



Measurement of Tau Neutrino Appearance in 8 Years of IceCube Data

and A Search for astrophysical neutrinos from the Local Universe

A thesis presented to the Faculty of Science in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Physics

Étienne Bourbeau

Supervisor: David Jason Koskinen

ABSTRACT

The mechanism by which neutrinos acquire mass is left unexplained by the Standard Model of particle physics. Several models attempting to explain this process imply the existence of yet undiscovered neutrino species, which would contradict the current paradigm that neutrino oscillations between the three known flavours preserve unitarity. Current constraints on unitarity are the weakest for the $U_{\tau 3}$ element of the neutrino mixing matrix, and can be improved by measuring the tau neutrino appearance channel $\nu_\mu \rightarrow \nu_\tau$.

An analysis aimed at measuring the appearance of tau neutrinos in eight years of data from the IceCube Neutrino Observatory is presented. Improvements on the previous event selection include the use of machine learning classifiers at final-level to improve both the signal efficiency and background rejection by 40%. A revised treatment of systematic uncertainties compared to previous measurements is presented, along with new studies that constrain even further the potential impact of detector noise mis-modeling on oscillation measurements. Results from a blind fit to data show significant pulls from a number of systematic parameters, indicating that the treatment of the DIS cross section uncertainty and DOM efficiency should be revised.

Neutrinos can be used as astrophysical messengers, and may point to the origins of the physical process by which ultra-high energy cosmic rays (UHECRs) are accelerated. A connection as yet to be made, however, between the astrophysical flux of neutrinos observed by IceCube and observations made by light-based astronomy. A search for correlations between clusters of neutrino in seven years of IceCube track events and the 2MRS catalog of galaxies is performed. No significant correlations are observed, leading to the derivation of upper limits on the density and luminosity of neutrino source populations located in the Local Universe (≈ 500 Mpc).

RÉSUMÉ

Standardmodellen i partikelfysik forklarer ikke, hvordan neutrinoer bliver udstyret med masse. Flere modeller prøver at forklare denne mekanisme med en ny slags af endnu uopdagede neutrinopartikler, som går imod det paradigme, hvor neutrinosvingning er en enhedprocess. Målinger af neutrinosvingningskanal $\nu_\mu \rightarrow \nu_\tau$ kan forbedre nuværende begrænsninger på elementet $U_{\tau 3}$ af blandingsmatrixet, som er det svageste.

En analyse rettet mod målingen af tau neutrino i otte års data fra IceCube Eksperimentet bliver præsenteret. Analysen forbedrer tidligere arbejde med neutrinodatasamling med nye maskinlæringsalgoritmer, der forhøjer såvel nøjagtigheden af neutrinotal og baggrundsafvigelse med 40%. Nye systematiske usikkerheder bliver inkluderet i analysen, og nye undersøgelser giver endnu stærkere grænse for, hvor meget detektorlarm påvirker måling af neutrinosvingning. Med resultater fra en blind-fit kan man se, at flere systematiske usikkerhedsparametre bliver trukket ud fra deres nominalværdi. Det betyder, at ydeligere arbejde med bl.a. DIS-tværtssnit og DOM-effektivitet behøves, før tau-neutrinosvingning kan blive målt.

Neutrinoer kan bruges som astrofysiske budbringere. Oplysninger fra astrofysiske neutrinoer kan pege mod retningen af de fysiske processer, hvor ultrahøje energi kosmiske stråler bliver accelereret. Men en forbindelse mangles mellem strømmen af astrofysiske neutrinoer målt i IceCube og observationer, der blive målt med lysbaseret astronomi. En søgning efter korrelationer mellem neutrino-multipla i syv års data fra IceCube og den 2MRS katalog af galakser præsenteres. Ingen væsentlig korrelationer var fundet. Med resultaterne kommer der nye grænser for lysstyrke og massefylde af neutrinokilder i det locale univers (≈ 500 Mpc).

ACKNOWLEDGEMENTS

Dear Reader,

You hold in your hands (or more likely on your screen) the result of a PhD, which is admittedly quite a nice achievement. What might not be apparent to you, however, is that the years of labour that went into this thesis have undoubtedly been the best years of my life thus far. That is arguably an even more impressive achievement, given the long hours at the office, the never-ending challenges inherent to experimental physics, uncooperative code, and the anxiety of the final months. To be able to still enjoy all of this requires the support of many people I wish to thank here, before you go on reading all about IceCube and neutrinos.

First, I would like to thank my advisor, Jason Koskinen, who decided that using roof-window money to hire some random engineer from Quebec would be a wise investment. I'm grateful for the trust he put in my work, his guidance and, above all, his honest and frank-spoken take on the physics problems we had to solve. This willingness to discuss critically, but also openly about things has led to some great academic conversations, and has made me do better research in the end.

A key reason as to why I enjoyed working on my PhD is certainly the collaborative nature of the analyses I've worked on, and the qualities of the people I had the pleasure of working with, whether they be postdoc, professors, PhD or master's students, both within IceCube and at the Niels Bohr Institute. I particularly want to thank Mohamed Rameez for his guidance through the correlation analysis, and Tom Stuttard for... well... so many things! Unofficial mentor, outstanding oscNext collaborator, and above all a great office mate to have around, whether it be to discuss fundamentally deep physics stuff, complain about clusters, icetrax and numba issues, or enjoy a few beers in a bar, in Copenhagen or elsewhere in the world. I wish him the best of career in the future, and will truly miss his daily company.

It is said that there is a world outside of physics, and that it can be quite fun. I agree, and want to acknowledge the many, many people who, via trips, beers, parties and evenings have contributed to keeping me sane-ish: my awesome roommate Tania, Rasmus, Violeta, my PhD-cousin and longtime erasmus friend from Aarhus Marina, along with dozens of friends and colleagues both in and out of academia, in Denmark or abroad. And above everybody here, I thank the four people responsible for 90% of the amount of happiness I have felt this last and particularly challenging year: Maman, Papa, Flo, and of course you, *Talpina* ;) . I love you very much.

Finally, if you are reading this right now, it is likely that we've shared some good times, so I thank you as well. If anything, we've shared the common experience of a pandemic, so you'll likely join me in wishing sars-cov-2 a swift death.

Bonne lecture!

CONTENTS

| | | |
|---|--|----|
| 1 | INTRODUCTION | 1 |
| I DETECTOR AND INSTRUMENTATION | | |
| 2 | NOTIONS IN PARTICLE PHYSICS | 4 |
| 2.1 | The Standard Model of Particle Physics | 4 |
| 2.2 | Chirality in particle physics | 5 |
| 2.3 | The Concept/Problem of Particle Generations | 6 |
| 2.4 | The Electroweak Framework | 6 |
| 2.4.1 | Separating the Weak and Electromagnetic Forces | 8 |
| 2.4.2 | Spontaneous Symmetry Breaking and Massive Particles | 10 |
| 2.5 | The role of neutrinos in all of this | 10 |
| 2.6 | Potential for the existence of HNL and see-saw mechanism | 12 |
| 3 | THE ICECUBE NEUTRINO OBSERVATORY | 13 |
| 3.1 | The Cherenkov Effect | 13 |
| 3.2 | Detector Overview | 14 |
| 3.2.1 | Light Detection | 16 |
| 3.2.2 | Triggering and Filtering | 16 |
| 3.2.3 | Pulse Reconstruction and Cleaning | 19 |
| 3.2.4 | Characterization of events in IceCube | 21 |
| 3.3 | The DeepCore Sub-Array | 22 |
| 4 | CHARACTERIZATION OF NOISE IN ICECUBE PMTS | 25 |
| 4.1 | Current Noise Model | 25 |
| 4.2 | Comparing log10dt distributions in simulation and data | 28 |
| 4.3 | Results of dark rate tests | 32 |
| 4.4 | Results of coincidence Tests | 34 |
| 4.5 | concluding remarks | 35 |
| II TAU APPEARANCE ANALYSIS | | |
| 5 | PROPERTIES OF THE NEUTRINO | 38 |
| 5.1 | Neutrino oscillations | 38 |
| 5.1.1 | Unitarity of a Matrix | 38 |
| 5.1.2 | Oscillations in vacuum | 39 |
| 5.2 | Unknown Properties of the neutrino | 43 |
| 5.2.1 | Neutrino Mass and Ordering | 43 |
| 5.2.2 | CP Violation | 44 |
| 5.2.3 | Majorana vs. Dirac | 44 |
| 5.2.4 | Disappearance vs. Appearance Experiments | 45 |
| 5.2.5 | Current experimental results | 47 |
| 5.2.6 | Constraints on unitarity | 49 |
| 6 | NEUTRINO OSCILLATIONS WITH ICECUBE | 53 |
| 6.1 | General Measurement Concept | 53 |
| 6.2 | Cosmic Rays and the Atmospheric Neutrino Flux | 55 |
| 6.2.1 | Neutrino propagation in matter | 58 |
| 6.3 | Neutrino Interactions with matter | 59 |
| 6.3.1 | Quasielastic Scattering | 61 |
| 6.3.2 | Resonance Production | 61 |
| 6.3.3 | Deep Inelastic Scattering | 61 |

| | | |
|--|---|-----|
| 6.3.4 | Additional considerations for neutrino-nucleus interactions | 62 |
| 6.4 | Experimental Flavour differentiation | 62 |
| 6.5 | Current results from previous experiments | 63 |
| 7 | EVENT SELECTION | 65 |
| 7.1 | Simulation Production | 65 |
| 7.2 | Selection Levels | 68 |
| 7.2.1 | Level 3: Low-level variables starting events | 69 |
| 7.2.2 | Level 4: Noise and Muon classifiers | 71 |
| 7.2.3 | Level 5: Muon corridor cuts | 74 |
| 7.2.4 | Level 6: Event Reconstruction | 76 |
| 7.2.5 | Level 7: Analysis-level cuts | 81 |
| 7.2.6 | Level 7: PID | 83 |
| 7.3 | Detector Data | 87 |
| 7.4 | Containment cut and final noise cleanup | 91 |
| 8 | TAU NEUTRINO APPEARANCE ANALYSIS | 93 |
| 8.1 | The PISA framework | 93 |
| 8.2 | Analysis binning and cut optimization | 94 |
| 8.3 | Systematic Uncertainties | 98 |
| 8.3.1 | Neutrino flux systematics | 99 |
| 8.3.2 | Neutrino Propagation | 100 |
| 8.3.3 | Cross section uncertainties | 101 |
| 8.3.4 | Detector effects | 106 |
| 8.3.5 | Muon Systematics | 108 |
| 8.3.6 | Assessing the need for noise parameter systematics | 109 |
| 8.4 | Likelihood Formulation | 111 |
| 8.4.1 | Choice of optimization metric | 111 |
| 8.5 | Final Choice of Systematics | 113 |
| 9 | RESULTS | 119 |
| 9.1 | Pre-unblinding Test results | 119 |
| 9.1.1 | Injection-Recovery Test | 119 |
| 9.1.2 | Sensitivity test | 120 |
| 9.1.3 | Ensemble Test | 122 |
| 9.2 | Blind fit agreement between data and simulation | 124 |
| 9.2.1 | Data-MC agreement in key variables | 125 |
| 9.2.2 | Systematic parameter fits and pulls | 129 |
| 9.3 | Potential roadmap to unblinding | 133 |
| 9.4 | Concluding Remarks | 134 |
| III 2MRS CORRELATION SEARCH | | |
| 10 | NEUTRINO ASTRONOMY | 136 |
| 10.1 | Potential Sources of Ultra-High Energy Cosmic Rays | 136 |
| 10.2 | Limitations of photons in astronomy | 137 |
| 10.3 | The neutrino / cosmic ray connection | 139 |
| 10.4 | The diffuse astrophysical Flux and searches for neutrino sources | 139 |
| 10.4.1 | The Seven-Year Point Source Search | 140 |
| 11 | SEARCH FOR CORRELATIONS BETWEEN NEUTRINOS AND THE LOCAL UNI- VERSE | 143 |
| 11.1 | Motivation | 143 |
| 11.2 | Analysis Method | 144 |
| 11.3 | Results | 147 |
| 11.4 | Concluding Remarks | 151 |

IV CONCLUSION AND FUTURE WORK

| | |
|---------------|-----|
| 12 CONCLUSION | 154 |
|---------------|-----|

V APPENDICES

| | |
|---|-----|
| A FULL RESULTS OF THE NOISE MEASUREMENT STUDY AT SNOLAB | 157 |
| B ADDITIONAL INFORMATION ABOUT THE TAU APPEARANCE ANALYSIS | 161 |
| B.1 Definition of the Veto-Identified Causal Hit region | 161 |
| B.2 Full list of systematic uncertainties | 161 |
| B.3 Definition of the extended Poissonian likelihood considered | 163 |
| B.3.1 Corrections specific to the generalized likelihood | 164 |
| B.3.2 Toy Monte Carlo study of the optimization metrics | 165 |
| B.4 Full set of systematic pull plots | 172 |
| B.4.1 Neutrino Flux Uncertainties | 173 |
| B.4.2 Neutrino Oscillation Uncertainties | 175 |
| B.4.3 Cross Section Uncertainties | 176 |
| B.4.4 Discrete Detector Uncertainties | 177 |
| B.4.5 Muon Uncertainty | 179 |
| B.5 Full results of the systematic impact test | 179 |
| B.5.1 Test where N_τ is set to 0.75 | 180 |
| B.5.2 Test where N_τ is set to 1.00 | 182 |
| B.5.3 Test where N_τ is set to 1.25 | 184 |
| B.6 Additional information about the Level 7 Muon Classifier | 185 |
| C ADDITIONAL MATERIAL FOR THE 2MRS CORRELATION SEARCH | 187 |
| C.1 Computing a limit on the density of neutrino sources | 187 |

LIST OF FIGURES

| | | |
|----|---|----|
| 1 | The Standard Model of particle physics | 5 |
| 2 | Fundamental neutrino interactions | 9 |
| 3 | Cherenkov radiation principle and detection | 14 |
| 4 | The IceCube South Pole Neutrino Observatory | 15 |
| 5 | Schematics of a photomultiplier tube | 16 |
| 6 | SPE distribution for a single DOM | 17 |
| 7 | Example of the waveform digitization process | 18 |
| 8 | Zoomed-in ATWD waveform | 19 |
| 9 | The IceCube Laboratory (ICL) | 20 |
| 10 | Steps of an SRT cleaning procedure | 21 |
| 11 | Typical event morphologies in IceCube. | 22 |
| 12 | Typical event morphologies in IceCube-DeepCore. | 24 |
| 13 | Characteristic noise timing plots | 26 |
| 14 | Experimental setup of the SNOLab noise measurements | 28 |
| 15 | Noise distributions measured at -10°C | 30 |
| 16 | ATWD Noise distributions from simulated DOMs operating at -10°C | 31 |
| 17 | FADC Noise distributions from simulated DOMs operating at -10°C | 31 |
| 18 | Fit of the raw time difference distribution to a Poisson-distribution. | 32 |
| 19 | Noise Rates of the four DOMs tested vs. temperature | 33 |
| 20 | Evidence of in-situ quieting since deployment | 34 |
| 21 | Schematic of the Coincidence Test experimental setup | 35 |
| 22 | The two mass ordering configurations: normal (left) and inverted (right) | 44 |
| 23 | Muon neutrino survival probabilities | 46 |
| 24 | Current Best-fit of all neutrino oscillation parameters. Source [48]. | 49 |
| 25 | Consistency of PMNS angle measurements with unitarity | 50 |
| 26 | Constraints on individual PMNS matrix elements | 51 |
| 27 | Analysis principle of an atmospheric neutrino oscillation measurement | 54 |
| 28 | The energy spectrum of the cosmic-ray flux. Source: [56] | 55 |
| 29 | Breakdown of atmospheric shower products | 57 |
| 30 | Energy spectrum of the cosmic-ray neutrino flux | 58 |
| 31 | Survival probability of a muon neutrino in the presence of matter | 59 |
| 32 | contributions to the neutrino-nucleus cross-sections | 60 |
| 33 | Diagram depiction of neutrino-nuclei interactions | 60 |
| 34 | Event morphologies of neutrino interactions. | 62 |
| 35 | Current tau neutrino appearance results | 64 |
| 36 | Absorption and scattering properties of the ice around DeepCore | 67 |
| 37 | Background reduction levels in the oscNext event selection. | 69 |
| 38 | Illustration of an unsimulated population of events in GRECO. | 71 |
| 39 | Performances of the level 4 Noise classifier. | 73 |
| 40 | Performances of the Level 4 atmospheric muon classifier. | 74 |
| 41 | Example of muon corridors. | 75 |
| 42 | Comparison of Retro to a traditional table reconstruction method. | 77 |
| 43 | Energy resolution of the events reconstructed with RetroReco. | 79 |
| 44 | Zenith resolution of the events reconstructed with RetroReco. | 80 |
| 45 | Example of a muon sneaking through one of the corridors in IceCube. | 81 |
| 46 | Performances of the Level 7 Muon classifier. | 82 |

| | | |
|----|--|-----|
| 47 | Particle-wise rates as a function of the Level 7 classifier cut | 83 |
| 48 | Performances of the PID classifier | 85 |
| 49 | PID output distributions according to true energy | 86 |
| 50 | seasonal rates at final level | 88 |
| 51 | Distribution of seasonal data in the analysis space | 89 |
| 52 | Seasonal consistency of reconstructed energy (chi2 metric) | 90 |
| 53 | Seasonal consistency of reconstructed energy (chi2 metric) | 91 |
| 54 | Schematic depiction of a neutrino pipeline in PISA. | 94 |
| 55 | First optimization study of the tau appearance analysis sensitivity. | 96 |
| 56 | Second optimization study of the tau appearance analysis sensitivity. | 97 |
| 57 | Distribution of events in the final analysis template. | 98 |
| 58 | Definition of the parameter space used in the analysis. | 99 |
| 59 | Pull plots showing the impact of the neutrino spectral index on the analysis. | 100 |
| 60 | Contributions of QE and RES events in the sample | 102 |
| 61 | Effect of the DIS systematic parameter on the total cross section | 104 |
| 62 | Differential DIS cross section corrections | 105 |
| 63 | Impact of the DIS systematic on the analysis | 106 |
| 64 | Angular acceptance curves simulated in the discrete systematic sets. | 108 |
| 65 | Distribution of the number of triggered DOMs in the noise side-band. | 109 |
| 66 | Outcome of the noise systematic study. | 110 |
| 67 | Bias of the four optimization metrics. | 113 |
| 68 | Coverage of the four optimization metrics. | 114 |
| 69 | Impact of individual systematic uncertainties on the analysis | 115 |
| 70 | Impact of individual systematic uncertainties on N_τ | 116 |
| 71 | Result of the impact test after parameter fixing | 117 |
| 72 | Result of the injection-recovery test on Asimov datasets. | 120 |
| 73 | Sensitivity of the tau neutrino appearance analysis | 121 |
| 74 | Comparison of the sensitivity with established measurements. | 122 |
| 75 | Test Statistic distribution of 200 pseudo-experiments. | 123 |
| 76 | Parameter fit and pull information for one of the systematics | 124 |
| 77 | Data-MC agreement at the best-fit point (reconstructed energy) | 125 |
| 78 | Data-MC agreement at the best-fit point (reconstructed $\cos(\theta_{zenith})$) | 126 |
| 79 | Data-MC agreement at the best-fit point (PID classifier) | 126 |
| 80 | Data-MC agreement at the best-fit point (L7 Muon classifier) | 127 |
| 81 | Data-MC agreement at the best-fit point (reconstructed vertex position) | 128 |
| 82 | 2D XY-vertex agreement between data and two DOM efficiency sets. | 129 |
| 83 | Goodness-of-fit of the analysis applied to real data | 130 |
| 84 | Fit value of all systematic parameters after blind fits. | 131 |
| 85 | Summary of the pulls on systematic uncertainties after blind fit. | 132 |
| 86 | Artist's depiction of an active galactic nuclei (AGN). | 137 |
| 87 | Interaction cross section between photons and two different targets | 138 |
| 88 | Observable cosmological distances of the universe | 139 |
| 89 | Energy spectrum of the HESE event selection. | 140 |
| 90 | Pre-trial significance map of the 7-years Point Source Search. | 141 |
| 91 | Distribution of galaxies coloured by redshift. Source: [136]. | 143 |
| 92 | Distribution of galaxy redshifts in the 2MRS catalog. | 144 |
| 93 | Optimization metric value as a function of multiplet selection cuts. | 145 |
| 94 | Distribution of 2MRS galaxies and locations of selected multiplets. | 147 |
| 95 | Test-statistic Distribution and fitted value for data. | 148 |
| 96 | Feldman-Cousins interval for the correlation analysis | 149 |
| 97 | Density-luminosity limits obtained from the correlation analysis result | 151 |

| | | |
|-----|--|-----|
| 98 | Noise distributions measured at -20°C | 157 |
| 99 | Noise distributions measured at -30°C | 158 |
| 100 | Noise distributions measured at -40°C | 158 |
| 101 | ATWD Noise distributions from simulated DOMs operating at -20°C | 159 |
| 102 | FADC Noise distributions from simulated DOMs operating at -20°C | 159 |
| 103 | ATWD Noise distributions from simulated DOMs operating at -30°C | 160 |
| 104 | FADC Noise distributions from simulated DOMs operating at -30°C | 160 |
| 105 | VICH variable causal region | 161 |
| 106 | Correlation matrix of systematic uncertainties. | 162 |
| 107 | Distribution of events in the 1D Toy Monte Carlo test case | 166 |
| 108 | Likelihood scans performed on the 1D toy Monte carlo test. | 167 |
| 109 | Special Toy MC case with very large data statistics (1000 points). | 168 |
| 110 | Bias in the toy MC likelihood study | 169 |
| 111 | Bias in the toy MC likelihood study | 170 |
| 112 | Test statistic distribution of 500 fits to pseudo-data trials. | 171 |
| 113 | Coverage test of the 1D toy Monte Carlo pseudo-data w. various levels of MC statistics. | 172 |
| 114 | Neutrinos spectral index $\Delta\gamma$ | 173 |
| 115 | Barr parameter Y (applied to kaons) | 173 |
| 116 | Barr parameter H | 173 |
| 117 | Barr parameter I | 174 |
| 118 | Barr parameter K | 174 |
| 119 | Barr parameter Y (applied to anti-kaons) | 174 |
| 120 | Barr parameter G | 175 |
| 121 | Barr parameter Z (applied to kaons) | 175 |
| 122 | Neutrino Mass splitting, Δm_{31}^2 | 175 |
| 123 | Neutrino mixing angle θ_{23} | 176 |
| 124 | Quasielastic axial mass, $M_{A\text{ }QE}$ | 176 |
| 125 | Resonant axial mass, $M_{A\text{ }RES}$ | 176 |
| 126 | CSMS Cross-section scaling (1: CSMS-like, 0:genie-like) | 177 |
| 127 | Hole ice acceptance parameter p_0 | 177 |
| 128 | Hole ice acceptance parameter p_1 | 177 |
| 129 | DOM efficiency | 178 |
| 130 | Bulk ice scattering coefficient | 178 |
| 131 | Bulk ice absorption coefficient | 178 |
| 132 | Neutrino effective area | 179 |
| 133 | Overall Muon rate (scaled effective area) | 179 |
| 134 | Positive pull, $N_{\tau}=0.75$ | 180 |
| 135 | Tau normalization offset for a positive pull, $N_{\tau}=0.75$ | 180 |
| 136 | Negative pull, $N_{\tau}=0.75$ | 181 |
| 137 | Tau normalization offset for a negative pull, $N_{\tau}=0.75$ | 181 |
| 138 | Positive pull, $N_{\tau}=1.00$ | 182 |
| 139 | Tau normalization offset for a positive pull, $N_{\tau}=1.00$ | 182 |
| 140 | Negative pull, $N_{\tau}=1.00$ | 183 |
| 141 | Tau normalization offset for a positive pull, $N_{\tau}=1.00$ | 183 |
| 142 | Positive pull, $N_{\tau}=1.25$ | 184 |
| 143 | Tau normalization offset for a positive pull, $N_{\tau}=1.25$ | 184 |
| 144 | Negative pull, $N_{\tau}=1.25$ | 185 |
| 145 | Tau normalization offset for a positive pull, $N_{\tau}=1.25$ | 185 |

LIST OF TABLES

| | | |
|----|---|-----|
| 1 | Definition of the Fiducial Volume of the DeepCore Filter | 23 |
| 2 | Result of the DOM-to-DOM noise coincidence rate | 35 |
| 3 | Fit parameter used in the hadronic energy scaling | 77 |
| 4 | Event rates at Level 7 | 92 |
| 5 | Final Analysis Binning | 97 |
| 6 | Fixed and unfixed oscillation parameters for this analysis (Normal Ordering) | 101 |
| 7 | Systematic uncertainties included in the tau neutrino appearance analysis . . | 118 |
| 8 | Outcome of additional fit configurations. | 133 |
| 9 | Systematic uncertainties considered in the analysis | 163 |
| 10 | Hyperparameters used in the training of the Level 7 Muon Classifier | 186 |

INTRODUCTION

The 20th century can probably be thought of as the most prolific era in the development of our understanding of matter. Beginning in the 1920's, an impressive synergy between quantum mechanics, classical field mechanics, group theory and relativity, combined with a large wealth of experimental evidence discovered from the 1950's onward, led to the development of what we know today as the Standard Model of particle physics. This model has had several predictive successes, yet it remains an incomplete framework, as it struggles to explain certain experimental measurements.

The most direct challenge to the Standard Model comes from neutrinos, light-weight neutral particles that were initially proposed by Wolfgang Pauli in 1930 and later discovered by the Cowan-Reines experiment in 1956 [1]. At the core of the problem is the phenomenon of neutrino oscillation, where neutrinos change flavour as they propagate through space. The confirmed experimental manifestation of oscillations implies the existence of massive neutrino states, and the Standard Model fails to provide an adequate mechanism for generating it.

A number of theories suggest that the standing problem of the neutrino mass can be resolved by invoking additional neutrino states beyond the three that have already been measured. This proposition has consequences on the assumed unitarity associated with the mathematical description of neutrino oscillation. The core of this thesis presents an analysis which aims to probe the unitarity of neutrino oscillation by measuring the appearance of tau neutrinos in a large data sample of neutrinos collected by the IceCube Neutrino Observatory, from 2012 to 2020.

Beyond oscillations, neutrinos also possess interesting properties that allow them to travel large distances through space unimpeded. As such, they have the potential for revealing yet unknown information about very energetic processes that occurred in distant regions of the Universe, information that could further the understanding of physics beyond what Earth-based colliders can do. This thesis also presents the result of a search for astrophysical neutrino sources in the local Universe, which was published in [2].

This thesis is divided in three parts. Part I presents background information relevant to the two analyses featured in the thesis, namely an introduction to the Standard Model and its relationship with neutrinos (Chapter 2), and a description of the IceCube neutrino detector (Chapter 3), including a detailed laboratory study of the instrumental noise in the IceCube optical sensors (Chapter 4). Part II is dedicated to the tau neutrino appearance analysis, and begins with the description of the mathematical framework used to describe three-flavour neutrino oscillations (Chapter 5); it continues with a presentation of the main principles of the analysis (Chapter 6) and the event selection process (Chapter 7); Chapter 8 covers how the analysis is practically implemented, while Chapter 9 presents the results of the pre-unblinding phase, where a series of tests are run prior to obtaining permission for looking at the real data. Finally, Part III of the thesis covers the neutrino astronomy search, starting with an overview of the field (Chapter 10) followed by a description of methods and results obtained in the correlation search (Chapter 11).

STATEMENT OF CONTRIBUTIONS

The work featured in this thesis is the result of a collaboration with multiple individuals. My contributions to this work is as follows:

All experimental methods, measurements and results presented in the noise measurement section of Part I are the result of my own work, with some logistical support from the SNOLAB personnel and additional support provided by many conversations with IceCube calibration experts.

The work that went into the tau neutrino appearance analysis was shared by several members of IceCube's oscillation working group. Many aspects of the analysis, namely the core principle and the methods used in the lower levels of the event selection, have been initially developed in a previous iteration of this analysis. My specific contributions in the improvement of the event selection consist of the training and implementation of the Level 7 Muon classifier, along with several contributions in the checks performed on the quality of the selection, from Level 4 onward: these concern data-MC agreement, all seasonal consistency checks and a few investigations on the efficiency of certain types of cuts associated with noise and muon rejection.

The core simulation software, and the software responsible for processing events up to Level 7 was established by fellow collaborators. The core analysis software (PISA) as well, though I have played a role in the development of a few core features related to the likelihood implementations. I was responsible for revisiting the electromagnetic-to-hadronic factor that goes into the energy estimation, but have not contributed to the reconstruction algorithm itself (RetroReco). The development of the scripts used to perform the various pre-analysis, blind fits, systematic checks and post-blind fit tests is the result of a close collaboration between Tom Stuttard (postdoc, NBI), Kayla Leonard (PhD student, UW-Madison) and myself: so were the scripts used to interpret and visualize the outcome of these tests. The specific choices made in the analysis pipeline are also the outcome of the same collaboration, though I have diverged from the collaborative work starting from Section 8.5.

The work done in the 2MRS correlation search was mostly done by myself, with guidance from Mohamed Rameez, who wrote the theoretical paper outlying the analysis principle, and helped me derive the limits features at the end of Chapter 11. I have written the core software used in this analysis to perform the multiplet selection and run the likelihood optimization fit on both scrambled and un-scrambled data. That software relies on several packages developed in python and in C by other people (such as GSL, Healpix, and other python libraries). Additional scrutiny by IceCube's neutrino source working group and by other members of the collaboration has helped the development of the analysis. The paper in which this analysis is featured was mostly co-authored by Steve Sclafani (PhD student, Drexel) and myself. Steve is the main analyzer of the *template analysis* featured in [2].

Part I

DETECTOR AND INSTRUMENTATION

NOTIONS IN PARTICLE PHYSICS

As this thesis covers both aspects of fundamental particle physics and astronomy, the first natural question that can occur in the reader's head might be: where to start?

The chosen approach in this work is to begin with an introduction of what is common to both topics: the theoretical context in which the properties of the neutrinos are understood, and the instrumentation used to acquire the data featured in the two analyses, namely the IceCube Neutrino Observatory.

2.1 THE STANDARD MODEL OF PARTICLE PHYSICS

Modern particle physics offers a coherent understanding of matter and the interactions that govern its behaviour, by reducing them into combinations of 17 particles¹ and four fundamental forces. Particles are divided into two main classes: *fermions*, which constitute the elementary building blocks of matter and have a spin of $1/2$, and *bosons*, which act as mediators of the fundamental forces and have integer spin numbers. The fundamental forces include the *electromagnetic force*, the *weak nuclear force* (or weak interaction), and the *strong nuclear force* (the fourth one, the *gravitational force*, has yet to be included in this framework). The theory encompassing these concepts is known as the Standard Model of particle physics, which is summarized graphically in Fig. 1: the leftmost figure presents the properties of all particles mentioned above, while the rightmost figure graphically represents how each particle type is linked to the fundamental interactions.

The Standard Model distinguishes two sub categories of fermions: *quarks* which are elementary constituents of *hadrons* (composite particles like pions, protons and neutrons), and *leptons* which include three types of electron-like particles with an electric charge of -1 , and three neutral particles called *neutrinos*. Particles of the lepton family interact only via the electromagnetic force, whose carrier is the *photon*, and the weak force, which is mediated by the Z and W^{+-} bosons. Quarks additionally interact via the strong force, whose carrier is the *gluon*. Lastly, the *Higgs boson* H^0 is the particle manifestation of the field responsible for generating the masses of all other elementary particles (save for possibly the neutrinos, which is addressed later in section 2.5).

In the theoretical framework of the Standard Model, particles and forces are represented as *fields* $\phi(x, y, z, t)$, which describe their properties and kinematics in both time and space. The mathematical formulation of these fields is subject to constraints derived from the laws of physics that are experimentally observed. The most fundamental of these constraints is that a field's *Lagrangian*, $\mathcal{L}(\phi, \delta_\mu \phi)$ must minimize the *action* defined in:

$$S = \int \mathcal{L}(\phi, \delta_\mu \phi) d^4x. \quad (1)$$

Where ϕ and $\delta_\mu \phi$ denote respectively the field and its first derivative, and d^4x denotes an integration over all of space-time. This equation is the fundamental principle behind classical field mechanics,

¹ Note that fermions additionally have corresponding antiparticles




Figure 1: The Standard Model of particle physics. Source: [3].

which is used to describe interactions between continuous fields such as classical electrodynamics [4]. An additional requirement for the proper description of the above quantities comes from special relativity: since the equations of motion derived in any inertial reference frame must be the same, the Lagrangian must be invariant under Lorentz transformations.

A third requirement is that the theoretical framework developed here must be *quantized*, to satisfy the requirements imposed by quantum mechanics. This is achieved through the *canonical*, or *second quantization*, which is the analogous process of deriving the quantum mechanical equations of motions for particles (also known as the first quantization), but this time with fields.

Lastly, groups of Lagrangian terms associated to the conservation of a particular quantum property (such as e.g. electric charge or spin) exhibit certain *symmetries*, a feature known as Noether's theorem. The association of fundamental interactions to groups of symmetries is a key concept of *gauge theory*, which further restricts the expression of the Lagrangian to representations that remain invariant under the transformations allowed by its associated symmetry group. As these transformations are dependent on space-time coordinates, one speaks of *local symmetries*. Symmetry groups are at the core of the search of a unified theory of physics, where the experimentally observed forces of nature could stem from a single, larger supersymmetric group of interactions. While this grand unification has yet to occur, this approach was quite successful in unifying two of the fundamental forces, electromagnetism and the weak interactions, into a coherent overarching framework called the *electroweak theory*.

Because the topics in this thesis involve neutrino interactions, the electroweak theory will be given further thought in section 2.4.

2.2 CHIRALITY IN PARTICLE PHYSICS

Chirality is usually defined as the projection of a particle's spin on its momentum vector, in the asymptotic ultra-relativistic limit. It is a fundamental property of particles that defines their behaviour under

the application of the parity operator, which describes the result of applying a point-reflection as:

$$\mathbf{P}\psi(\vec{x}, t) = \psi(-\vec{x}, t). \quad (2)$$

Where $\psi(\vec{x}, t)$ represents the wavefunction of a particle. If the resulting wavefunction of a particle is antisymmetric under the operation (meaning $\psi(\vec{x}) = -\psi(-\vec{x}, t)$), the interaction is known to be chiral: this means that we can divide particles into two categories based on their *handedness*. This separation leads to the notion that particles from the Standard Model can be *left-handed* or *right-handed*.

This distinction is of great importance for the neutrinos, as electroweak interactions are known to be chiral, and work differently on left-handed and right-handed fermions.

2.3 THE CONCEPT/PROBLEM OF PARTICLE GENERATIONS

A defining feature of the Standard Model is the fact that fermions appear to exist in three sets of *generations*, which differ by their masses and their flavour quantum number. The charged leptons, for instance, exist as either *electrons* (e , first generation), *muons* (μ , second generation) or *tauon, or tau lepton* (τ , third generation). While these three particles have different masses, all of them possess the same electric charge of -1. In a similar manner, three generations of *up-type quarks* (Up u , Strange s and Top t , with an electric charge of +2/3) and *down-type quarks* (Down d , Charm c and Bottom b , with an electric charge of -1/3) also exist.

To this day, the origin of particle generations is still an unsolved problem in particle physics. In the Standard model, generations are addressed theoretically by invoking a set of *ad hoc* constants called *Yukawa couplings*, giving each generation of fermions different masses during the spontaneous symmetry breaking mechanism (see section 2.4.2). The value of these constants are however not the result of a theoretical prediction, and are essentially derived from the measured masses of the particles in the laboratory.

