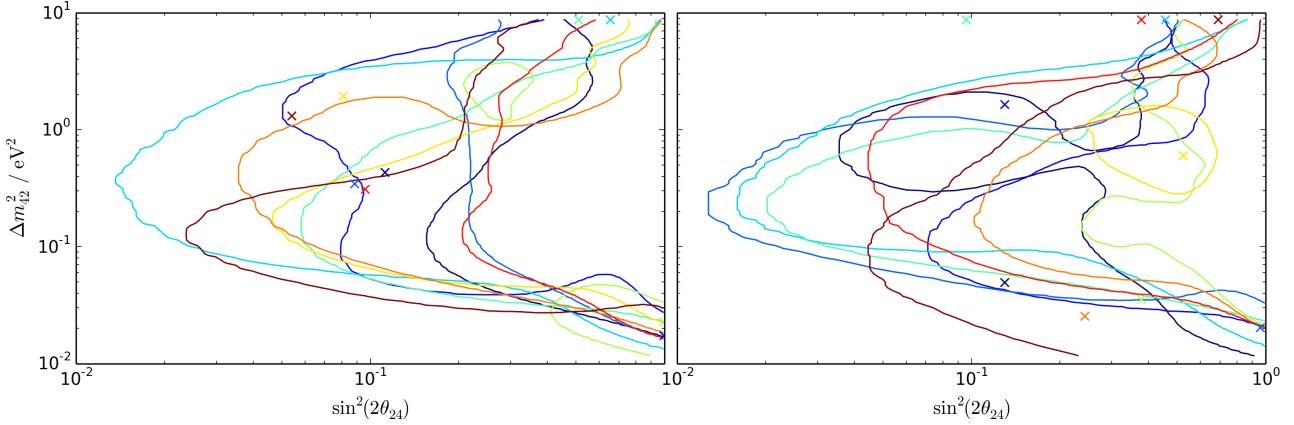


Figure G.1: Contours of twenty (2×10) randomly selected pseudo experiments for the null hypothesis, with their points of best fit.

```

for all pseudo experiment contours do
    for all  $\Delta m_{42}^2$  do
        let  $s$  be the 1d spline of  $\Lambda - \Lambda_{90\% \text{ C.L.}}$  vs.  $\theta_{24}$ 
        let  $t$  be the roots of  $s$ 
        if length of  $t$  is 1 then
            save  $t[0]$  with weight 1
        else if length of  $t$  is 0 then
            if  $s > 0$  then
                save min.  $\theta_{24}$  (i.e.,  $1^\circ$ ) with weight  $w$ 
            else
                save max.  $\theta_{24}$  (i.e.,  $42.5^\circ$ ) with weight 1
        else
            save the largest root in  $t$  and every second root from there on
            with equal weights that sum up to 1
    for all  $\Delta m_{42}^2$  do
        compute median and percentiles of saved roots, taking into account weights

```

The algorithm only counts excluding crossings (under the fairly weak assumption that the maximum mixing angle is always excluded in sensitive regions), but is still somewhat unphysical in that it effectively considers regions as homogeneously fractionally excluded while in fact they are fully excluded in some regions and not at all in others.

The weight w governs whether contours that exclude null are largely ignored or not. If it is 1 (its default value), they are taken into account normally; since usually more than 2.5% of the null experiments erroneously exclude null at the 90% C.L., this leads to the light blue area encompassing the null hypothesis. Setting w to much smaller values can solve this perceived problem, but does not make much sense physically. A cleaner method is to fully discard experiments that exclude the null hypothesis, but this does

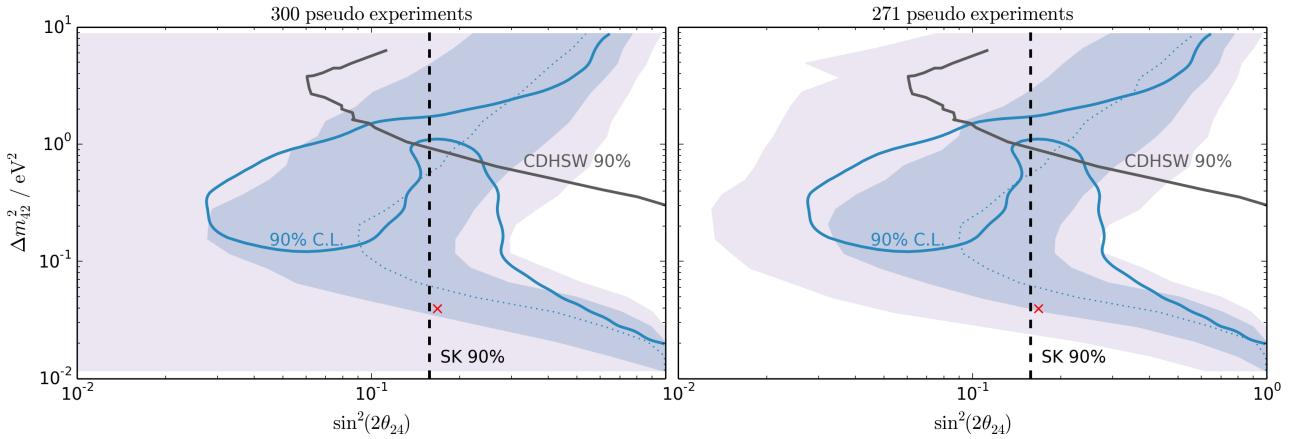


Figure G.2: Brazilian-flag plots of the final result, created by the 1d algorithm. **Left:** full ensemble; **right:** null results only.

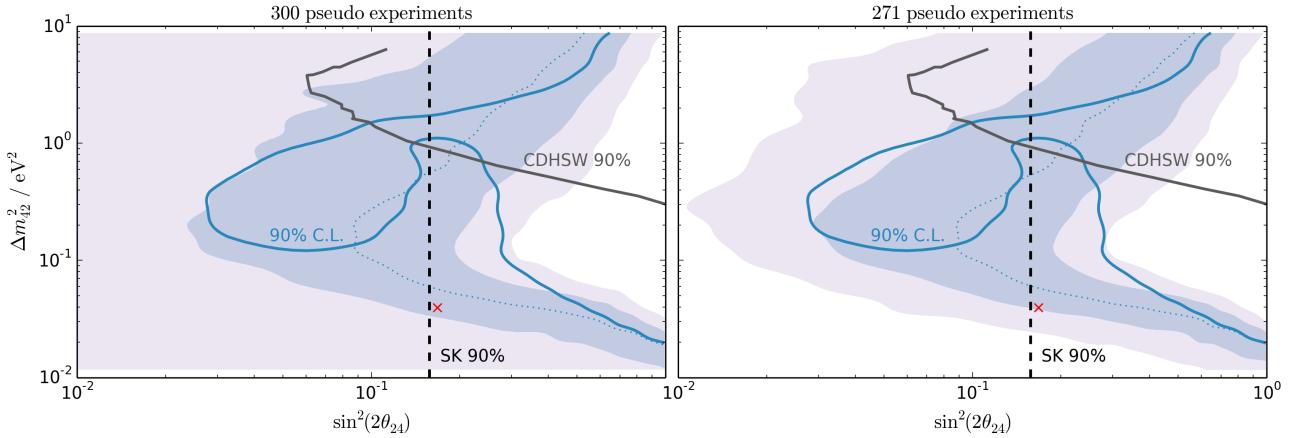


Figure G.3: Brazilian-flag plots of the final result, created by the final 2d algorithm. **Left:** full ensemble (same as figure 7.8); **right:** null results only.

not represent the full pseudo ensemble, but rather is an a-posteriori decision to visualize where contours lie if they do not exclude the null hypothesis.

Figure G.2 shows the Brazilian-flag plots analog to figure 7.8 created by this algorithm.

A better algorithm can be formulated by taking a two-dimensional approach, based on the binning of the LLH scan: The LLH landscapes of all experiments are shifted to be zero at their lowest point (including the null hypothesis). One then simply calculates the medians and percentiles over all experiments for every bin, and computes contours as lines of equal heights (e.g., using matplotlib [177]).

This algorithm is simpler and less ambiguous than the one-dimensional version and can naturally handle very complex contours, which is why it was selected to create the final Brazilian-flag plots, shown in figures 7.8 and G.3.

Acknowledgments

While writing this thesis, I was supported by many people whom I now want to thank, hoping not to forget anybody.