2.4 THE ELECTROWEAK FRAMEWORK

The electroweak theory is defined by the symmetry group $SU(2)_L \times U(1)$ [5]. This means that all interactions under this force originate from the following local symmetry groups:

- $SU(2)$, which is a symmetry group that can be represented by Pauli matrices,
- $U(1)$, which denotes phase transformations $\exp -i\theta$, where $\theta \in [0, 2\pi]$.

Also essential is the notion that the electroweak interaction is *chiral*, as the $SU(2)$ transformations only apply to left-handed particles (hence the subscript L); the $U(1)$ transformations meanwhile apply to both right-handed and left-handed particles. The most general Lagrangian satisfying this gauge

group² is:

$$\begin{aligned}
\mathcal{L} = & i\overline{L}_L \not{D} L_L + i\overline{Q}_L \not{D} Q_L + \sum_{f \in [e, u, d]} i\overline{f}_R \not{D} f_R \\
& - \frac{1}{4} \underline{A}_{\mu\nu} A^{\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} \\
& + (D_\mu \Phi)^\dagger (D^\mu \Phi) - \mu^2 * \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2 \\
& - y^e (\overline{L}_L \Phi e_R + \overline{e}_R \Phi^\dagger L_L) \\
& - y^d (\overline{Q}_L \Phi d_R + \overline{d}_R \Phi^\dagger Q_L) \\
& - y^u (\overline{Q}_L \tilde{\Phi} u_R + \overline{u}_R \tilde{\Phi}^\dagger Q_L).
\end{aligned} \tag{3}$$

Each term has been colour-coded based on the types of interaction it leads to: terms involving neutrinos directly are coloured in blue, while terms in red represent are involved in neutrino-matter interaction. The orange terms, meanwhile, are explored in more details in section 2.4.2. Equation 3 summarizes the possible types of interactions that can occur in the electroweak theoretical framework. As can be seen above, the latter prompts interactions between two different sets of left-handed, fermion *doublets*:

$$\begin{aligned}
L_L & \in \left(\begin{bmatrix} \nu_e \\ e_L \end{bmatrix}, \begin{bmatrix} \nu_\mu \\ \mu_L \end{bmatrix}, \begin{bmatrix} \nu_\tau \\ \tau_L \end{bmatrix} \right) \\
Q_L & \in \left(\begin{bmatrix} u_L \\ d_L \end{bmatrix}, \begin{bmatrix} c_L \\ s_L \end{bmatrix}, \begin{bmatrix} t_L \\ b_L \end{bmatrix} \right),
\end{aligned} \tag{4}$$

and three sets of right-handed, fermion *singlets*:

$$f_R \in (e_R, u_R, d_R), (\mu_R, c_R, s_R), (\tau_R, t_R, b_R). \tag{5}$$

Were it not for the bottom four lines of this Lagrangian, all fermions would be massless, and neutrinos would interact via the fields included in the *covariant derivative* $\not{D} = \gamma^\mu D_\mu$ ³, defined as:

$$\gamma^\mu D_\mu = \gamma^\mu \left(\delta_\mu + ig \underline{A}_\mu \cdot \underline{I} + ig' B_\mu \frac{Y}{2} \right), \tag{6}$$

Where : $\underline{A}^\mu = (A_1^\mu, A_2^\mu, A_3^\mu)$ and $\underline{I} = (I_1, I_2, I_3)$.

In the above equation, B^μ and the elements of operator \underline{A}^μ are vector potentials, respectively associated with the weak hypercharge Y and weak isospin I quantum numbers, the quantities conserved in this gauge group. Also introduced in Eq. 6 are two coupling constants intervening in the electroweak force, g and g' . Finally, components of operator \underline{A}^μ are associated with a set of generator matrices \underline{I} , which constitute an orthonormal basis capable of describing the symmetry operations of the SU(2) group. One representation of these matrices is known as the set of *Pauli matrices* defined as:

$$I_1 = \frac{1}{2} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad I_2 = \frac{1}{2} \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad I_3 = \frac{1}{2} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{7}$$

² Note that this section is only derived for a single generation of fermions.

³ In this notation, γ^μ represents *gamma matrices*, the group of matrices capable of generating all the possible Lorentz transformations.

2.4.1 Separating the Weak and Electromagnetic Forces

The hypercharge Y operator yields a different value of hypercharge depending on the fermion to which it is applied. For the lepton fields, one obtains:

$$YL_L = -L_L, \quad Y\alpha_R = -2\alpha_R, \quad \text{Where } \alpha \in [e, \mu, \tau]. \quad (8)$$

With this equation and the other relationships described earlier, it is possible to expand the interaction Lagrangian term involving a neutrino and lepton flavour α , which takes the form⁴:

$$i\bar{L}_L \not{D} L_L = -\gamma^\mu (\bar{\nu}_{\alpha,L} \quad \bar{\alpha}_L) \cdot \left(g \frac{A_1}{2} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + g \frac{A_2}{2} \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} + g \frac{A_3}{2} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} - g' \frac{B}{2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \cdot \begin{pmatrix} \nu_{\alpha,L} \\ \alpha_L \end{pmatrix}. \quad (9)$$

Combining the 2x2 matrices into a single one, we get:

$$i\bar{L}_L \not{D} L_L = -\frac{1}{2} \gamma^\mu (\bar{\nu}_{\alpha,L} \quad \bar{\alpha}_L) \cdot \begin{bmatrix} gA_3 - g'B & g(A_1 - iA_2) \\ g(A_1 + iA_2) & -(gA_3 + g'B) \end{bmatrix} \cdot \begin{pmatrix} \nu_{\alpha,L} \\ \alpha_L \end{pmatrix}. \quad (10)$$

As can be seen from Eq. 10, the interactions of the lepton doublet with the electroweak bosons fall into two different categories. From the non-diagonal terms, it is possible to derive the *charged-current interaction* Lagrangian, which allows for the transformation of a left-handed lepton α_L into its associated neutrino $\nu_{\alpha,L}$:

$$\mathcal{L}_{CC} = -\frac{g}{2} \gamma^\mu [\bar{\nu}_{\alpha,L} (A_1^\mu - iA_2^\mu) \alpha_L + \bar{\alpha}_L (A_1^\mu + iA_2^\mu) \nu_{\alpha,L}]. \quad (11)$$

From the diagonal terms of Eq. 10, it is possible to define the left-handed component of the *neutral-current interaction* which, for the neutrino, is defined as:

$$\mathcal{L}_{NC} = -\frac{1}{2} \gamma^\mu \bar{\nu}_{\alpha,L} (gA_3 - g'B) \nu_{\alpha,L}. \quad (12)$$

Moreover, the pure electroweak boson fields can also be re-expressed as a more familiar combinations using eqs. 13, 14 and 15.

$$W^\mu = \frac{A_1^\mu - iA_2^\mu}{\sqrt{2}}, \quad (13)$$

$$Z^\mu = \cos(\theta_W) A_3^\mu - \sin(\theta_W) B^\mu, \quad (14)$$

$$A^\mu = \sin(\theta_W) A_3^\mu + \cos(\theta_W) B^\mu. \quad (15)$$

These equations represent the definition of the weak bosons Z^0 and W^{+-} and the electromagnetic vector potential A^μ associated with the photon A^μ . This choice of representation is made to mirror the observations that we experimentally detect photons and weak bosons. The angle θ_W is called the Weinberg angle, and it represents the mixing of the original electroweak bosons and their neutral manifestation as Z^0 s and γ s.

⁴ Here, the purely kinetic term δ_μ is ignored.

At the end of these transformation, it is possible to defined the neutrino's **weak charged current interaction** as:

$$\mathcal{L}_{CC} = -\frac{g}{\sqrt{2}} \bar{\nu}_\alpha \gamma^\mu (1 - \gamma^5) \nu_\alpha W_\mu, \quad (16)$$

where the term $(1 - \gamma^5)$ is an operator that allows one to obtain the left-handed component of a particle field. The **neutral current interaction**, meanwhile, can be defined as a combination of two terms: one involving weak interactions, and one involving electromagnetic interactions.

$$\mathcal{L}_{NC} = -\frac{1}{2} \bar{\nu}_\alpha \gamma^\mu [(g \cos(\theta_W) + g' \sin(\theta_W)) Z_\mu + (g \sin(\theta_W) - g' \cos(\theta_W)) A^\mu]. \quad (17)$$

Since it is observed experimentally that neutrinos do not interact electromagnetically, the second term is equal to zero. This means that $g \sin(\theta_W) = g' \cos(\theta_W)$, which allows one to re-express the coupling from Eq. 17 with the following manipulation:

$$\begin{aligned} g \cos(\theta_W) + g' \sin(\theta_W) &= g \cos(\theta_W) + g \frac{\sin^2(\theta_W)}{\cos(\theta_W)}, \\ &= g \left[\cos(\theta_W) + \frac{1}{\cos(\theta_W)} (1 - \cos^2(\theta_W)) \right], \\ &= \frac{g}{\cos(\theta_W)}. \end{aligned} \quad (18)$$

The weak neutral current Lagrangian can therefore be written as:

$$\mathcal{L}_{NC} = -\frac{g}{2 \cos \theta_W} \bar{\nu}_\alpha \gamma^\mu (1 - \gamma^5) \nu_\alpha Z_\mu. \quad (19)$$

Equations 16 and 19 can be visually represented with the use of *Feynman diagrams*: Figure 2 illustrates the fundamental interaction vertices involving neutrinos.




Figure 2: The fundamental Feynman vertices of neutrino interactions: charged (left) and neutral (right) current interactions

2.4.2 Spontaneous Symmetry Breaking and Massive Particles

The bottom four lines of Eq. 3 introduce further complexity to the (so far) relatively simple form of the electroweak interactions described earlier. This complexity stems from the scalar potential Φ associated with the *Higgs boson*. The Higgs potential defined by the last two terms of the third line can be re-written as:

$$\begin{aligned} -V(\Phi) &= \mu^2(\Phi^\dagger\Phi) + \lambda(\Phi^\dagger\Phi)^2, \\ &\approx \lambda \left(\Phi^\dagger\Phi - \frac{v^2}{2} \right)^2. \end{aligned} \quad (20)$$

This potential has a non-zero minimum at a value $\Phi^\dagger\Phi = v^2/2$. This value is known as the *vacuum expectation value* (VEV), which introduces massive bosons W^+ and Z when expanding the first term of the same Lagrangian line. These masses are related to the VEV, but also the Weinberg angle θ_W via:

$$m_W = \frac{gv}{2}, \quad (21)$$

$$m_Z = \frac{gv}{2\cos(\theta_W)}. \quad (22)$$

The Higgs potential interaction with the weak force bosons is thus responsible for their masses. Moreover, the Higgs can also interact with fermions via Yukawa interactions such as the ones written in the last three lines of Eq. 3. Particularly, we have for the lepton doublets:

$$-y^e(\bar{L}_L\Phi e_R + \bar{e}_R\Phi^\dagger L_L). \quad (23)$$

Once again, the presence of the non-zero VEV value v in the Higgs potential will lead to mass terms involving pairs of right and left-handed lepton counterparts:

$$m_{l_\alpha} = y_\alpha^l \left(\frac{v}{\sqrt{2}} \right). \quad (24)$$

This means that the mass of a charged lepton of flavour α is proportional to the VEV, as well as to a constant y^α called the *Yukawa coupling*. This constant is actually part of a larger matrix, Y^L , which plays a role in the process of mixing, which is given further thought in the next section.

2.5 THE ROLE OF NEUTRINOS IN ALL OF THIS

In many ways, the theories described above are challenged by what is known today about neutrinos. One of these problems is the absence of right-handed, neutrino singlets, $|\nu_R\rangle$, which prevents the existence of mass terms of the form:

$$-y^e\bar{L}_L\Phi\nu_R. \quad (25)$$

This lack of right-handed neutrinos means that under the Standard model, *neutrinos cannot be massive*. This affirmation is in direct contradiction with the confirmed existence of **neutrino oscillations**, which implies the existence of at least three neutrino states with differentiated masses (see Sec. 5.1 for more details). This is the most direct experimental evidence proving that the prevailing theory in

particle physics remains incomplete.

Another challenge posed by neutrinos, and more generally by the Yukawa coupling scheme introduced in the last section, is the concept of *mixing*. In Eq. 23, the choice of Yukawa coupling has been made to make the mass eigenstate of this part of the Lagrangian coincide with the flavour eigenstates of the weak interaction Lagrangians, Eq. 16 and Eq. 19. However, this choice of parametrization is not unique: the Yukawa coupling matrix Y^l can be any arbitrary choice of unitary matrix, which can later be diagonalized by a set of unitary matrices V_R^l and V_L^l such that $V_R^{l\dagger} Y^{l'} V_L^l = \text{diag}(y_e, y_\mu, y_\tau)$.

This transformation has an impact on the way weak interactions are expressed, as the flavour states must now be expressed as:

$$\alpha_L = V_L l_{\alpha, L}. \quad (26)$$

Inserting that transformation in Eq. 16 for example gives the following:

$$-\frac{g_W}{2\sqrt{2}} \bar{\nu}_\alpha \gamma^\mu (1 - \gamma^5) V_L l_{\alpha, L} W_\mu. \quad (27)$$

Now, because the neutrino has no Yukawa mass term, a transformation of the field by V_L^+ can be made so as to cancel the one made on the charged lepton field⁵. This silent transformation means that *neutrino flavours are actually defined by their associated charged lepton in CC interactions, and not by their intrinsic properties*.

Meanwhile, in the quark sector, the existence of two Yukawa couplings (between Q_L and two right-handed fermions, d_R and u_R) means that the corresponding CC Lagrangian is modified by the introduction of the following transformations:

$$U_L = V_L^U u_L, \quad U_R = V_R^U u_L, \quad D_L = V_L^D d_L, \quad D_R = V_R^D d_R, \quad (28)$$

Where the uppercase states denote the flavor states, and the lowercase states represent the mass eigenstates. As there are two additional transformation matrices involved in this case, and all particles involved have right-handed components, no transformation can hide the change that occurred in the CC Lagrangian, which is altered by the matrix product:

$$U_{CKM} := V_L^{U\dagger} V_L^D. \quad (29)$$

This matrix is known as the Cabbibo-Kobayashi-Maskawa (CKM) matrix, which encodes the fact that mass-eigenstates of the up-like and down-like quarks are not the same as the weak interaction eigenstates. This matrix, and the concept of mixing, are quite interesting because *they are directly applicable to neutrino oscillation* (a similar matrix, the PMNS matrix is used to explained oscillation in section 5.1). Neutrinos, as they are currently understood, thus pose an inconsistency challenge to the Standard Model: they should not have mass and yet, they behave exactly as expected of a mixtures of massives fermion states.

⁵ In a more complicated way, this happens as well in the NC Lagrangian[5], but this will not be shown here.

2.6 POTENTIAL FOR THE EXISTENCE OF HNL AND SEE-SAW MECHANISM

Many theories have been proposed to explain neutrino oscillation and its corollary, the existence of massive neutrinos. A number of them begin with the realization that the inconsistency mentioned above could be resolved with the introduction of a right-handed singlet component to the neutrino, ν_R . This is at the basis of the *see-saw mechanism*, which adds an additional Yukawa coupling to account for this additional neutrino state.

The consequence of adding this terms implies the existence of one, or several, generations of massive neutrino states. Such a right-handed counterpart of the neutrino is called a heavy neutral lepton (HNL). These neutrinos would mix with the other three massive states, but cannot take part in the weak interaction, as the number of participating generation of leptons has already been strictly constrained [6]. These so-called *sterile* neutrinos could be detected in neutrino oscillation experiments as apparent inconsistencies between the standard 3-neutrinos expectation and the observed data; one of these inconsistencies is a breach of *unitarity*, a property of the mixing matrix that is further described in section 5.1.1.

THE ICECUBE NEUTRINO OBSERVATORY

The detection of neutrinos is a challenging endeavour, given their absence of electric charge and weak interaction cross section. Any experiment designed to do so must overcome these obstacles by suppressing large backgrounds, and accurately reconstructing the direction, energy and morphology of neutrino interactions occurring inside the detection volume.

The following section gives a description of the main components of the IceCube Neutrino Observatory, and details the basic physical mechanisms relevant for the detection of neutrino in large Cherenkov detectors. Initially envisioned in the 1960's, the concept for this class of experiment was successfully concretized into a prototype detector called AMANDA[7] in 1997, followed by ANTARES[8] in 2008 and finally, IceCube[9] in 2010.

3.1 THE CHERENKOV EFFECT

When highly energetic charged particles move through a medium, a coherent wavefront of electromagnetic radiation can be emitted along the particle's path, provided that the velocity of the particle is faster than the phase velocity of light in the aforementioned medium. This is known as Cherenkov radiation[10][11], which is the core principle behind neutrino detection in IceCube.

Since the moving particle has an electric charge, its displacement causes a disruption in the field of the polarized surroundings. This leads to the emission of radiation which, due to the refractive index of the medium, travels more slowly than the particle that causes it. As is illustrated in Fig. 3a, this leads to a buildup of constructively interfering wavefronts at a fixed angle with respect to the particle's trajectory, the Cherenkov angle θ_C , which depends on the refraction index n and particle velocity v :

$$\cos(\theta_C) = \frac{1}{n\beta} = \frac{c}{nv}. \quad (30)$$

The observational consequence of Cherenkov radiation is that the fast particle acts as a moving light cone that, if viewed head-on, will be detectable as well-defined circles such as the one that can be seen in Fig. 3b, which shows an example event from Super Kamiokande.

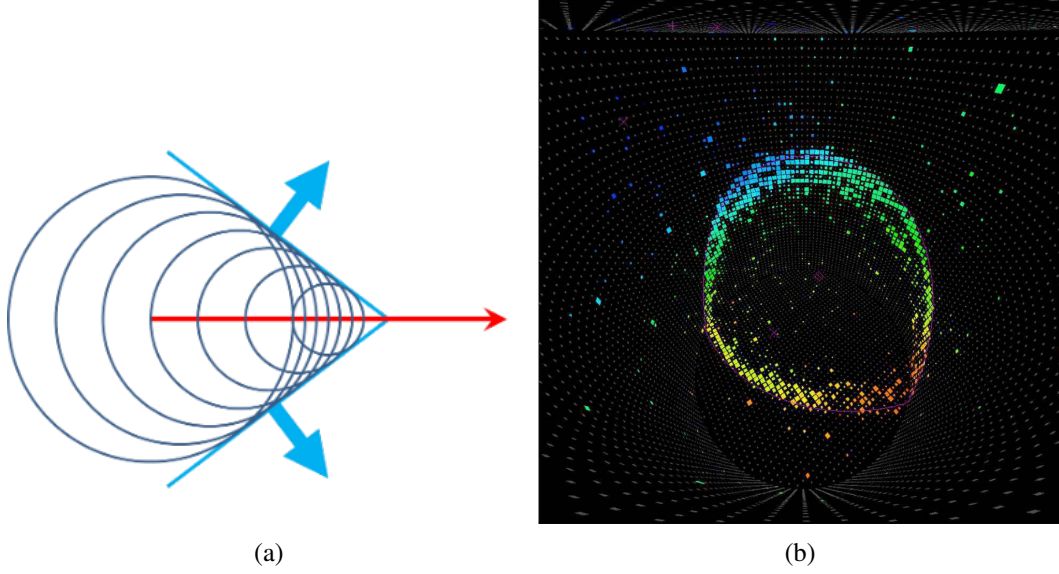



Figure 3: Cherenkov radiation principle and detection. a) Schematic representation of the Cherenkov effect: a particle moving along the red arrow induces radiation depicted by the circles in the drawing. Since these circles propagate outward more slowly than the speed at which they are created, a cone of constructively interfering electromagnetic waves appears, leading to light emission in the directions of the blue arrows. b) Example of a Cherenkov ring detected by Super-Kamiokande. The width of the ring is associated with the time during which a particle satisfies the Cherenkov threshold, once it enters the detector. (source: [12])

Any particle physics experiment using a transparent dielectric medium as an interaction volume can make use of Cherenkov radiation to characterize the morphology and energy of an interaction. Imaging Atmospheric Cherenkov Telescopes (IACTs) use, for example, the atmosphere as a detection medium, and use the Cherenkov light from particle showers to reconstruct very-high-energy gamma rays. Neutrino detectors, on the other hand, make use of large volumes of water, either liquid or frozen, to observe Cherenkov light produced out of neutrino-nucleus interactions.

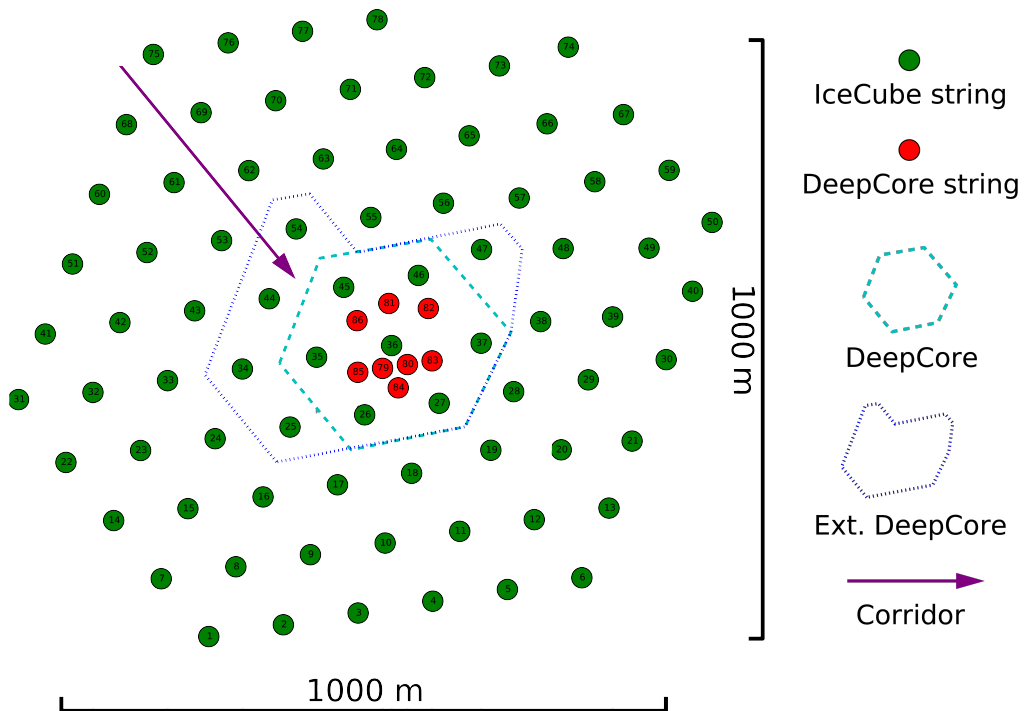
3.2 DETECTOR OVERVIEW

The IceCube Neutrino Observatory, represented in Fig. 4, is a detector made from 1 km^3 of instrumented ice buried 1500 m to 2500 m underground, near the Amundsen-Scott South Pole Station in Antarctica. The detector consists of 86 strings of Digital Optical Module (DOM), which were drilled into the ice and constitute a tri-dimensionnal grid capable of collecting Cherenkov light emitted by particle showers produced in the ice. IceCube is complemented at the surface by IceTop, an ensemble of 81 surface stations used to perform cosmic-ray studies¹.

¹ The data from IceTop is not used in this thesis.



(a)



(b)

Figure 4: The IceCube South Pole Neutrino Observatory. a) Three-dimensional view of the array, showing the location of the strings with respect to the ICL and the depth of instrumented ice. b) Top view of the strings position, including the location of DeepCore strings at the center of the detector.

3.2.1 Light Detection

The elementary unit of detection in IceCube is a photomultiplier tube or PMT. Originally invented in the 1930's[13], photomultipliers have been used extensively in particle physics as a mean to convert single photons into a measurable electrical current[13]. Figure 5 presents the schematics of a typical PMT: it consists of a vacuum-sealed glass bulb, with a sensor end covered by a thin deposited layer of photosensitive material, which is capable of releasing an electron when a photon hits this material, via the photoelectric effect. Under normal operation, a large electric potential is applied between a photocathode and the anode, such that the original photo-electron produced at the cathode can be accelerated toward the other end of the device, where a set of dynodes are positioned to amplify the electronic signal via secondary emission. The end of the dynode amplification stage is then connected the anode, out of which a measurable electronic signal can be converted into a digital waveform.

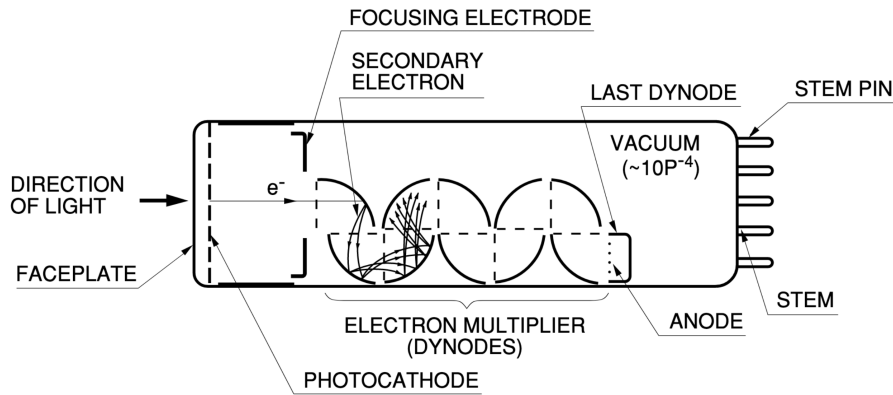


Figure 5: Schematics of a photomultiplier tube's inner workings. Source: [13]

Every PMT in IceCube is integrated into a self-contained digital optical module called a DOM[9]. A DOM consists of a photomultiplier tube, digitizer board, and spherical pressure vessel, along with a set of 12 LED's that are used for calibrating the detector. DOMs convert the electric signal from the PMT into a digitized waveform with the use of both a Fast Analog-to-Digital Converter (FADC), and two Analog Transient Waveform Digitizers (ATWD). These devices, embedded in the DOM's main board, differ in their resolution and timing ranges and provide complimentary information about the events recorded; the FADC covering a longer time range (256 samples at 25 ns/sample), the ATWDs providing higher resolution waveforms to process (128 samples at 3.3 ns/sample). The digitized information is sent back at the surface via one of the many coaxial cables that constitute a string of the detector.

At the surface, the response of all 5160 DOMs is monitored by the experiment's triggering and filtering software, which is responsible for constructing the individual detector events that are part of a particular data sample.

3.2.2 Triggering and Filtering

The most basic level of trigger that exists in IceCube is the discriminator threshold, which is the voltage level above which a particular DOM will begin recording waveform information, once the electronic signal surpasses the digitizer baseline voltage. The absolute value of that threshold is unique to each DOM, and is calibrated each year to correspond to an integrated charge of 0.2 photo-




Figure 6: The Single Photo-Electron (SPE) distribution for a DOM. The histogram shows the distribution of single-photon pulses in data as a function of the charge the photon leaves in the DOM. The pe scale is set by fitting the peak and shape of this distribution to a template, which is shown in blue in the figure. Note that the single photon template is extracted out of a fit that includes additional contributions from multi-photon contamination. Source: [14]

electron units (pe).

Because of uncertainties in the amplification process of a PMT, the absolute charge collected from a single photoelectron emitted at the photocathode will not be exactly the same each time a photon is detected. These statistical fluctuations need to be calibrated for, which is done by making a selection of high-quality, single photon events in the data, and modeling the shape of the Single Photoelectron (SPE) spectrum. Figure 6 shows a typical SPE distribution for an IceCube DOM. The peak of this distribution corresponds to the most likely charge that a single photon will produce in the PMT, which is the quantity used to normalize the charge scale in units of photo-electrons (pe). Another feature of this distribution is the presence of an exponential distribution at very low charge levels. This background contribution includes the contribution of basic electronic noise to the signal coming out of the PMT. The discriminator threshold is chosen to limit the intrusion of this noise background into the waveforms that are recorded.

The next level of hardware trigger is called the *local coincidence* criteria, and requires each DOM to constantly monitor the activity of their immediate neighbors. Each DOM pair is connected to each other with a local coincidence cable. Once a DOM in the pair triggers its discriminator threshold, it will send a signal to the other, at which point a local coincidence will be declared if the other DOM has also triggered its discriminator within $\pm 1\mu\text{s}$ of being notified. The reasoning behind this trigger logic is simple: if two closely spaced DOMs detect a photon within the same time interval, it is likely that the light detected comes from a physical event occurring in the ice.

The local coincidence trigger is of primordial importance, because it determines the type of waveform information that will be recorded about the event. Events triggering a DOM's discriminator threshold, but not a local coincidence will produce Soft Local Coincidence waveforms (SLC), which only convey information about the three bins around the peak FADC waveform. In the other case, the event will be recorded as a Hard Local Coincidence (HLC), and the full waveform information from both the FADC and the ATWD will be recorded and sent to the surface. Figure 7 illustrates the typical information obtained from an HLC event, including the instrumental deadtime which is represented by gray boxes. A zoomed-in view of the first triggered event in this particular event is shown in Fig.

8, where it is possible to compare the resolution and time range of the two digitizers. From Fig. 7 it is also possible to see how the use of multiple digitization channels (particularly, a second ATWD channel) can reduce the long deadtime associated with the digitization of an ATWD waveform. the local coincidence trigger is also a method for reducing deadtime, as it offers the possibility to abort an ATWD digitization process, when it comes from a (likely background) SLC waveform.

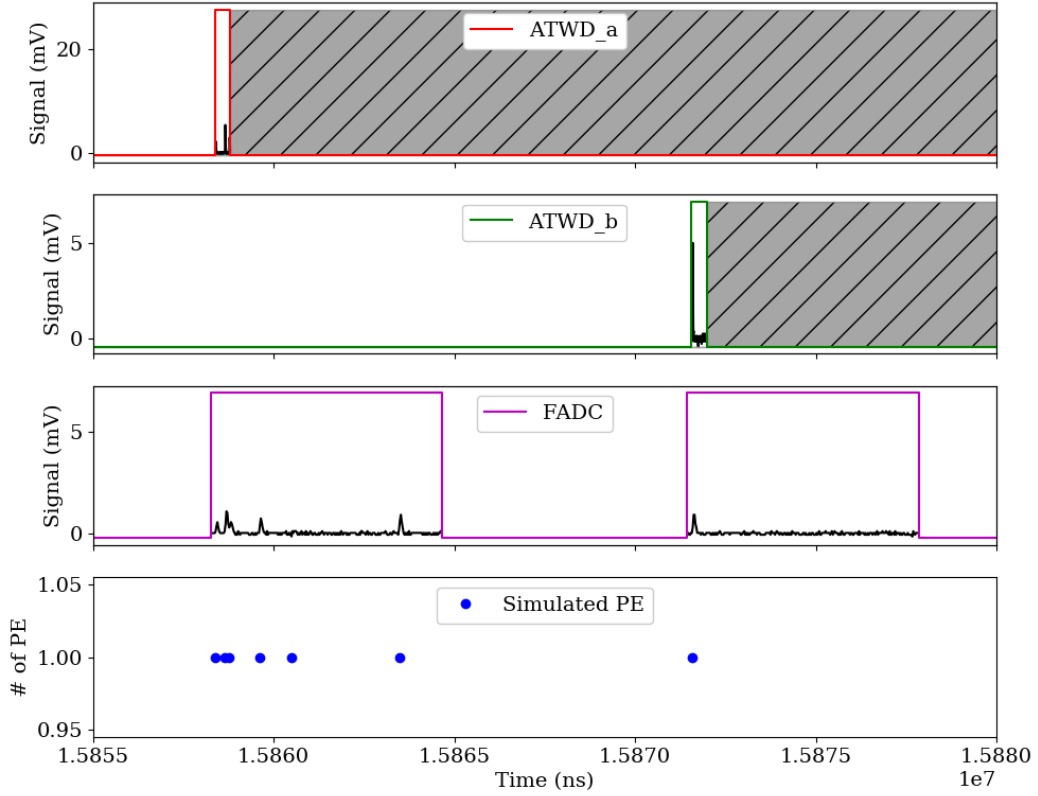


Figure 7: Example of the waveform digitization process of simulated photoelectrons, in an HLC-triggered DOM. The top three boxes show the waveform signal recorded by every digitizing device onboard a single DOM in a simulated event. The bottom box shows the timing of the photons produced in the simulation as a reference. Grayed-out areas indicate deadtime periods associated with the digitization and rearming process of the equipment.

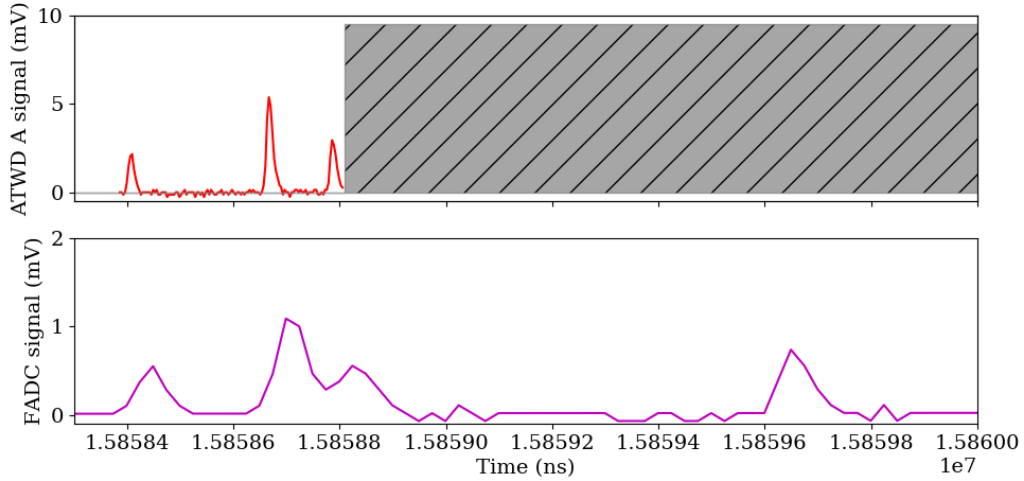


Figure 8: Zoomed-in view of an ATWD waveform and its corresponding FADC trace, for the same HLC hit.

Discriminator thresholds and local coincidences are the only instances of hardware trigger that are encoded into the firmware of a DOM's main board. As such, these types of triggers are usually kept the same for the entirety of the experiment's lifetime, as it is difficult and risky to change them once the detector is powered up. The higher levels of data processing however can be easily changed from year to year, as they pertain to the software that is run online at the IceCube Laboratory (ICL, shown in Fig. 9), the surface building that hosts the servers responsible for receiving and packaging the information from all DOMs into event formats understandable by the entire collaboration. This next level of processing is called filtering, and consists of a series of rules that define which combination of DOM triggers are deemed useful for physics analyses.

Filters vary wildly in configurations, each one being specifically designed for a particular analysis type. A filter can consist of a combination of software triggers: an example of such trigger is the Single Multiplicity Trigger (SMT), which keeps all events with a minimum number of local coincidence instances in the detector. Other filters can also include time window searches, incorporate information from IceTop-triggered events, or look for a particular distribution of charge within the detector. The High-Energy Starting Event filters (HESE) is an example of the latter: it looks for events with at least 1500 p.e. of charge deposited in the detector, with less than 6 p.e. deposited on the boundary DOMs. Section 3.3 describes the particular filter used in oscillation analysis.