First of all, I want to thank my advisor and first referee Prof. Christopher Wiebusch. He kindled my interest in astroparticle physics and invited me into his workgroup and the IceCube collaboration for my diploma thesis and subsequently this work. He was always available for invaluable advice and constructive discussions, and without his mentorship this thesis would never have come to pass.

I also want to thank Prof. Thomas Hebbeker for finding the time to be my second referee. His lectures at the *Schüleruni* program strongly influenced my decision to study physics, and I greatly enjoyed his lectures in the course of my studies. I am happy neither Christopher nor you took offense at me possibly falling slightly asleep during lectures every now and then.

A very important contribution to this work was the creation of state-of-the-art Monte Carlo data for the IC-59 detector configuration, which I could not have done without the help from Michael Larson and Markus Vehring. Many thanks!

I would like to thank Martin Leuermann and Sebastian Schoenen for sharing their KDE code, as well as Martin Bissok, Christian Haack, Leif Rädel, René Reimann, Anne Schukraft, and the rest of the IceCube Aachen workgroup for the many fruitful discussions we had during all stages of my work. The workgroup and especially the Kinderzimmer provide a great friendly and professional environment, with the only downside that lunch is sometimes objectively eaten too early.

Thanks also go to the rest of the IceCube collaboration. It was a great experience to work together constructively with people from all over the world, and I am grateful that I was able to contribute a little bit to the efforts of keeping the detector running and improving the data quality even further. Very special thanks go to Janet for teaching me that rhetoric and psychology can be important in scientific discussions.

The RWTH Aachen University IT Center uncomplicatedly provided over 40 CPU years for my analysis code, to which I am very grateful. I also want to thank the many authors

ACKNOWLEDGMENTS

of Python and the Python packages that were used for this work, most prominently NumPy, SciPy, Matplotlib, and ColorPy [178], as well as Anaconda, Inc. for providing a free system-independent Python distribution [179].

Last but not least I would like to thank my girlfriend Saskia Philippen for proofreading, as well as her and my family for their support and encouragement during the last years. Sadly, my mother will no longer be able to read these words; may she rest in peace.

Declaration

Marius Wallraff hereby declares that this dissertation and the contents contained therein are his own and were generated independently, as a result of his own original research.

I hereby declare under oath:

1. This work has been fully or mainly produced during the period of doctoral student of this faculty and university;
2. If any component of this dissertation has been used for an academic degree or another qualification at this or another institution, this was clearly indicated;
3. When any other own or third-party publications were used, these were clearly named;
4. When any other own or third-party publications were cited, the source was always specified. This dissertation is entirely my own work, with the exception of such citations;
5. All major sources of support have been designated;
6. Whenever a part of this dissertation is based on collaboration with others, I clearly indicated what was developed by others and what by myself;
7. Parts of this work have been previously published in:
 - Dissertation of Anne Schukraft [114] (see chapter 3)
 - Proceedings of the 33rd Internat. Cosmic Ray Conference [120] (see chapter 7)
 - Proceedings of the 34rd Internat. Cosmic Ray Conference [121] (see chapter 7)

Aachen, October 30, 2019, _____

DECLARATION

Eidesstattliche Erklärung

Marius Wallraff erklärt hiermit, dass diese Dissertation und die darin dargelegten Inhalte die eigenen sind und selbstständig, als Ergebnis der eigenen originären Forschung, generiert wurden.

Hiermit erkläre ich an Eides statt:

1. Diese Arbeit wurde vollständig oder größtenteils in der Phase als Doktorand dieser Fakultät und Universität angefertigt;
2. Sofern irgendein Bestandteil dieser Dissertation zuvor für einen akademischen Abschluss oder eine andere Qualifikation an dieser oder einer anderen Institution verwendet wurde, wurde dies klar angezeigt;
3. Wenn immer andere eigene oder Veröffentlichungen Dritter herangezogen wurden, wurden diese klar benannt;
4. Wenn aus anderen eigenen oder Veröffentlichungen Dritter zitiert wurde, wurde stets die Quelle hierfür angegeben. Diese Dissertation ist vollständig meine eigene Arbeit, mit der Ausnahme solcher Zitate;
5. Alle wesentlichen Quellen von Unterstützung wurden benannt;
6. Wenn immer ein Teil dieser Dissertation auf der Zusammenarbeit mit anderen basiert, wurde von mir klar gekennzeichnet, was von anderen und was von mir selbst erarbeitet wurde;
7. Teile dieser Arbeit wurden zuvor veröffentlicht und zwar in:
 - Dissertation von Anne Schukraft [114] (siehe Kap. 3)
 - Proceedings of the 33rd Internat. Cosmic Ray Conference [120] (siehe Kap. 7)
 - Proceedings of the 34rd Internat. Cosmic Ray Conference [121] (siehe Kap. 7)

Aachen, 30. Oktober 2019, _____

EIDESSTATTLICHE ERKLÄRUNG

Bibliography

- [1] Sun Tzu and Cleary, T.: *The Art of War: Complete Texts and Commentaries*. Shambhala, **2005**. ISBN 9780834827301.
URL https://books.google.de/books?id=trpSpjQjg_cC
- [2] Chadwick, J.: “Intensitätsverteilung im magnetischen Spectrum der β -Strahlen von radium B + C”. *Verhandl. Dtsc. Phys. Ges.*, vol. 16, p. 383, **1914**.
URL <http://cds.cern.ch/record/262756>
- [3] Cowan, C. L., Reines, F., Harrison, F. B., Kruse, H. W., and McGuire, A. D.: “Detection of the Free Neutrino: a Confirmation”. *Science*, vol. 124(3212), pp. 103–104, **1956**. doi:10.1126/science.124.3212.103.
URL <http://www.sciencemag.org/content/124/3212/103.short>
- [4] Danby, G., Gaillard, J. M., Goulianos, K. A., Lederman, L. M., Mistry, N. B., Schwartz, M., and Steinberger, J.: “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. *Phys. Rev. Lett.*, vol. 9, pp. 36–44, **1962**. doi:10.1103/PhysRevLett.9.36.
URL <https://www.physik.uni-bielefeld.de/~yorks/pro13/PhysRevLett.9.36.pdf>
- [5] Kodama, K. *et al.*: “Observation of tau neutrino interactions”. *Physics Letters B*, vol. 504(3), pp. 218 – 224, **2001**. ISSN 0370-2693. doi:10.1016/S0370-2693(01)00307-0.
URL <http://arxiv.org/abs/hep-ex/0012035>
- [6] Anicin, I. V.: “The Neutrino – Its Past, Present and Future”. *SFIN (Institute of Physics, Belgrade) year XV, Series A: Conferences*, vol. No. A2, pp. 3–59, **2002**.
URL <http://arxiv.org/abs/physics/0503172>
- [7] Olive, K. A. *et al.*: “Review of Particle Physics”. *Chin. Phys.*, vol. C38, p. 090001, **2014**. doi:10.1088/1674-1137/38/9/090001.
URL <http://pdg.lbl.gov/>
- [8] Wurm, M.: “Solar Neutrino Spectroscopy”. *Phys. Rept.*, vol. 685, pp. 1–52, **2017**. doi:10.1016/j.physrep.2017.04.002.
URL <https://arxiv.org/abs/1704.06331>

BIBLIOGRAPHY

[9] Agostini, M. *et al.*: “Spectroscopy of geoneutrinos from 2056 days of Borexino data”. *Phys. Rev.*, vol. D92(3), p. 031101, **2015**. doi:10.1103/PhysRevD.92.031101.
URL <https://arxiv.org/abs/1506.04610>

[10] L’Annunziata, M.: *Handbook of Radioactivity Analysis*. Elsevier Science, **2012**. ISBN 9780123848741.
URL <https://books.google.com.au/books?id=S4FrejzJy0cC>

[11] Mirizzi, A., Tamborra, I., Janka, H.-T., Saviano, N., Scholberg, K., Bollig, R., Hudepohl, L., and Chakraborty, S.: “Supernova Neutrinos: Production, Oscillations and Detection”. *Riv. Nuovo Cim.*, vol. 39(1-2), pp. 1–112, **2016**. doi:10.1393/ncr/i2016-10120-8.
URL <https://arxiv.org/abs/1508.00785>

[12] Aartsen, M. G. *et al.*: “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. *Science*, vol. 361(6398), p. eaat1378, **2018**. doi:10.1126/science.aat1378.
URL <https://arxiv.org/abs/1807.08816>

[13] Rädel, L. C.: *Measurement of High-Energy Muon Neutrinos with the IceCube Neutrino Observatory*. Ph.D. thesis, RWTH Aachen University, **November 2017**.
URL http://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaaaaayrjby

[14] Schoenen, S. J.: *Discovery and Characterization of a Diffuse Astrophysical Muon Neutrino Flux with the IceCube Neutrino Observatory*. Ph.D. thesis, RWTH Aachen University, **Februar 2017**.
URL http://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaaaaawaarf

[15] Blennow, M., Edsjo, J., and Ohlsson, T.: “Neutrinos from WIMP annihilations using a full three-flavor Monte Carlo”. *JCAP*, vol. 0801, p. 021, **2008**. doi:10.1088/1475-7516/2008/01/021.
URL <https://arxiv.org/abs/0709.3898>

[16] Li, Y.-F.: “Detection prospects of the cosmic neutrino background”. *International Journal of Modern Physics A*, vol. 30(12), p. 1530031, **2015**. doi:10.1142/S0217751X15300318.
URL <http://arxiv.org/abs/1504.03966>

[17] Faessler, A., Hodak, R., Kovalenko, S., and Šimkovic, F.: “Search for the Cosmic Neutrino Background and KATRIN”. *Rom. J. Phys.*, vol. 58(9-10), pp. 1221–1231, **2013**.
URL <https://arxiv.org/abs/1304.5632>

- [18] Anchordoqui, L. A.: “Ultra-High-Energy Cosmic Rays”. *arXiv astro-ph*, **2018**.
URL <https://arxiv.org/abs/1807.09645>
- [19] Dembinski, H. P., Engel, R., Fedynitch, A., Gaisser, T., Riehn, F., and Stanev, T.: “Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10^{11} GeV”. *PoS*, vol. ICRC2017, p. 533, **2018**. doi:10.22323/1.301.0533.
URL <https://arxiv.org/abs/1711.11432>
- [20] Ave, M., Boyle, P. J., Gahbauer, F., Hoppner, C., Horandel, J. R., Ichimura, M., Muller, D., and Romero-Wolf, A.: “Composition of Primary Cosmic-Ray Nuclei at High Energies”. *Astrophys. J.*, vol. 678, p. 262, **2008**. doi:10.1086/529424.
URL <http://arxiv.org/abs/0801.0582>
- [21] Fowler, J. W., Fortson, L. F., Jui, C. C. H., Kieda, D. B., Ong, R. A., Pryke, C. L., and Sommers, P.: “A Measurement of the cosmic ray spectrum and composition at the knee”. *Astropart. Phys.*, vol. 15, pp. 49–64, **2001**. doi:10.1016/S0927-6505(00)00139-0.
URL <http://arxiv.org/abs/astro-ph/0003190>
- [22] Bertaina, M. *et al.*: “The cosmic ray energy spectrum in the range 10^{16} – 10^{18} eV measured by KASCADE-Grande”. *Astrophysics and Space Sciences Transactions*, vol. 7(2), pp. 229–234, **2011**. doi:10.5194/asta-7-229-2011.
URL <http://www.astrophys-space-sci-trans.net/7/229/2011/>
- [23] Aab, A. *et al.*: “Measurement of the cosmic ray spectrum above 4×10^{18} eV using inclined events detected with the Pierre Auger Observatory”. *JCAP*, vol. 1508, p. 049, **2015**. doi:10.1088/1475-7516/2015/08/049.
URL <http://arxiv.org/abs/1503.07786>
- [24] Ruffini, R., Vereshchagin, G. V., and Xue, S. S.: “Cosmic absorption of ultra high energy particles”. *Astrophys. Space Sci.*, vol. 361, p. 82, **2016**. doi:10.1007/s10509-016-2668-5.
URL <http://arxiv.org/abs/1503.07749>
- [25] Dermer, C. D. and Atoyan, A.: “Ultrahigh energy cosmic rays, cascade gamma-rays, and high-energy neutrinos from gamma-ray bursts”. *New J. Phys.*, vol. 8, p. 122, **2006**. doi:10.1088/1367-2630/8/7/122.
URL <http://arxiv.org/abs/astro-ph/0606629>
- [26] Gaisser, T. K. and Honda, M.: “Flux of atmospheric neutrinos”. *Ann. Rev. Nucl. Part. Sci.*, vol. 52, pp. 153–199, **2002**. doi:10.1146/annurev.nucl.52.050102.090645.
URL <http://arxiv.org/abs/hep-ph/0203272>
- [27] Gaisser, T. K. and Stanev, T.: “Path length distributions of atmospheric neutrinos”. *Phys. Rev. D*, vol. 57, pp. 1977–1982, **Feb 1998**. doi:10.1103/PhysRevD.57.1977.
URL <http://arxiv.org/pdf/astro-ph/9708146.pdf>

BIBLIOGRAPHY

[28] Gaisser, T.: *Cosmic Rays and Particle Physics*. Cambridge University Press, **1991**. ISBN 9780521339315.
URL <https://books.google.de/books?id=5JMmJI7Mm30C>

[29] Battistoni, G., Ferrari, A., Lipari, P., Montaruli, T., Sala, P. R., and Rancati, T.: “A Three-dimensional calculation of atmospheric neutrino flux”. *Astropart. Phys.*, vol. 12, pp. 315–333, **2000**. doi:10.1016/S0927-6505(99)00110-3.
URL <http://arxiv.org/abs/hep-ph/9907408>

[30] Abbasi, R. *et al.*: “Observation of Anisotropy in the Galactic Cosmic-Ray Arrival Directions at 400 TeV with IceCube”. *The Astrophysical Journal*, vol. 746(1), p. 33, **2012**.
URL <http://stacks.iop.org/0004-637X/746/i=1/a=33>

[31] Halzen, F. and Martin, A. D.: *Quark & Leptons: An Introductory Course In Modern Particle Physics*. John Wiley & Sons, **2008**.
URL http://web-docs.gsi.de/~gkalicy/books/Quark_Leptons.pdf

[32] Honda, M., Kajita, T., Kasahara, K., Midorikawa, S., and Sanuki, T.: “Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data”. *Phys. Rev.*, vol. D75, p. 043006, **2007**. doi:10.1103/PhysRevD.75.043006.
URL <http://arxiv.org/abs/astro-ph/0611418>

[33] Honda, M., Kajita, T., Kasahara, K., Midorikawa, S., and Sanuki, T.: “HKKMS2006 flux tables corresponding to astro-ph/0611418”.
URL <http://www.icrr.u-tokyo.ac.jp/~mhonda/nflx2006/>

[34] Coppi, P.: “How Do We know Antimater is Absent?” *eConf*, vol. C040802, p. L017, **2004**.
URL <http://www.slac.stanford.edu/econf/C040802/papers/L017.PDF>

[35] Adamson, P. *et al.*: “Measurements of atmospheric neutrinos and antineutrinos in the MINOS Far Detector”. *Phys. Rev.*, vol. D86, p. 052007, **2012**. doi:10.1103/PhysRevD.86.052007.
URL <http://arxiv.org/abs/1208.2915>

[36] Khachatryan, V. *et al.*: “Measurement of the charge ratio of atmospheric muons with the CMS detector”. *Phys. Lett.*, vol. B692, pp. 83–104, **2010**. doi:10.1016/j.physletb.2010.07.033.
URL <http://arxiv.org/abs/1005.5332>

[37] Levy, J.-M.: “Kinematics of an off axis neutrino beam”. *arXiv hep-ex*, **2010**.
URL <http://arxiv.org/abs/1005.0574>

[38] Barr, G. D., Gaisser, T. K., Robbins, S., and Stanev, T.: “Uncertainties in Atmospheric Neutrino Fluxes”. *Phys. Rev.*, vol. D74, p. 094009, **2006**. doi: 10.1103/PhysRevD.74.094009.
URL <http://arxiv.org/abs/astro-ph/0611266>

[39] Cohen, A. G., Glashow, S. L., and Ligeti, Z.: “Disentangling Neutrino Oscillations”. *Phys. Lett.*, vol. B678, pp. 191–196, **2009**. doi:10.1016/j.physletb.2009.06.020.
URL <http://arxiv.org/abs/0810.4602>