In addition to selecting specific DOM trigger configurations, filters may also run some low-level reconstructions algorithms to better select events of interest. These reconstructions are often used offline to perform initial background reduction cuts on an event selection (see Chapter 7). All events passing at least one of the active list of filters are sent via satellite, within minutes to hours, to the main database where they can be viewed and processed by the entire collaboration.

3.2.3 Pulse Reconstruction and Cleaning

Following digitization, the information contained in the FADC and ATWD waveforms is calibrated to subtract specific traits associated with the operational gain and baseline electronics noise of individual DOMs. The calibrated waveforms are then fed through a deconvolution algorithm, which attempts to estimate the distribution of underlying photoelectrons that produced the observed signal. The output of this algorithm is a *pulse* which returns the estimated time and number of photoelectrons



Figure 9: The IceCube Laboratory (ICL)

deposited by the detected photons.

Information from an event in the detector is organized as a *pulse series*, which is a sequence, one per each DOM, containing the arrival time, the charge and the HLC/SLC status of the light pulses detected, following a proper calibration of the raw waveforms. In its initial state, a pulse series will usually be contaminated by several spurious pulses coming from unrelated noise. Since the presence of these pulses can have a negative impact on the reconstruction accuracy of the event, a cleaning algorithm is run to select the subset of pulses that are most likely to be physical.

The cleaning algorithm used in oscillation analyses is a combination of a Seeded Radius-Timing (SRT) and a Time Window (TW) cleaning steps. SRT cleaning uses radius and time information to determine the list of DOMs that are likely to take part in a physics readout. The algorithm starts with a pre-selected list of HLC-triggered DOMs. Iterating over each sensor in the list, the algorithm determines if additional SLC hits were produced within a radius of 150 m from the current DOM. Next, the arrival times of the pulses in the current DOM is compared to the arrival times of the found SLC hits: if the timing of pulses in the latter coincide to the timing of the former within a time window of $1\mu\text{s}$, the SLC pulse information is added to the cleaned pulse series. Figure 10 illustrates the SRT pulse selection process.

The SRT procedure is repeated on the new list of DOMs that now includes a first set of SLC-triggered DOMs. This iterative process is repeated for up to three iterations, or once no more DOMs meet the RT requirements. Following the SRT cleaning, a static time-window cleaning then removes all pulses occurring outside a time window of $[-5, +4]\mu\text{s}$ around the event's trigger time. The pulse series obtained at the end of the time window cleaning is defined as the final, *cleaned* pulse series, and it will be used at analysis-level to reconstruct particle interactions.

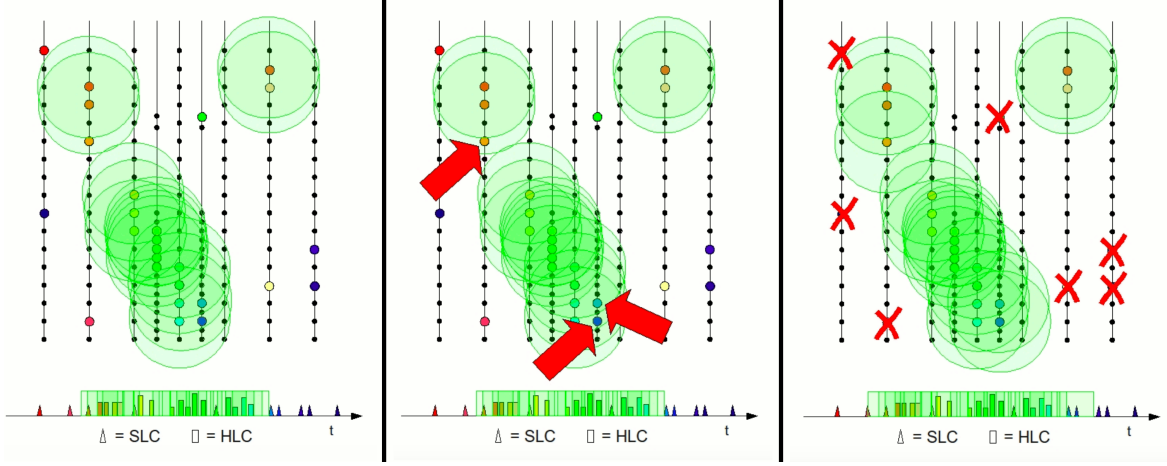


Figure 10: The three steps of an SRT pulse cleaning procedure. In these figures, the colored beads indicate DOMs that have received some light in the event. The green circles indicate the radial area of allowed spatial coincidence around the HLC-triggered DOMs of the event. An associated coincidence time window is also indicated by green boxes around the pulses created by HLC DOMs, on the time axis located at the bottom of the plot. DOMs identified by the red arrows on the middle plot correspond to SLC hits that are reintegrated into the cleaned pulse series, while the red crosses on the right plot indicate SLC hits that are removed by the SRT algorithm.

Note that beyond pulse cleaning and event selection, individual pulse units are rarely used to quantify the physics of an event: this is due to the fact that the deconvolutional algorithm has a tendency to overestimate the number of pulses contained in a waveform with multi-pe charges, a phenomenon is known *pulse splitting*. Instead of relying on pulses, the quantification of physics events is usually based on the numbers of triggered DOMs (a quantity defined as a *hit* in this thesis) or the charge deposited in a DOM.

3.2.4 Characterization of events in IceCube

A critical aspect of particle detectors is their ability to distinguish the energy deposition patterns from different types of events: the better the resolution, the easier it is to identify the interaction vertices and quantify the kinematics of the outgoing particles, which gives information that can be used to identify the types of particle involved and the nature of the interaction. Unlike smaller, but more densely instrumented neutrino detectors, the sparse grid of DOMs available in IceCube do not allow one to see the fine details of a neutrino interaction. Rather, the characterization of an event has historically been limited to distinguishing between two broad categories of events:

- **Cascades**, shown in Fig. 11a are characterized by spherical light patterns outgoing from the interaction vertex. They are produced by hadronic and electromagnetic cascades that accompany nuclei breakups.
- **Tracks**, shown in Fig. 11b form cylindrical light patterns along the path taken by muons. Because the lifetime of a relativistic muon is several microseconds long, it can survive over distances much larger than the detector, thus enabling the array of DOMs to resolve its trajectory.

While every type of neutrino interaction leads to slightly different morphologies (see Section 6.4 for more details), they both fall into these general categories. Cascade-like events are associated with

ν_e CC, most ν_τ CC, as well as all neutral current interactions. Meanwhile, tracks can be produced by background cosmic-ray muons, ν_μ CC and a fraction of ν_τ CC interactions, as they must involve the production of a muon.

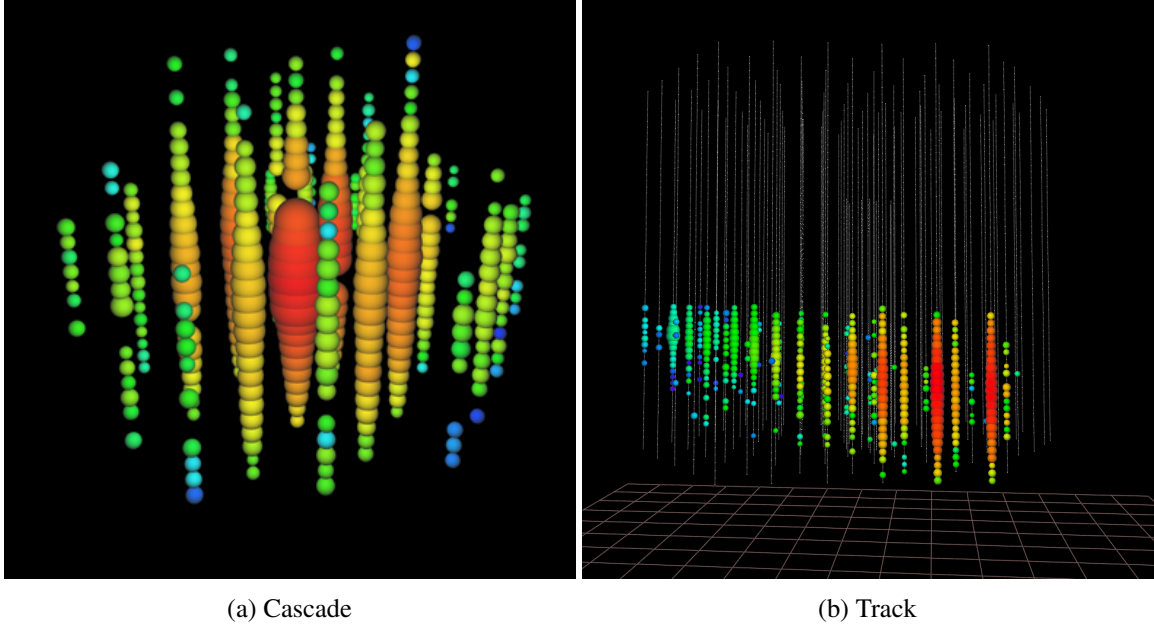


Figure 11: Typical event morphologies in IceCube.

In recent years, the development of machine learning classification algorithms has allowed for a finer characterization of events along a continuous spectrum between the two types of events. These algorithms are particularly useful to characterized events with hybrid or unclear morphologies such as low-energy events used in oscillation analyses (see Section 7.2.6 of Chapter 7).

3.3 THE DEEPCORE SUB-ARRAY

The geometry spacing of the DOMs in IceCube is quite large: 78 of its strings are distant from one another by 125 m on average, and DOMs are positioned at intervals of 17 m on their data cable. This spacing is well adapted for detecting high-energy astrophysical neutrinos with energies of $\mathcal{O}(\text{TeV})$ to $\mathcal{O}(\text{PeV})$, as these interactions are expected to generate copious amounts of light in the ice. Atmospheric neutrinos in the oscillation regime ($\mathcal{O}(\text{GeV})$) produce a lot less Cherenkov light however, meaning that the sensitivity of the IceCube array to events of less than 100 GeV is very limited. The DeepCore sub-array was designed to overcome this problem, and lower the energy threshold of the experiment to energies that would permit atmospheric neutrino oscillation analysis and searches for physics beyond the standard model, like dark matter annihilation searches[15].

DeepCore consists of a set of 8 additional strings located at the center of the array (the red circles in Fig. 4b), and located in the deepest portion of the larger IceCube volume. Strings have an average spacing of 72 m between each other, while DOMs are positioned 7 m apart on each string. Most of the DOMs used in DeepCore were also fit with higher quantum efficiency photomultipliers, which, together with the tighter spacing allowed a reduction of the energy threshold to ≈ 5 GeV.

Associated with DeepCore is a specific filter designed to select events that have the potential to come from low-energy neutrino interactions. The filter looks for triggered events with more than

Table 1: Definition of the Fiducial Volume of the DeepCore Filter

| String type | String number | DOM numbers |
|-------------|-------------------------|-------------|
| DeepCore | 79–86 | 11–60 |
| IceCube | 25–27, 34–37, 44–47, 54 | 39–60 |

3 HLC-triggered DOMs in a fiducial volume defined in Table 1, within a time window of $2.5 \mu\text{s}$ centered around the event trigger. It then applies a dedicated SRT cleaning on the raw pulse series, using a slightly different configuration than the collaboration-wide SRT cleaning (it namely keeps all HLC hits rather than just the core ones). Information from this pulse series is then split according to the fiducial volume and its complementary veto region that includes every DOM outside the fiducial region. The veto is used to eliminate most events associated with atmospheric muons (tracks that deposit light before reaching the fiducial volume).

The DeepCore veto algorithm takes in as input the timing of the first pulse recorded by a hit (i.e. a triggered DOM). The algorithm does the following:

- First, the center-of gravity (COG) of the hits occurring inside the fiducial volume is calculated. Hits occurring more than one standard deviation away from the mean timing distribution are excluded from the COG computation.
- Second, an interaction time is estimated by computing the light travel time between each hit contributing to the COG and the COG itself.
- Third, the time of a hit in the veto region is compared to the estimated interaction time by computing the apparent speed of a particle that would travel from the COG to the hit DOM. If the derived speed of any veto hit is contained within $[0.25, 0.4] \text{ m/ns}$, the event is discarded on the basis that the veto hit could be causally related to a muon crossing the detector.

Figures 12a and 12b illustrate typical cascade and track events at the typical oscillation energy regime. As is evident from these images, the visual distinctions that demarcated tracks from cascades at higher energies are much more difficult to perceive in DeepCore events, which is illustrative of the main experimental challenge faced when one wishes to categorize neutrino interactions in the context of atmospheric oscillations.

Following the application of the DeepCore filter, the remaining events are further processed using a chain of selection steps which collectively are known as the oscNext event selection.


(a) 47 GeV ν_e event classified as a cascade(b) 45 GeV ν_μ event classified as a track

Figure 12: Typical event morphologies in IceCube-DeepCore. The lower energies involved in oscillation analyses means that particle interactions generally deposit less energy in the detector, which means that events reconstructed in this regime are dimmer and have less well-defined shapes.

CHARACTERIZATION OF NOISE IN ICECUBE PMTS

Noise plays an important role in the realm of low-energy physics analyses in DeepCore. Noise in IceCube refers to the sporadic emission of pulses inside individual DOMs, in the absence of external light sources. The major source of this noise is suspected to come from radioactive decays within the glass of the DOMs; recent studies comparing the result of laboratory measurements to simulated scintillation light from radioactive decays corroborate this hypothesis, though it has yet to explain the behaviour of the noise over its full timescale [16].


Noise can be characterized in-situ with a special type of datastream called *HitSpool* [17]. As its name indicates, this stream spools all raw hits detected by all DOMs around a circular buffer, capable of storing several days of data. It is mainly used in supernova-triggered searches, where the raw hit information is critical to understand the evolution of a supernova's neutrino burst¹. Upon detection of a supernova, a request is sent to save dedicated periods of HitSpool data directly to disk; such requests can also be performed manually on demand, for example in order to provide calibration data for IceCube's noise model.

The noise model that is currently used by IceCube was developed in [18] to simulate the noise component of the event pulses detected. It relied on HitSpool calibration data from 2014, and due to limitations of the DOM digitizing software, it needed to extrapolate the behaviour of the noise in timescales shorter than $\approx \mathcal{O}(\mu\text{s})$. This section describes a series of laboratory measurements performed on spare DOMs that were meant to test the validity of IceCube's current noise model over very short timescales.


4.1 CURRENT NOISE MODEL

Initial modeling of the dark rates inside IceCube DOMs assumed that the probability of a noise hit followed a simple Poissonian distribution of parameter λ , denoting the average noise hit rate recorded by a single PMT. However, further investigation into the timing structure of these pulses showed that the noise mechanism was more complex than that. Figure 13a shows the arrival time of four successions of 50 hits, within HitSpool data recorded for a single DOM. As can be seen in the figure, discrete jumps occur whereby several hits are generated within a timescale that is much shorter than the timescale expected from a purely Poissonian, uncorrelated process (represented by the black line). The presence of *correlated noise* can also be visualized in logarithmic space, such as the one shown in Figure 13b. This plot illustrates that the distribution of time intervals between successive noise hits is characterized by a complex structure, spanning all timescales from seconds to micro seconds.

¹ A neutrino burst from a supernova is detectable in IceCube as a momentary increase of global DOM trigger rates.



(a)



(b)

Figure 13: Characteristic representations of noise pulse timing in IceCube. a) Timing of a sequence of noise pulses (labelled as hits in the plot) in a DOM, folded in intervals of 50 ms. The black line indicates a pure Poisson expectation of 220 Hz. b) Distribution of time intervals between successive noise pulses, in base-10 logarithmic scale (known as a log- Δt distribution). Plot taken from [9].

Histograms of time intervals between pulses like the one in fig. 13b, known colloquially as log- Δt distributions, constitute the main data used to model noise properties inside the IceCube DOMs,

which is done in a software module call *vuvuzela*. Vuvuzela attempts to mimic the $\log\text{-}\Delta t$ distribution of each individual DOM, by modeling the timing distribution of the noise pulses as a combination of two probability density functions: an *uncorrelated* one to model the standard Poissonian component, and a *correlated* function to model the other component. The timing of photons produced by the uncorrelated component is Poisson-distributed, with a rate defined by λ_{uncorr} . Meanwhile, two variables are needed to describe the correlated component: one is the number of photons produced per each instance of a correlated event, and the other is the expected distribution of arrival times between these hits, which is modeled by a log-normal distribution of parameters $\mu_{\text{corr},t}$ and $\sigma_{\text{corr},t}$. Adding to this a distinct Poisson rate characterizing the likelihood of a correlated event happening, we end up with a list of five parameters to fully describe the noise behaviour in IceCube DOMs:

- λ_{uncorr} : The expectation rate of the uncorrelated noise component.
- λ_{corr} : The expectation rate of correlated component (akin to the radioactive decay rate).
- η_{corr} : The expected number of photons emitted per each instance of correlated events.
- $\mu_{\text{corr},t}$: The mean time interval between the emission of two correlated photon hits.
- $\sigma_{\text{corr},t}$: The width of the timing distribution.

These parameters were fitted for each of the ≈ 5000 DOMs of the array by performing a likelihood fit of the model to HitSpool data collected in the 2014 season. An important feature of figure 13b is its abrupt cutoff at $\approx \mu\text{s}$ timescales. The reason for this cutoff has to do with the limitations of the digitizing equipment installed onboard the IceCube DOM: for hits which do not satisfy HLC triggering conditions, digitization of the ATWD waveform is abandoned by the DAQ, and only the charge stamp from the FADC is kept in the datastream. Given that a vast majority of noise hits are SLC hits, that all digitizers suffer from $\approx \mu\text{s}$ deatime after triggering, pulse pairs emitted with a $\log_{10}(\Delta t) < -5$ are not included in the data used to fit the noise parameters.

This lack of data in the low $\log\text{-}\Delta t$ region could in theory fail to describe this parameter space, as there is no experimental data to assess the validity of the model over these timescales. This could potentially be problematic in oscillation analyses that rely on low-energy (hence dim) events: the latter are more likely to consist of small bursts of correlated pulses, which is a kind of signature that noise could in mimic.

For the reason mentioned above, a laboratory experiment was designed and run at the SNOLab underground laboratory in Sudbury, Canada to measure the shape of the low $\log\text{-}\Delta t$ region. Given its location, 2100m underneath the surface, and its large overburden of roughly 6000 meter water-equivalent[19], SNOLab was an optimal location for performing noise measurements while being shielded from cosmic rays, which mimics to some degree an experimental setup similar to that of IceCube. Four spare DOMs stripped of their mainboards were shipped to SNOLab, where the signal cable of their PMT was modify to allow for a direct measurement of pulses using a high duty-cycle oscilloscope. Bypassing the original digitizing equipment allowed much finer measurements of noise pulses, because the high resolution of the oscilloscope could easily probe timescales of less than 1 ns. The DOMs were placed inside a commercial freezer, which could reproduce the range of temperatures at which the sensors operate in the real array ($T \in [-8, -35]^{\circ}\text{C}$).

Figures 14a, 14b and 14c show photographs of the experimental setup used to measure the dark noise. First, the high voltage supply of each DOM needed to be adjusted to level the PMT gain between all three DOMs. This was done by placing the sensor opposite an LED flasher originally developed to calibrate the PMTs of the VERITAS telescope [20]. Based on SPE distributions obtained from flasher runs at multiple operational voltage, it was possible to determine the voltage required to

operate each DOM at a gain of $10^{7.5}$. Next, each DOM was covered with electrical tape to avoid measuring photons that would internally reflect back into the PMT at the glass / air boundary². DOMs were then optically isolated from one another and placed inside of the freezer, kept in darkness, while the oscilloscope triggered on pulses produced by noise pulses.




Figure 14: Experimental setup of the SNOLab noise measurements. a) One of the tested DOMs in its calibration setup, inside the freezer. b) Test DOM with electrical tape for the noise measurement. c) Three of the tested DOMs at the end of a noise measurement run.

4.2 COMPARING LOG10DT DISTRIBUTIONS IN SIMULATION AND DATA

Figure 15 shows the $\log-\Delta t$ distributions obtained during the dark noise measurement runs described above, for an operating temperature of -10°C (Appendix A presents the same plots for the other temperatures). The results show that the low $\log-\Delta t$ end of the experimental distribution does extend beyond the range of HitSpool data, as expected. Furthermore, the shape of the short timescale tail ap-

² Since in-situ DOMs are surrounded by ice which has a similar refractive index as glass, such internal total reflection is not expected to occur[21].

pears to be consistent accross all four DOMs used in the measurements, and across all temperatures tested.

To compare the shape of this distribution to the modeling used in vuvuzela, raw photon-level simulation files of pure noise events were generated using the standard production software. However, in order to have access to the short timescale information from these files, the usual detector response software had to be modified so that all hits of a target DOM would be registered as HLC. This allowed for the full waveform information to be kept in the simulation file, which is a way of by-passing the hard microsecond threshold limitations described earlier. Note that in these HLC-only noise datasets, two types of waveforms become available (the ATWD waveform and the FADC waveform). For the purpose of this comparison, pulse timing information from both waveforms were compared separately to the laboratory data.

Figure 16 and 17 show the $\log\text{-}\Delta t$ distributions obtained if one looks respectively at the ATWD or FADC waveform. These present the distributions of nine simulated DOMs with operational temperatures close to the ones used in the laboratory test. Because of the complexity of the DOM's dual digitizer architecture, a quantitative comparison of the two curves was not possible, since no single channel could encompass the behaviour of the noise over the full timescale as was done in the lab. Nevertheless, a qualitative agreement can be observed between the shape of the distribution at low $\log\text{-}\Delta t$ in laboratory settings and a combination of the distributions seen in both channels. For instance, the presence of the small bump at the $\mathcal{O}(10)$ ns timescale in the lab measurements, which is likely the result of either PMT prepulses or early afterpulses [13], can also be observed in the $\log\text{-}\Delta t$ distributions of the ATWD channel. Meanwhile, the relatively continuous prolongation of the correlated component of the $\log\text{-}\Delta t$ histogram, from $2.45\text{ }\mu\text{s}$ down, matches observations from the FADC-derived $\log\text{-}\Delta t$ distribution. One can further add to these observations the fact that the location of the main afterpulsing peak (around a $\log\text{-}\Delta t$ of -5) and the main uncorrelated peak, the relative size of all components of the distribution, and the total span of the distribution are globally consistent between the SNOLAB and the combined ATWD-FADC histograms.

Taken together, the results of these tests indicate that the current model of noise timing properties is not too far off from what is observed directly out the anode of the PMT. The only caveat to this would be the overall scale of the rates obtained in the lab, which are noticeably higher than those obtained in simulation. That topic is further explored in section 4.3.




Figure 15: Distributions of hit times between successive noise pulses in the four tested DOMs, at -10°C .




Figure 16: Distributions of ATWD hit times between successive noise pulses in simulated DOMs, at -10°C .




Figure 17: Distributions of FADC hit times between successive noise pulses in simulated DOMs, at -10°C .

It should be noted that while it was not deemed necessary to pursue the noise structure comparison to a more precise extent, the technique devised in this study to get rid of the hard $\log\text{-}\Delta t$ cutoff at $2.45\text{ }\mu\text{s}$ can also be applied to HitSpool data from the in-situ DOMs. Much like in simulation, the local coincidence bypass trick mentioned above can be configured into a DOM's firmware to extend the timescale of the recorded noise pulses into the low $\log\text{-}\Delta t$ region. This study has therefore laid out a path for a finer characterization of the noise properties of individual DOMs, if it ever becomes apparent that further precision modeling of this mechanism becomes necessary for future physics results.

4.3 RESULTS OF DARK RATE TESTS

Using the same data, the total noise rate of each tested DOM was calculated at four different temperatures ($[-10, -20, -30, -40]^\circ\text{C}$). The correlated and uncorrelated noise rates were then extracted from the raw histogram of time intervals between successive pulses, like the one in Fig. 18. Uncorrelated rates were obtained by fitting the rate of a Poisson-distributed hit sequence to the long-timescale data of the histogram (as uncorrelated noise is dominant in the 0.1 to 1.0 ms range). The rate of the correlated component is then estimated by subtracting the uncorrelated component from the total rate.

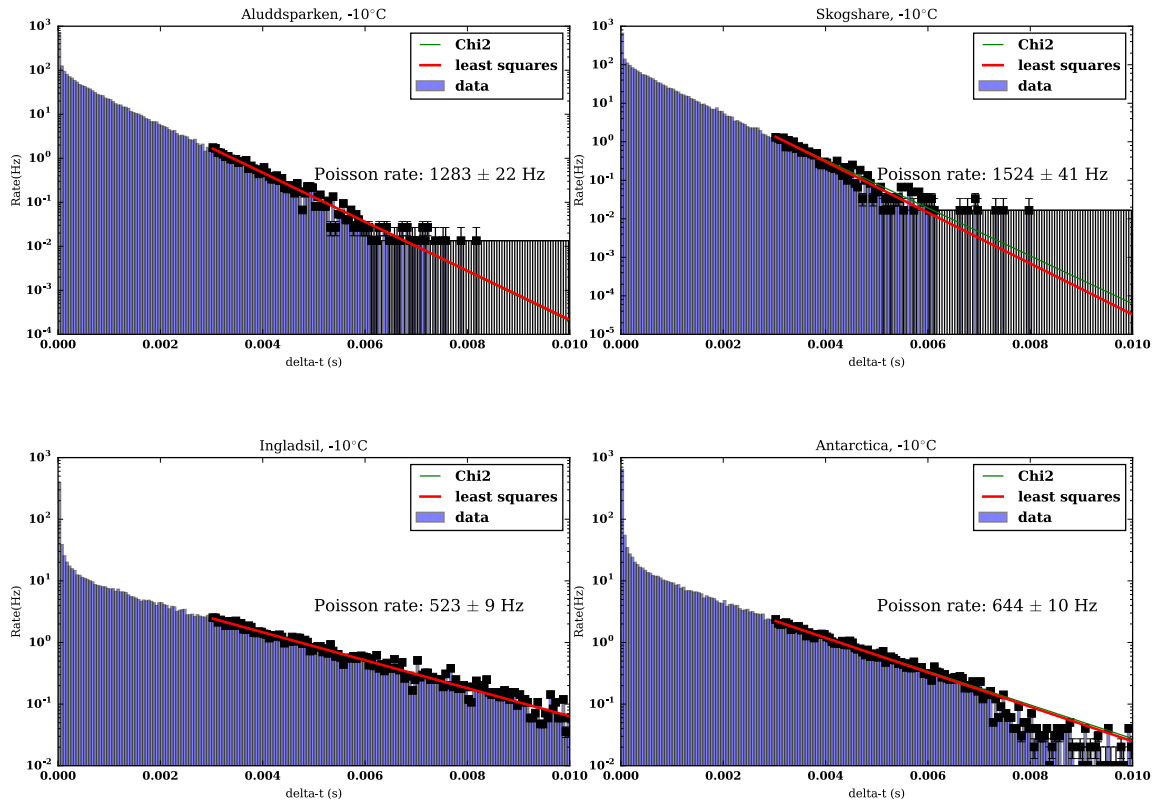


Figure 18: Fit of the raw time difference distribution to a Poisson-distribution.

Figure 19 summarizes the rates measured as a function of the temperature. Several observations were surprising from these plots. First, the large variability between the four DOMs tested, particularly with respect to the uncorrelated component rate, seemed to indicate that the choice of the DOMs used for this study might not have been ideal: the latter were taken out of a pile of initially rejected DOMs that had never been fully characterized, meaning that the DOM used for this setup might not have been representative of the ones used in IceCube. However, a careful look at the calibration logs

of these DOMs did not indicate that the reason for which each one had been rejected was connected to the PMT³. Second, the overall rate of noise pulse recorded as a function of temperature averaged a value roughly twice as high as the average rate measured in-situ [9], even taking into account the standard deviation of the measurements across all DOMs.

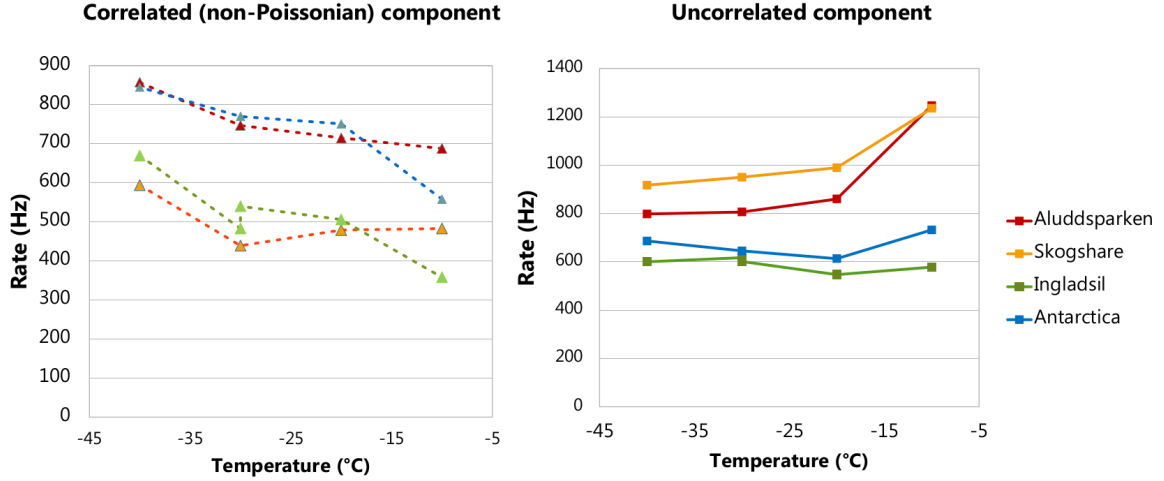


Figure 19: Correlated and uncorrelated components of the noise rates for the four DOMs tested, as a function of temperature.

Larger noise rates have also been observed by other surface lab measurements [21], [16] (see Fig. 20b), which might indicate that *cosmic ray activation* might play a role in the production of isotopes inside the glass. This phenomenon has been documented in the case of dark matter experiments, for instance in silicon [22] and germanium [23] detectors. Several types of construction material (such as copper and titanium) have also been a source of undesired isotopes, activated from exposure to cosmic rays at the surface [24]. These materials used in the construction of these detectors must usually spend an extended period of time underground before being cleared to use.

³ All DOMs used were actually rejected because of mainboard problems, which were in any case not used in the SNOLAB tests.




Figure 20: a) Evidence of DOM quieting on six DOMs of String 1. The total trigger rate of these DOMs was measured at frequent intervals following the initial deployment of the string, uncovering an exponential decrease in the observed rate. b) Total noise rates measured in different laboratory settings: at SNOLab in the study described in this chapter (black), at UW-Madison during the Final Assembly Testing phase prior to deployment (yellow), and in Chiba during calibration studies for the IceCube Upgrade (orange). The nominal noise rate measured in-situ is shown in blue.

Previous studies of the in-situ noise rates have also demonstrated a certain *quieting* of the DOMs over the years following deployment. Figure 20a shows the raw DAQ trigger rates for six DOMs of String 1 as a function of time. A clear exponentially decreasing trend is visible, leading to a factor of ≈ 2 drop in the total rate over the course of the first 3 years of operations. Meanwhile, Fig. 20b shows temperature-dependent rate measurements performed in four different measurement campaign, which each measure higher rates than in-situ data, at different time since the DOM manufacturing. A coherent explanation of these two plots could be that once buried deeply underground, the DOM glass vessel ceases to be exposed to the large cosmic-ray radiation that are present during surface measurements. Given that the DOMs used in SNOLab have now remained underground for at least two years, it would be possible to verify this theory by performing a second round of measurements in the future.

4.4 RESULTS OF COINCIDENCE TESTS

Another test performed at SNOLab was to assess the probability that scintillation light from one DOM reaches a neighbouring DOM. While this probability is negligibly small in IceCube and even DeepCore, where DOMs are separated by several meters, this assumption must be revisited with the upcoming deployment of the IceCube Upgrade[25]. First, because the separation between DOMs is much smaller (approximately 2 meters), and second, because the new instruments being deployed will have PMTs oriented upward, leaving pairs of sensors directly facing each other.

To evaluate the probability of coincidence, two DOMs were placed facing each other inside the freezer, at two distances (0.6 m and 1.2 m), as pictured in Fig. 21. The oscilloscope triggered on one DOM, and recorded the signal from both DOMs to disk. After the experiment, the data was analyzed offline: for a time period of 40 ns around the triggering time of the first DOM, the waveform of the second DOM was searched for any traces of pulses occurring within that window.

Table 2: Result of the DOM-to-DOM noise coincidence rate

| Test condition | Coincidence Rate |
|----------------|------------------|
| Control | 0.2 Hz |
| D = 63 cm | 15.8 Hz |
| D = 113 cm | 5.2 Hz |




Figure 21: Schematic of the Coincidence Test experimental setup. Two DOMs are positioned in the freezer facing each other, and the coaxial cable from both PMTs are monitored on a high dutycycle oscilloscope. The latter is set to trigger on light pulses occurring in one of the DOMs: if a light pulse is detected in the other one within a short time window around the trigger, the event is considered a coincident noise event.

The result of this test is shown in Table 2, where the rate of coincidences is reported for a control run (where both DOMs are taped and optically isolated from one another) and for the two distances at which the effect was measured. Assuming an $1/R^2$ dependence on this phenomenon, the expected rate of coincidence between two DOMs in DeepCore is expected to be negligible. Furthermore, since this measurement has been made in air, this can be considered an upper bound on the expected rate in experimental conditions, as the scattering length is much smaller in glacial ice.

At distance of 2 m, which is the planned IceCube Upgrade inter-DOM distance, that same upper bound leads to a manageable rate of ≈ 3 Hz. These results must however be interpreted with caution, as the glass of the new optical modules is known to be intrinsically more noisy than the one used in the original detector. Additionally, the fact that the Upgrade will integrate multiple sensors means that coincidence events might be detected at both ends of an optical module pair. In order to properly account for this potentially new and important source of correlated noise, the coincidence test performed at SNOLAB should definitely be run on Upgrade modules.

4.5 CONCLUDING REMARKS

The results of the multiple measurements conducted at SNOLab can be summarized as follow:

- **Low delta-T distributions:** A comparison between the shape of the $\log\text{-}\Delta t$ distributions, for both the experimentally measured DOMs and simulated DOMs with similar temperature conditions, shows that the two qualitatively agree with each other. A strategy exists to further quantify the agreement using in-situ data, if future physics analyses show a greater sensitivity to the detailed time structure of noise pulses.
- **Correlated and Uncorrelated rates vs. Temperature:** The noise rates inside the four test DOMs was measured at several temperatures. The trends observed differ between the four DOMs, though all four DOMs seem to have higher rates of correlated events as the temperature is lowered, in agreement with prior observations [21] [26] [16]. Given the remarkable coherence of the $\log\text{-}\Delta t$ results between all DOMs, it is unlikely that the variability seen is the result of faulty DOM. Rather, these results could show that this variability is intrinsic to the DOMs population (with a caveat regarding the absolute rate scale). Given the small number of DOMs tested, and a relatively high uncertainty regarding certain parameters of the experiments (namely, the PMT gain), another round of measurement with higher precision equipment should be probably made to confirm the validity of these results.
- **Absolute scale of the noise rates:** It was found that the total noise rate of the tested DOMs about a factor of 2 higher than the currently measured rates inside the detector (500 Hz). This was at first surprising, but after verification is in line with the rates observed in other laboratory tests performed within the collaboration. A proposed hypothesis for this behaviour would involve cosmic ray activation of the glass, but could not be tested due to a lack of long-term measurements on DOMs located in an underground facility.
- **Coincidence Test:** The expected rate of coincidence, whereby a DOM would be triggered by scintillation on its bottom neighbor, is negligible with the sensor spacing used in IceCube and DeepCore. This test should however be repeated for the new IceCube Upgrade sensors, which are made from intrinsically noisier silica glass.