[40] Esposito, S.: “On the indistinguishability of Majorana from Dirac neutrino propagation in a stellar medium”. *Nuovo Cim.*, vol. B111, pp. 1449–1459, **1996**. doi: 10.1007/BF02741484.
URL <http://arxiv.org/abs/hep-ph/9607371>

[41] Kuo, T. K. and Pantaleone, J.: “Neutrino oscillations in matter”. *Rev. Mod. Phys.*, vol. 61, pp. 937–979, **Oct 1989**. doi:10.1103/RevModPhys.61.937.
URL <http://www.nikhef.nl/~{}h84/matterosc.pdf>

[42] Schmitz, N.: *Neutrinophysik*. Teubner Studienbücher Physik. Vieweg+Teubner Verlag, **2013**. ISBN 9783322801142.
URL <https://books.google.de/books?id=Ud2SxVnGIgIC>

[43] Nunokawa, H., Parke, S. J., and Valle, J. W. F.: “CP Violation and Neutrino Oscillations”. *Prog. Part. Nucl. Phys.*, vol. 60, pp. 338–402, **2008**. doi:10.1016/j.ppnp.2007.10.001.
URL <https://arxiv.org/abs/0710.0554v2>

[44] Gonzalez-Garcia, M. C., Maltoni, M., and Schwetz, T.: “Updated fit to three neutrino mixing: status of leptonic CP violation”. *JHEP*, vol. 11, p. 052, **2014**. doi: 10.1007/JHEP11(2014)052.
URL <http://arxiv.org/abs/1409.5439>

[45] Hooper, D., Morgan, D., and Winstanley, E.: “Probing quantum decoherence with high-energy neutrinos”. *Phys. Lett.*, vol. B609, pp. 206–211, **2005**. doi:10.1016/j.physletb.2005.01.034.
URL <https://arxiv.org/abs/hep-ph/0410094>

[46] Wolfenstein, L.: “Matter effects on neutrino propagation”. In *Neutrinos from the Lab, the Sun, and the cosmos: Proceedings, 2000 SLAC Summer Institute on Particle Physics (SSI 00): Stanford, USA, 14-25 Aug 2000*, p. ch05, **2000**.
URL <http://www.slac.stanford.edu/gen/meeting/ssi/2000b/writeups/wolfenstein/wolfenstein.pdf>

[47] Giunti, C. and Kim, C.: *Fundamentals of Neutrino Physics and Astrophysics*. OUP Oxford, **2007**. ISBN 9780198508717.
URL <https://books.google.de/books?id=HioSDAAAQBAJ>

BIBLIOGRAPHY

[48] Dziewonski, A. M. and Anderson, D. L.: “Preliminary reference Earth model”. *Physics of the Earth and Planetary Interiors*, vol. 25(4), pp. 297 – 356, **1981**. ISSN 0031-9201. doi:[http://dx.doi.org/10.1016/0031-9201\(81\)90046-7](http://dx.doi.org/10.1016/0031-9201(81)90046-7).
URL <http://www.sciencedirect.com/science/article/pii/0031920181900467>

[49] Grotzinger, J., Jordan, T., and Schweizer, V.: *Press/Siever Allgemeine Geologie*. Springer Berlin Heidelberg, **2016**. ISBN 9783662483428.
URL <https://books.google.de/books?id=9fvcDAAAQBAJ>

[50] Chu, S. Y. F., Ekström, L. P., and Firestone, R. B.: “The Lund/LBNL Nuclear Data Search”, **Feb 1999**.
URL <http://nucleardata.nuclear.lu.se/toi/>

[51] Choubey, S.: “Signature of sterile species in atmospheric neutrino data at neutrino telescopes”. *JHEP*, vol. 12, p. 014, **2007**. doi:[10.1088/1126-6708/2007/12/014](https://doi.org/10.1088/1126-6708/2007/12/014).
URL <https://arxiv.org/abs/0709.1937>

[52] Akhmedov, E. K.: “Parametric resonance in neutrino oscillations in matter”. *Pramana*, vol. 54, pp. 47–63, **2000**. doi:[10.1007/s12043-000-0006-4](https://doi.org/10.1007/s12043-000-0006-4).
URL <https://arxiv.org/abs/hep-ph/9907435>

[53] Merfeld, K. M. and Latimer, D. C.: “Parametric enhancement of flavor oscillation in a three-neutrino framework”. *Phys. Rev.*, vol. C90(6), p. 065502, **2014**. doi:[10.1103/PhysRevC.90.065502](https://doi.org/10.1103/PhysRevC.90.065502).
URL <https://arxiv.org/abs/1412.2728>

[54] Schael, S. *et al.*: “Precision electroweak measurements on the Z resonance”. *Phys. Rept.*, vol. 427, pp. 257–454, **2006**. doi:[10.1016/j.physrep.2005.12.006](https://doi.org/10.1016/j.physrep.2005.12.006).
URL <https://arxiv.org/abs/hep-ex/0509008>

[55] Aseev, V. N. *et al.*: “An upper limit on electron antineutrino mass from Troitsk experiment”. *Phys. Rev.*, vol. D84, p. 112003, **2011**. doi:[10.1103/PhysRevD.84.112003](https://doi.org/10.1103/PhysRevD.84.112003).
URL <https://arxiv.org/abs/1108.5034>

[56] Kraus, C. *et al.*: “Final results from phase II of the Mainz neutrino mass search in tritium beta decay”. *Eur. Phys. J.*, vol. C40, pp. 447–468, **2005**. doi:[10.1140/epjc/s2005-02139-7](https://doi.org/10.1140/epjc/s2005-02139-7).
URL <https://arxiv.org/abs/hep-ex/0412056>

[57] Abazajian, K. N. *et al.*: “Light Sterile Neutrinos: A White Paper”. *arXiv hep-ex*, **2012**.
URL <https://arxiv.org/abs/1204.5379>

[58] Parke, S. and Ross-Lonergan, M.: “Unitarity and the three flavor neutrino mixing matrix”. *Phys. Rev.*, vol. D93(11), p. 113009, **2016**. doi:[10.1103/PhysRevD.93.113009](https://doi.org/10.1103/PhysRevD.93.113009).

113009. And corresponding NuFact 2016 proceedings.
URL <https://arxiv.org/abs/1508.05095>

[59] Antusch, S. and Fischer, O.: “Probing the non-unitarity of the leptonic mixing matrix at the CEPC”. *International Journal of Modern Physics A*, vol. 31(33), p. 1644006, **2016**.
URL <https://arxiv.org/abs/1604.00208>

[60] Mohapatra, R. and Pal, P.: *Massive Neutrinos in Physics and Astrophysics*. World Scientific Lecture Notes in Physics. World Scientific Publishing Co. Pte. Ltd., **2004**. ISBN 9789814488129.
URL <https://books.google.de/books?id=HsbUCgAAQBAJ>

[61] Aguilar-Arevalo, A. *et al.*: “Evidence for neutrino oscillations from the observation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam”. *Phys. Rev.*, vol. D64, p. 112007, **2001**. doi:10.1103/PhysRevD.64.112007.
URL <https://arxiv.org/abs/hep-ex/0104049>

[62] Aguilar-Arevalo, A. A. *et al.*: “Improved Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment”. *Phys. Rev. Lett.*, vol. 110, p. 161801, **2013**. doi:10.1103/PhysRevLett.110.161801.
URL <https://arxiv.org/abs/1303.2588>

[63] Giunti, C. and Laveder, M.: “Statistical Significance of the Gallium Anomaly”. *Phys. Rev.*, vol. C83, p. 065504, **2011**. doi:10.1103/PhysRevC.83.065504.
URL <https://arxiv.org/abs/1006.3244>

[64] Mention, G., Fechner, M., Lasserre, T., Mueller, T. A., Lhuillier, D., Cribier, M., and Letourneau, A.: “The Reactor Antineutrino Anomaly”. *Phys. Rev.*, vol. D83, p. 073006, **2011**. doi:10.1103/PhysRevD.83.073006.
URL <https://arxiv.org/abs/1101.2755>

[65] Gariazzo, S., Giunti, C., Laveder, M., and Li, Y. F.: “Updated Global 3+1 Analysis of Short-BaseLine Neutrino Oscillations”. *JHEP*, vol. 06, p. 135, **2017**. doi:10.1007/JHEP06(2017)135.
URL <https://arxiv.org/abs/1703.00860>