Part II

TAU APPEARANCE ANALYSIS

PROPERTIES OF THE NEUTRINO

Since their original discovery in 1956 by Cowan and Reines [1], significant efforts have been made to better understand neutrinos. It is now known, for example, that there exists at least three flavours of neutrinos, mirroring the generations of charged leptons in the Standard Model [27][28]. It was also shown that due to an asymmetry with respect to how the electroweak force treats particles based on chirality, neutrinos seem to only exist as left-handed fermions.

The most intriguing and surprising property of neutrinos is their ability to change flavours, the phenomenon known as *neutrino oscillations*[29]. As discussed in section 2.5, the experimental observation of this effect throws a wrench in the electroweak theory, by proving the existence of *massive* neutrinos. Under the currently accepted theory, this shouldn't happen, as there exists no right-handed neutrino singlet capable of producing massive particles the same way massive charged leptons are created (see section 2.5).

This chapter describes the known properties of the neutrino, starting with the phenomenology and current experimental overview of neutrino oscillation. A particular emphasis will be made on the concept of *unitarity*, which provides the underlying assumptions used in the standard treatment of three-flavours oscillations; it is also the main phenomenological property that can be probed by the tau appearance analysis. Note that unless specified, all equations derived in this section use natural units ($c = \hbar = 1$).

5.1 NEUTRINO OSCILLATIONS

5.1.1 Unitarity of a Matrix

In order to understand the derivation of neutrino oscillation phenomenology presented in this work, a definition of *unitarity* and its resulting mathematical relationships is necessary. Let M be a matrix of the form¹:

$$M = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{bmatrix}. \quad (31)$$

This matrix is said to be unitary if it fulfills the following mathematical relationship:

$$M^\dagger M = I_3. \quad (32)$$

This relationship states that the product of the matrix with its transpose conjugate must be equal to the 3x3 identity matrix I_3 . Writing down explicitly this product, and the ensuing equalities allows one to derive some properties of the matrix; for example, requiring the diagonal terms of the product

¹ The elements $U_{\alpha i}$ of this matrix will later refer to neutrino oscillation mixing elements later on.

$M^\dagger M$ to be equal to 1 gives us the following unitarity condition:

$$\begin{aligned} |U_{e1}|^2 + |U_{\mu1}|^2 + |U_{\tau1}|^2 &= 1, \\ |U_{e2}|^2 + |U_{\mu2}|^2 + |U_{\tau2}|^2 &= 1, \\ |U_{e3}|^2 + |U_{\mu3}|^2 + |U_{\tau3}|^2 &= 1. \end{aligned} \quad (33)$$

Which states that the complex product of each matrix element in a given column must sum up to 1. Taking the fact that for a unitary matrix, the commutation relationship $MM^\dagger = M^\dagger M$ must hold, we can derive similar constraints for the sum of each squared elements in any row of the matrix:

$$\begin{aligned} |U_{e1}|^2 + |U_{e2}|^2 + |U_{e3}|^2 &= 1, \\ |U_{\mu1}|^2 + |U_{\mu2}|^2 + |U_{\mu3}|^2 &= 1, \\ |U_{\tau1}|^2 + |U_{\tau2}|^2 + |U_{\tau3}|^2 &= 1. \end{aligned} \quad (34)$$

Finally, six additional conditions can be imposed by developing the expressions that must go to zero in the upper and lower triangles of the matrix, which are known as the unitarity triangle constraints:

$$\begin{aligned} &\text{Requiring that } U^\dagger U = I_3 : \\ &U_{e1}^* U_{e2} + U_{\mu1}^* U_{\mu2} + U_{\tau1}^* U_{\tau2} = 0 \\ &U_{e1}^* U_{e3} + U_{\mu1}^* U_{\mu3} + U_{\tau1}^* U_{\tau3} = 0 \\ &U_{e2}^* U_{e3} + U_{\mu2}^* U_{\mu3} + U_{\tau2}^* U_{\tau3} = 0 \end{aligned} \quad (35)$$

$$\begin{aligned} &\text{Requiring that } UU^\dagger = I_3 : \\ &U_{e1}^* U_{\mu1} + U_{e2}^* U_{\mu2} + U_{e3}^* U_{\mu3} = 0 \\ &U_{e1}^* U_{\tau1} + U_{e2}^* U_{\tau2} + U_{e3}^* U_{\tau3} = 0 \\ &U_{\mu1}^* U_{\tau1} + U_{\mu2}^* U_{\tau2} + U_{\mu3}^* U_{\tau3} = 0 \end{aligned}$$

These relationships make up the unitarity constraints that are imposed in the PMNS oscillation framework described in the next section.

5.1.2 Oscillations in vacuum

The phenomenon of neutrino oscillation, as it is presently understood, emerges from the fact that eigenstates of neutrino propagation differ from the eigenstates of the charged-current interaction defined in Eq. 16, which is the basis for the determination of a neutrino's flavour. This means that flavour neutrinos can be represented as superpositions of mass eigenstates: a neutrino of flavour α can be written down as:

$$|\nu_\alpha\rangle = \sum_{k \in 1,2,3} U_{\alpha k}^* |\nu_k\rangle, \quad (36)$$

where the terms $U_{\alpha k}$ are elements of the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [30][31] which defines the mixing strength between each combination of flavour and mass states:

$$M_{PMNS} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{bmatrix}. \quad (37)$$

From an experimental perspective, the observable quantity from neutrino oscillation is the measurement of a change in the flavour composition of some initial flux of neutrinos. That quantity directly depends on the probability of a neutrino to oscillate from flavour α to flavour β , $P_{\alpha \rightarrow \beta}$, which is canonically derived by considering the plane-wave solution to the neutrino propagation hamiltonian \hat{H} , which satisfies the time-dependent Schrödinger equation:

$$i \frac{d}{dt} |\nu_k(t)\rangle = \hat{H} |\nu_k(t)\rangle. \quad (38)$$

Since the propagation eigenstates are the mass states, $k \in [1, 2, 3]$. The solution to Eq. 38 is a wave function of the form:

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle. \quad (39)$$

Using the relationship given by Eq. 36, that solution can be written as:

$$|\nu_\alpha(t)\rangle = \sum_{k \in (1,2,3)} U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle, \quad (40)$$

$$= \sum_{k \in (1,2,3)} U_{\alpha k}^* e^{-iE_k t} \sum_{\beta \in (e,\mu,\tau)} U_{\beta k} |\nu_\beta\rangle. \quad (41)$$

Where Eq. 41, the mass eigenstate $|\nu_k\rangle$ are re-expressed in terms of a superposition of flavour states:

$$|\nu_k\rangle = \frac{1}{\sqrt{U_{\alpha\alpha} U_{\alpha\alpha}^*}} \sum_{\beta \in (e,\mu,\tau)} U_{\beta k} |\nu_\beta\rangle. \quad (42)$$

From mathematical constraints imposed by unitarity (see Eq.33), the normalization factor $\frac{1}{\sqrt{U_{\alpha\alpha} U_{\alpha\alpha}^*}} = 1$. It is then possible to define $P_{\alpha \rightarrow \beta}$ as the quantum mechanical probability to observe a transition from state α at $t_0 = 0$ to state β at later time t , a quantity given by the square of the transition amplitude:

$$A_{\alpha \rightarrow \beta} = \langle \nu_\beta(t) | \nu_\alpha(0) \rangle. \quad (43)$$

Expanding this product we get:

$$\begin{aligned} \langle \nu_\beta(t) | \nu_\alpha(0) \rangle &= \left[\sum_{k \in (1,2,3)} U_{\beta k} e^{-iE_k t} (U_{ek} |\nu_e\rangle + U_{\mu k} |\nu_\mu\rangle + U_{\tau k} |\nu_\tau\rangle) \right]^* \\ &\cdot \left[\sum_{k \in (1,2,3)} U_{\alpha k} (U_{ek} |\nu_e\rangle + U_{\mu k} |\nu_\mu\rangle + U_{\tau k} |\nu_\tau\rangle) \right], \end{aligned} \quad (44)$$

The expression above contains several terms which need to be multiplied together. However, most of these products are greatly simplified, due once again to the unitary nature of the PMNS matrix. For example, computing the first term of each square bracket together, we get the following:

$$\begin{aligned} U_{\beta 1}^* e^{iE_1 t} (U_{e1}|\nu_e\rangle + U_{\mu 1}|\nu_\mu\rangle + U_{\tau 1}|\nu_\tau\rangle)^* \cdot U_{\alpha 1} (U_{e1}|\nu_e\rangle + U_{\mu 1}|\nu_\mu\rangle + U_{\tau 1}|\nu_\tau\rangle) \\ = U_{\beta 1}^* U_{\alpha 1} e^{iE_1 t} (U_{e1}^* U_{e1} + U_{\mu 1}^* U_{\mu 1} + U_{\tau 1}^* U_{\tau 1}), \end{aligned} \quad (45)$$

Which is obtained by considering the fact that the flavour states $(|\nu_e\rangle, |\nu_\mu\rangle, |\nu_\tau\rangle)$ form an orthonormal basis (see Eq.34):

$$\langle \nu_\alpha | \nu_\beta \rangle = \delta_{\alpha\beta}. \quad (46)$$

In addition, the unitary conditions (see Eq. 35 of section 5.1.1) further require the cross-terms of Eq.44 to cancel, as can be seen from rearranging the first of these cross terms:

$$\begin{aligned} U_{\beta 1}^* U_{\alpha 2} e^{iE_1 t} (U_{e1}^*|\nu_e\rangle + U_{\mu 1}^*|\nu_\mu\rangle + U_{\tau 1}^*|\nu_\tau\rangle) \cdot (U_{e2}|\nu_e\rangle + U_{\mu 2}|\nu_\mu\rangle + U_{\tau 2}|\nu_\tau\rangle) \\ = U_{\beta 1}^* U_{\alpha 2} e^{iE_1 t} (U_{e1}^* U_{e2} + U_{\mu 1}^* U_{\mu 2} + U_{\tau 1}^* U_{\tau 2}), \\ = 0. \end{aligned} \quad (47)$$

Bringing everything together, the transition amplitude simplifies to:

$$\langle \nu_\beta(t) | \nu_\alpha(0) \rangle = \sum_{k \in (1,2,3)} U_{\alpha k}^* U_{\beta k} e^{-iE_k t}. \quad (48)$$

Expanding the product $|\langle \nu_\beta(t) | \nu_\alpha(0) \rangle|^2$ yields the general relationship describing oscillation probabilities in vacuum:

$$\begin{aligned} |\langle \nu_\beta(t) | \nu_\alpha(0) \rangle|^2 &= U_{\alpha 1}^* U_{\beta 1} U_{\alpha 1} U_{\beta 1}^* + U_{\alpha 2}^* U_{\beta 2} U_{\alpha 2} U_{\beta 2}^* + U_{\alpha 3}^* U_{\beta 3} U_{\alpha 3} U_{\beta 3}^* \\ &+ U_{\alpha 2}^* U_{\beta 2} U_{\alpha 1} U_{\beta 1}^* e^{-i(E_2 - E_1)t} + U_{\alpha 1}^* U_{\beta 1} U_{\alpha 2} U_{\beta 2}^* e^{-i(E_1 - E_2)t} \\ &+ U_{\alpha 3}^* U_{\beta 3} U_{\alpha 1} U_{\beta 1}^* e^{-i(E_3 - E_1)t} + U_{\alpha 1}^* U_{\beta 1} U_{\alpha 3} U_{\beta 3}^* e^{-i(E_1 - E_3)t} \\ &+ U_{\alpha 3}^* U_{\beta 3} U_{\alpha 2} U_{\beta 2}^* e^{-i(E_3 - E_2)t} + U_{\alpha 2}^* U_{\beta 2} U_{\alpha 3} U_{\beta 3}^* e^{-i(E_2 - E_3)t}, \end{aligned} \quad (49)$$

$$= \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t} \quad (50)$$

The above expression is composed of a sum of diagonal matrix elements $(U_{\alpha k}^* U_{\beta k} U_{\alpha k} U_{\beta k}^*,$ the first row of Eq. 49) and sums of mixed quartic products of the form $U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^*,$ where $j \neq k$. These last terms are complex numbers, and are multiplied by phase terms involving the mass difference between any two pair of states jk . Since each pair of mass indices has a complex conjugate, it is convenient to use Euler's identity and complex number algebra to explicitly reveal the oscillatory nature of the probability. Making the following substitutions:

$$U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* := A_{jk}, \quad (E_j - E_k)t := \phi, \quad (51)$$

We can represent the sum of any two pair of quartic product as a sum of the following terms:

$$U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* e^{-i(E_k - E_j)t} + U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_j - E_k)t} \\ = [\Re(A_{jk}) + i\Im(A_{jk})] e^{-i\phi} + [\Re(A_{jk}) - i\Im(A_{jk})] e^{i\phi}. \quad (52)$$

where the real (\Re) and imaginary (\Im) symbols have been introduced to distinguish the real and complex part of A_{jk} . A couple of algebraic manipulation can then be applied to obtain the terms below:

$$[\Re(A_{jk}) + i\Im(A_{jk})] e^{-i\phi} + \\ [\Re(A_{jk}) - i\Im(A_{jk})] e^{i\phi} = 2\Re(A_{jk}) \left[\frac{e^{i\phi} + e^{-i\phi}}{2} \right] + 2\Im(A_{jk}) \left[\frac{e^{i\phi} - e^{-i\phi}}{2i} \right] \quad (53)$$

$$= 2\Re(A_{jk}) \cos(\phi) + 2\Im(A_{jk}) \sin(\phi) \quad (54)$$

$$= 2\Re(A_{jk}) \left[1 - 2\sin^2\left(\frac{\phi}{2}\right) \right] + 2\Im(A_{jk}) \sin(\phi), \quad (55)$$

Where the trigonometric identity $\cos(2\theta) = 1 - 2\sin^2(\theta)$ was used to convert the cosine into a sine. Re-inserting this result into Eq. 50 leads to the expression:

$$|\langle \nu_\beta(t) | \nu_\alpha(0) \rangle|^2 = \sum_{i=0}^3 U_{\alpha i}^* U_{\beta i} U_{\alpha i} U_{\beta i}^* \\ + 2 \sum_{k>j} \Re \left(U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \right) \cdot \left[1 - 2\sin^2 \left(\frac{(E_k - E_j)t}{2} \right) \right] \\ + 2 \sum_{k>j} \Im \left(U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \right) \cdot \sin((E_k - E_j)t). \quad (56)$$

As it is written above, the probability in Eq. 56 depends on the difference in total energy between pairs of mass eigenstates; if one assumes that all eigenstates possess the same energy as the initially created flavour neutrinos, that difference is dependent on the masses of the two states $|\nu_k\rangle$ and $|\nu_j\rangle$. This can be made evident by writing the energy-momentum dispersion relation, in the case of ultra-relativistic particles. Defining \vec{p} as the momentum of one of the state, and $E_{tot,j}$ its total energy:

$$E_{tot,j} = \sqrt{\vec{p}^2 + m_j^2} \quad (57)$$

$$= |\vec{p}| \sqrt{1 + \frac{m_j^2}{\vec{p}^2}} \quad (58)$$

$$\approx |\vec{p}| \left(1 + \frac{m_j^2}{2\vec{p}^2} \right) \quad (59)$$

$$\approx E_{tot,j} + \frac{m_j^2}{2E_{tot,j}}. \quad (60)$$

In the ultra-relativistic limit, $\frac{m_j^2}{\vec{p}^2} \ll 1$, allowing a first-order Taylor expansion of the square root in Eq. 57, leading to the approximation of Eq. 59. The same ultra-relativistic limit also means that $E_{tot,j} \approx |\vec{p}| = E$, which leads to Eq. 60. Putting back this approximation into the energy difference

term of Eq. 56 leads to the equality $E_k - E_j = (m_k^2 - m_j^2)/2E$. Moreover, the distance traveled by an ultra-relativistic neutrino can be approximated as:

$$L \approx c \cdot t. \quad (61)$$

Defining the symbol $\Delta m_{kj}^2 = m_k^2 - m_j^2$ it is possible to rewrite the oscillation probability expression as:

$$\begin{aligned} P_{\alpha \rightarrow \beta}(t) = & \sum_{i=0}^3 \Re(U_{\alpha i}^* U_{\beta i} U_{\alpha i} U_{\beta i}^*) \\ & + 2 \sum_{k>j} \Re(U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*) \left[1 - 2 \sin^2 \left(\frac{\Delta m_{jk}^2 L}{4E} \right) \right] \\ & + 2 \sum_{k>j} \Im(U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*) \cdot \sin \left(\frac{\Delta m_{jk}^2 L}{2E} \right). \end{aligned} \quad (62)$$

It is important to remember that two approximations were made to derive the oscillation picture described by Eq. 62. One of these comes from Eq. 61, where neutrinos are assumed to have a velocity equal to c , which is not true. Another assumption is that all mass states composing a flavour state $|\nu_\alpha\rangle$ have the same momentum and energy. In reality, Heisenberg's uncertainty principle implies that the momentum and energy of these constituents are subject to intrinsic fluctuations, which brings concerns regarding the coherence of the mass states during propagation. To account for this aspect, the neutrino should be treated as a wave packet, something that has been done in [32]. Fortunately enough, the phenomenological outcome of this treatment does not differ significantly from the one derived above.

5.2 UNKNOWN PROPERTIES OF THE NEUTRINO

Prior to discussing neutrino appearance and disappearance phenomenology, it is important to acknowledge that current unknowns regarding the properties of the neutrino itself may impact the outcome of the oscillation probabilities defined earlier. This section briefly describes three aspects of the neutrino that remain a mystery for modern particle physics, and how they might interfere in neutrino oscillation measurements.

5.2.1 Neutrino Mass and Ordering

As was explained in section 5.1, the existence of neutrino oscillation implies the existence of massive neutrinos. Unfortunately, oscillation experiments are only sensitive to *mass differences* between the three neutrino states, which leaves the question of the neutrino's absolute mass scale unanswered. This is known as the *mass ordering problem*, which can be illustrated by figure 22. Given the known mass splittings, an ambiguity is left as to whether the third mass eigenstate is heavier or lighter than the other two states; one then speaks of *normal* or *inverted* ordering respectively.

Current constraints from dedicated experiments place the effective mass of the electron-neutrino to be smaller than 1.1 eV [33]. Indirect constraints from cosmology, meanwhile, argue for a limit on the sum of the neutrino masses to be smaller than 0.12 or 0.54 eV, depending on the choice of constraints [34].




Figure 22: The two mass ordering configurations: normal (left) and inverted (right)

The choice of mass ordering may have an impact on the sign (positive or negative) of the mass differences. This has consequences on the oscillation probability, which includes a sine term (see the last term of Eq. 62) that is sensitive to a change of sign; it also has consequences on the sign of the matter effects that come into play when performing oscillation measurements through the Earth (the origin of these effects is explained in Chapter 6). In its current form however, the IceCube DeepCore detector is not sensitive enough to determine the mass ordering on its own, as evidenced by the results published in [35]. This property is therefore not relevant to this work. Normal ordering is assumed throughout this thesis.

5.2.2 CP Violation

The combination of the charge conjugation operator C and the parity operator P yield what is known as the CP transformation. A major breakthrough in particle physics was the discovery that this symmetry transformation was violated in the quark sector [36]. The violation of CP symmetry is often associated with the problem of baryon asymmetry of the Universe, which is the reason why matter and anti-matter in the Big Bang did not completely annihilate, leaving a possibility for us to exist and do physics.

In the PMNS matrix, CP violation in the neutrino sector is parametrized as a complex phase, δ_{CP} . The magnitude of that phase, which determines the scale of CP violation in the neutrino sector, can be observed by comparing the neutrino and anti-neutrino channels of an accelerator experiment [37] [38]. Since IceCube cannot distinguish between neutrinos and anti-neutrinos, this property is not relevant for the work described further. The following equation will therefore assume a CP phase of 0 to simplify derivations.

5.2.3 Majorana vs. Dirac

All known fermions must have a wave function that satisfies Dirac's equations:

$$(i\gamma^\mu \delta_\mu - m)\psi = 0, \quad (63)$$

This is done by defining ψ as a two-dimensional *spinor*, with both a left-handed and a right-handed component:

$$\psi = \begin{bmatrix} \psi_L \\ \psi_R \end{bmatrix}. \quad (64)$$

In 1937, Ettore Majorana proposed a similar equation that involved a particle field ψ and its charge conjugate ψ_C [5]:

$$i\gamma^\mu \delta_\mu \psi - m\psi_C = 0 \quad (65)$$

This equation holds like Dirac's equation is and only if the particle described by ψ is neutral, and its own antiparticle. This is known as the *Majorana neutrino* hypothesis, and it has several consequences on the representation of the particle in both the CC and NC interaction lagrangians, and the PMNS matrix. In the latter case, two additional phases are added such that

$$U_{PMNS}^{Majorana} = U_{PMNS}^{Dirac} \cdot \begin{bmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (66)$$

It can be proven, however, that oscillation experiments are not sensitive to these phases. The modified Majorana mixing matrix, $M_{PMNS, Majorana}$, is then:

$$U_{PMNS}^{Majorana} = U_{PMNS}^{Dirac} \cdot \begin{bmatrix} U_{e1}e^{i\alpha_1/2} & U_{e2}e^{i\alpha_2/2} & U_{e3} \\ U_{\mu1}e^{i\alpha_1/2} & U_{\mu2}e^{i\alpha_2/2} & U_{\mu3} \\ U_{\tau1}e^{i\alpha_1/2} & U_{\tau2}e^{i\alpha_2/2} & U_{\tau3} \end{bmatrix}. \quad (67)$$

Thus, for any quartic product of matrix element $U_{\alpha j}^* U_{\alpha k} U_{\beta j} U_{\beta k}^*$, the majorana phases will be cancelled out due to the presence of a conjugated pair. For this reason, oscillation experiments cannot be used to determine the nature of the neutrino, which is why this property of the neutrino is no relevant for this work.

5.2.4 Disappearance vs. Appearance Experiments

As can be seen from Eq. 62, the probability for a neutrino to transition between two flavour states depends on three aspects:

- The value of the PMNS matrix elements $U_{\alpha k}$.
- The value of the neutrino mass splittings Δm_{kj}^2 .
- The ratio of baseline-to-energy, L/E , which is experiment dependent.

The experimental observation of neutrino oscillation therefore depends on the energy range and flavour composition of the initial neutrino flux, and on the distance between production and detection. The phenomenology of oscillation can also be further categorized based on whether the final flavour state is the same as the original flavour (one then speaks of *disappearance experiment*) or if it is different (*appearance experiment*).

Figure 23 illustrates how a muon neutrino disappearance probability, $P_{\nu_\mu \rightarrow \nu_\mu}$, evolves as a function of an experiment's baseline, for three different neutrino energies in a scenario in which there is no CP violation. In this particular case, Eq. 62 is simplified by the use of the same flavour index: the quartic product $U_{\mu j}^* U_{\mu j} U_{\mu k} U_{\mu k}^* = |U_{\mu j}|^2 |U_{\mu k}|^2$ is a real quantity, so the sine component vanishes,

leaving:

$$P_{\nu_\mu \rightarrow \nu_\mu} = |U_{\mu 1}|^4 + |U_{\mu 2}|^4 + |U_{\mu 3}|^4 + 2|U_{\mu 1}|^2|U_{\mu 2}|^2 + 2|U_{\mu 1}|^3|U_{\mu 3}|^2 + 2|U_{\mu 2}|^2|U_{\mu 3}|^2 - 2 \left[U_{\mu 1}^2 U_{\mu 2}^2 \sin^2 \left(\frac{\Delta m_{12}^2 L}{4E} \right) + U_{\mu 1}^2 U_{\mu 3}^2 \sin^2 \left(\frac{\Delta m_{13}^2 L}{4E} \right) + U_{\mu 2}^2 U_{\mu 3}^2 \sin^2 \left(\frac{\Delta m_{23}^2 L}{4E} \right) \right]. \quad (68)$$




Figure 23: Survival probability (also known as disappearance probability) of muon neutrinos of three different energies, as a function of the baseline at which it is detected. Vertical lines show the baselines of three existing oscillation experiments. Note that unlike accelerator experiments, IceCube is sensitive to all baselines included in the shaded gray area.

Using once again unitarity constraints (specifically, Eq. 34), and noting that the terms in the first row of Eq. 68 can be re-written as:

$$\begin{aligned} & |U_{\mu 1}|^4 + |U_{\mu 2}|^4 + |U_{\mu 3}|^4 + 2|U_{\mu 1}|^2|U_{\mu 2}|^2 + 2|U_{\mu 1}|^3|U_{\mu 3}|^2 + 2|U_{\mu 2}|^2|U_{\mu 3}|^2 \\ &= (|U_{\mu 1}|^2 + |U_{\mu 2}|^2 + |U_{\mu 3}|^2)^2 \\ &= 1, \end{aligned} \quad (69)$$

the muon disappearance probability is therefore equal to:

$$P_{\nu_\mu \rightarrow \nu_\mu} = 1 - 2 \left[U_{\mu 1}^2 U_{\mu 2}^2 \sin^2 \left(\frac{\Delta m_{12}^2 L}{4E} \right) + U_{\mu 1}^2 U_{\mu 3}^2 \sin^2 \left(\frac{\Delta m_{13}^2 L}{4E} \right) + U_{\mu 2}^2 U_{\mu 3}^2 \sin^2 \left(\frac{\Delta m_{23}^2 L}{4E} \right) \right]. \quad (70)$$

As they only involve a single flavour of neutrino, muon disappearance experiments are only sensitive to mixing elements of the second row: $U_{\mu 1}$, $U_{\mu 3}$, and $U_{\mu 2}$. This is different from the neutrino appearance case, where the oscillation probability involves terms from two of the mixing matrix rows.

In the specific case of tau neutrino appearance, we have that²:

$$\begin{aligned}
P_{\nu_\mu \rightarrow \nu_\tau} = & |U_{\mu 1}^* U_{\tau 1}|^2 + |U_{\mu 2}^* U_{\tau 2}|^2 + |U_{\mu 3}^* U_{\tau 3}|^2 \\
& + 2U_{\mu 1}^* U_{\tau 1} U_{\mu 2} U_{\tau 2}^* + 2U_{\mu 1}^* U_{\tau 1} U_{\mu 3} U_{\tau 3}^* + 2U_{\mu 2}^* U_{\tau 2} U_{\mu 3} U_{\tau 3}^* \\
& - 4 \left[U_{\mu 1}^* U_{\tau 1} U_{\mu 2} U_{\tau 2}^* \sin^2 \left(\frac{\Delta m_{12}^2 L}{4E} \right) \right. \\
& + U_{\mu 1}^* U_{\tau 1} U_{\mu 3} U_{\tau 3}^* \sin^2 \left(\frac{\Delta m_{13}^2 L}{4E} \right) \\
& \left. + U_{\mu 2}^* U_{\tau 2} U_{\mu 3} U_{\tau 3}^* \sin^2 \left(\frac{\Delta m_{23}^2 L}{4E} \right) \right], \tag{71}
\end{aligned}$$

Where, as in the disappearance case, it is possible to use the unitarity angle constraints (this time, Eq. 35) to show that:

$$\begin{aligned}
& |U_{\mu 1}^* U_{\tau 1}|^2 + |U_{\mu 2}^* U_{\tau 2}|^2 + |U_{\mu 3}^* U_{\tau 3}|^2 + 2U_{\mu 1}^* U_{\tau 1} U_{\mu 2} U_{\tau 2}^* \\
& + 2U_{\mu 1}^* U_{\tau 1} U_{\mu 3} U_{\tau 3}^* + 2U_{\mu 2}^* U_{\tau 2} U_{\mu 3} U_{\tau 3}^* \\
& = (U_{\mu 1}^* U_{\tau 1} + U_{\mu 2}^* U_{\tau 2} + U_{\mu 3}^* U_{\tau 3})^2 = 0. \tag{72}
\end{aligned}$$

The tau appearance probability is therefore equal to:

$$\begin{aligned}
P_{\nu_\mu \rightarrow \nu_\tau} = & -4 \left[U_{\mu 1} U_{\tau 1} U_{\mu 2} U_{\tau 2} \sin^2 \left(\frac{\Delta m_{12}^2 L}{4E} \right) \right. \\
& + U_{\mu 1} U_{\tau 1} U_{\mu 3} U_{\tau 3} \sin^2 \left(\frac{\Delta m_{13}^2 L}{4E} \right) \\
& \left. + U_{\mu 2} U_{\tau 2} U_{\mu 3} U_{\tau 3} \sin^2 \left(\frac{\Delta m_{23}^2 L}{4E} \right) \right]. \tag{73}
\end{aligned}$$

The expression obtained in Eq.73 is slightly more complicated than in the disappearance case. Nevertheless, given the current measured values of the mass splittings and the baseline/energy ranges considered by an experiment like IceCube, it is possible to simplify the probability by neglecting the solar contribution (Δm_{12}^2), which leaves the following:

$$P_{\nu_\mu \rightarrow \nu_\tau} \approx -4U_{\mu 3} U_{\tau 3} \cdot \left[U_{\mu 1} U_{\tau 1} \sin^2 \left(\frac{\Delta m_{13}^2 L}{4E} \right) + U_{\mu 2} U_{\tau 2} \sin^2 \left(\frac{\Delta m_{23}^2 L}{4E} \right) \right]. \tag{74}$$

One can note that whereas a disappearance channel only probes mixing elements from a single row, appearance experiments are sensitive to at least two rows. This is of crucial importance for constraining the values of the PMNS matrix elements of the tau sector, for which no disappearance experiment is possible³.

5.2.5 Current experimental results

A wide range of neutrino appearance and disappearance experiments have contributed to the measurements of the neutrino mass splittings and the individual mixing elements of the PMNS matrix. These cover a wide range of L/E parameter space and use one of the following sources of neutrino

² Note that here again we assume that there is no CP violation

³ There is currently no experiment capable of measuring tau-neutrino disappearance.

as a starting flux:

- **Solar neutrinos:** These consist of MeV neutrinos produced in the fusion reactions occurring inside the Sun. These neutrinos are exclusively produced as ν_e 's, which oscillate across their propagation through the solar system. The Sudbury Neutrino Observatory (SNO)[39] and Borexino[40] experiments performed oscillation measurements using this neutrino flux.
- **Reactor Neutrinos:** These come from a man-made MeV anti- ν_e flux coming from the decay of uranium fission products in commercial nuclear reactors. Oscillation experiments in this regime measuring those are concerned with identifying the maxima of ν_e disappearance, which occurs relatively close to the emission source (10m to 1000m). KamLand[41], Daya Bay[42], RENO[43], and Double Chooz[44] are examples of reactor neutrino experiments.
- **Accelerator Neutrinos:** These consist of neutrinos produced in particle accelerators aimed at a fixed target: they are generated from the decay of mesons produced when high-energy protons collide with the heavy nuclei of the target. This flux of neutrinos can consist of ν_e 's and ν_μ 's, with the possibility to reverse the accelerator polarity and choosing instead antineutrino fluxes, to probe CP violation. Accelerator neutrino flux have energies of $\mathcal{O}(1 \text{ GeV})$, and are used in both short ($\mathcal{O}(100\text{m})$) and long ($\mathcal{O}(100\text{km})$) baseline experiments. MicroBooNE is an example of short baseline, while T2K[45] and NO ν A [46] are examples of long baseline experiments.
- **Atmospheric Neutrinos:** These consist of neutrinos originating from meson decay products in atmospheric cosmic-ray showers. While the flavour content of these neutrinos resembles a lot that of accelerator experiments, the energies involved are much higher ($\mathcal{O}(1\text{-}100 \text{ GeV})$) than other oscillation experiments. The set of baselines needed to observe the oscillation maximum is thus correspondingly larger ($\mathcal{O}(10000\text{km})$). Super-Kamiokande[47], ANTARES[8] and IceCube are both probing this type of neutrino flux.

To simplify the expression of oscillation probabilities in an experimental framework, the PMNS matrix is usually re-parametrized into three distinct matrices:

$$M_{PMNS} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & s_{23} \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (75)$$

Where : $c_{ij} = \cos(\theta_{ij})$, $s_{ij} = \sin(\theta_{ij})$.

In this representation, the free parameters of the PMNS matrix are re-arranged into three mixing angles (θ_{12} , θ_{23} , θ_{13}), and a complex phase quantifying the CP-violation in the neutrino sector (δ_{CP}). This separation of the PMNS matrix into three distinct matrices allows one to isolate the pertinent oscillation parameter for each of the types of neutrino flux measured: the first matrix is associated with atmospheric neutrinos, the second one to reactor experiments, and the third one to solar neutrino measurements. Table 24 lists the current best-fit values of all these parameters, as obtained from a global fit on the most recent neutrino oscillation data[48], for the case of normal mass ordering.

| | | Normal Ordering (best fit) | | Inverted Ordering ($\Delta\chi^2 = 2.7$) | |
|-----------------------------|---|---------------------------------|-------------------------------|--|-------------------------------|
| | | | | | |
| | | bfp $\pm 1\sigma$ | 3σ range | bfp $\pm 1\sigma$ | 3σ range |
| without SK atmospheric data | $\sin^2 \theta_{12}$ | $0.304^{+0.013}_{-0.012}$ | $0.269 \rightarrow 0.343$ | $0.304^{+0.013}_{-0.012}$ | $0.269 \rightarrow 0.343$ |
| | $\theta_{12}/^\circ$ | $33.44^{+0.78}_{-0.75}$ | $31.27 \rightarrow 35.86$ | $33.45^{+0.78}_{-0.75}$ | $31.27 \rightarrow 35.87$ |
| | $\sin^2 \theta_{23}$ | $0.570^{+0.018}_{-0.024}$ | $0.407 \rightarrow 0.618$ | $0.575^{+0.017}_{-0.021}$ | $0.411 \rightarrow 0.621$ |
| | $\theta_{23}/^\circ$ | $49.0^{+1.1}_{-1.4}$ | $39.6 \rightarrow 51.8$ | $49.3^{+1.0}_{-1.2}$ | $39.9 \rightarrow 52.0$ |
| | $\sin^2 \theta_{13}$ | $0.02221^{+0.00068}_{-0.00062}$ | $0.02034 \rightarrow 0.02430$ | $0.02240^{+0.00062}_{-0.00062}$ | $0.02053 \rightarrow 0.02436$ |
| | $\theta_{13}/^\circ$ | $8.57^{+0.13}_{-0.12}$ | $8.20 \rightarrow 8.97$ | $8.61^{+0.12}_{-0.12}$ | $8.24 \rightarrow 8.98$ |
| | $\delta_{CP}/^\circ$ | 195^{+51}_{-25} | $107 \rightarrow 403$ | 286^{+27}_{-32} | $192 \rightarrow 360$ |
| | $\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$ | $7.42^{+0.21}_{-0.20}$ | $6.82 \rightarrow 8.04$ | $7.42^{+0.21}_{-0.20}$ | $6.82 \rightarrow 8.04$ |
| | $\frac{\Delta m_{3\ell}^2}{10^{-3} \text{ eV}^2}$ | $+2.514^{+0.028}_{-0.027}$ | $+2.431 \rightarrow +2.598$ | $-2.497^{+0.028}_{-0.028}$ | $-2.583 \rightarrow -2.412$ |

Figure 24: Current Best-fit of all neutrino oscillation parameters. Source [48].

5.2.6 Constraints on unitarity

In standard oscillation experiments, measurements of the oscillation parameters (see Table 24) are calculated using the assumption that the PMNS matrix is unitary; however, the current precision of these experiments does not rule out the possibility that this assumption might be false. In a method described in [49], the consistency of the current oscillation parameters with PMNS unitarity is evaluated two ways.