[66] Collin, G. H., Argüelles, C. A., Conrad, J. M., and Shaevitz, M. H.: “Sterile Neutrino Fits to Short Baseline Data”. *Nucl. Phys.*, vol. B908, pp. 354–365, **2016**. doi:10.1016/j.nuclphysb.2016.02.024.
URL <https://arxiv.org/abs/1602.00671>

[67] Munroe, R.: “xkcd”, **2006**.
URL <https://xkcd.com/>

BIBLIOGRAPHY

[68] Aguilar-Arevalo, A. A. *et al.*: “Significant Excess of ElectronLike Events in the MiniBooNE Short-Baseline Neutrino Experiment”. *Phys. Rev. Lett.*, vol. 121(22), p. 221801, **2018**. doi:10.1103/PhysRevLett.121.221801.
URL <https://arxiv.org/abs/1805.12028>

[69] An, F. P. *et al.*: “Evolution of the Reactor Antineutrino Flux and Spectrum at Daya Bay”. *Phys. Rev. Lett.*, vol. 118(25), p. 251801, **2017**. doi:10.1103/PhysRevLett.118.251801.
URL <https://arxiv.org/abs/1704.01082>

[70] Dentler, M., Hernández-Cabezudo, I., Kopp, J., Maltoni, M., and Schwetz, T.: “Sterile neutrinos or flux uncertainties? — Status of the reactor anti-neutrino anomaly”. *JHEP*, vol. 11, p. 099, **2017**. doi:10.1007/JHEP11(2017)099.
URL <https://arxiv.org/abs/1709.04294>

[71] Hayes, A. C., Jungman, G., McCutchan, E. A., Sonzogni, A. A., Garvey, G. T., and Wang, X.: “Analysis of the Daya Bay Reactor Antineutrino Flux Changes with Fuel Burnup”. *Phys. Rev. Lett.*, vol. 120(2), p. 022503, **2018**. doi:10.1103/PhysRevLett.120.022503.
URL <https://arxiv.org/abs/1707.07728>

[72] Mention, G., Vivier, M., Gaffiot, J., Lasserre, T., Letourneau, A., and Materna, T.: “Reactor antineutrino shoulder explained by energy scale nonlinearities?” *Phys. Lett.*, vol. B773, pp. 307–312, **2017**. doi:10.1016/j.physletb.2017.08.035.
URL <https://arxiv.org/abs/1705.09434>

[73] Ho, C. M. and Scherrer, R. J.: “Sterile Neutrinos and Light Dark Matter Save Each Other”. *Phys. Rev.*, vol. D87(6), p. 065016, **2013**. doi:10.1103/PhysRevD.87.065016.
URL <https://arxiv.org/abs/1212.1689>

[74] Chu, X., Dasgupta, B., Dentler, M., Kopp, J., and Saviano, N.: “Sterile Neutrinos with Secret Interactions – Cosmological Discord?” *arXiv hep-ph*, **2018**.
URL <https://arxiv.org/abs/1806.10629>

[75] Rott, C.: “Status and Prospects of the JSNS2 Experiment”. In *39th International Conference on High Energy Physics (ICHEP 2018) Seoul, Gangnam-Gu, Korea, Republic of, July 4-11, 2018*, **2018**.
URL <https://arxiv.org/abs/1811.03321>

[76] Mertens, S.: “Status of the KATRIN Experiment and Prospects to Search for keV-mass Sterile Neutrinos in Tritium β -decay”. *Phys. Procedia*, vol. 61, pp. 267–273, **2015**. doi:10.1016/j.phpro.2014.12.043.
URL http://inspirehep.net/record/1357702/files/10.1016_j.phpro.2014.12.043.pdf

[77] Aartsen, M. G. *et al.*: “The IceCube Neutrino Observatory: Instrumentation and Online Systems”. *JINST*, vol. 12(03), p. P03012, **2017**. doi:10.1088/1748-0221/12/03/P03012.
URL <https://arxiv.org/abs/1612.05093>

[78] Andres, E. *et al.*: “Observation of high-energy neutrinos using Cherenkov detectors embedded deep in Antarctic ice”. *Nature*, vol. 410, pp. 441–443, **2001**. doi:10.1038/35068509.
URL <https://www.nature.com/articles/35068509>

[79] Abe, K. *et al.*: “Hyper-Kamiokande Design Report”. *arXiv physics.ins-det*, **2018**.
URL <https://arxiv.org/abs/1805.04163>

[80] Aslanides, E. *et al.*: “A deep sea telescope for high-energy neutrinos”. *arXiv astro-ph*, **1999**.
URL <https://arxiv.org/abs/astro-ph/9907432>

[81] Feynman, R., Leighton, R., and Sands, M.: *The Feynman Lectures on Physics: Mainly mechanics, radiation, and heat*. Basic Books. Basic Books, **2011**. ISBN 9780465024933.
URL <https://books.google.de/books?id=bDF-uoUmttUC>

[82] Čerenkov, P. A.: “Visible Radiation Produced by Electrons Moving in a Medium with Velocities Exceeding that of Light”. *Phys. Rev.*, vol. 52, pp. 378–379, **Aug 1937**. doi:10.1103/PhysRev.52.378.
URL <http://link.aps.org/doi/10.1103/PhysRev.52.378>

[83] Wallraff, M.: “Design, Implementation and Test of a New Feature Extractor for the IceCube Neutrino Observatory”, **2010**. Diplomarbeit.
URL https://www.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaakaiil

[84] Frank, I. M. and Tamm, I.: “Coherent visible radiation of fast electrons passing through matter”. *C. R. Acad. Sci. URSS*, vol. 14, pp. 109–114, **1937**.
URL http://link.springer.com/chapter/10.1007%2F978-3-642-74626-0_2

[85] Warren, S. G.: “Optical constants of ice from the ultraviolet to the microwave”. *Appl. Opt.*, vol. 23(8), pp. 1206–1225, **Apr 1984**. doi:10.1364/AO.23.001206.
URL <http://ao.osa.org/abstract.cfm?URI=ao-23-8-1206>

[86] Aartsen, M. G. *et al.*: “Measurement of South Pole ice transparency with the IceCube LED calibration system”. *Nucl. Instrum. Meth.*, vol. A711, pp. 73–89, **2013**. doi:10.1016/j.nima.2013.01.054.
URL <https://arxiv.org/abs/1301.5361>

[87] Formaggio, J. A. and Zeller, G. P.: “From eV to EeV: Neutrino Cross Sections Across Energy Scales”. *Rev. Mod. Phys.*, vol. 84, p. 1307, **2012**. doi:10.1103/RevModPhys.

84.1307.
URL <https://arxiv.org/abs/1305.7513>

[88] Glashow, S. L.: “Resonant Scattering of Antineutrinos”. *Phys. Rev.*, vol. 118, pp. 316–317, **Apr 1960**. doi:10.1103/PhysRev.118.316.
URL <http://link.aps.org/doi/10.1103/PhysRev.118.316>

[89] Gandhi, R., Quigg, C., Reno, M. H., and Sarcevic, I.: “Ultrahigh-energy neutrino interactions”. *Astropart. Phys.*, vol. 5, pp. 81–110, **1996**. doi:10.1016/0927-6505(96)00008-4.
URL <https://arxiv.org/abs/hep-ph/9512364>

[90] Learned, J. G. and Mannheim, K.: “High-energy neutrino astrophysics”. *Annual Review of Nuclear and Particle Science*, vol. 50(1), pp. 679–749, **2000**. doi:10.1146/annurev.nucl.50.1.679
URL <http://dx.doi.org/10.1146/annurev.nucl.50.1.679>

[91] Groom, D. E.: “Degradation of resolution in a homogeneous dual readout hadronic calorimeter”. *Nucl. Instrum. Meth.*, vol. A705, pp. 24–31, **2013**. doi:10.1016/j.nima.2012.12.080.
URL <https://arxiv.org/abs/1210.2334v2>

[92] Kowalski, M.: “On the Čerenkov light emission of hadronic and electro-magnetic cascades”. *internal report AMANDA-IR/20020803*, DESY-Zeuthen, **Aug 2002**.
URL <https://internal.icecube.wisc.edu/reports/amanda/data/20020803-track.pdf>

[93] Voigt, B.: *Sensitivity of the IceCube detector for ultra-high energy electron-neutrino events*. Ph.D. thesis, Humboldt-Universität zu Berlin, Mathematisch-Naturwissenschaftliche Fakultät I, **Nov 2008**.
URL <http://edoc.hu-berlin.de/docviews/abstract.php?id=29421>