First, the results of multiple types of experiments that are sensitive to θ_{12} , θ_{23} , and θ_{13} are overlaid Fig. 25. In this figure, the constraints from various experiments in the θ_{12}/θ_{13} and θ_{13}/θ_{23} planes are drawn on top of another; these will take varying shapes, based on the combinations of ν_e , ν_μ , and ν_τ disappearance / appearance channels being probed by any particular experiment. In a unitarity scenario, it should be expected that all contours should overlap in the parameter space. As can be seen by the black circles visible in both plots, this condition is satisfied for current measurements, which means that all are independently consistent with a unitarity scenario.




Figure 25: Current constraints from individual oscillation experiments on the PMNS matrix mixing angles. Constraints are shown in the θ_{13} , θ_{12} plane (a), and in the θ_{23} , θ_{13} plane (b). If the three-flavour oscillation paradigm holds and the PMNS matrix is unitary, there should exist a region in these plots where all measurements overlap. That region is plotted as a black circle. Figures taken from [49].

Another method for probing unitarity is to re-derive the oscillation probabilities for disappearance and appearance experiments, using an *agnostic* parametrization of the PMNS matrix:

$$M_{agnostic} = \begin{bmatrix} |U_{e1}| & |U_{e2}|e^{i\phi_{e2}} & |U_{e3}|e^{i\phi_{e3}} \\ |U_{\mu1}| & |U_{\mu2}| & |U_{\mu3}| \\ |U_{\tau1}| & |U_{\tau2}|e^{i\phi_{\tau2}} & |U_{\tau3}|e^{i\phi_{\tau3}} \end{bmatrix}. \quad (76)$$

In this treatment, the matrix is now described by 15 independent parameters (9 amplitudes $U_{\alpha j}$ and 4 phases ϕ_j) that are no longer subject to the constraints defined in Eqs. 34, 33, and 35. Similarly, the description of flavour states $|\nu_\alpha\rangle$ now need to be normalized by the sum of the row elements:

$$|\nu_\alpha(t)\rangle = \frac{1}{\sqrt{N_\alpha}} \sum_{i=1}^3 U_{\alpha i}^* e^{-iE_i t} |\nu_i(t)\rangle. \quad (77)$$

These changes have consequences on the analytical form of Eqs.68 and 71, which now incorporate additional normalization terms (which, unlike Eqs.34 and 33 are $\neq 1$) and unitarity triangle products (which, unlike Eq.35, are $\neq 0$). The likelihood of obtaining a given set of matrix elements under this agnostic framework can be compared with the likelihood of getting the same result in a standard unitary scenario. Figure 26 shows the results of such a comparison. It displays the constraints, in units of $\Delta\chi^2$, on the individual amplitudes of the matrix described in Eq.76: blue lines show these constraints given the currently observed oscillation values, whereas the red line show the expected constraints on unitarity one could achieve with a set of future experiments: T2HK[50], DUNE[51], JUNO[52], and a prospective new result from IceCube.




Figure 26: Constraints on the value of the individual PMNS mixing matrix elements, as derived in [49]. While elements of the first row are constrained to sub-percent precision, elements of the bottom row have much larger uncertainties. A new measurement by IceCube could particularly improve constraints on $U_{\tau 3}$.

For this work, the most important item from Fig.26 is the bottom right plot, indicating the current and future constraints on the mixing element $U_{\tau 3}$ of the oscillation matrix. As pointed out in 5.1.1 (particularly by Eq. 74), the tau appearance channel is the only one capable of probing this particular element, whose amplitude remains poorly constrained compared for example to the level of precision at which the CKM mixing elements are known ($O(0.1\%)$, see [53]), or to the mixing elements involving electron neutrino measurements ($O(0.1 - 1\%)$). The future IceCube measurement is here the most important contributor to the improvements that could be made in constraining that element in the future (the red curve on the plot), based on the findings in [49]. The only thing that is assumed about this future measurement is that it reaches a precision of order $O(10\%)$ on the tau neutrino normalization constant, which is precisely the order of magnitude reached by the sensitivity of this analysis (see Chapter 9).

Note that these conclusions still hold in the specific sub-case where unitarity is broken by the existence of a larger, $m \times n$ unitary matrix describing oscillations with additional sterile neutrino states,

which is a scenario that has been previously explored in [54]. It is also worth noting that relaxing unitarity assumptions also has an effect on the constraints of other elements of the PMNs matrix such as the CP phases, and the components of the unitarity triangles defined by Eq. 35. The constraints on these quantities from both current and future experiments have been calculated in [49] and in [55], but in these cases, the achievable gain from a new IceCube measurement is not significant.

NEUTRINO OSCILLATIONS WITH ICECUBE

Neutrino oscillation analyses in IceCube belong to the category of experiments probing *atmospheric neutrino oscillations*. Contrary to particle accelerator experiments, which rely on a fixed baseline and a very narrow energy spectrum, atmospheric neutrino experiments rely on a source of neutrino that is abundant, yet slightly more chaotic: the flux that is produced by the collision of *cosmic rays* with atoms of the upper atmosphere.

The following chapter describes the means of neutrino production, the physics of neutrino interactions, and how a typical atmospheric neutrino analysis is conducted.

6.1 GENERAL MEASUREMENT CONCEPT

The concept of an atmospheric neutrino experiment is depicted in Fig. 27a: From their production site in the atmosphere, neutrinos start traveling through the Earth. While they propagate, these neutrinos will oscillate mostly as they would in vacuum, following the probability expressions derived in Chapter 5. Upon reaching a detector, a fraction of these neutrinos will interact with the surrounding matter, producing a signal that can then be registered as an event. The frequency at which this interaction happens depends on the cross section of neutrino-nuclei interaction, which is described in Section 6.3.

A key feature of atmospheric neutrino interaction is the relatively high energies at which they can be detected ($O(1 - 100 \text{ GeV})$) compared to other sources of neutrinos. This means that the energy deposition in the detector is sufficient for IceCube to correctly reconstruct the original neutrino's direction, which is often expressed in terms of a zenith angle θ_{zenith} . This quantity can be understood as a proxy for determining the particular baseline of an event which, together with its energy, constitute the fundamental observables necessary for measuring oscillation. Since neutrinos are produced all around the Earth, it is therefore possible to map out neutrino oscillation probabilities into two-dimensional *oscillograms*, such as the one featured in Fig. 27b. For any given set of mixing angles and mass splittings, it is possible to measure, for instance, the disappearance of muon neutrinos by looking for a signal deficit in the shape of concentric arcs, in zenith and energy space.

Now, this ideal experimental picture naturally becomes more complex when considering the following limitations:

- A limited knowledge about the composition of the neutrino flux,
- Distorsions induced by the propagation of neutrinos in matter,
- A limited knowledge about the neutrino interaction cross section,
- Uncertainties related to the reconstructed energies, direction, and flavour of the detected events.

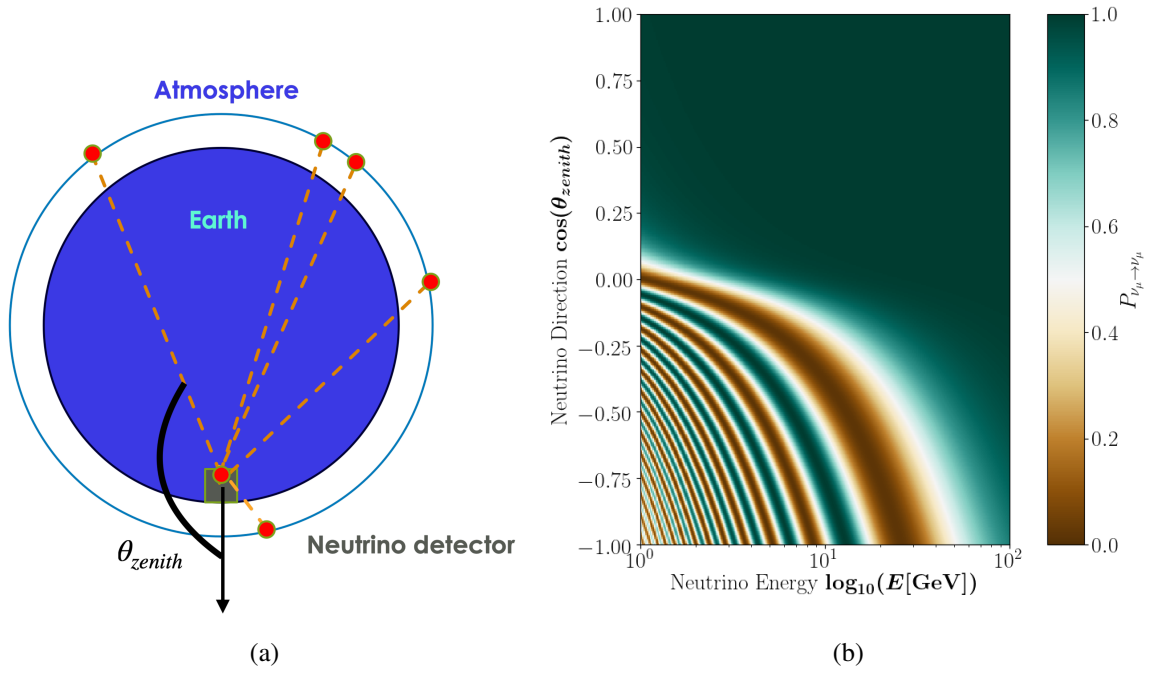


Figure 27: a) Analysis principle behind an atmospheric neutrino oscillation measurements. Neutrinos of various energy are produced in the atmosphere, and travel through the Earth to interact inside the detector. The distance a neutrino travelled is encoded in the event's zenith information, θ_{zenith} . b) *Oscillogram* depicting the probability for a muon neutrino to be detected as a muon neutrino in IceCube. That probability depends on both the mixing angle (θ_{23}) that controls the strength of the phenomena, and the mass splitting (Δm_{23}^2) which determines the location of the oscillation bands.

The following sections cover the first three limitations below, while the fourth one is described in detail in the event selection description from Chapter 7.

6.2 COSMIC RAYS AND THE ATMOSPHERIC NEUTRINO FLUX

Cosmic rays refer to a source of radiation observable on Earth whose intensity increases with altitude. This radiation comes from a flux of charged particles of extra-terrestrial origin, and spans an enormous range of energies ranging from MeVs all the way up EeVs[56]. The energy spectrum of this flux, shown in Fig. 28, has been studied by many experiments both on the ground as well as in the atmosphere, low-Earth orbit and beyond¹.




Figure 28: The energy spectrum of the cosmic-ray flux. Source: [56]

Upon reaching the Earth's atmosphere, cosmic-ray primaries interact inelastically with the nuclei of air molecules, triggering the breakdown of the initial nucleons into a particle shower of light and unstable mesons that will eventually decay into secondary products. Figure 29 shows the distribution of the most common mesons in the showers, as a function of altitude. These consist mostly of neutral (π^0) and charged (π^+ and π^-) pions, whose principal decay processes are:

$$\pi^0 \rightarrow \gamma + \gamma \quad (98.82\% \text{ of the time}), \quad (78)$$

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (99.99\% \text{ of the time}), \quad (79)$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \quad (99.99\% \text{ of the time}). \quad (80)$$

¹ Fun fact: Beyond includes exotic locations such as the surface of the Moon [57] and interstellar space [58].

Muons produced by these decays also contribute to the neutrino flux by decaying themselves into electrons:

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \quad (\approx 100\% \text{ of the time}). \quad (81)$$

For most of the energy range pertaining to oscillations (1-100 GeV), muon neutrinos are the dominant by-products of cosmic-ray interactions. As the energy of the primary particle increases, contributions to the neutrino flux from heavier kaons (the strange-quark equivalent of the pion meson) become important. The latter decay in the following fashion²:

$$K^+ \rightarrow \mu^+ + \nu_\mu \quad (63\% \text{ of the time}), \quad (82)$$

With a small fraction of kaons decaying through the mode:

$$K^+ \rightarrow \pi^0 + e^+ + \nu_e \quad (5\% \text{ of the time}). \quad (83)$$

That last decay is responsible for a non-negligible production of electron neutrinos in the initial flux. Finally, higher energies also allow for the production of charm (D) mesons, which have a sufficient rest mass to decay into a tau lepton. Therefore, tau neutrinos are also expected to contribute to the overall neutrino flux, via the main decay modes of the tau lepton:

$$\tau^- \rightarrow \nu_\tau + \pi^- + \pi^0 \quad (\approx 25\% \text{ of the time}), \quad (84)$$

$$\tau^- \rightarrow \nu_\tau + \mu^- + \bar{\nu}_\mu \quad (\approx 17\% \text{ of the time}), \quad (85)$$

$$\tau^- \rightarrow \nu_\tau + e^- + \bar{\nu}_e \quad (\approx 18\% \text{ of the time}), \quad (86)$$

$$\tau^- \rightarrow \nu_\tau + \pi^- \quad (\approx 11\% \text{ of the time}), \quad (87)$$

$$\tau^- \rightarrow \nu_\tau + \pi^- + 2\pi^0 \quad (\approx 9\% \text{ of the time}), \quad (88)$$

$$\tau^- \rightarrow \nu_\tau + \pi^- + \pi^- + \pi^+ \quad (\approx 9\% \text{ of the time}). \quad (89)$$

This contribution of tau neutrino to the atmospheric flux is known as the *prompt neutrino flux*, and its flux is several order of magnitudes smaller than the other two channels[59]; so much so, in fact, that it has not yet been measured. That contribution is therefore neglected in the atmospheric flux treatment of this analysis.

² K^- decay into the charged conjugate modes of Eqs. 82 and 83.

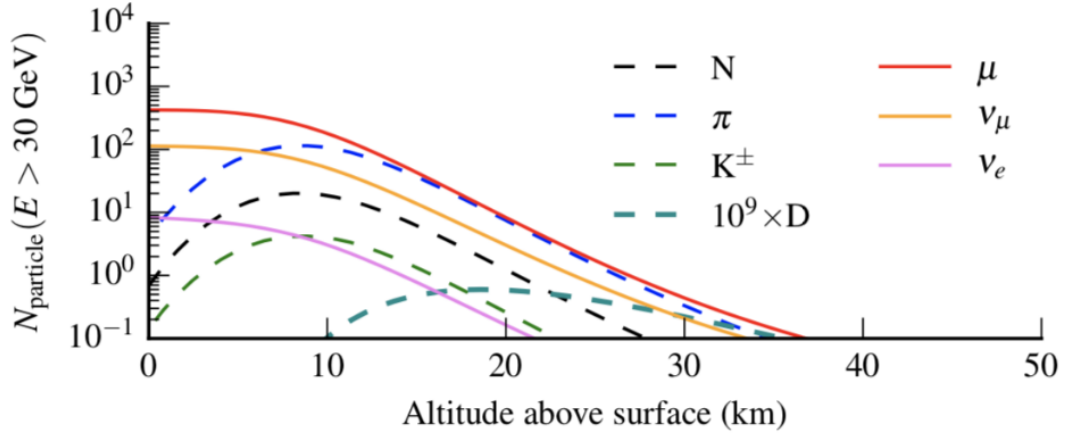


Figure 29: Breakdown of the column density of atmospheric shower by-products as a function of height. Pions and kaons are the main contributor to neutrino flux production mechanism. Source : [60]

The total flux of each neutrino flavour, shown in Fig. 30, is obtained by solving the cascade equations, which describe the change in particle fluxes across time, altitude and direction of observation. This flux can be affected at energies below 10 GeV by geomagnetic fields of the Earth, which change its azimuthal composition[60]. These computations are also subject to uncertainties in the hadronic interaction model used to model collisions. This has an impact on the outgoing flux of cosmic-ray secondaries, which is taken into account in the systematic uncertainty treatment described in Section 8.3.




Figure 30: Energy spectrum of the neutrino flux produced by atmospheric cosmic-ray showers, divided into individual flavour contributions. The values presented here correspond to expected flux values at the surface of IceCube, and were calculated using the MCEq software. Source: [60]

Finally, another important by-product of these atmospheric showers is the production of atmospheric *muons* at almost every step of the meson decay paths. Given their relativistic speeds and long decay time, these muons are still abundant on the Earth's surface, and can penetrate several kilometers into the ground before decaying. Muons constitute the main source of background in IceCube and are the target of multiple selection cuts featured in Chapter 7.

6.2.1 Neutrino propagation in matter

An experimental reality that is not taken into account by Eq.62 is the presence of matter (namely, the Earth³) along the propagation path of atmospheric neutrinos. The process of *coherent neutrino scattering* induces an additional perturbative potential in the Hamiltonian, the strength of which depends on the density of free electrons (N_e) the neutrino encounters:

$$V_{eff} = \sqrt{2}G_F N_e, \quad (90)$$

where V_{eff} is the effective potential and G_F is the Fermi coupling constant. The consequence of this matter potential depends on the properties of the medium and its associated density. In the Sun, where N_e is very large and changes smoothly from the core to the photosphere, the matter potential will change the effective mass eigenstates of the neutrinos propagating out to space, leading to a

³ also known as \oplus , home or *over here*.

matter effect known as the Mikheyev–Smirnov–Wolfenstein (MSW) effect[61].

On Earth, smaller and non-varying densities of electrons mean that no MSW effect is observed when atmospheric neutrinos propagate. However, additional *parametric resonances* occur at the boundaries of the Earth’s layers, where the electron density suddenly changes at the interface, the most important of which being the core/mantle boundary. Upon encounter of this wall-shaped potential, the propagation of neutrinos is shifted by a constant phase, inducing a *delay* in the particle’s transition amplitude[62][63]. Figure 31 shows the consequence of parametric resonances on the 2D oscillation probability observable by an atmospheric neutrino detector like IceCube.




Figure 31: Survival probability of a muon neutrino in the presence of matter, shown for the normal ordering. a) Parametric resonances due to the matter-induced potential introduces slight distortions in the oscillogram. b) Ratio of the muon neutrino survival probability in matter vs. a pure vacuum expectation.

As can be seen in Fig. 31, parametric resonances mostly affect lower energy events below 5 GeV, which is close to the DeepCore detection threshold. This means that this effect introduces mostly sub-dominant distortions to the oscillogram, which limits its impact on the overall tau appearance analysis.

6.3 NEUTRINO INTERACTIONS WITH MATTER

As was mentioned in Section 2.1, neutrinos are neutral fermions which interact with matter solely via the *charged current interaction* (with the Lagrangian defined by Eq.91) and the *neutral current interaction* (with the Lagrangian defined by 92).

$$-\frac{g}{2\sqrt{2}}\bar{\nu}_\alpha\gamma^\mu(1-\gamma^5)\alpha W_\mu \quad (91)$$

$$-\frac{g}{2\cos\theta_W}\bar{\nu}_\alpha\gamma^\mu(1-\gamma^5)\nu_\alpha Z_\mu \quad (92)$$

As it was also seen earlier, neutrinos and quarks can interact together, as they both carry electroweak quantum numbers. However, quarks cannot exist as single entities, meaning that neutrino-matter interactions must always involve the nucleus and underlying nucleons of an atom. As such, the proper modeling of the cross sections depends heavily on the elements of the structures that become affected by the scattering, which itself depends on the energies involved and, most importantly, on the Lorentz-invariant momentum transfer Q^2 that occurs between an incoming neutrino and the target nucleon. Depending of the neutrino's kinetic energy, interactions with an atomic nuclei can be further classified into three categories, each of them probing different realms of physics occurring within the nuclei themselves:

- Quasielastic Scattering (abbreviated QE),
- Resonant pion (and other meson) production (RES),
- Deep Inelastic Scattering (DIS).

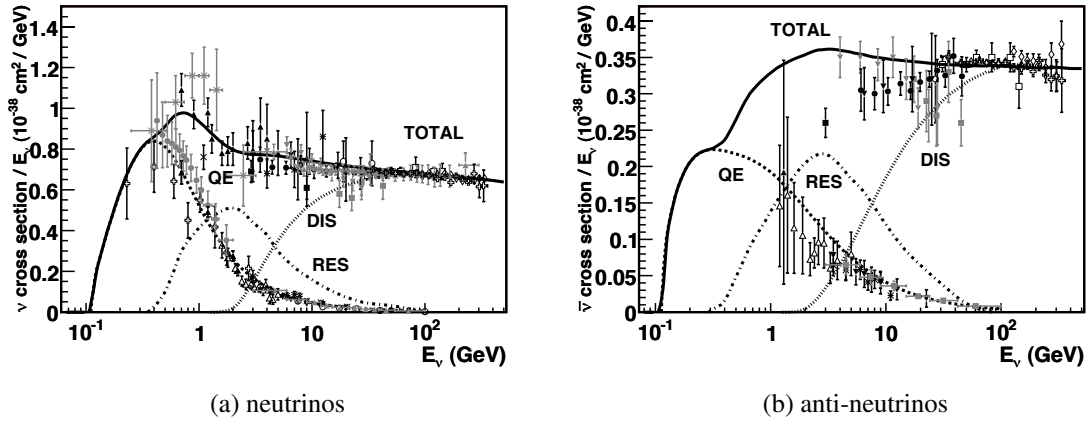


Figure 32: Contributions of the different neutrino-nucleus interactions to the total CC cross section of neutrinos (left) and anti-neutrinos (right). Source: [64]

Figure 32 depicts the contributions of these effects in the total cross section measurements for both neutrinos and anti-neutrinos.




Figure 33: Schematic depiction of the three kinds of neutrino-nuclei interactions. a) Quasielastic scattering b) Resonance production c) Deep Inelastic Scattering. Diagrams adapted from [65].

6.3.1 Quasielastic Scattering

Quasielastic scattering, depicted in Fig. 33a, dominates the neutrino interaction cross-section below 1 GeV. At this energy range, the neutrino deposits little energy compared to the rest mass of the nucleus, which can be treated as a single particle (neutron or proton) [64][66]. The QE cross section is usually expressed as a function of the system's invariant mass transfer, Q^2 as:

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2 M^2 |V_{ud}|^2}{8\pi E_\nu^2} \cdot \left[A \pm \frac{(4ME_\nu - Q^2 - m^2)}{M^2} B + \frac{(4ME_\nu - Q^2 - m^2)^2}{M^4} C \right], \quad (93)$$

In this equation, G_F is the Fermi coupling constant, M is the nucleon mass, m the outgoing lepton mass, V_{ud} is the CKM matrix element joining up and down quarks together and E_ν the incoming neutrino's energy. Meanwhile, the multiplicative factors A , B and C encompass additional combinations of scaling factors called *form factors*, which are empirically defined parameters used to model the distribution of charge within the target nucleon. For example, the factor B is defined by:

$$B = \frac{Q^2}{M} F_A (F_1 + F_2) \quad (94)$$

Where F_1 and F_2 are called the Dirac and Pauli *vector form factors*, and F_A is the *axial form factor*. Under an assumption called the *conserved vector current hypothesis* (CVC), the formers can be related to their electromagnetic counterparts, meaning that their values can be constrained using electron-nucleus scattering data [67]. The axial component, on the other hand, must be obtained from neutrino scattering data. This form factor is modeled as a dipole defined by:

$$F_A = \frac{g_A}{1 + M_{A, QE}^2 / Q^2}, \quad (95)$$

Where g_A is the axial coupling constant and M_A is an empirical parameter called the *axial mass*, which has a nominal value of 0.95 GeV in this analysis; the value comes from a combined analysis of several neutrino-nucleon target experiments[68].

6.3.2 Resonance Production

As the energy transferred to the nucleus approaches the rest mass of pion mesons, a resonant production of the latter can occur in the interaction: this is depicted by Fig. 33b. Once produced, the pion can either re-interact with the neighbouring constituents of the nucleus, or propagate coherently outside of it. The modeling of these interactions most commonly used was developed in [69]. It also contains an axial form factor which has the same shape as Eq.95, and the nominal value of its axial mass component, $M_{A, RES}$, is determined to be 1.12 GeV in [68].

6.3.3 Deep Inelastic Scattering

At energies above 10 GeV, the dominating neutrino interaction mechanism leads to a total breakdown of the target nucleon, a process called Deep Inelastic Scattering. In this regime, the neutrino interacts directly with individual quarks inside the nucleon, and leave behind a group of loosely defined hadronic states which can be modelled as *partons* (see the diagram in Fig. 33c). Modern techniques to calculate the DIS cross sections involve calculating the predicted parton density function from perturbative quantum chromo-dynamics (QCD) as a function of both the transfer momentum Q^2 and

a variable called Bjorken's x variable [70], which quantifies the momentum transfer between the original nucleon and the hit quark that produces the hardonic products[71].

6.3.4 Additional considerations for neutrino-nucleus interactions

Since the measurement and modeling of neutrino cross section is a very active domain of research, many models and computation methods exist to predict the total interaction cross section[56]. With the help of new accelerator measurements, several methods try to incorporate corrections that must be made to account for many effects not taken into account in the original derivation of the cross sections. Most of these are nuclear effects, and relate to the fact that nucleons interacting with the neutrinos are not free, especially in heavy target experiments. Another type of adjustment is the use of radiative correction to reduce uncertainties related to the absolute value of the weak coupling constant G_F . Finally, whereas most cross section measurements assume lepton universality, there exist a significant transition region where the mass of the tau lepton has a non-negligible impact in the kinematic suppression of the cross-section[72]. This could an uncertainty in the experimental treatment of the tau neutrino interaction.

6.4 EXPERIMENTAL FLAVOUR DIFFERENTIATION

In order to perform a muon neutrino disappearance measurement (and a tau neutrino appearance experiment), one must be able to distinguish the experimental signature left by ν_μ from that of the other neutrino species. Figure 34 shows the event morphologies left by all four types of neutrino interactions that can be detected in IceCube. All interactions leave a cascade-like imprint, which can be *hadronic* or *electromagnetic* in nature. That cascade can also be accompanied by a track in ν_μ -CC and some ν_τ -CC events generate a track-like morphology. Tau neutrinos can also generate so-called *double-cascade* events where a first cascade is initiated by the nuclei breakdown of the target, and a second one appears from the decay of the tau lepton generated in the CC interaction. Finally, part of the energy in an interaction can be carried away by the neutrino left intact in neutral current interaction, or by a neutrino emitted in the decay of the tau lepton (for ν_τ -CC interactions).




Figure 34: Event morphologies of the different kind of neutrino interactions: ν_e -CC (top left), ν_μ -CC (top right), ν_τ -CC (bottom left), and ν -NC (bottom right). The legend indicates the experimental signal associated with each symbol. Note that ν_τ -CC interactions can have different experimental signatures, depending on the fate of the outgoing tau lepton.

With the large and well-defined track they leave behind, muon neutrinos are the most easily recognizable events in IceCube, which is why they constitute the target "signal" population of events for

the training of a particle ID classifier algorithm in section 7.2.6. Meanwhile, due to both experimental and simulation⁴ constraints, all other neutrino interactions (ν_e -CC, ν_τ -CC and all NC events) cannot be properly distinguished from one another. For example, given the very short lifetime of the tau lepton (290×10^{-15} s, corresponding to a decay length of $87.03 \mu\text{m}$ at rest[56]), both cascades from tau-neutrino CC interactions are unresolvable by the current detector geometry. Also, the particular event hypothesis templates used to reconstruct cascade events do not make several key distinction between the hadronic and electromagnetic components of the light emitted by electron and tau neutrino interactions, which is further explained in section 7.2.4.

This brings about an important point regarding the nature of the tau-neutrino appearance in IceCube: unlike other experiments that can identify individual tau neutrino flavoured interactions, the appearance measurement in this work is a *statistical* one, as tau-neutrino fractions are solely determined by their expected global contribution to the 3-flavour oscillation expectation template.

6.5 CURRENT RESULTS FROM PREVIOUS EXPERIMENTS

After modeling the many intricacies related to the production, propagation, and detection of atmospheric neutrinos, one can obtain a simulation template describing the expected number of events in any set of PID, zenith, and energy bins. This prediction is influenced by the value of the neutrino oscillation parameters used to simulate the oscillation process. Using a minimization algorithm, it is therefore possible to tune the simulation expectations to best match the observed count in experimental data. For atmospheric neutrino experiments, such minimization is usually sensitive to two parameters:

- θ_{23} , the atmospheric mixing angle,
- Δm_{23}^2 , the mass splitting between the second and third mass eigenstates.

A typical muon disappearance search in IceCube will involve a fit of these two physics parameters, along with a series of systematic nuisance parameters that are described in Chapter 8. If one is interested in probing the unitarity of oscillations seen in IceCube, a tau appearance search can be made by including, in addition to the above mixing parameters, a tau normalization constant N_τ , whose main effect on the analysis is that it scales up or down the total contribution of tau neutrinos to the overall simulation template. This is the core idea behind the analysis presented in this thesis.

Figure 35 shows all measurements of the tau neutrino normalization performed to this day. Besides IceCube, two other experiments have had sensitivity to the measurement of the tau neutrino appearance channel. The first experiment to do so was the Oscillation Project with Emulsion-tRacking Apparatus, or OPERA [73], which was located at the Gran Sasso Underground laboratory in Italy. The goal of the experiment was to detect the appearance of tau neutrinos from an oscillated flux of muon neutrinos created by the proton beam of the Super Proton Synchrotron (SPS), which serves as the final injector to CERN's Large Hadron Collider. To do so, an array of mm-thick nuclear emulsion films was used to accurately identify the track signal left by the tau lepton produced from the neutrino's CC interaction. OPERA detected a total of five tau neutrinos, which was used to set a first and only direct measurement of the tau normalization factor.

⁴ The instances where a ν_τ decays into a muon for instance, are not tagged in simulation, preventing a proper characterization of this subclass of tau neutrino events.




Figure 35: Currently measured values of the tau neutrino normalization N_τ .

The second experiment to have performed a tau neutrino appearance measurement was Super-Kamiokande [74]. This megaton, water-cherenkov detector located in Japan is sensitive to atmospheric neutrinos, and used a neural-network approach to identify tau-neutrino candidate in 7.6 years of data. Lastly, the current best-measurement of the tau neutrino normalization has been performed by IceCube in [75]; this analysis, which used three years of IceCube data, pioneered many techniques that are still used in this analysis, particularly regarding the selection of events.

Two analyses are featured in [75]: a principal analysis (\mathcal{A}) with 2.8 years of data livetime, and a confirmatory analysis (\mathcal{B}) using a slightly smaller dataset, but a mostly similar treatment of systematic uncertainties; both resulted in values of N_τ that are consistent with each other, and with the 3-flavour neutrino hypothesis of 1.0 (to within 1-1.5 σ). It is to note that two fit results were reported: one in which the total contribution of tau-neutrinos (NC+CC) is weighted by N_τ , and one for which the normalization factor is only applied to the CC components of the tau neutrinos detected (CC only). While only the first result makes physical sense in the context of standard neutrino oscillation⁵, the CC result nevertheless allows for a more direct comparison to the results of OPERA and Super-K, which could only detect CC interactions in their respective experiment.

The next sections of this thesis describe an IceCube analysis that shares many similarities with the one performed in [75]. It differs mainly by the treatment of systematic uncertainties, and by the event selection used to isolate neutrinos in the data, which is presented in the next chapter.

⁵ No out-the-shelf particle physics scenario predicts a non-unitary result visible in CC, but not NC interactions.

EVENT SELECTION

The following sections give an overview of **oscNext**, the latest event selection that has been developed to perform neutrino oscillation physics analyses with data from DeepCore.

In IceCube, the trigger level has a seasonally averaged rate of 2.7 kHz, 99% of which consisting of background atmospheric muon events. Meanwhile, the high-level reconstruction algorithm used in the oscillation analyses takes an average of 1 minute to process a single event, making this the bottleneck of the whole simulation production flow. The purpose of these steps is therefore to reduce the number of background events being fed unnecessarily into this step, up to the point where the data flow becomes computationally manageable.

A typical event selection includes *cuts* using variables that have good feature discrimination between a desired signal, and an unwanted background. Such variables are often developed using a *nominal simulation set*, which constitutes the best-guess approximation of what the actual detector data will look like¹. In the case of oscNext, seven levels of cuts and processing have been used to obtain the detector data and simulation set that are used in the analysis fit. The following sections give an overview of each step involved in the selection process.

Note that any discussion of neutrinos in this chapter (including the figures) is meant to be representative of both neutrinos and anti-neutrinos, as the two types of particles leave an identical trace in the detector.

7.1 SIMULATION PRODUCTION

Reproducing particle interactions in IceCube requires the use of an *event generator* to simulate the initial interaction of the primary cosmic ray, and the number, type and kinematics of the resulting products. The latter then have to be *propagated* through the surrounding matter and their energy losses properly parametrized; the photons emitted from these losses must be propagated, until they either get absorbed by the ice or reach a DOM. Finally, the *detector response simulation* must accurately reproduce the efficiency at which the equipment is capable of converting these photons into waveforms, and the likelihood of triggering on a particular event.

Three classes of "particles" have been simulated for this event selection: neutrinos, atmospheric muons, and noise events. The neutrinos simulation sets contain an effective livetime of 70 years, to ensure enough statistics with respect to the livetime used in the data.

Neutrino-nucleus interactions described in section 6.3 are simulated using the GENIE event generator [76], a publicly available software that is widely used in the experimental neutrino physics

¹ Note that this represents the best guess *prior to performing a likelihood fit*, so the rates displayed in this section do not take systematic uncertainties into account.

community. GENIE simulates both the kinematics of said neutrino-nuclei interactions, and the nucleon decay processes resulting from the initial collision, over energies ranging from a few MeV up to 5 TeV. The GENIE software can be tuned to the needs of a particular experiment, and includes re-weighting options to calculate systematic uncertainties related to the various cross section measurements. For this particular analysis, a tuning package developed by the Km3Net collaboration [77] was used to extend the energy range of the neutrinos produced to 10 TeV.

GENIE requires as input a choice of interaction volume and initial neutrino flux; in this sample, the volume is chosen to be cylindrical, and the flux is set to be a single power law of index -2.0. To optimize computation time, the simulated volume of events in oscNext is chosen to vary as a function of the neutrino's initial energy. At lower energies, only events occurring in or very near DeepCore are generated, as anything occurring outside this region will likely not pass the DeepCore filter. At higher energies, the volume is extended to account for events that may not be fully contained inside DeepCore, but may still pass the coarse veto filtering algorithm. The neutrino cross sections used come from computations made by the GENIE collaboration; the DIS component is later re-weighted to the more recent CSMS model described in section 8.3.3.

Muons are produced using a bespoke software from IceCube called MuonGun[78]. As its name indicates, this algorithm "fires" muons across a user-defined cylinder volumes, and simulates the propagation of the muon through that volume using the PROPOSAL code. The simulated events can then be re-weighted according to a particular flux model (in this case, the H4a cosmic ray flux model [79]), to appropriately represent the probability of seeing a muon from a particular direction in the sky. The advantage of this software is that it is computationally much more efficient than full scale cosmic ray simulation, because all the muons simulated end up interacting in the detector. In addition, a kernel density estimator (KDE) technique is used to weight the generation volume and spectrum of muons[80]. This KDE weights the probability for muons to be generated in such a way as to enhance muon production in areas of the parameter space where muons were found in the final level of the previous event selection. This ensures that there will be enough background statistics in the areas of the detector where the rejection efficiency is low.

The products of the simulated interactions (outgoing lepton, mesons and hadronic showers from GENIE, or muons from MuonGun) are propagated through the surroundings using two different software: PROPOSAL [81] is used on ν_μ interactions to propagate the daughter muon particle, while GEANT4 [82] is used to propagate the products of hadronization. These two software determine the regions in space-time where photons are emitted, which serves as the input for the next step of the simulation chain: *photon propagation*.

Photons are propagated using a GPU-based algorithm named CLSIM [83]. Photons propagating in a medium can be either absorbed or scattered in another direction. In order to accurately propagate photons in IceCube, a model of the optical properties of the glacial ice is a critical component of the experiment. The model used in simulation relies on calibration data from LED flashers installed on the main board of each DOM, as well as constraints from ice core studies, and scattering measurements obtained in six of IceCube's holes during deployment [84]. In a calibration run, one or more LEDs of a DOM emit light flashes that are monitored by the neighboring DOMs. The arrival times of the photons that manage to reach the monitor is compared with the expected arrival times in simulation, where the bulk ice optical properties of the ice are fit in layers of 10m, using a likelihood minimization technique [85]. Optical scattering and absorption coefficients have additional dependences on photon wavelength, temperature and depth, which have previously been established and constrained by IceCube's pathfinding predecessor, AMANDA [86].




Figure 36: Absorption (red) and scattering (blue) properties of the ice as a function of depth, as incorporated in the ice model used for photon propagation. The DeepCore fiducial region includes DOMs located within depths of approximately -150 m to -500 m (in IceCube coordinates), below the main dust layer. The dots drawn on top of the plot show the z-position of the DOMs on the main DeepCore string, as well as on the central string #36 which has the larger, IceCube spacing configuration. Green dots denote normal quantum efficiency DOMs, whereas black DOM denote High Quantum Efficiency (HQE) DOMs. Red dots indicate DOMs that are faulty and thus not considered in the simulation.

Figure 36 illustrates the scattering and absorption properties of the bulk ice surrounding DeepCore, as defined by the ice model. This information is encoded in the depth-wise coefficients, which determines in CLSIM the probability for a photon to vanish in the ice, or change its direction. At the end of this stage, a photon surviving the propagation to reach the surface of a DOM is recorded, and will serve as an input to the detector response simulation.

Noise events are simulated using the Vuvuzela model [18]. This software code uses the information about the fitted noise parameter described in section 4.1 to produce simulated photoelectrons being generated in the PMT from radioactive decays scintillating in the glass. Noise is incorporated in two ways in the simulation. First, the Vuvuzela module is used in the neutrino and muon simulation chains to incorporate spurious noise pulses within physical events. Second, it is also used to generate noise-only simulation, since a lot of these events are capable of passing the basic trigger conditions of DeepCore. In order to best capture the features of triggerable noise events, simulated noise events must be generated over much longer time windows than the other simulated events (100 ms instead of $\approx 1\mu\text{s}$).

All three simulation types are processed through a processing pipeline that simulates every step of the detector's response. The experimental signal from the PMT is simulated as an experimental waveform that includes effects that were measured in laboratory settings, such as afterpulsing and jitter. The resulting waveform is sent through a simulation of the digitization process, which includes effects like electronic noise and the shape of the SPE distributions. After this point, both simulated events and detector data events are treated the same, using the filtering process described in section 3.2.2.

7.2 SELECTION LEVELS

Section 3.2.2 covered the general triggering and filtering scheme used by the entire IceCube collaboration to produce a low-level selection of events that are deemed to have a physics potential. In the following sections, the other levels of event selection and processing used for neutrino oscillation analyses will be described, starting with level 3. The objective associated with each of these levels is to reduce the number of events that need to be sent to computationally intensive reconstructions, while eliminating as many background events as possible. This process can be summarized graphically by Fig. 37, which plots the event rates of each type of simulated particle at each step of the selection.




Figure 37: Summary of the background removal procedure across all levels of the oscNext event selection.

7.2.1 Level 3: Low-level variables starting events

The selection process specific to the oscNext event selection begins with the Level 2 DeepCore filter, which was described in 3.3. Among the events satisfying this filter condition, a set of preliminary cuts on low-level variables has been made to efficiently remove noise and muon events. These cuts are applied mostly on the number and distributions of *hits* (defined as the number of triggered DOMs in an event) and *pulses* (defined as the reconstructed pulse entries in a DOM's pulse series): events with a lot of hits² concentrated in DeepCore will be kept, whereas events with a low number of hits, a distribution of hits significantly spread out in time or incorporating a lot of hits in the IceCube veto defined in section 3.3 will likely be cut out from the selection. The Level 3 cuts applied to the selection are the following:

- **NchCleaned**: The number of hits in the cleaned pulse series. Required to be > 6
- **NoiseEngine**: Output from an algorithm that projects all possible pairs of hits onto a sphere, rejecting events with very diffuse projections.
- **NAbove200Hits**: The number of pulses above a depth of -200 m, which roughly corresponds to the top of DeepCore. Required to be < 10 .
- **MicroCount hits**: The number of hits inside the busiest 300 ns time window of an event. The time window considered is restricted to a period covering $4 \mu\text{s}$ before and $5 \mu\text{s}$ after the trigger time. Required to be > 2 .
- **VertexGuess.z**: The depth at which the first clean pulse of an event occurs. Required to be < -120 m.
- **DCFiducialHits**: The number of hits inside the DeepCore fiducial volume. Required to be > 2 .

² This logic also applies to pulses.

- **Veto/Fiducial ratio:** The ratio of hits in the veto region w.r.t hits in the fiducial volume. Required to be < 1.5 .
- **CausalVetoHits:** Using the list of hits passing the causality condition of the DeepCore filter, determine how many are located within the veto region. Required to be < 7 .
- **C2HR6:** The fraction of DOMs that are triggered within 600 ns of the third clean pulse of the event (thus ignoring the first 2). Required to be > 0.37 .
- **CleanedTimeLength:** The time interval between the first and last pulse of the cleaned pulse series. Required to be $< 5 \mu\text{s}$.
- **UncleanedTimeLength:** Same as the previous quantity, but for the uncleaned pulse series. Required to be $< 13 \mu\text{s}$.
- **RTVeto:** The number of hits found by the radius-time (RT) veto algorithm. This number depends on the number of hits in the fiducial region.

All variables developed at Level 3 consists of re-optimized elements originally used in past oscillation analyses. One of the key difference with previous event samples is the absence of charge information in the selection variables. Because of potential mismodeling of the charge in reconstructed pulse (due to the pulse splitting issue mentioned in 3.2.3), it was found that hits were a more robust quantity to use as criteria for an event selection. This choice has the advantage of protecting the selection process from potential data-simulation disagreements.

Two additional cuts are performed on the event selection, to prevent mismodeling that was found when looking at data-Monte Carlo agreement in the previous version of the event selection. The first of these cuts is associated with rare events where certain DOMs would spontaneously flare for an extended period of time, producing an excess of extremely bright events in the selection [80]. The second cut was added after the discovery of a sub-population of events present in data with no counterpart in simulation, which had distinctly bad goodness-of-fit values in their final-level reconstructions. Figure 38 shows the location of this population in the reconstructed energy / reduced-log-likelihood (rLLH) plane. Further investigation on these particular events revealed that they came from long trigger windows from the Slow Particle (SLOP) filter, which is designed to look for slow-moving monopoles inside IceCube. Level 3 therefore includes a cut to remove events passing this particular filter, because it is not included in the conventional IceCube simulation.




Figure 38: Scatter plot showing simulation (blue) and data (black) events, based on the logarithm of their reconstructed energy and their reduced likelihood value (rLLH). This value corresponds to the absolute value of the likelihood obtained during the reconstruction fit, divided by the number of degrees of freedom. This variable can be thought of a weak proxy for the quality of an event’s reconstruction. A distinct population of data events with no simulated counterpart can be seen above the red line (the latter is only a visual guide).

7.2.2 Level 4: Noise and Muon classifiers

After an initial removal of background events with low-level cuts, the data / Monte Carlo agreement of the sample becomes decent enough to use simulation-driven machine learning tools to further reduce the background noise and muon events. Classifier algorithms such as Boosted Decision Trees (BDT)[87][88][89] have been used successfully in many particle physics experiments³. Given a large enough training sample, BDTs have the capacity to efficiently recognize complex relationships between a restricted set of input variables, and their categorization as signal or background.

Level 4 of the event selection consists of two separate BDT classification steps, each of them specifically designed to target either noise or muons events. Figure 39a and 39b show the performances of the noise classifier, which was trained on five input variables:

- **NchCleaned**: The number of hits in the cleaned pulse series (already used in Level 3).
- **STW_m3500p4000.DTW200**: Number of hits within a combination of a static and dynamic time windows. First, a static window of $[-3.5, +4.0] \mu\text{s}$ is centered around the event’s trigger time. Within this time interval, a dynamic time window of 200 ns in length is slid across, and positioned as to maximize the number of hits occurring within it. The number of hits within

³ See for example the following, non-exhaustive list of reference: [38], [90], [75] and [91]

that optimized time window becomes a good measure of the pulse clustering in the event, with a small clustering being typically associated with self-triggered noise events.

- **iLineFit.speed:** A low-level, track hypothesis is fitted to the event using the algorithm described in [92]. The resulting speed estimate of the particle is, in the case of a noise event, more often extreme (either very small or very large) than the other types of simulated events.
- **Fill ratio from mean:** The best discriminatory variable against noise. The fill ratio is defined by calculating the ratio of triggered to untriggered DOMs within a certain radius of a vertex hypothesis. For this variable, the radius is defined as the mean distance between all hits and the specified vertex. The latter is defined to be the first HLC hit occurring within the event.
- **FullTimelengthRatio:** Take the ratio of the event duration calculated using the cleaned and the uncleaned pulse series. The event duration is defined as the difference between the arrival time of the first and last pulse in the pulse series. For noise events, the uniform spread in pulse timings will generally yield higher ratios than physics events.

Figure 39a is an example of a classifier output score distribution, which shows the output score attributed by the BDT to the Level 3 simulation events. These events are given a true label (in this particular case, noise and neutrinos making up the signal), and the datasets are split into two samples. The first one is dedicated for *training*, where the algorithm learns to recognize and weight significant relationships between input variables that optimize classification. The second sample, meanwhile, is fed into the already trained BDT in a *testing* phase where one can verify that the features highlighted by the algorithm do not depend on very specific features from the events used in the training, a phenomenon called *overtraining*⁴. Overtraining can happen for instance when the provided training data is sparse: in these cases, statistical fluctuations between individual events might be wrongly perceived as defining features of a class.

As can be seen in the figure, the presence of sharp peaks at output values of 1.0 and 0.0 is a sign of good classifier performance, indicating that the trained algorithm has a very high separation power. Additionally, the fact that the training and testing sample distributions almost entirely overlap is a good indication that the BDT was not overtrained. Figure 39b shows another view of the algorithm's performance. In this plot, the rate of events surviving a cut on the classifier output score is shown in cyan for noise, and in red for neutrinos. This plot shows that the noise rate can drop significantly even for small values of classifier output, while retaining almost 100% of the signal neutrino events. Given these performances, a first Level 4 cut was devised by removing events with a score of less than 0.5.

⁴ This is also known as *overfitting* in the literature.




Figure 39: Performances of the Level 4 noise classifier. a) Normalized distribution of noise and neutrino events based on the output value returned by the classifier (also referred to as the neutrino probability). The presence of two sharp peaks at 0 and 1 are a sign that the BDT achieves good discrimination power between the two kinds of events. b) Rates of neutrino and noise events remaining in the selection, as a function of the threshold cut applied to the classifier's output. As the cut is moved toward higher neutrino probability values, the neutrino rate remains essentially the same, while the noise rate drops sharply by four orders of magnitudes. This shows that the trained classifier is capable of keeping a maximum amount of signal (high efficiency) while cutting away a maximum amount of background (high purity).

Following the cut on the noise classifier, the remainder of the sample was used to train another BDT, this time to remove muon events from the selection. The classifier was trained using nine input variables, which are listed below:

- **ICVetoHits**: The number of hits in the IceCube veto region defined at Level 3.
- **RTVeto250Hits**: The number of hits in the radius-time region defined at Level 3.
- **NAbove200Hits**: Number of pulses observed above a depth of -200 m, defined at Level 3.
- **L4_VICH_nch**: Number of Veto-Identified Causal Hits (VICH). The arrival time of every pulse is compared to the time of the pulse closest to the event trigger time. This variable counts the number of DOMs that have at least one pulse with a distance and timing outside an approximate causality region with respect to the trigger (see Fig. 105 in Appendix B for the region definition).
- **L4_accumulated_time**: Time at which 75% of the event's total charge has been detected.
- **L4_first_hlc_rho**: Radial distance between String 36 and the first hit satisfying an HLC condition.
- **SRTTWOOfflinePulsesDCHitStatistics.cog.z**: Depth of the center-of-gravity constructed from all hits in the cleaned pulse series.

- **SRTTWOOfflinePulsesDCHitStatistics.z_travel**: Range of depth covered by the hits of the cleaned pulse series.
- **SRTTWOOfflinePulsesDCHitStatistics.z_sigma**: Width of the distribution of all hits z-positions.

Contrary to the noise classifier, the Level 4 atmospheric muon BDT classifier was trained directly on detector data rather than simulation. This is done because the event statistics obtained from data, which remains dominated by atmospheric muons at 99%, is larger than what can be obtained from the available Muongun simulation, as the latter is optimized for the final level parameter space. In order to average over seasonal variations in the atmospheric muon flux in the training, a set of three full runs were taken at different periods of each detector season. Figure 40a shows the output score distribution for events tagged as muons in cyan and purple, and for the signal neutrinos in red and green. Once again, the trained algorithm is capable of achieving a very good event separation and little overtraining. The efficiency at which muons are eliminated while neutrinos are kept is illustrated in Figure 40b.

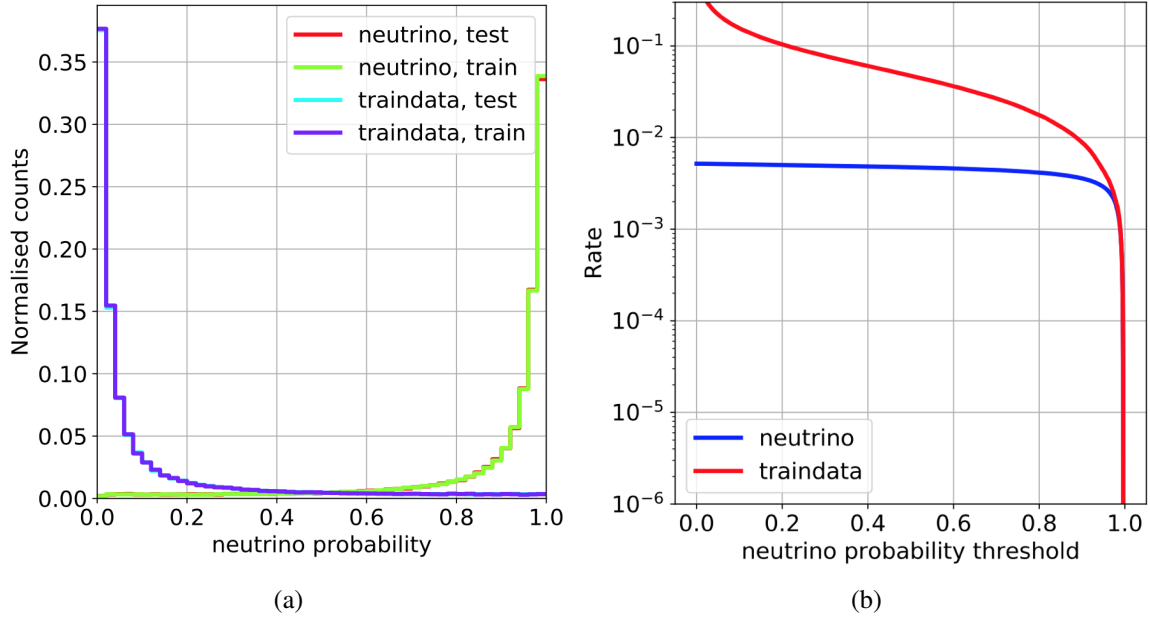


Figure 40: Performances of the Level 4 atmospheric muon classifier. a) Normalized distributions of neutrino and detector data events (used as a proxy for atmospheric muons) as a function of their classifier output score (also referred to as a neutrino probability). The distributions are once again well distinguishable, and there is good agreement between test and train samples. b) Rates of neutrinos and training data as a function of the choice of threshold used to the classifier's output.

7.2.3 Level 5: Muon corridor cuts

Following the previous event selection stage, almost all the noise events, and a majority of the muon events are eliminated from the sample. The remaining muons, however, are difficult to remove, because they usually consist of tracks passing through one of the *corridors* of IceCube: because of its regular grid shape and large spacing between strings, the standard IceCube array, which is used as a veto for DeepCore events, includes directionnal corridors of uninstrumented ice where muons can sneak through while still depositing light in DeepCore; examples of such corridors are identified in Fig. 41. To mitigate this problem, specific directionnal cuts have been made based on the observed morphology of the corridor muons in the event viewer. In the end, these cuts eliminate 97% of the




Figure 41: Example of muon corridors.

muons passing Level 4, at the expense of cutting away 48 % of the neutrinos.

Potential corridor events are identified by running a dedicated *corridor cut module*. This algorithm first defines the location of all pre-identified corridors; it then computes the event's center-of-gravity using the uncleaned pulse serie, and proceeds to find the DeepCore string located closest to that CoG. From there, every corridor that might be oriented toward that string is checked for hits occurring within a certain radius of 250 m and a time window of $[-1.0, 1.0] \mu s$ (this is similar to the way the RT cut is computed in Level 4).

Level 5 also introduces a low-level reconstruction that was used in early muon disappearance analyses of DeepCore [93]. These reconstructions only work for events that have a lot of direct light (ie very few pulses diverging from the ideal expectation of a single Cherenkov track), which usually corresponds to very up-going tracks. This reconstruction, called SANTA, can provide better directional information about the event than the lower level reconstructions run so far. Therefore, certain cuts at Level 5 are only applied if there is agreement between SANTA-reconstructed directions and those obtained from lower-level reconstruction. These cuts are listed below:

- **Radial Containment:** If any of the three following vertex guesses lie further than 150 m from String 36, the event is removed from the selection: the position of the first HLC hit (`first_hlc_rho`), the location of the brightness string, or the Level 3 vertex guess.
- **Vertical Containment:** Events lying outside the depth interval $[-499, 200]$ m are removed (according to the Level 3 vertex guess and the first HLC position)
- **Corridor Cut:** Remove events that have more than 2 hits along pre-identified directions that are known to correspond to corridors. This condition is lifted if the angular separation between the corridor and the SANTA reconstructed track is smaller than 0.7 rad.
- **Harsher classifier cuts:** the threshold for both the noise and muon classifier is set respectively to 0.85 and 0.9

At this point in the processing, the overall rates of events finally become manageable for the reconstruction software, which is described in the next section.

7.2.4 Level 6: Event Reconstruction

The sixth level of the event selection corresponds to the step at which the high-quality reconstruction, called RetroReco [94], is performed. Whereas previous reconstructions so far relied on simple geometric computations roughly approximating the shape of a track, this algorithm is based on *reconstruction tables*, which compare an event hypothesis with detailed templates describing the expected light imprint of the interaction in the detector. These templates are encoded in large, multi-dimensional tables built from simulating a particular light source a large number of times, at many locations and orientations in the array. Given a particular hypothesis, tables are interpolated to obtain a probability density function of photon arrival time at each DOM inside IceCube, which is then fed into the likelihood defined by the equation below [95]:

$$LLH_{reco} = - \sum_i^{N_{DOM}} \sum_j^{T_i} [N_{ij} \cdot \ln(\mu_{ij} - \mu_{ij} + \ln(\Gamma(N_{ij} + 1)))] . \quad (96)$$

The fundamental probability unit of this likelihood is the probability for observing N_{ij} photons in DOM i in time bin j , given a corresponding expectation μ_{ij} obtained from interpolating a photon table. That probability is assumed to be Poisson-distributed, and allows for non-integer photon numbers by using the gamma function instead of a pure factorial component. Summing over all time bins and all DOMs, a likelihood value is obtained, which is used to determine the optimal event hypothesis matching the data during minimization.

The photon expectation μ_{ij} for an event hypothesis occurring at space-time coordinate (x, y, z, t) and direction (θ, ϕ) are obtained using splines that were precomputed from the simulated tables. These have been generated with an assumption of azimuthal symmetry: a light source is simulated at position $(x = 0, y = 0, z, \theta, \phi, t = 0)$, and the resulting probability map is translated in x, y coordinates and in time, when the event hypothesis is off-axis. Separate tables are created for the two kinds of events detectable in the array: point-like light sources are used for cascades, while moving cherenkov light cones of constant intensity are used for tracks.

One of the defining features of RetroReco is the method by which tables are simulated: as the name suggests, the time axis of the simulation is reversed, such that photons are simulated as propagating out of the DOMs. The advantage of this technique can be seen in Fig. 42, which compares the basic principles of RetroReco with that of a traditional photon propagation code: whereas the latter is wasting a lot of computer resources by propagating a large quantity of photons that never end up being detected by the array, all photons simulated by RetroReco are used in the construction of the probability tables.

In order to fit the energy of an event, the expected photon counts μ_{ij} are initially derived in the photon tables from simulated electromagnetic cascades of 1 GeV, which are then scaled linearly up or down depending on the overall brightness of the event. If the hypothesis involves a Cherenkov-emitting track, the photon expectation of the latter is derived from a different set of tables created using small track segment, assumed to be minimum ionizing, instead of isotropic light source. The reconstructed length of the final track is then obtained from iteratively adding segments until the likelihood starts worsening. In order to account for the energy-dependent light emission in track-like particles, the final reconstructed energy of the track is interpolated from tables estimates in [96] using the reconstructed length as input.

An additional scaling must be applied to the reconstructed energy of the cascade, which is assumed to come from purely electromagnetic particle showers. In reality, cascade events from neutrino interactions are never purely electromagnetic (EM for short), as inelastic scattering breaks apart the




Figure 42: Schematic comparison of Retro Reco with respect to a traditional table method. In the traditional tables (upper figure), photons are generated from a light source and propagate outward, leading to wasteful GPU processing on all photons that fail to reach a DOM. In RetroReco (lower figure), tables are generated by propagating photons backward in time, starting at the DOM. This means that every photon simulated will be used to map out the likelihood space.

Table 3: Fit parameter used in the hadronic energy scaling

| Parameter | Value |
|-----------|------------|
| E_o | 0.18791678 |
| E_{th} | 0.2 |
| f_o | 0.30974123 |
| m | 0.16267529 |

nuclei, generating a hadronic component to the cascade. Since the hadrons produced are more massive and slower than the electrons and positrons created during a pure EM shower, less Cherenkov light is expected from hadronic cascades (HD for short), meaning that the table-derived energy E_{EM} is likely underestimating the true energy of the event.

To fix this, a correction factor is being applied based on GEANT4 simulation of EM and HD cascades in ice, performed in [97] and [98]. In these simulations, proton and electron primaries of equal energies were simulated to interact in the ice, and the total light yield of the cascade was taken to be the sum of the track length of all decay products with a velocity larger than the Cherenkov threshold. From that, a ratio of light yield was computed between EM and HD cascades at various energies, and the dependence on the true energy of the interaction was fit with the following:

$$F(E) = \frac{T_{HD}(E)}{T_{EM}(E)} = 1 + \left(\frac{E}{E_o} \right)^{-m} \cdot (f_o - 1). \quad (97)$$

Where $T_{HD}(E)$ and $T_{EM}(E)$ correspond to the respective Cherenkov track length (light yield) of the hadronic and electromagnetic cascade simulated with a primary of energy E . The functional form of this relationship has been heuristically derived in [99], and assumes that the scaling is due to a “leakage” of neutral pions from the hadronic channel: since the latter decay into photons, they generate EM cascades that contribute more and more to the total light yield of the shower, as you move up in energy. The parameters E_o , m , and f_o are empirical factors that are fit to the points obtained in the GENIE simulation.

The use of this factor was revisited in the present analysis. Table 3 lists the values of the fit parameters used in the hadronic scaling factor for oscNext. In essence, they are the same value as the ones used in analyses \mathcal{A} and \mathcal{B} of [75] (the values used in these analyses can be found in [100]), except for the value of a threshold parameter E_{th} . That parameter can be interpreted as a critical energy threshold below which the neutral pion leakage stops happening, and it is used in the processing script to apply a cutoff to the scaling at low energies. Given a lack of documentation justifying the choice that had been made earlier, it was decided to change that threshold to roughly correspond to the pion production threshold of 0.2 GeV. This choice was motivated by two reasons:

- Given that the simulation used to fit these values had a lower bound of 10 GeV, the constraints on the low-energy tail of the functional form are very weak.
- The previous value of E_o had the technical disadvantage of introducing a slight kink in the reconstructed energy distribution at ≈ 2 GeV, which could affect events close to the energy threshold of this analysis.

Lastly, a pre-cleaning of the pulse series was found to improve the quality of the reconstruction, as well as the data-Monte Carlo agreement of the fundamental charge and hit distributions. The steps of this pre-cleaning are the following:

- Events with less than 8 pulses in the detector are discarded,
- Events which fail to reconstruct are discarded,
- Reconstructions ending after very few iterations are discarded, as they fail to properly estimate the energy of the event.

The performance of the reconstruction can be seen in Figures 43 and 44, where the resolutions of the reconstructed energies and zenith angles are compared to their true values. These two variables are crucial to any oscillation analysis, as they are the experimental handles used to determine the baseline-to-energy ratio L/E that appears in oscillation probabilities.

From the energy resolution plots, it is possible to note that RetroReco performs best - meaning that the median reconstructed energy most closely match the 1:1 line - on ν_e CC and ν_μ CC events, as these constitute ideal cascade and track morphologies. In the case of ν_τ CC and neutral current interactions, energy estimation is impeded by the presence of a neutrino as outgoing product of the interaction. The fractional energy carried by this hidden participant, a quantity referred to as *inelasticity*, has a flux-averaged value of 0.5 in the energy region of interest (100 GeV)[101]. This means that a sizeable portion of the initial neutrino energy would leave no trace in the detector, leading to consistently underestimated energy values, as show in Fig. 43d. In ν_τ CC interaction, the energy carried away by the neutrino (from the decay of the tau lepton) accounts for a smaller fraction of the total energy budget, which results in a smaller bias than in the neutral current case.

In addition, biases in the energy reconstructions are visible at both ends of the energy spectrum. In the low energy range ($E_{True} < 5$ GeV), that bias comes from the fundamental limitations of DeepCore in its ability to resolve events, as only a limited number of DOMs are triggered. In the high energy range ($E_{True} > 300$ GeV), a bias occurs in track-like events (ν_μ CC and a fraction of ν_τ CC) from the fact that the track hypothesis used in RetroReco does not account for stochastic losses of energy along the muon's path. Finally, another bias is caused by the fact that the size of the light imprint associatd with these events is large, which means that a significant fraction of these are not contained within DeepCore.




Figure 43: Energy resolution of neutrino events reconstructed with RetroReco. The colored scale of each histogram denotes the unweighted number of simulated events that falls within a particular bin. The black solid line represents the median energy resolution, while the black dashed lines represent the $\pm 1\sigma$ spread of the distributions.




Figure 44: Zenith resolution of neutrino events reconstructed with RetroReco, which is defined as the difference $\theta_{reco} - \theta_{True}$. The colored scale of each histogram denotes the unweighted number of simulated events that falls within a particular bin. The black solid line represents the median resolution, while the black dashed lines represent the $\pm 1\sigma$ spread of the distributions.

7.2.5 Level 7: Analysis-level cuts

With reconstructed quantities that are now refined at Level 6, a final cleaning of the sample can be performed to remove the muons still present. Figure 45 shows such muon event. The event in question sneaks into DeepCore by passing through one of the corridors where IceCube DOMs have difficulty catching the light emitted by the muon track prior to reaching DeepCore. In these leftover backgrounds, even the Level 5 corridor cut fail to identify them.




Figure 45: Example of an atmospheric muon sneaking through one of the corridors in IceCube. Along these directions, the light emitted by the track is not seen by the outer IceCube DOMs, resulting in a background event in DeepCore.

Another BDT classifier was trained, where the input variables used are derived from the RetroReco event fits. Additionally, another iteration of the corridor cut module from Level 5 is run, this time using narrower settings for the corridor definitions: the radius is set to 75 m, and the time window restricted to $[-150, 250]$ ns. After several initial rounds of training, the final selected set of 5 input variables to the classifier are:

- **L4_MuonClassifier_Data_ProbNu**: The Level 4 muon classifier output, which still has some muon discrimination power.
- **L7_reconstructed_vertex_z**: The reconstructed vertex z-position from RetroReco.
- **L7_reconstructed_rho_36**: The reconstructed radial distance between the RetroReco vertex and String 36.
- **L7_NarrowCorridor_z_min**: Deepest hit recorded in the ice that matches the narrow corridor hypothesis.
- **L7_santa_retro_angles_diff**: Angle between the SANTA and RetroReco reconstructed directions.

The selection of these particular variables comes from several considerations. First, the variables needed to show good data-MC agreement at Level 6 to make sure that the event selection would behave as predicted on data. Second, a preference was given to choose input variables that showed little correlations between each other, such that an eventual problem in one of the input variable could

be isolated by removing that variable from the classifier. Finally, the selected inputs needed to have large values of *feature importance*, a quantity used to describe both the frequency at which an input is used to decide an event’s classification, and the strength at which the same feature has discrimination power over individual events. The feature importance of an input variable is estimated using the SHAP metrics [102], which applies concepts of game theory to establish the relative importance of each inputs. The classifier was trained on 60% of the baseline muongun set and 50% of the baseline neutrino setss using the `LightGBM` [89] algorithm (same as for the Level 4 classifiers). Table 10 (in Appendix B) lists the chosen hyperparameters used during the BDT training.

The performance of the Level 7 classifier is shown in Fig. 46a (output score distribution) and Fig. 46b. The latter is a third way of representing the effectiveness of the algorithm in removing background whilst keeping a maximal amount of signal event. The yellow line shows the rates of both muons and neutrinos as the choice of classifier output cut is moved from 0.0 to 1.0, a cut of 0.0 corresponding to the highest rate values on the plot. The blue star corresponds to the rates obtained on OscNext if one were to apply the GRECO⁵-style straight cuts on DeepCore containment only, while the red star denotes the signal and background rates obtained at final level on the GRECO sample (column “L7 LE” of [75]). As can be seen, the fact that both these points lie on the left and above the yellow line mean that the classifier saves more neutrinos for an equal choice of atmospheric muon rate, and removes more background events for an equal choice of neutrino rate. Finally, like the other classifiers trained in earlier levels, there appears to be little overtraining given the good overlap between the test and training sample curves.

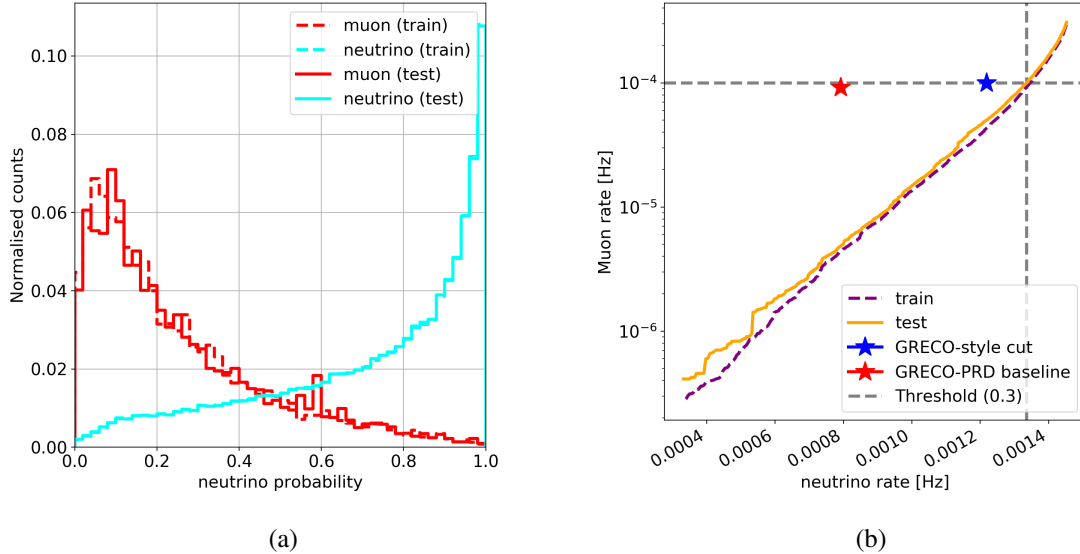


Figure 46: Performances of the level 7 Muon classifier. (a): Output value of the classifier for both neutrino and muon classes. (b): Signal and background rates as a function of the cut applied to the classifier score, for the training (purple dashed line) and test (yellow solid line) data samples. The two stars correspond to rates obtained by the previous GRECO event selection (red) in [75], and the rates one would obtain GRECO-like cuts to the oscNext event selection (blue).

For the tau neutrino appearance analysis, the choice of classifier cut must also take into account its effect on the rate of individual neutrino flavours. This can be illustrated by the plots in Figure 47 which shows, for both the training (left) and testing (right) samples, the rates of ν_e , ν_μ , and ν_τ as a function of the choice of cut. Section 8.2 details the choice of cut made to the Level7 output classi-

⁵ GRECO is the colloquial name used to described analysis \mathcal{A} of [75] internally within IceCube




Figure 47: Particle-wise rates as a function of the Level 7 classifier cut
Rates of all particle types as a function of the cut applied to the classifier score.

fier in the tau appearance analysis, a choice that was made to balance sensitivity with good data-MC agreement.

In addition to the cut on the Level 7 classifier, a closer scrutiny of the events led to the discovery of a small subpopulation of *coincident events*, characterized by the presence of two distinct light patterns (one from a neutrino and one from an atmospheric muon) at different locations in DeepCore. In order to eliminate these events from the selection, the following straight cuts were applied:

- **z.travel.top15:** The average z-direction of hit DOMs located in the 15 top-most layers of IceCube. Negative values are indicative of downgoing events, so these events are removed from the sample.
- **n.outter:** The number of hits in the outermost strings of IceCube. Events with more than 8 hits are removed from the sample.

7.2.6 Level 7: PID

At the last stage of Level 7, a particle identification (PID) classification is made based on the morphology of the detected events, to separate the sample in subsets with different neutrino flavour compositions. It is crucial that ν_μ events be isolated from the other types of neutrinos, because they dominate the composition of the full sample: any attempt at performing a ν_τ appearance analysis on all events at once would be virtually impossible, because the expected signal from ν_τ 's would be of the same order as the statistical fluctuations in the ν_μ signal. In previous event selections, PID was determined based on either the reconstructed length of the track fit, or the likelihood ratio between a track+cascade and cascade-only hypothesis [75]. Machine learning can be used here as well to refine this categorization, by feeding both these variables (and a few others) into a boosted decision tree. That BDT was trained to distinguish between tracks and non-tracks (i.e. cascades or badly reconstructed track) events. Its inputs are the following variables:

- **Track Length:** The reconstructed length of the track hypothesis. That length corresponds to the sum of the track segments added to the cascade hypothesis during minimization, in RetroReco. Longer tracks are mostly associated with ν_μ CC interactions.
- **Cascade Energy:** The reconstructed energy of the cascade part of the reconstruction hypothesis.
- **Zero dLLH:** The difference, in log-likelihood space, between the cascade-only and the cascade+track fits made by RetroReco. Lower values mean that the addition of the track component does not really change the goodness-of-fit, which is evidence that an event morphology is a cascade.
- **Zenith:** The reconstructed zenith angle of the particle's direction. This variable helps identify vertically upgoing events, which are usually better resolved events due to the tighter vertical spacing of the DOMs along strings. For this subclass of event, it is easier to distinguish tracks from the other types of events.
- **Zenith Uncertainty:** An estimate of the zenith uncertainty based on the minimization steps. Tracks have usually better pointing resolutions than cascades, which translates into smaller uncertainties.