[94] Baier, V. N. and Katkov, V. M.: “Concept of formation length in radiation theory”. *Phys. Rept.*, vol. 409, pp. 261–359, **2005**. doi:10.1016/j.physrep.2004.11.003.
URL <https://arxiv.org/abs/hep-ph/0309211>

[95] Chirkin, D. and Rhode, W.: “Muon Monte Carlo: A High-precision tool for muon propagation through matter”. *arXiv hep-ph*, **2004**.
URL <https://arxiv.org/abs/hep-ph/0407075>

[96] Lynch, G. R. and Dahl, O. I.: “Approximations to multiple Coulomb scattering”. *Nucl. Instrum. Methods B*, vol. 58(1), pp. 6 – 10, **1991**. ISSN 0168-583X. doi: [http://dx.doi.org/10.1016/0168-583X\(91\)95671-Y](http://dx.doi.org/10.1016/0168-583X(91)95671-Y).
URL <http://www.sciencedirect.com/science/article/pii/0168583X9195671Y>

[97] Dutta, S. I., Huang, Y., and Reno, M. H.: “Tau neutrino propagation and tau energy loss”. *Phys. Rev.*, vol. D72, p. 013005, **2005**. doi:10.1103/PhysRevD.72.013005.
URL <https://arxiv.org/abs/hep-ph/0504208>

[98] Hallen, P.: *On the Measurement of High-Energy Tau Neutrinos with IceCube*. Master's thesis, RWTH Aachen, III. Physikalisches Institut B, **Nov 2013**. URL https://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaapwhzq

[99] Abbasi, R. *et al.*: “The Design and Performance of IceCube DeepCore”. *Astropart. Phys.*, vol. 35, pp. 615–624, **2012**. doi:10.1016/j.astropartphys.2012.01.004. URL <https://arxiv.org/abs/1109.6096>

[100] Погода и климат: “Weather and Climate – The Climate of Amundsen–Scott”, **2017**. URL <http://www.pogodaiklimat.ru/climate2/89009.htm>

[101] The IceCube Collaboration: “Anatomy of a DOM”. IceCube Public Gallery, **Nov 2006**. URL <http://icecube.wisc.edu/gallery/view/140>

[102] Abbasi, R. *et al.*: “The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping”. *Nucl. Instrum. Meth.*, vol. A601, pp. 294–316, **2009**. doi:10.1016/j.nima.2009.01.001. URL <https://arxiv.org/abs/0810.4930>

[103] Ackermann, M. *et al.*: “Optical properties of deep glacial ice at the South Pole”. *J. Geophys. Res. Atmos.*, vol. 111(D13), p. D13203, **2006**. doi:10.1029/2005JD006687. URL <http://at-web.physik.uni-wuppertal.de/~kampert/Publications-PDF/Ackermann-JGR-06.pdf>

[104] Price, P. B., Woschnagg, K., and Chirkin, D.: “Age *vs* depth of glacial ice at South Pole”. *Geophysical Research Letters*, vol. 27(14), pp. 2129–2132, **2000**. ISSN 1944-8007. doi:10.1029/2000GL011351. URL <http://onlinelibrary.wiley.com/doi/10.1029/2000GL011351/abstract>

[105] Chirkin, D.: “Evidence of optical anisotropy of the South Pole ice”. In *Proceedings, 33rd International Cosmic Ray Conference (ICRC2013): Rio de Janeiro, Brazil*, p. 0580, **July 2013**. URL <https://galprop.stanford.edu/elibrary/icrc/2013/papers/icrc2013-0580.pdf>

[106] Chirkin, D.: “Photon Propagation and Ice Properties”. *Icecube bootcamp*, University of Wisconsin Madison, **2012**. URL <http://icecube.wisc.edu/~dima/work/BKP/DCS/BOOTCAMP12/ppaip.ppt>

[107] Merck, M., Chirkin, D., Vélez, J. C. D., and Skarupka, H.: “IceCubes GPGPU’s cluster for extensive MC production”. *Journal of Physics: Conference Series*, vol. 396(2), p. 022046, **2012**. URL <http://stacks.iop.org/1742-6596/396/i=2/a=022046>

BIBLIOGRAPHY

[108] Kopper, C. *et al.*: “An OpenCL-based photon-tracking simulation using a (source-based) ray tracing algorithm modeling scattering and absorption of light in the deep glacial ice at the South Pole or Mediterranean sea water”. GitHub, **Jan 2011**.
URL <https://github.com/claudiok/clsim>

[109] Rongen, M.: “Measuring the optical properties of IceCube drill holes”. *EPJ Web of Conferences*, vol. 116, p. 06011, **2016**. doi:10.1051/epjconf/201611606011.
URL <http://dx.doi.org/10.1051/epjconf/201611606011>

[110] Ahrens, J. *et al.*: “Muon track reconstruction and data selection techniques in AMANDA”. *Nucl. Instrum. Meth.*, vol. A524, pp. 169–194, **2004**. doi:10.1016/j.nima.2004.01.065.
URL <https://arxiv.org/abs/astro-ph/0407044>

[111] Aartsen, M. G. *et al.*: “Energy Reconstruction Methods in the IceCube Neutrino Telescope”. *JINST*, vol. 9, p. P03009, **2014**. doi:10.1088/1748-0221/9/03/P03009.
URL <https://arxiv.org/abs/1311.4767v3>

[112] Larson, M. J.: *Simulation and Identification of Non-Poissonian Noise Triggers in the IceCube Neutrino Detector*. Master’s thesis, Graduate School of the University of Alabama, Department of Physics, **2013**.
URL https://docushare.icecube.wisc.edu/dsweb/Get/Document-68303/LarsonThesis_final.pdf

[113] Zoll, M. C. R.: “Preparations for the next solar WIMP Analysis with IceCube: Advances in simulation, filtering, event topology identification and analysis approach”, **2014**.
URL <http://www.diva-portal.org/smash/record.jsf?pid=diva2:699510&dswid=305>

[114] Schukraft, A.: *Search for a diffuse flux of extragalactic neutrinos with the IceCube Neutrino Observatory*. Ph.D. thesis, RWTH Aachen University, **June 2013**.
URL http://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaapwiuw

[115] Abbasi, R. *et al.*: “An improved method for measuring muon energy using the truncated mean of dE/dx”. *Nuclear Instruments and Methods in Physics Research A*, vol. 703, pp. 190–198, **March 2013**. doi:10.1016/j.nima.2012.11.081.
URL <https://arxiv.org/abs/1208.3430>

[116] Heck, D., Schatz, G., Thouw, T., Knapp, J., and Capdevielle, J. N.: “CORSIKA: A Monte Carlo code to simulate extensive air showers”, **February 1998**.
URL <https://www.ikp.kit.edu/corsika/70.php>

[117] Gazizov, A. and Kowalski, M. P.: “ANIS: High energy neutrino generator for neutrino telescopes”. *Comput. Phys. Commun.*, vol. 172, pp. 203–213, **2005**. doi:

10.1016/j.cpc.2005.03.113.
URL <https://arxiv.org/abs/astro-ph/0406439>

[118] Chirkin, D. and Rhode, W.: “Muon Monte Carlo: A High-precision tool for muon propagation through matter”. *arXiv hep-ph*, **2004**.
URL <https://arxiv.org/abs/hep-ph/0407075>

[119] Allison, J. *et al.*: “Recent developments in Geant4”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 835, pp. 186 – 225, **2016**. ISSN 0168-9002. doi: <http://doi.org/10.1016/j.nima.2016.06.125>.
URL <http://www.sciencedirect.com/science/article/pii/S0168900216306957>

[120] Wallraff, M.: “Search for sterile neutrinos with the IceCube Neutrino Observatory”. In *Proceedings, 33rd International Cosmic Ray Conference (ICRC2013): Rio de Janeiro, Brazil, July 2-9, 2013*, p. 0455, **July 2013**.
URL <https://galprop.stanford.edu/elibrary/icrc/2013/papers/icrc2013-0455.pdf>

[121] Wallraff, M. and Wiebusch, C.: “Search for sterile neutrinos with the IceCube Neutrino Observatory”. *PoS*, vol. ICRC2015, p. 1100, **2016**. doi:10.22323/1.236.1100.
URL <https://pos.sissa.it/236/1100/pdf>