The training of the BDT was performed using XGBoost[88], and the events were re-weighted to an unoscillated atmospheric flux to avoid biasing the classifier toward specific oscillation templates. Figure 48a shows the output score of the trained PID classifier, where a value of 1.0 is taken to represent a track-like event. Contrary to the previous BDT performance plots shown, the ν_μ CC distribution (in blue) is characterized by a very narrow peak around values of 1.0, but also a wide bump roughly centered around an output score of 0.5. The meaning of this structure is that while there is a significant population of events that can easily be recognized as tracks in the detector, there is also a lot of events for which the classification becomes ambiguous: a value of 0.5 in the classifier essentially means the events are indistinguishable.




Figure 48: Performances of the Level 7 PID classifier. a) Classifier output score for ν_μ CC events (blue), compared to the score of other neutrino types (red) and background events (muons and noise, in orange). Vertical lines show the boundaries of the chosen PID bins in the analysis. b) ROC curve for the PID classifier (shown in purple for the training set, and in yellow for the testing set), compared with the PID identifier used in analysis \mathcal{A} (the GRECO label) and in analysis \mathcal{B} (the DRAGON label) of [75].

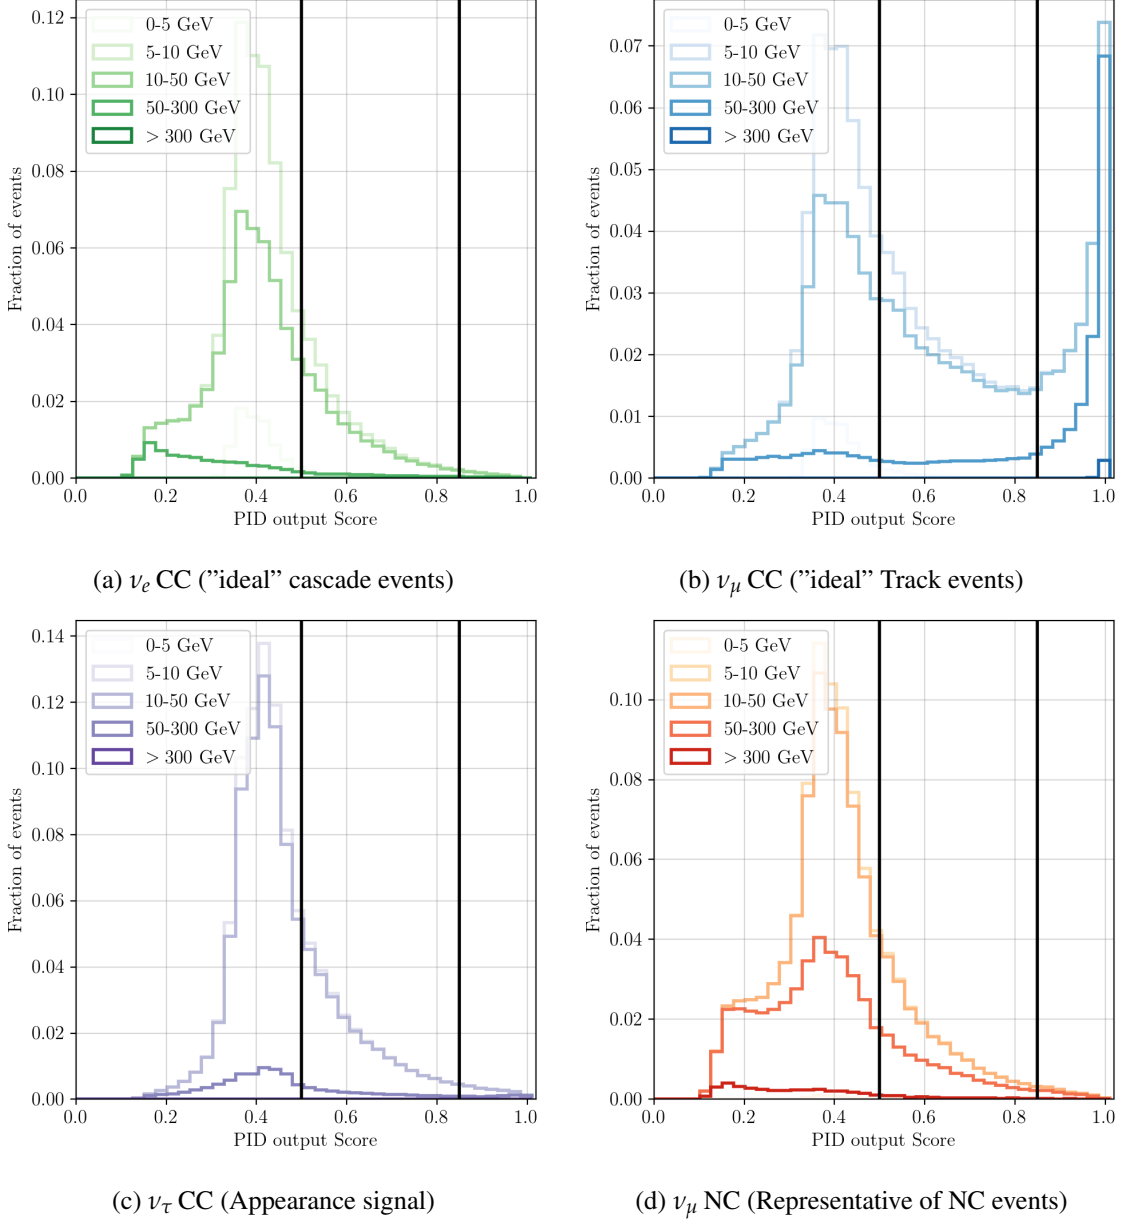


Figure 49: Selected PID distributions as a function of the true energy of the simulated neutrinos, showing how the algorithm performs on different types of events. A common feature of these plots is the improvement in classification power for higher energy events, as the latter tend to migrate toward the extreme scores of 1 (track-like) and 0. (cascade-like).

The performances of the PID classifier can further be explained by subdividing the neutrino events into their respective interaction type, as well as by their true energy, which is done in Fig. 49. As can be seen in these plot, part of the middle bump can be explained by the fact that low-energy events likely do not deposit enough light in DeepCore to be differentiated as tracks or cascades. At higher energies, on the other hand, events that constitute ideal cascades (neutral current and ν_e CC events), tend to be better classified with values slightly closer to 0.0 than the main confusion peak of 0.5. Figure 49c also shows that ν_τ CC events are difficult to classify, as they lead to multiple morphologies that do not completely resemble tracks nor cascades (see Fig.34).

Figure 48b displays the receiving operating characteristic (ROC) curve, which is similar to the curve represented in Fig. 46b, except that the background and signal rates are normalized over the

total data rate to express *efficiencies* at selecting signal and rejecting background. The larger is the area under the curve, the more efficient is the classification. In this plot, the PID classifier performances can be compared to the performances obtained using the previous two classification methods: a likelihood-ratio cut (which was used in analysis \mathcal{B}^6 of [75]) and a cut on track length (which was used in analysis \mathcal{A}). One can observe that the BDT only marginally improves upon the past techniques, and remain relatively close to the 1:1 line (which is the performance obtained by classifying cascades and track at random). The reason for this is that there exist fundamental limitations to the power of DeepCore in distinguishing between events, as its coarse and sparse geometry does not allow for a good characterization of events below ≈ 10 GeV. This limits the improvement in PID recognition that can be made with the current set of experimental observables, and is an area where a future experiment like the IceCube Upgrade could greatly improve oscillation studies.

7.3 DETECTOR DATA

Events passing the collaboration's list of online filters are sent via satellite and stored in a common database where they can be analyzed. The recorded events are saved in multiple files which cover a pre-defined interval of data-taking known as a *run*. The nominal duration of a run is 8 hours, but can be shorter if problems occurred during data taking. Each run is vetted by the operations team and fellow collaborators during monitoring shifts, to make sure that the data recorded is of good quality. Cuts are made on occasion to the usable livetime of a run, such that the data used in the physics analyses only contains data from periods of stable detector conditions. The detector data considered for use in this analysis includes everything recorded by IceCube since it began stable operation of its full, 86-string configuration. This corresponds to a time period extending from August 2011 to October 2020.

The flow of data is subject to changes in the online filter configuration that can have a sizeable impact on the event selection over time. These changes occur once a year, and define periods of uniform detector operation known as *seasons*. Therefore, before running the analysis on data, the properties of the event selection for every season must be evaluated to ensure that they do not vary significantly over time. Note that because the date at which the filter transition occurs can vary by a few months, a season does not generally cover a period of 1 year of detector operation. Thus, for the remainder of this section, references to a year number (for example "2016") will designate a detector season, and not a calendar year.

The first data consistency check is to compute the monthly rates accross all the runs used in the analysis, at final analysis level; this is shown in figure 50. Vertical lines have been included to distinguish between various moments in time. The black dashed lines correspond to the 1st of january of each year. The colored lines, meanwhile, represent the official start dates of a detector season.

⁶ Analysis \mathcal{B} is known colloquially as the DRAGON analysis in IceCube.




Figure 50: seasonal rates at final level

As can be seen in Fig. 50, the overall data rates accross time are stable, and statistically consistent with one another: A linear fit through the data points shown yield an average rate of 1.08 mHz, with a fitted slope that is consistent with 0.0. For reasons that will become apparent in the next paragraphs, data from the 2011 season was excluded from the analysis, as it showed a significant degree of incompatibility with the rest of the dataset. The data concerned by this exclusion is identified by a shaded region in Fig. 50.

The next check performed to ensure seasonal consistencies was to look at the 1D distribution of all three analysis variables (reconstructed energy, coszen and PID classifier output), in addition to a set of variables that were used in the lower levels of the sample selection. Fig 51 shows how these distributions look like for the case of the analysis variables. The bottom section of the plot shows the ratio between every season with respect to 2014 as a reference⁷.

⁷ Ratios with respect to each season have also been made, but they are not shown here.




Figure 51: 1D distributions of the detector data on a seasonal basis, in the three variables used in the analysis: energy (top left), $\cos(\theta_{zenith})$ (top right) and PID (bottom). The bottom section of each plot shows the ratio of expected event counts for each detector season with respect to the 2014 season (which was used to model detector properties in simulation). These counts have been re-scaled to take into account the different livetime of each season.

At first glance, all distributions look reasonably similar, and have flat ratios that seem to be close to 1.0. Given the small statistical error bars however, this check must be followed by a more quantitative description of the inter-season agreement. One method that was used was to look at the sum of the bin-wise chi-square defined by Eq. 98.

$$\chi_{ref,y}^2 = \sum_i \left(\frac{(N_{y,ref}^i - N_{ref}^i)^2}{N_{y,ref}^i + N_{ref}^i} \right). \quad (98)$$

$$\text{where} \quad N_{y,ref}^i = N_y^i \cdot L_{ref} / L_y$$

Where the χ^2 is summed over every bin i of the histogram, N^i refers to the bin count of either the reference (ref) or the compared season (y), and L represents the livetime of a season. The number of bins included in the sum (ie. the number of degree of freedom) is defined by excluding all empty bins of the reference histogram, and all bins of the comparison year with less than 10 events (to avoid biasing the chi2 with very low bin counts). This metric turned out to be sensitive to trends that affect the total rates of a season y compared to a reference ref . The comparison matrix on Fig. 52

shows the reduced chi2 value between every pair of season. This test was quite effective at detecting specific season outliers, as shown in Fig. 52. In this plot, 2011 stands out as having χ^2 values that are much worse than the other seasons. This gave the first indication that including this data in the analysis might be problematic.




Figure 52: Compatibility of individual seasons with one another, evaluated using the χ^2 defined in Eq. 98. The number and color in each square is a measure of the agreement between a reference and a comparison season, over the 1D distribution of the reconstructed energy of their respective events. The value of the chi2 shown has been divided by the number of degrees of freedom considered.

Another method used to compare seasonal distributions was the use of a 2-sample Kolmogorov-Smirnov (KS) test. This test statistic is more sensitive to changes in the shape of the distributions, which is useful for spotting disagreements which do not relate to the overall data rates. Since shaped-based disagreements cannot easily be accounted for without producing season-specific simulation sets⁸, the KS test was used as the principle mean of decision for seasonal consistency. The p-values from these 2-sample tests are shown in Fig. 53 for the same analysis variables.

⁸ An overall rate issue, meanwhile, could be accounted for by adding an overall normalization factor in the analysis




Figure 53: Compatibility of individual seasons with one another, evaluated using the KS-test metric. The number in each square represents the p-value from a 2 sample KS test performed on the cumulative distribution function of a pair of seasons. Large p-values indicate good agreement between the distribution shapes.

As can be seen from Fig. 53, the p-values obtained from these KS test are large, indicating that one cannot reject the hypothesis that the distributions from each pair of season are the same. An exception to this are the sections of the reconstructed energy plot associated with 2011, which show several p-values below 10% (namely, three p-values of 1% or less out of 45 independent comparison performed). This constitutes the second evidence justifying the exclusion of the 2011 season from the analysis.

7.4 CONTAINMENT CUT AND FINAL NOISE CLEANUP

After examining the reconstructed vertex distributions in both data and Monte Carlo, a noticeable disagreement was seen whereby events with their vertex reconstructed significantly outside the Deep-Core fiducial volume were occurring more frequently in data than in simulation. To eliminate this problem, the following straight cuts were applied on the selection:

- Reconstructed vertex z-position > -500 m
- Reconstructed vertex z-position < -200 m

Table 4: Event rates at Level 7

| | DC Filter [mHz] | Level 3 [mHz] | Level 4 [mHz] | Level 5 [mHz] | Level 6 [mHz] | Level 7 [mHz] |
|-------------------------|---------------------------|-------------------------|-------------------------|-------------------------|-------------------------|-------------------------|
| ν_e CC ⁹ | 1.61 | 0.95 | 0.84 | 0.48 | 0.43 | 0.26 |
| ν_μ CC | 6.16 | 3.77 | 3.11 | 1.39 | 1.31 | 0.71 |
| ν_τ CC | 0.193 | 0.129 | 0.119 | 0.071 | 0.066 | 0.049 |
| ν NC | 0.86 | 0.53 | 0.46 | 0.23 | 0.21 | 0.12 |
| Atm. μ | 7273 | 505 | 28.1 | 1.00 | 0.97 | 0.048 |
| Noise | 6621 | 36.6 | 0.28 | 0.07 | 0.06 | 0.007 |
| Total MC | 13903 | 547 | 32.9 | 3.16 | 3.05 | 1.19 |
| Data | 16168 | 582 | 31.6 | 3.15 | 2.93 | 1.08 |

- Reconstructed vertex $\rho_{36} < 300$ m

Additionally, harsher cuts on the Level7 muon classifier, and on the Level4 noise classifier were applied in order to reduce the noise rates to negligible levels, and to further reduce muon contamination in the sample. These cuts are:

- Level7 muon classifier output > 0.4
- Level4 noise classifier output > 0.95

Note that these cuts were not developed in the chronological order followed in this thesis. The issues that required the cuts were discovered after performing the binning optimization of Section 8.2, hence the difference in classifier threshold cuts.

Table 4 summarizes the process of event selection in oscNext by showing the rates of various particle types across each level of cuts described in this chapter. This table includes the final Level 7 rates after applying the coincidence, containment and noise reduction cuts, as well as by performing cuts associated with the choice of binning that will be described in Chapter 8. Despite the fact that quality cuts applied late in the selection led to a significant drop in neutrino retention efficiency, that efficiency is slightly higher than the previous IceCube event selection, leading to overall larger event statistics than what was expected of purely extrapolating the longer livetime of detector data available.

⁹ Rates for neutrinos include anti-neutrino contributions.

TAU NEUTRINO APPEARANCE ANALYSIS

This chapter describes how the analysis concepts introduced in Chapter 6 are implemented in practice on the `oscNext` event selection to perform a tau neutrino appearance analysis. As it was mentioned there, the goal of this analysis is to assess the compatibility of the IceCube DeepCore data to the expectation from a standard three-flavour neutrino oscillation theory. This is done by fitting a tau normalization factor N_τ to a template simulating the expected event count obtained in data, all the while accounting for systematic uncertainties that have a measurable impact on the expected event count. The template is three dimensional, and each dimension provides information about an observable of atmospheric neutrino oscillations: reconstructed energy estimates the true neutrino energy E , reconstructed $\cos(\theta_{zenith})$ is a proxy for the oscillation baseline L , and the PID label offers a basic categorization of neutrino flavours.

8.1 THE PISA FRAMEWORK

The fit to data is performed with the PISA software[103], which follows a modular approach to perform template fits on data while incorporating re-weighting schemes for the treatment of systematics. The core principle of this framework is to define *event pipelines* that process and reweight MC events in a succession of *stages*, each responsible for implementing part of the physics that impacts a given particle type, from production to propagation and detection. A separate pipeline is defined for each group of particles that share a common set of reweighting stages: in the context of this analysis, this means that one pipeline is used for all flavors of neutrino and anti-neutrinos, whereas atmospheric muons are treated by a second pipeline. Given their very small contribution to the event sample at final level, noise events are not included in the analysis¹.

Figure 54 schematizes the steps that make up a typical pipeline: each stage module is responsible for changing the distribution of events obtained from the `oscNext` selection according to a particular set of fit parameters. The flux stage, for example, is responsible for reweighting the event distribution according to a neutrino flux model, for which the spectral index $\Delta\gamma_{sys}$ might vary from the nominal value $\Delta\gamma_0$ during a fit. The red distribution in Fig. 54 illustrates how such a change could impact the event distribution. Similarly the oscillation stage, which is responsible for reweighting neutrino events based on their oscillation probabilities, will yield slightly different distribution depending on the mass splitting chosen; this is illustrated by the green histogram in the same figure. It should be noted that a reweighting stage can be applied on dimensions that differ from the reconstructed dimensions of the final template; the number of dimensions involved in the reweighting also vary from stage to stage.

¹ While noise-only events are negligible, the potential effect of noise pulses within physics event has been investigated further in Section 8.3.6

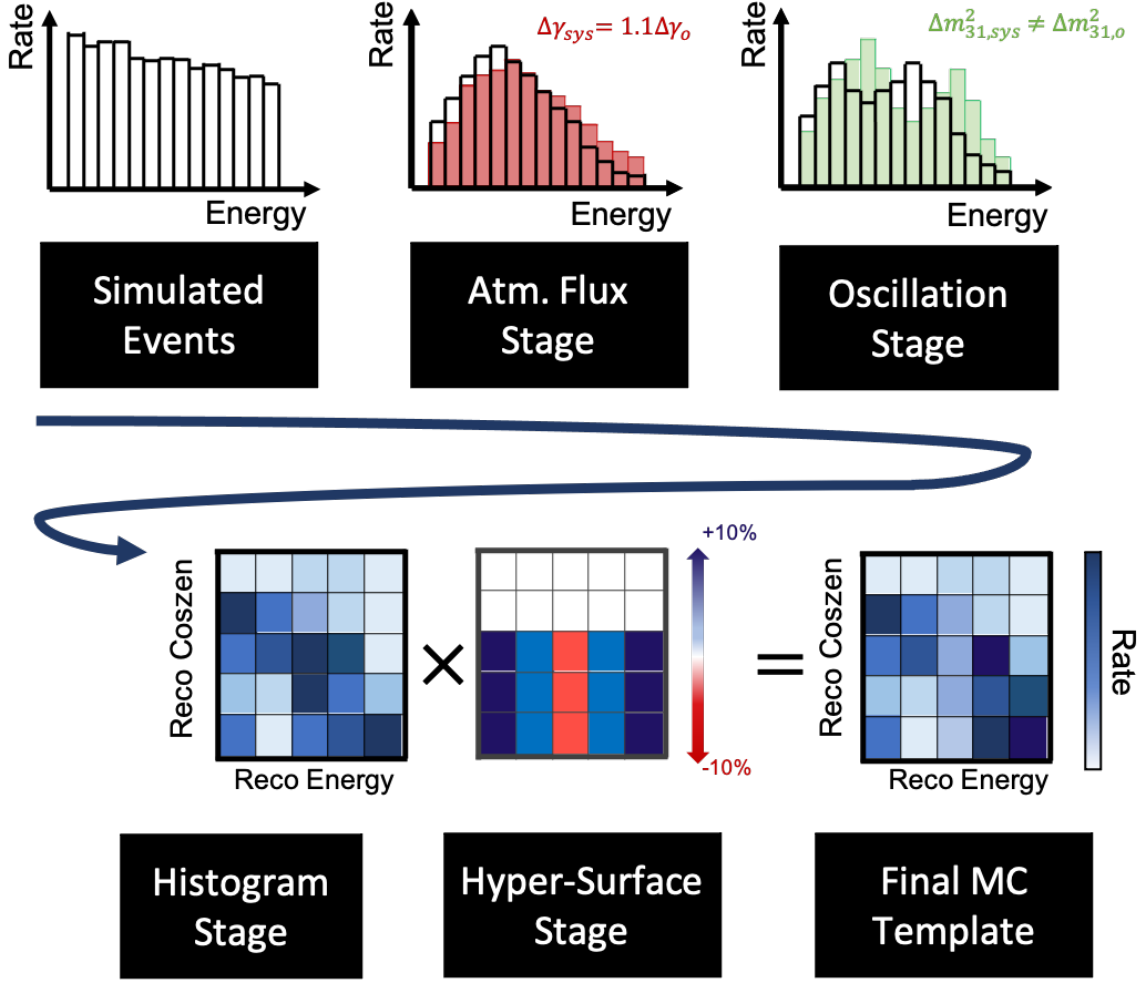


Figure 54: Schematic depiction of a neutrino pipeline in PISA. Starting with the oscNext events passing the Level 7 criteria, the simulated neutrino events are successively reweighted by each stage.

In addition to event-by-event reweighting functionalities, PISA also allows binned reweighting of the MC template, which is illustrated in the bottom row of Fig. 54. With the help of *hypersurfaces*, described in more details further down, the histogrammed simulation events can be smoothly reweighted without unphysical corrections by a template with the same binning dimensions, such as the one illustrated in the middle of the bottom row of the same figure. All parameters used in all stages, no matter how they are applied, can therefore be varied on a continuous scale in the course of an optimization process.

8.2 ANALYSIS BINNING AND CUT OPTIMIZATION

The particular choice of binning must be chosen to balance the need for high per-bin statistics with the need to have a high sensitivity to the predicted oscillation signal. An increased sensitivity can be obtained by using a finner binning to better capture the shape of the oscillation signal², but it can also be improved by modifying the relative contributions of various neutrino species in the PID categories used in the third dimension of the template. Guided by previous experience from past analyses, a choice of 12 energy bins, and 10 zenith bins was deemed appropriate given the expected statistics of

² This is valid up to the resolution limit of the reconstruction algorithm.

the event sample; meanwhile, the total number of PID bins was subject to a more systematic study, as will be shown below.

Several practical requirements were also integrated in the optimization process, particularly regarding the upper and lower limits of the binning. The first of these had to do with a more conservative treatment of the down-going region in $\cos(\theta_{zenith})$ space, to avoid a potential mismodeling of the *atmospheric self-veto effect*. This effect describes cases where the muon and the neutrino of a single shower both interact in the detector, creating an event that would have a much higher chance of being vetoed by the selection process. Given that muons and neutrinos were simulated separately, the impact of this effect could not be properly assessed with the available simulation, which meant that there was a possibility that the predicted rate of neutrino in the most down-going region was overestimated. To deal with this issue, three scenarios of binning cuts in zenith were tested:

- **Scenario A:** Full sky (no cuts),
- **Scenario B:** Mild down-going event cut (keep $\cos \theta_z < 0.6$),
- **Scenario C:** Harsh down-going event cut (keep $\cos \theta_z < 0.3$).

In addition to zenith angle cuts, three choices of cuts on the Level 7 muon classifier output variable were tested:

- **Scenario 1:** Full statistical power (keep classifier output score > 0.3),
- **Scenario 2:** Medium statistical power (keep classifier output score > 0.6),
- **Scenario 3:** Low statistical power (maximum purity) (keep classifier output score > 0.9).

The different scenarios of muon classifier cuts were made to explore the cost/benefits of including more events, and thus a higher number of background events into the analysis. All combinations of zenith and classifier cut scenarios were combined together, and evaluated using the outcome of a sensitivity test on the MC template (see section 9.1.2 for a description). The outcome of this optimization study is illustrated in Fig. 55, where the width of the likelihood scan at the 68% and 90% confidence levels are projected on the lower part of the plot. Whereas Scenario A1 gave the smallest confidence interval of all the scenarios tested, a preference for harsher zenith cuts (still to limit exposure to self-veto issues) led to the choice of scenario B1 as the most optimal zenith cutoff.




Figure 55: Sensitivity to the tau normalization for multiple choices of cuts in the final event selection. When testing for zenith angle and classifier cut configurations, a preference is seen for a less restrictive muon rejection, and a full-sky analysis (A1). Scenario B1 was chosen in the end, to mitigate the potential issues arising from unsimulated self-veto effects in the oscNext simulation files.

The second practical requirement of the binning analysis was that it should be common to both the muon disappearance analysis and the tau neutrino appearance analysis, to facilitate comparisons and cross-checks between the two. A third optimization was performed on the disappearance study, whereby four different PID binning configurations were explored:

- **Scenario 2:** Two PID bins only (boundaries: [0, 0.5, 1.0])
- **Scenario 3a:** Three PID bins option A (boundaries: [0., 0.35, 0.5, 1.0])
- **Scenario 3b:** Three PID bins option B (boundaries: [0., 0.5, 0.85, 1.0])
- **Scenario 4:** Four PID bins (boundaries: [0., 0.35, 0.5, 0.85, 1.0])

To ensure that the choice in PID binning of the disappearance analysis did not significantly hinder the performance of the appearance analysis, a sensitivity test was ran for each of these configurations, the outcome of which is shown in Fig. 56. This test showed that unlike the first optimization process, the choice of PID binning only marginally impacted the ν_τ appearance analysis, which meant that any configuration favored by the ν_μ disappearance analysis (scenario 3a) could also be used in this work.

Table 5: Final Analysis Binning

| Variable | Lower Bound | Upper Bound | # of bins |
|---|--------------|----------------|-----------|
| Coszen | -1. | 0.3 | 10 |
| Log10(E) | 0.69 (5 GeV) | 2.48 (300 GeV) | 12 |
| PID Bin Definitions (based on Classifier output) | | | |
| Cascade | [0., 0.5) | | |
| Mixed | [0.5, 0.85) | | |
| Tracks | [0.85, 1.0] | | |




Figure 56: Sensitivity to the tau normalization for multiple PID binning configuration. It was showed that the analysis was not too sensitive to these four choices of configurations, which meant that the final chosen PID binning (scenario 3a) could be made to optimize the muon disappearance analysis.

To complete the analysis binning optimization, the reconstructed energy range considered was chosen to include events solely within the range [5., 300] GeV. The lower bound was chosen to reflect minimal energy resolution of the detector and reconstruction algorithm, while the upper limit was chosen to limit to contribution of misreconstructed high-energy events to the selection. These limits completed the final choice of analysis binning, which is summarized in Table 5.

Figure 57 presents the MC event distribution in the chosen analysis binning. Note that the event counts showed in this figure correspond to the *nominal* MC expectation. In reality, the fit of several systematic uncertainty parameters is expected to change the final event counts in each bin, after the fit to data.




Figure 57: Distribution of events in the final analysis template, in the nominal simulation set.

8.3 SYSTEMATIC UNCERTAINTIES

In this analysis, systematic uncertainties are treated as nuisance parameters that are allowed to vary freely in the course of a fit to data. The validity of the physics results is conditional on obtaining reasonable values for these parameters, which are included to absorb the effect of known systematic uncertainties.

It is important to note that nuisance parameters do not quantify the real scale of a particular systematic uncertainty in IceCube. While a strong pull of a parameter in a fit might indicate that the analysis is indeed affected by the systematic studied, the absence of a pull in a parameter does not necessarily mean that the value chosen reflects the true scale of the effect. This must be determined independently using dedicated calibration studies.

8.3.1 Neutrino flux systematics

Systematics Included:

- **Barr Parameters (11 parameters)**
- π^+ to π^- ratio
- **Neutrino flux spectral index, $\Delta\gamma_\nu$**

The flux of neutrinos produced at the South Pole by cosmic rays is interpolated, as a function of energy, zenith and azimuth, from a tabulated model calculated in [104] (hereby referred to as the "Honda" model). The state-of-the art characterization of the uncertainties inherent to the atmospheric neutrino flux comes from [105]. Components of the flux are subdivided into categories that depend on the type of mesons responsible for production, and the variable x_{lab} , defined as the ratio between the energy carried by secondary shower particles to the energy of the initial cosmic-ray primary. Figure 58 provides a definition of the regions of parameter space that have dedicated parameters quantifying the uncertainty on the neutrino flux. Each parameter highlighted in green corresponds to a Barr parameter that was considered as a nuisance parameter in the analysis (the other ones were fixed at their nominal, after it was shown that they didn't not impact the analysis significantly).




Figure 58: Definition of the parameter space associated with each Barr parameter, along with their respective uncertainties. Parameters which are fit freely in the analysis are highlighted in green, while the other parameters were fixed at their nominal value. This figure has been modified from [105].

Proper calculation of the atmospheric neutrino flux is performed using the Matrix Cascade Equation software (MCEq)[106], which solves the one-dimensional evolution of particle densities in the development of atmospheric showers. Given a primary cosmic ray model (in this case, the GlobalSplineFitBeta model from [107]), and a model describing hadronic interactions occurring during particle collisions (in this case, SYBILL 2.3c [108]), MCEq calculates the expected flux of conventional and prompt leptons, including neutrinos, as a function of energy, altitude, and direction of observation. Alternatively, the software can derive predictions about the impact of each Barr parameter on the output neutrino yield. The impact of a Barr parameter on an event's weight in PISA is therefore defined as:

$$w_b = w_i + \left(b \cdot \frac{d\Phi_{nom}}{db} \right), \quad (99)$$

where b is the value of a particular Barr parameter and $\frac{d\Phi_{nom}}{db}$ is a gradient splined from MCEq which describes how the nominal flux Φ_{nom} changes as a function of parameter b . That correction is ap-

plied to an initial weighting w_i that is defined as:

$$w_i = \Phi_{\text{nom}} \cdot \left(\frac{E}{[24 \text{ GeV}]} \right)^{\Delta\gamma_\nu}, \quad (100)$$

Where Φ_{nom} is the baseline Honda flux, E the energy of the neutrino, and $\Delta\gamma_\nu$ is a parameterized change in the neutrino flux spectral index. The latter is allowed to vary the shape of the energy distribution of events, to account for uncertainties in the cosmic ray spectral index.

The impact of varying the neutrino spectral index by $+1\sigma$ on the final reconstructed analysis space can be observed in a *pull plot* such as the one in Fig. 59. This kind of plot shows the bin-by-bin impact of any systematic uncertainty on the final Monte Carlo template. That impact is quantified by the *bin pull*, defined as:

$$\text{Pull}_{\text{bin } i} = \frac{(N_{\text{pull}} - N_{\text{nominal}})}{\sqrt{N_{\text{nominal}}}}, \quad (101)$$

where N_{pull} is the expected count in an analysis bin when the spectral index is increased by 1σ , while N_{nominal} is the expected count when the parameter is set to its nominal value of 0.0. The full set of pull plots from the systematics included in the analysis is included in Appendix B.




Figure 59: Binwise pulls observed in the analysis bins when the neutrino spectral index $\Delta\gamma_\nu$ is increased by 1σ . Blue-coloured regions show bins where the expected event count is lower than in the nominal MC template, whereas red-coloured regions indicated areas where the pulled systematic leads to an excess in events w.r.t. the nominal template. One can also note that pulls occur only across the energy dimension of the histograms as is intended by the spectral index.

From a pull plot such as Fig. 59 one can see that an increase in the neutrino spectral index will lead to an exponential rescaling of the expected number of events as a function of energy, which translates into a roughly linear pull as a function of the reconstructed logarithmic energy. From this plot it is also possible to see the energy pivot of the spectral index shift, which is characterized by the white boundaries between the blue and red regions of the histograms. Pull plots for every systematic uncertainty are included in appendix B.

8.3.2 Neutrino Propagation

Systematics Included:

- **Atmospheric neutrino mixing angle θ_{23}**
- **Neutrino mass splitting Δm_{31}^2**

The propagation of neutrinos and the calculation of their oscillation probabilities across their path is done with the Prob3 Software package [109], which relies on theoretical computations from [110]. The original code was integrated into the PISA framework as a self-standing stage that assigns, based on a neutrino's origin and direction, a probability of conversion into one of the three standard flavours.

The oscillation stage takes into account the parametric resonance effect described in section 6.2.1 by propagating the neutrinos through successive layers of various electron densities. The details of the Earth layer structure is contained in a file providing the boundaries of a 12-layer modeling of the Earth's structure, called the Preliminary Reference Earth Model (PREM) [111]. The output probability will be affected by the choice of oscillation parameters, which can be either fixed at their current best-fit values or left free to fit. Table 6 lists the values used for the other oscillation parameters in the analysis; most nominal values have been taken from version 4.0 of the global neutrino fit results, using atmospheric neutrino data from Super-Kamiokande [112], except for θ_{23} and Δm_{31}^2 which were taken from the previous muon neutrino disappearance analysis performed on IceCube [113].

Table 6: Fixed and unfixed oscillation parameters for this analysis (Normal Ordering)

| Parameter | Nominal value | Fixed? | Prior |
|-------------------|------------------------------------|--------|-------------------------------------|
| θ_{12} | 33.82° | Yes | N/A |
| θ_{13} | $8.61 \pm 0.13^\circ$ | Yes | N/A |
| δ_{CP} | 217° | Yes | N/A |
| Δm_{12}^2 | $7.39 \times 10^{-5} \text{eV}^2$ | Yes | N/A |
| θ_{23} | 49.7° | No | $[30, 60]^\circ$ |
| Δm_{31}^2 | $2.525 \times 10^{-3} \text{eV}^2$ | No | $[1, 7] \times 10^{-3} \text{eV}^2$ |

Given that the tau appearance involves mainly the $\nu_\mu \rightarrow \nu_\tau$ channel, only θ_{23} and Δm_{31}^2 (or Δm_{23}^2 ³) or are directly impacting this analysis⁴: these are therefore treated as nuisance parameter that are expected to have values that are compatible with the result of the ν_μ disappearance analysis. The reasons for not fixing these to the outcome of the muon disappearance analysis is to ensure that the fitted tau neutrino normalization lies in a region of the parameter space that is compatible with the disappearance result.

8.3.3 Cross section uncertainties

Systematics Included:

- **QE axial Mass, $M_{A, QE}$**
- **RES axial Mass, $M_{A, RES}$**
- **DIS cross section uncertainty**
- **Effective area scaling, $A_{eff \nu}$**

³ Since IceCube is sensitive to neither the sign of Δm_{23}^2 , nor the solar mass splitting which is fixed in the analysis, the largest mass splittings of the normal ordering, Δm_{31}^2 and Δm_{23}^2 only differ by a constant factor. Both can thus be used interchangeably as fit variables. This analysis uses Δm_{31}^2 .

⁴ As an extra precautions, the impact of varying θ_{13} and δ_{CP} is assessed later in Section 8.5.