[122] Aartsen, M. G. *et al.*: “The IceCube Neutrino Observatory Part V: Neutrino Oscillations and Supernova Searches”. In *Proceedings, 33rd International Cosmic Ray Conference (ICRC2013): Rio de Janeiro, Brazil, July 2-9, 2013*, **2013**.
URL <https://arxiv.org/abs/1309.7008>

[123] Aartsen, M. G. *et al.*: “The IceCube Neutrino Observatory - Contributions to ICRC 2015 Part V: Neutrino Oscillations and Supernova Searches”. In *Proceedings, 34th International Cosmic Ray Conference (ICRC 2015): The Hague, The Netherlands, July 30-August 6, 2015*, **2015**.
URL <https://arxiv.org/abs/1510.05227>

[124] Hellwig, D.: *Berechnung von Oszillationseffekten atmosphärischer Neutrinos mit zusätzlichen sterilen Massenzuständen*. Bachelor’s thesis, RWTH Aachen, III. Physikalisches Institut B, **Aug 2012**.
URL https://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaaaaaakaigg

[125] Razzaque, S. and Smirnov, A. Yu.: “Searches for sterile neutrinos with IceCube DeepCore”. *Phys. Rev.*, vol. D85, p. 093010, **2012**. doi:10.1103/PhysRevD.85.093010.
URL <https://arxiv.org/abs/1203.5406>

BIBLIOGRAPHY

[126] Lyons, L.: *Statistics for Nuclear and Particle Physicists*. Cambridge University Press, **1986**. ISBN 9780521255400.
URL <https://books.google.de/books?id=40NgQgAACAAJ>

[127] Rüschenendorf, L.: *Mathematische Statistik*. Springer Spektrum, **2014**. ISBN 978-3-642-41996-6.
URL <https://dx.doi.org/10.1007/978-3-642-41997-3>

[128] Chirkin, D.: “Likelihood description for comparing data with simulation of limited statistics”. *arXiv astro-ph*, **April 2013**.
URL <https://arxiv.org/abs/1304.0735>

[129] Brent, R.: *Algorithms for Minimization Without Derivatives*. Dover Books on Mathematics. Dover Publications, **1972/2013**. ISBN 9780486143682.
URL <https://books.google.de/books?id=AITCAgAAQBAJ>

[130] Jones, E., Oliphant, T., Peterson, P. *et al.*: “SciPy: Open source scientific tools for Python”, **2001**–. [Online; accessed 2015-01-04].
URL <http://www.scipy.org/>

[131] Wilks, S. S.: “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. *Ann. Math. Statist.*, vol. 9(1), pp. 60–62, **03 1938**. doi: 10.1214/aoms/1177732360.
URL <http://dx.doi.org/10.1214/aoms/1177732360>

[132] Koob, A.: *On the estimation of neutrino oscillation parameters with global fits of the IceCube atmospheric neutrino data*. Master’s thesis, RWTH Aachen, III. Physikalisches Institut B, **Sep 2014**.
URL http://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaaaaapwicy

[133] Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J.: “Algorithm 778: L-BFGS-B: Fortran Subroutines for Large-scale Bound-constrained Optimization”. *ACM Trans. Math. Softw.*, vol. 23(4), pp. 550–560, **December 1997**. ISSN 0098-3500. doi:10.1145/279232.279236.
URL <http://doi.acm.org/10.1145/279232.279236>

[134] James, F. and Roos, M.: “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations”. *Comput. Phys. Commun.*, vol. 10, pp. 343–367, **1975**. doi:10.1016/0010-4655(75)90039-9.
URL [https://doi.org/10.1016/0010-4655\(75\)90039-9](https://doi.org/10.1016/0010-4655(75)90039-9)

[135] Montaruli, T.: “The NeutrinoFlux C++ class: a service fluxes for weighting”, **2005**–. [Online; accessed 2017-08-18].
URL http://icecube.wisc.edu/~tmontaruli/neutrinoflux/NeutrinoFlux_Teresa.html

[136] Gaisser, T. K.: “Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio”. *Astropart. Phys.*, vol. 35, pp. 801–806, **2012**. doi:10.1016/j.astropartphys.2012.02.010.
URL <https://arxiv.org/abs/1111.6675>

[137] Elbert, J. W.: “Multiple muons produced by cosmic ray interactions”. In *DUMAND Summer Workshop, Volume 2 - UHE Interactions, Neutrino Astronomy*, vol. 2 (edited by Roberts, A.), pp. 101–121, **1979**.
URL <http://adsabs.harvard.edu/abs/1979duma....2..101E>

[138] Gaisser, T. K.: “Atmospheric Lepton Fluxes”. *EPJ Web Conf.*, vol. 99, p. 05002, **2015**. doi:10.1051/epjconf/20159905002.
URL <https://arxiv.org/abs/1412.6424>

[139] Haack, C., Schukraft, A., Zilles, A., and Wiebusch, C.: “Re-weighting atmospheric neutrino fluxes to account for the cosmic-ray knee with neutrino flux”. *internal report icecube/201301002-v2*, RWTH Aachen University, **Sep 2013**.
URL https://internal.icecube.wisc.edu/reports/data/icecube/2013/01/002/icecube_201301002_v2.pdf

[140] Fedynitch, A., Engel, R., Gaisser, T. K., Riehn, F., and Stanev, T.: “Calculation of conventional and prompt lepton fluxes at very high energy”. *EPJ Web Conf.*, vol. 99, p. 08001, **2015**. doi:10.1051/epjconf/20159908001.
URL <https://arxiv.org/abs/1503.00544>

[141] Collin, G.: “An estimation of systematics for up-going atmospheric muon neutrino flux at the south pole”. *Tech. Rep.*, Massachusetts Institute of Technology, **Aug 2015**.
URL <http://hdl.handle.net/1721.1/98078>

[142] Gaisser, T. K. and Honda, M.: “Flux of atmospheric neutrinos”. *Ann. Rev. Nucl. Part. Sci.*, vol. 52, pp. 153–199, **2002**. doi:10.1146/annurev.nucl.52.050102.090645.
URL <https://arxiv.org/abs/hep-ph/0203272>

[143] Fedynitch, A., Becker Tjus, J., and Desiati, P.: “Influence of hadronic interaction models and the cosmic ray spectrum on the high energy atmospheric muon and neutrino flux”. *Phys. Rev. D*, vol. 86(11), 114024, **December 2012**. doi:10.1103/PhysRevD.86.114024.
URL <https://arxiv.org/abs/1206.6710>

[144] Hörandel, J. R.: “On the knee in the energy spectrum of cosmic rays”. *Astropart. Phys.*, vol. 19, pp. 193–220, **2003**. doi:10.1016/S0927-6505(02)00198-6.
URL <https://arxiv.org/abs/astro-ph/0210453>

[145] Zatsepin, V. I. and Sokolskaya, N. V.: “Three component model of cosmic ray spectra from 100-gev up to 100-pev”. *Astron. Astrophys.*, vol. 458, pp. 1–5, **2006**.

BIBLIOGRAPHY

doi:10.1051/0004-6361:20065108.
URL <https://arxiv.org/abs/astro-ph/0601475>

[146] Adriani, O. *et al.*: “PAMELA Measurements of Cosmic-Ray Proton and Helium Spectra”. *Science*, vol. 332, p. 69, **April 2011**. doi:10.1126/science.1199172.
URL <https://arxiv.org/abs/1103.4055>

[147] Auffenberg, J., Haack, C., Hellwig, D., Rädel, L., Schoenen, S., and Wiebusch, C.: “On the Correction of Atmospheric Temperature Variations in Atmospheric Neutrino Flux Simulations”. *internal report icecube/201505001-v3*, RWTH Aachen University, **Aug 2015**.
URL https://internal.icecube.wisc.edu/reports/data/icecube/2015/05/001/icecube_201505001_v3.pdf

[148] Parkinson, C., Graham, S., and Przyborski, P.: “Aqua Earth-observing satellite mission”, **2002**–. [Online; accessed 2017-07-02].
URL <https://aqua.nasa.gov/>

[149] United States Committee on Extension of the Standard Atmosphere (COESA), Dubin, M. *et al.*: “U.S. Standard Atmosphere”. *technical report 19770009539*, NASA, **Oct 1976**.
URL <https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19770009539.pdf>

[150] Hellwig, D.: *Untersuchung der Produktionsraten von Mesonen mit Strangeness über die zeitliche Variation des mit IceCube gemessenen Neutrinoflusses*. Master’s thesis, RWTH Aachen, III. Physikalisches Institut B, **Sep 2014**.
URL http://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaaaaapwien

[151] Miyamoto, H.: “IceCube GoldenPMT Calibration”, **Mar 2008**. [Online; accessed 2017-09-02].
URL http://www.ppl.phys.chiba-u.jp/research/IceCube/pmt/GoldenPMT/error_summary/index.html

[152] Aartsen, M. G. *et al.*: “Searches for Sterile Neutrinos with the IceCube Detector”. *Phys. Rev. Lett.*, vol. 117(7), p. 071801, **2016**. doi:10.1103/PhysRevLett.117.071801.
URL <https://arxiv.org/abs/1605.01990>

[153] Asimov, I.: “Franchise”. In *If: Worlds of Science Fiction*. Quinn Publishing, **Aug 1955**.
URL <https://archive.org/details/195508IF0000>

[154] Cowan, G., Cranmer, K., Gross, E., and Vitells, O.: “Asymptotic formulae for likelihood-based tests of new physics”. *Eur. Phys. J.*, vol. C71, p. 1554, **2011**. doi:10.1140/epjc/s10052-011-1554-0, 10.1140/epjc/s10052-013-2501-z. [Erratum: *Eur.*

Phys. J.C73,2501(2013)].
URL <https://arxiv.org/abs/1007.1727>

[155] Abe, K. *et al.*: “Limits on sterile neutrino mixing using atmospheric neutrinos in Super-Kamiokande”. *Phys. Rev.*, vol. D91, p. 052019, **2015**. doi:10.1103/PhysRevD.91.052019.
URL <https://arxiv.org/abs/1410.2008>

[156] Feldman, G. J. and Cousins, R. D.: “A Unified approach to the classical statistical analysis of small signals”. *Phys. Rev.*, vol. D57, pp. 3873–3889, **1998**. doi:10.1103/PhysRevD.57.3873.
URL <https://arxiv.org/abs/physics/9711021>

[157] Cowan, G.: *Statistical Data Analysis*. Oxford University Press, **1998**.
URL http://www.sherrytowers.com/cowan_statistical_data_analysis.pdf

[158] Knuth, D.: *Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Pearson Education, **2014**. ISBN 9780321635761.
URL <https://books.google.de/books?id=Zu-HAwAAQBAJ>

[159] Sousa, A. B.: “MINOS Search for Sterile Neutrinos”. In *Proceedings, 13th International Workshop on Neutrino Factories, Superbeams and Beta beams (NuFact11): Geneva, Switzerland, August 1-6, 2011*, **2011**. doi:10.1088/1742-6596/408/1/012026. [J. Phys. Conf. Ser.408,012026(2013)].
URL <https://arxiv.org/abs/1110.3455>

[160] Lindner, M., Rodejohann, W., and Xu, X.-J.: “Sterile neutrinos in the light of IceCube”. *JHEP*, vol. 01, p. 124, **2016**. doi:10.1007/JHEP01(2016)124.
URL <https://arxiv.org/abs/1510.00666>

[161] Vehring, M.: *Measurement of Neutrino Oscillations with IceCube-DeepCore*. Ph.D. thesis, RWTH Aachen University, **March 2017**.
URL http://www.institut3b.physik.rwth-aachen.de/global/show_document.asp?id=aaaaaaaaaaaxjhfk

[162] Aartsen, M. G. *et al.*: “Search for sterile neutrino mixing using three years of IceCube DeepCore data”. *Phys. Rev.*, vol. D95(11), p. 112002, **2017**. doi:10.1103/PhysRevD.95.112002.
URL <https://arxiv.org/abs/1702.05160>

[163] Zuber, K.: *Neutrino Physics, Second Edition*. Series in High Energy Physics, Cosmology and Gravitation. Taylor & Francis, **2011**. ISBN 9781420064711.
URL <https://books.google.de/books?id=0DRaQiiINvkJ>

[164] Couchot, F., Henrot-Versillé, S., Perdereau, O., Plaszczynski, S., Rouillé d’Orfeuil, B., Spinelli, M., and Tristram, M.: “Cosmological constraints on the neutrino mass

BIBLIOGRAPHY

including systematic uncertainties”. *Astron. Astrophys.*, vol. 606, p. A104, **2017**. doi:10.1051/0004-6361/201730927.
URL <https://arxiv.org/abs/1703.10829>

[165] Oldengott, I. M., Barenboim, G., Kahlen, S., Salvado, J., and Schwarz, D. J.: “How to relax the cosmological neutrino mass bound”. *JCAP*, vol. 1904(04), p. 049, **2019**. doi:10.1088/1475-7516/2019/04/049.
URL <https://arxiv.org/abs/1901.04352>

[166] Wallraff, M. and Wiebusch, C.: “Calculation of oscillation probabilities of atmospheric neutrinos using nuCraft”. *Comput. Phys. Commun.*, vol. 197, pp. 185–189, **2015**. doi:10.1016/j.cpc.2015.07.010.
URL <https://arxiv.org/abs/1409.1387v2>

[167] Brown, P. N., Byrne, G. D., and Hindmarsh, A. C.: “VODE: a variable-coefficient ODE solver”. *SIAM J. Sci. Stat. Comput.*, vol. 10(5), pp. 1038–1051, **September 1989**. ISSN 0196-5204. doi:10.1137/0910062.
URL <http://dx.doi.org/10.1137/0910062>

[168] Wendell, R.: “Prob3++ software for computing three flavor neutrino oscillation probabilities”, **2012**–.
URL <http://www.phy.duke.edu/~raw22/public/Prob3++>

[169] Aartsen, M. G. *et al.*: “Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption”. *Nature*, vol. 551, pp. 596–600, **2017**. doi:10.1038/nature24459.
URL <https://arxiv.org/abs/1711.08119>

[170] C. A. Arguelles Delgado, J. S., B. Jones: “IC86 Sterile Search Brief Update”, **Oct 2015**.
URL https://docushare.icecube.wisc.edu/dsweb/Get/Document-75671/28oct_sterile_talk.pdf

[171] C. A. Arguelles Delgado, C. N. W., J. Salvado: “ ν -SQuIDS”, **September 2014**.
URL <https://github.com/arguelles/nuSQuIDS>

[172] Rosenblatt, M.: “Remarks on Some Nonparametric Estimates of a Density Function”. *Ann. Math. Statist.*, vol. 27(3), pp. 832–837, **09 1956**. doi:10.1214/aoms/1177728190.
URL <https://doi.org/10.1214/aoms/1177728190>

[173] Parzen, E.: “On Estimation of a Probability Density Function and Mode”. *Ann. Math. Statist.*, vol. 33(3), pp. 1065–1076, **09 1962**. doi:10.1214/aoms/1177704472.
URL <https://doi.org/10.1214/aoms/1177704472>

[174] Gramacki, A.: *Nonparametric Kernel Density Estimation and Its Computational Aspects*. Studies in Big Data. Springer International Publishing, **2017**. ISBN 9783319716886.
URL <https://books.google.de/books?id=PCpEDwAAQBAJ>

[175] Silverman, B.: *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, **1986**. ISBN 9780412246203.
URL <https://books.google.de/books?id=e-xsrjsL7WkC>

[176] Efron, B., Rogosa, D., and Tibshirani, R.: “Resampling Methods of Estimation”. In *International Encyclopedia of the Social & Behavioral Sciences (Second Edition)* (edited by Wright, J. D.), pp. 492 – 495. Elsevier, Oxford, second edition edn., **2015**. ISBN 978-0-08-097087-5. doi:<https://doi.org/10.1016/B978-0-08-097086-8.42165-3>.
URL <http://www.sciencedirect.com/science/article/pii/B9780080970868421653>

[177] Hunter, J. D.: “Matplotlib: A 2D graphics environment”. *Computing In Science & Engineering*, vol. 9(3), pp. 90–95, **2007**. doi:[10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).
URL <https://matplotlib.org/>

[178] Kness, M.: “ColorPy – A Python package for handling physical descriptions of color and light spectra”, **2009**.
URL <http://markness.net/colorpy/ColorPy.html>

[179] Anaconda, Inc.: “Anaconda Software Distribution”, **2012**.
URL <http://www.anaconda.com>