Multiple systematic uncertainties are used to model the uncertainties on the various contributions to the total neutrino-nucleon cross section. At lower energies, the GENIE event generator can reweight individual events obtained from QE and RES interactions for different values of axial masses ($M_{A,QE}$, $M_{A,RES}$), the empirical parameters described in Section 6.3. The simulation produced for this analysis had GENIE compute a weight for five different mass values, which were fit using a second degree polynomial. In the analysis pipeline, the value of the axial mass parameters is then allowed to vary continuously along this interpolated curve, within a prior of $[-2\sigma, 2\sigma]$. The treatment of these two systematic uncertainties is a conservative choice in regards to the relatively modest contributions of quasielastic and resonant events to the final event sample: as can be seen in Fig. 60, RES events account for 10% of the total number of neutrinos (15% of anti-neutrinos), whereas QE events only account for 5% of the total number of neutrinos (11% of anti-neutrinos).




Figure 60: Distribution of QE, RES, and DIS events at final level in the event selection, as a function of the neutrino’s true energy. Dark colors represent neutrino events, whereas light colors represent anti-neutrino events. The contributions of QE and RES events in the sample is limited to $O(10\%)$ of the total number of events, and concerns only neutrinos with the lowest energies.

In the DIS regime, a systematic parameter is used to accommodate for theoretical uncertainties between the nominal cross section computed by the GENIE[76] and CSMS [71] models. The main difference between the two models is the choice of parton density functions, and the method used to calculate the DIS contribution to the cross section. GENIE uses a leading-order approximation derived in [114] and the GRV98 parton density function (PDF)[115]. Meanwhile, CSMS uses a next-to-leading order approximation in their calculations, and a PDF derived from more recent data at the HERA experiment[116].

Since it relies on better calculation methods and uses state-of-the-art PDFs, the CSMS model is deemed to be the most accurate of the two; results from a measurement of inelasticity in IceCube's very-high energy neutrino sample also seems to support this claim[101]. However, the cross section calculations provided by CSMS can only be trusted above neutrino energies of 100 GeV[117], leaving the bulk of this analysis' energy range uncovered by CSMS predictions.

Following a detailed study of the differences between the two cross section models [117], a systematic uncertainty parameter was created to accomodate fits to both models. That parameter scales the strength of two corrections made to convert the GENIE model to a CSMS-like shape: one on the total cross section, and one on the differential cross section $d\sigma/dE dy$, which depends on both the energy E of an event and its inelasticity (y). Both corrections are applied by finding the CSMS weight of an event from interpolated splines, defined as:

$$w_{CSMS,tot}(E) = \frac{\sigma_{CSMS,tot}(E)}{\sigma_{GENIE,tot}(E)} - 1, \quad (102)$$


for the total cross section, and:

$$w_{CSMS,diff}(E, y) = \frac{\sigma_{CSMS,diff}(E, y)}{\sigma_{GENIE,diff}(E, y)} - 1, \quad (103)$$


or the differential cross section. The final weight w_f calculated for the event is then:

$$w_f = w_i \cdot (1 + w_{CSMS,tot}(E) * P_{DIS}) (1 + w_{CSMS,diff}(E, y) * P_{DIS}), \quad (104)$$

where w_i is the initial weight of the event, and P_{DIS} is the scaling introduced by the DIS systematic parameter. The effect of this parameter on the total cross section can be visualized in Fig. 61, while Fig. 62 illustrates the shape of the differential corrections in the true energy/inelasticity parameter space where they are defined. Note that in order to provide a continuous systematic to lower energies, the ratio of GENIE to CSMS cross section is extrapolated as a constant for energies below 100 GeV.



(a) neutrinos



(b) anti-neutrinos

Figure 61: Effect of the DIS systematic parameter P_{DIS} on the total cross section for a) neutrinos and b) anti-neutrinos. The red dashed line (GENIE_SUM) corresponds to a nominal cross section model, based on the GENIE prediction (where the systematic parameter is set to 0). The light green and dark red dashed line represent cross section values with the parameter set to respectively -1 and 1. The range covered by these model uncertainties can be compared to the available cross section data points. Figures taken from [117].




Figure 62: Corrections applied to the DIS differential cross section. For each type of neutrino, and each type of weak interactions, the 2D ratio between the CSMS and GENIE calculations is computed, which is the quantity plotted in this figure. The differential correction is obtained by subtracting 1.0 to the maps shown here, and scaling the effect strength by the DIS parameter P_{DIS} .

The resulting impact of the DIS parameter in the analysis space is quite complex, given the multi-dimensional nature of the corrections involved and selection effects in the final-level oscNext sample. To help visualize this effect, Figure 63 shows the individual contributions of both the total and differential cross section effect, in a scenario where P_{DIS} is set to a value of +1.




Figure 63: Impact of the DIS systematic parameter P_{DIS} on the analysis template, divided into its total and differential component: the total cross section correction (top row), the differential correction (middle row), and the combined effect of both corrections (bottom row). In each case, the parameter governing the shift is changed to a value of 1.0

Finally, the total number of neutrinos present in the sample can be scaled up or down by an effective area parameter $A_{eff\nu}$, which tries to compensate potential misestimation of total neutrino rate in the data, irrespective of their flavour or interaction type.

8.3.4 Detector effects

Systematics Included:

- Bulk ice absorption coefficient
- Bulk ice scattering coefficient
- Hole ice angular acceptance parameters (p_0 and p_1)
- DOM optical efficiency

In addition to continuous systematics, a series of discrete simulation sets have been created to take into account uncertainties associated with the detector itself. In these sets, one or several properties of the detector are modified from their nominal values, and the full event production chain is run. The count expectations from these discrete systematic sets are used to generate multi-dimensional *hypersurfaces*, so that the fit of those systematic parameter to data can be interpolated over a continuous range of values. Gradients can also be derived for each of the parameter dimensions included in the hypersurface, which provide an estimate of bin-by bin evolution of the expected count in the template after an incremental change in the value of the systematic parameters.

For a given analysis bin, the scaled count N_{sys} associated with a hypersurface $f(p_1, p_2, \dots, p_N)$ built out of N systematic parameters will be equal to:

$$\begin{aligned} N_{sys} &= N_{nominal} \cdot f(p_1, p_2, \dots, p_N), \\ &= N_{nominal} \cdot \left[c + \sum_{i=1}^N f_i(\Delta p_i, m_i) \right], \end{aligned} \quad (105)$$

where c is an overall offset constant common to all surfaces, and $f_i(\Delta p_i)$ is a function that relates the variation in expected events to the parameter gradient Δp_i . That function varies for each systematic, and takes the following form:

$$f_i(\Delta p_i) = \begin{cases} \exp[m_i \cdot (\Delta p_i)] - 1 & \text{for DOM efficiency in muons} \\ m_i \cdot (\Delta p_i) & \text{for all other systematics} \end{cases}. \quad (106)$$

Where m_i is a parameter that is fit using the discrete systematic sets. From the above treatment it is therefore possible to attribute, for continuous values of detector systematics, a definite change in the bin count expectation, which allows these uncertainties to be treated like any other in the fitting procedure. No cross terms between the various hypersurface parameters are included in this formulation.

Three detector effects, encompassing 5 parameters, are included in the hypersurface treatment. The first one is the average value of the bulk ice absorption and scattering coefficient, whose nominal values as a function of detector depth are shown in Fig. 36 of Section 7.1. Discrete systematic uncertainty sets of simulated neutrino and muon interactions have been generated for values of absorption and scattering ranging between +10% and -10% with respect to their baseline values, which is the prior range of uncertainty recommended by the Calibration Working Group following the characterization of the likelihood space around the best fit point of the LED flasher fits to the ice model.

The next detector effect taken into account concerns the optical properties of the refrozen ice around each IceCube string, which are expressed in terms of angular acceptance curves for photons that reach the DOMs in simulation. The probability for a photon arriving with an angle η with respect to the face of the PMT has been measured in-situ, with special calibration runs using the LED flashers. To model the uncertainties associated with these measurements, the angular acceptance of photons in simulation can be parametrized using a curve defined by two parameters, p_0 and p_1 . Figure 64 illustrates the shape taken by this curve as a function of various combinations of parameter values used to produce systematic uncertainty sets (the red curve shows the nominal acceptance model). As can be seen, the systematic uncertainty on the photon angular acceptance is particularly high for photons reaching the DOM from a vertically upgoing angle $\eta \approx 0$ because of a lack of flasher data to constrain this section of the curve, as no LEDs are oriented directly upward on the DOM board.

The origin of the true angular acceptance shape is unknown. Anecdotal evidence taken from a calibration camera positioned along one of the strings hints at the existence of a bubble column roughly aligned with the forwardmost portion of the DOM PMTs. This hypothesis has never been properly confirmed however, despite several attempts at calibration measurements to better visualize it.

The last sets of discrete systematic uncertainties have had their *overall optical efficiency* shifted between 90% and 110% of its baseline value. Optical efficiency is a scaling factor to scale the probability of photons to be converted into photoelectrons. This factor is derived from calibration studies on a sample of minimum ionizing muons: since these muons lose energy at a constant rate along their path, the Cherenkov light from their passage can be used as a standard candle. An event selection




Figure 64: Angular acceptance curves simulated in the discrete systematic sets, for various values of *hole ice* parameters p_0 and p_1 . Angular acceptance is defined as a function of $\cos(\eta)$, the cosine of the angle between a photon’s arrival direction and the surface of the PMT (which is assumed to be pointing downward). The acceptance curve used in the nominal Monte Carlo simulation is drawn in red.

of high purity minimum ionizing muons, designed in [118] and later refined in [119], is used in both data and simulation sets of varying optical efficiencies. The average charge collected in data is then compared to the one obtained in simulation, which yields a value of DOM efficiency that is used to scale the new physics simulations. The value of this optical parameter is subject to some uncertainty. Since the simulation propagates photons using an idealized modeling of the light scattering properties, any localized impurities or nuances in the structure of the ice sheets could lead to modify the overall light collection efficiency at the DOMs. For this reason, optical efficiency is included as a nuisance parameter in the analysis.

8.3.5 Muon Systematics

- **Spectral index, $\Delta\gamma_\nu$**
- **Effective area scaling**

Two other systematic parameters are applied to the muon pipeline of the analysis, to account for both the uncertainty in the muon production mechanism in atmospheric showers (similarly to the neutrino spectral index) and the uncertainty in IceCube’s effective area. The latter parameter is simply a normalization factor that scales up or down the muon contribution to the Monte Carlo template.

8.3.6 Assessing the need for noise parameter systematics

In previous analyses, uncertainties related to the Vuvuzela parameters of each DOM were assumed to have a negligible impact on the analysis. As this analysis now cumulates 8 years of detector data live-time, this assumption needed to be revisited, as lower levels of the event selection showed evidence of a time-based change in the number of uncleaned pulses recorded by each DOM. Despite the small contribution of noise-only events to the final event sample, a change in the Vuvuzela parametrization could possibly have an effect on the time structure of the noise pulses inserted into simulated physics event, which is the main reason why this systematic uncertainty was studied.

The choice of parameter to vary was based on the assumption that the correlated component of the distribution is the more likely to sneak into real physics event, because of its short timescale structure. Systematic uncertainty sets where the correlated noise rate was varied by $\pm 5\%$ and $\pm 10\%$, and the impact on the final event selection was estimated by selecting a data side-band that showed a) good data-MC agreement with the nominal MC, and b) was overwhelmingly dominated by noise events.

To construct the side-band, the distribution of the low-level variable "NChCleaned" (the number of triggered DOMs in the clean pulse series of an event) was constructed using events that a) failed the Level 4 noise classifier cut, and b) failed the Level 4 muon classification as well. This distribution is plotted in Fig. 65.




Figure 65: Distribution of the number of triggered DOMs, for events that satisfy the criteria of the noise side-band. The side-band is defined to correspond to the two first non-empty bins of the distribution, as these are dominated by noise events. Counts correspond to the nominal expected number of events for 8 years of livetime.

The event count in the two bins of the side-band were compared in data and the various Monte Carlo sets; this is shown in Fig. 66a. Note that the shaded grey bands represent the uncertainties in the rates obtained in data, for three groups of seasons: 2011/2012, 2013 to 2015, and 2016 to 2020. These groups corresponds to seasons in which small configuration changes in the online filters

might have resulted in variations of the noise rates⁵. To evaluate the need for systematic set, a χ^2 was constructed where the rates obtained in the first two bins of the data distribution are compared to each of the simulated noise-only simulation sets. That is to say:

$$\chi^2 = \frac{(N_{data} - N_{MC\ set\ i})^2}{\sigma_{data}^2 + \sigma_{MC}^2}. \quad (107)$$

The error used for the data count is taken to be the systematic uncertainty associated with season-to-season variations in the rate (the gray bands). MC uncertainty, meanwhile, is evaluated as the sum of the event weights squared, $\sum w_i^2$.

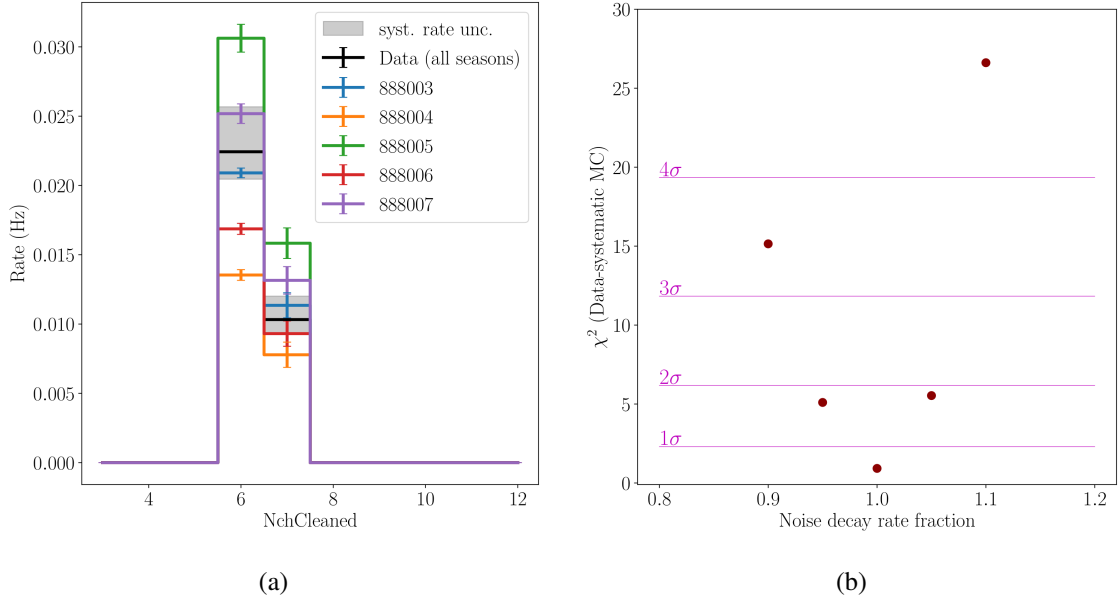


Figure 66: Outcome of the noise systematic uncertainty study. a) Event rates in the noise side-band, in both data and the various Monte Carlo sets. The shaded gray area represents the systematic uncertainty associated with the yearly variation of noise rates across the eight seasons of data used. The colored histogram represent the MC expectation associated with each of the systematic noise sets. b) Variation of the χ^2 defined in Eq. 107, as a function of the systematic variation in λ_{corr} . Purple lines denote the usual significance thresholds for a χ^2 with d.o.f. = 2.

Figure 66b shows the χ^2 metric as the correlated noise rates in the MC set increases, along with the confidence levels associated with this particular side-band (two bins are used, therefore the d.o.f. is 2). What this figure shows is that if the correlated noise rates had been consistently off by 5%, one could expect that the side band would be significantly disagreeing with the data count. This means that even fluctuations as little as 5% can safely be ruled out as a reasonable choice of fluctuation, thereby eliminating the need for a noise systematic of this sort.

⁵ These variations would have been subtle however, as all seasons have been reprocessed with the same offline filter processing.

8.4 LIKELIHOOD FORMULATION

The statistical significance of a unitarity-violating signal is evaluated by using a binned likelihood as the optimization metric. From this likelihood, a test statistic is defined as the following ratio:

$$TS = -2 \ln \left[\frac{LH_{best-fit}}{LH_{N_\tau=1}} \right], \quad (108)$$

where $LH_{best-fit}$ is the minimized likelihood obtained from the fit to data, and $LH_{N_\tau=1}$ is the likelihood value obtained when the tau normalization value is fixed to 1.0. If N_τ deviates significantly from a value of 1.0 (by more than 3σ), the result could then be interpreted as evidence of a tension between data and the standard three-neutrino oscillation paradigm.

A simple formulation of a likelihood function is the poisson likelihood, which is defined as

$$\begin{aligned} L &= \sum_i^{N_{bins}} Poi(k_{data,i}; \lambda), \\ &= \sum_i^{N_{bins}} \frac{\lambda^{k_{data,i}} \exp^{-\lambda}}{k_{data,i}!}. \end{aligned} \quad (109)$$

In the poisson likelihood, the probability of obtaining data count $k_{data,i}$ in bin i is assumed to be poisson-distributed around expectation value λ which, in the context of Monte Carlo simulation, is given by the sum of all the weighted MC events w_i falling into a bin:

$$\lambda = \sum_i^{k_{mc}} w_i. \quad (110)$$

A potential issue with this likelihood formulation is that it assumes that the expectation obtained from Monte Carlo is known to infinite precision, which is rarely the case in reality. Monte Carlo simulation production is an inherently stochastic process, which means that a finite simulation set should be considered as a random sampling of the true distribution of Monte Carlo weights. This problem, raised in [120] and [121], has been revisited in recent years ([85], [122], [123], [124]) which has led to the development of new likelihood formulations to better take into account the finite statistics involved in a simulation data set. In order to deal with potential problems stemming from finite MC statistics, a study of the performance of the available likelihood formulations was conducted prior to choosing the optimization metric for the analysis.

8.4.1 Choice of optimization metric

In addition to the standard Poisson likelihood defined in Eq. 109, three other types of metrics including finite MC statistics corrections were considered for the analysis. These include a *modified* χ^2 , defined as:

$$\chi_{mod,i}^2 = \frac{(k_{data,i} - k_{MC,i})^2}{k_{data,i} + \sum_i w_i^2}, \quad (111)$$

where $k_{MC,i}$ is the sum of the Monte Carlo weights in a bin i . Three types of extensions to the Poissonian likelihood were also considered, which have the form:

$$L_{bin\ i} = \int_0^\infty Poi(k_{data,i}; \lambda_i) \cdot P(\lambda_i) d\lambda_i, \quad (112)$$

where $Poi(k_{data,i})$ is the standard Poisson expression in Eq.109, with the exception that these likelihood integrate (or marginalize) over all possible expectation values λ_i that a finite MC template could yield. This is a bayesian construct that includes a prior $P(\lambda_i)$, encoding the probability for a given set of MC weights to sum up to expectation λ_i . The three extended likelihoods differ from each other by the type of prior used. These likelihoods:

- the *Effective likelihood* (mcllh_eff)
- the *Mean likelihood* (mcllh_mean)
- the *Generalized Poisson likelihood* (gen_llh)

Details about the difference in priors used in the three methods are available in Appendix B. The choice of metric for this analysis was based upon two criteria: the absence of a *bias* when fitting the tau normalization on statistically fluctuated datasets, and the proper *coverage* of the test statistic distribution. All available optimization metrics were evaluated using a simple 1D toy MC example, where the peak of a gaussian signal is fit on Poisson-fluctuated data using templates made from varying levels of MC statistic⁶. They were then tested on the full, multivariate fit of the ν_τ appearance analysis.

Bias was evaluated by determining the fit value returned for N_τ by a particular likelihood minimization, for the median of 100 pseudo-experiments, where pseudo-data comes the statistically fluctuated MC template. Figure 67 shows the median sensitivity for each of these metrics for the determination of the tau normalization⁷. As can be seen from this plot, all likelihood implementations showed no significant biases in the determination of the tau neutrino normalization.

Coverage was evaluated by looking at the test statistic distribution obtained on the previous pseudo-trials at the nominal value of N_τ . Since $LH_{best-fit}$ and $LH_{N_\tau=1}$ are *nested hypotheses* (in other words, the best-fit likelihood is simply an extended version of the null hypothesis with 1 additional parameter), Wilk's theorem states that the likelihood ratio of the two values should be χ^2 -distributed, with a degree of freedom equal to 1. Coverage is therefore evaluated by comparing the percentile intervals of the experimental TS distribution with those of a χ^2 statistic. The result of this comparison is shown in Fig. 68 where the expected percentiles of a χ^2 are plotted on the x-axis. These percentiles corresponds to well-defined values of test statistics; the observed fraction of test statistics corresponding to these values can then be plotted on the y-axis. Curves lying on the 1:1 line of Fig. 68 show that the percentiles of the observed and expected distributions are similar, which means that the TS distribution obtained from these metrics satisfy Wilk's theorem. This means that trustworthy confidence levels can be derived by comparing directly the values of likelihood ratio (or modified χ^2) obtained to the known significance threshold of an appropriate χ^2 function, without having to brute-force confidence intervals by running the fit on thousands of trials.

⁶ See Appendix B for more details on the toy MC.

⁷ The generalized likelihood implementation could not converge properly to a fit, and was therefore excluded from the study at this point.




Figure 67: Result of the sensitivity test scan applied on 100 pseudo-data trials of a full tau appearance analysis chain, for each of the optimization metric tested. The circle denote the location of the median fit value of N_τ , while the horizontal bars indicate the 68% and 90% confidence interval obtained from the pseudo-data trials. Every metric fits back a median N_τ value that is within 1-2% of 1.0, which does not constitute a statistically significant bias.

These results led to the conclusion that despite their theoretical advantages, novel metrics taking into account finite MC statistics were not required for this analysis. Instead, the regular poisson likelihood was chosen to the test statistic of this analysis, since it was clear from this study that the nominal MC templates had sufficiently high statistics.

8.5 FINAL CHOICE OF SYSTEMATICS

Despite providing useful parametrizations of known effects in the experiment, systematic uncertainties could also add unnecessary degrees of freedom in the analysis if it turns out that the latter is not sensitive to the parametrized effect. In order to keep in line with Occam's principle a *systematic impact test* is conducted. In this test, pseudo-data is generated from a Monte Carlo template in which a single systematic parameter is pulled away from its nominal expectation. The pulled pseudo-data is then fit with a Monte Carlo template in which the same parameter is fixed to the nominal value, and not allowed to vary during the fit.

The impact of not including that particular systematic parameter is measured in terms of test statistic difference, that is to say:

$$I_s = 2(LLH_{s \text{ fixed}, N_\tau \text{ free}} - LLH_{s \text{ fixed}, N_\tau \text{ fixed}}) - 2(LLH_{s \text{ free}, N_\tau \text{ free}} - LLH_{s \text{ free}, N_\tau \text{ fixed}}), \quad (113)$$

where I_s is the impact of systematic parameter s , N_τ is the tau neutrino normalization, and the labels *free* and *fixed* refer to the status of parameter s in the likelihood fit. The threshold for fixing a systematic uncertainty is set such that its associated impact must be smaller than 0.1σ , where σ refers to the width of the expected test statistic distribution of the null hypothesis. In this analysis, that distribution should be χ^2 -distributed with a degree of freedom of 1.0, which means that the numerical value of the threshold is set to 0.02.

The scale of the pulls used in this test depend on the type of prior assigned to the uncertainty: if the prior is gaussian, the pull is defined as that gaussian's standard deviation, otherwise the pull is taken




Figure 68: Coverage plot showing the performance of the four optimization metric applied to the full tau appearance analysis chain. In this case, the TS values were obtained by comparing the multidimensional fit with and without fixing the tau normalization to one. Every metric shows good coverage properties, which is indicative that the oscNext MC event templates do not suffer from finite Monte Carlo statistics.

to be half of the parameter fit range. Because of the expected asymmetric effect of some parameters⁸, the test was ran on both positive and negative pulls. In addition, the test was ran with pseudo-data created with tau neutrino normalizations of 0.75 and 1.25, corresponding roughly to the expected 2σ uncertainty range of the analysis. The purpose of running the test at these off-nominal point is to evaluate the impact of a pull on the determination of the true uncertainty bounds of the analysis.

Figure 69 shows the maximum impact obtained in this test, out of the six pull scenarios tested (2 signs of pulls \times 3 tau normalization values). In certain cases the parameter pull impact is negative, which can happen since the fits performed are not Asimov fits⁹. The vertical dashed line delimits the decision threshold of 0.02, whereby parameters with an absolute value of I_s lower than it have been fixed to their nominal values for the final analysis.




Figure 69: Impact of individual systematics on the tau neutrino appearance analysis. The value associated with each parameter corresponds to the largest pull value I_s recorded out of the six pull tests (negative/positive for three different tau normalization values) performed. Red values indicate that the impact value was negative. The vertical dashed line shows the threshold delimiting each parameter's fate: parameters with an absolute I_s score below the threshold are fixed in the final analysis.

To further ensure that the physics result remain unchanged by a pulled systematic uncertainty, the fitted values of N_τ associated with the tests shown in Fig.70 were plotted in terms of offsets with respect to the true value of the parameter in the test pseudo-data. The bottom bar of Fig. 70 indicates, for instance, that the θ_{13} pull that led to the largest impact score in Fig. 69 fitted back a tau normalization very close to 1.25, which was the true value used in the pseudo-data for that test. The

⁸ A decrease in DOM efficiency, for example, might lead to a smaller change than an increase.

⁹ In other words, the nominal MC template in every fit is not exactly the one used as pseudo-data, meaning that the likelihood values associated with the best-fit point are likely different from the ones from the Asimov case.

important result from Fig. 70 is that none of the parameters selected to be fixed led to a large offset in the physics parameter of interest, thus confirming that they could be fixed without causing any harm.




Figure 70: Impact of individual systematics on the fitted tau neutrino normalization N_τ . The quantity plotted alongside each parameter name relates to the fitted value of N_τ recorded in the test that produced the highest impact I_s in Fig. 69. That quantity measures the offset between the fitted normalization and the true value that was used in the pseudo-data, in units of standard deviation (here assumed to be 10%). Any offsets smaller than 0.1σ indicates that a systematic uncertainty can be fixed, as long as it also satisfies the previous I_s threshold. Also note that σ in this plot has a different meaning than in Fig.69.

Finally, given the existence of some correlations between some parameters, the impact test was re-run on every fixed parameter, using the newly reduced list of systematics in the Monte Carlo template. This test could have indicated whether a parameter's impact on the analysis had increased now that correlated partners were also fixed. As can be seen in Fig. 71, no significant change in the test outcome was visible. Table 7 lists the full list of systematic uncertainties retained for the analysis, while Table 9 in Appendix B summarizes the information regarding all uncertainties considered in this chapter.




Figure 71: Impact score of the newly fixed systematics on the final analysis. The impact of the pull systematic was assessed on the MC template with the reduced list of free-floating systematics, to make sure that potential correlations between pairs of fixed parameters might affect the result of the fit. All impact values are still smaller than the decision threshold (black vertical dashed line), indicating that the chosen parameters can remain fixed.

Table 7: Systematic uncertainties included in the tau neutrino appearance analysis

| Category | Parameter Name | Nominal Value | Prior | Full range |
|-------------|-----------------------------------|--------------------------|------------------|----------------|
| Flux | Spectral index $\Delta\gamma_\nu$ | 0.0 | gaussian (0.1) | [-0.5, 0.5] |
| | Barr Y (K) | 0.0 | gaussian (0.3) | [-1.5, 1.5] |
| | Barr H (π) | 0.0 | gaussian (0.15) | [-0.75, 0.75] |
| | Barr I (π) | 0.0 | gaussian (0.122) | [-0.61, 0.61] |
| | Barr W (K) | 0.0 | gaussian (0.4) | [-2, 2] |
| | Barr Y (\bar{K}) | 0.0 | gaussian (0.3) | [-1.5, 1.5] |
| | Barr G (π) | 0.0 | gaussian (0.3) | [-1.5, 1.5] |
| | Barr Z (K) | 0.0 | gaussian (0.122) | [-0.61, 0.61] |
| Propagation | Δm_{31}^2 | 2.38e-3 GeV ² | uniform | [0.001, 0.007] |
| | θ_{23} | 45.6 | uniform | [0, 90] |
| Interaction | M_A (QE) | 0.99 GeV | gaussian (0.198) | [-40%, 40%] |
| | M_A (RES) | 1.12 GeV | gaussian (0.224) | [-40%, +40%] |
| | DIS CSMS | 0.0 | gaussian (0.1) | [-3, 3] |
| Detector | Hole ice p_0 | 0.101569 | uniform | [-2, 1] |
| | Hole ice p_1 | -0.049344 | uniform | [-0.2, 0.2] |
| | DOM Efficiency | 1.0 | gaussian (0.1) | [0.8, 1.2] |
| | Bulk ice sca | 1.0 | uniform | [0.95, 1.15] |
| | Bulk ice abs | 1.0 | uniform | [0.9, 1.1] |
| | Eff. Area $A_{eff\nu}$ | 1.0 | uniform | [0, 3] |
| Muons | Eff. Area $A_{eff\mu}$ | 1.0 | uniform | [0, 3] |
| Physics | N_τ | 1.00 | uniform | [0, 3] |

RESULTS

As is the norm for particle physics analyses, the fit of the simulation template to data follows a rigorous blinding procedure, where every component of the analysis pipeline must be carefully vetted to ensure that the analysis does what it is meant to do. This is necessary because physicists can unconsciously bias their analyses to match the result they expect. To avoid this, an ensemble of tests and checks have been performed to evaluate the resiliency of the analysis pipeline to systematic uncertainties and potential mismodeling, the proper performances of the minimizing routine and the ability of the analysis to converge to the correct physics values.

Since an analysis like the tau neutrino appearance is a massively complex, collaborative undertaking, the path toward unblinding is akin to a winding road full of potholes, sketchy turns and off-road improvisation. While the official IceCube analysis is currently undergoing final review, I have proceeded separately to partly unblind the data, given the advanced maturity of the analysis and my graduation timeline. Therefore, the results shown here diverge from the official IceCube analysis and serve the purpose of providing an object of discussion in the thesis.

9.1 PRE-UNBLINDING TEST RESULTS

The pre-unblinding test phase defines the series of checks performed prior to fitting the MC template to real data. They are meant to evaluate the potential of the analysis for new discoveries, and its capacity to run smoothly in the idealized cases where the data is directly derived from Monte Carlo.

9.1.1 *Injection-Recovery Test*

- **Purpose:** Make sure that the minimizer is capable of converging toward a known true template.
- **Test description:** The nominal Monte Carlo template is weighted to specific values of the physics parameter, and provided as pseudo-data for the fit. Starting with the same Monte Carlo template, but at random initial guesses for the parameter, the minimization is performed to check whether the recovered values match the injected truth.

This test is performed using the un-modified, nominal Monte Carlo template as an *Asimov Dataset*, which means that one assumes in this test that the particular instance of simulation generated is absolutely representative of the data¹. This dataset is devoid of statistical fluctuations, meaning that a fit of the same template should return exactly the same parameter values, up to the numerical precision of the minimizer. Figure 72 shows the result of an inject-recovery test for various injected values of the tau normalization parameter. The value of the true normalization in the pseudo-data is changed five times, over values ranging from 0 to 2. For every injected true value (displayed on the x-axis),

¹ The term was coined in reference to Isaac Asimov's Franchise short story [125], where a single individual is deemed to be the most representative voter.

the fit procedure is run: starting with the nominal MC template, where $N_\tau = 1.0$, the tau neutrino normalization is allowed to vary, along with the values of all systematic nuisance parameters. The outcome of the fit is shown on the y-axis, where one can see that the analysis is capable of recovering the truth at every injection point, as all yellow dots fall exactly on the gray 1:1 line of Fig. 72, to within the precision error tolerated by the minimizer.



Figure 72: Result of the injection-recovery test on Asimov datasets. For each of the scan points shown in yellow, a fit was performed where the pseudo-data is generated with a different true tau normalization value. The purple dash line represents where the initial tau normalization value of the MC template, prior to the fit.

9.1.2 Sensitivity test

- **Purpose:** Quantify the ability of the analysis to reject the null hypothesis.
- **Test description:** Scan points are chosen whereby pseudo-data is generated using a range of tau normalizations (from 0 to 2.0). At every scan point, two fits are performed: one in which all parameters are allowed to vary (called a *free fit*), and one in which the tau normalization parameter is frozen to 1.0 (the *null hypothesis fit*).

Following these fits, the test statistic value of every data point is calculated as: $2(LLH_{free\ fit} - LLH_{null})$, and plotted in Fig. 73. The values of likelihood ratios can be converted into significance value, as Wilks' theorem applies in this situation (as the null hypothesis is nested within the free fit hypothesis); these values are shown as horizontal purple lines in the figure. From the shape of the curve it is possible to see that as one attempts to fit the nominal MC template with pseudo-data generated with $N_\tau \neq 1.0$, the value of the likelihood ratio increases, leading to smaller p-values and thus a stronger rejection of the null hypothesis.

This sensitivity scan can be performed using either Asimov datasets as in the injection-recovery test, or statistically fluctuated trials. While the former allows for a quick estimation of the sensitivity that can be used as a feedback during the development of the analysis (see for example the plots from Section 8.2), trials-based sensitivity is a more accurate estimation of the true performance of the analysis. Trials can also be useful to evaluate likelihood biases (see Section 8.4.1) and construct

expected test statistic distributions to evaluate the goodness of the fit to real data (see the ensemble test in next section).

On Fig. 73, the outcome of the sensitivity test on both types of pseudo-data is shown: the red line corresponds to the test performed with an Asimov dataset, while the shaded bands correspond to regions where 68% (dark orange) and 90% (light orange) of the trial parabolas end up after the test. The outcome of the trial-based sensitivity test can be summarized by looking at the median sensitivity, which is shown as a dark dashed line on the same figure.




Figure 73: Sensitivity of the tau neutrino appearance analysis. a) The red dashed line represents the sensitivity obtained from a single Asimov pseudo-dataset, while the black dashed line shows the median sensitivity of 400 statistically fluctuated, pseudo-data trials. The dark orange band illustrates the 68% confidence interval of the trials, while the light orange band shows the 90% confidence.

Since the median curve is well centered at the nominal tau normalization value, and agrees well with the outcome of the Asimov test, it is possible to conclude that the reported confidence intervals are robust against statistical fluctuations. These intervals can be constructed by determining the width of the median sensitivity parabola at the point where it intersects the 1σ and 2σ confidence levels of a χ^2 -distributed likelihood ratio, which are indicated by the horizontal purple lines in the

figure. These confidence intervals are projected on a horizontal line in the bottom section of the plot.

The same 1σ interval drawn at the bottom of Fig. 73 can also be compared to the currently reported measurements of the tau neutrino normalization. This is done in Fig. 74, where one can see that this analysis has the potential to constraint N_τ at least a factor of 4 better than the last IceCube result, reaching for the first time an order of 10% precision.




Figure 74: The 68% confidence interval, taken with respect to the width of the Asimov sensitivity curve, is compared with the previous results published in the literature: the previous IceCube analysis [75], OPERA [73] and Super Kamiokande [74].

9.1.3 Ensemble Test

- **Purpose:** Ensure the sanity of fit when the optimization is ran on multiple statistically fluctuated data.
- **Test description:** A number of pseudo-experiment datasets are created by statistically fluctuating the bin counts of the nominal simulation template². Each of these pseudo-experiments is fitted using the analysis procedure, and several ensemble distributions are plotted from the result of these fits. Note that for this test, only a free fit is performed at the $N_\tau = 1$ nominal value; unlike the sensitivity test, the null hypothesis fit is not included in the test.

The first plot made of this ensemble of pseudo-experiments is the distribution of the optimized likelihood value, shown in Fig. 75. This distribution can be useful in two ways. First, it can serve as a check that the minimization process is well-behaved, whereby all likelihood values distribute themselves in a unimodal and smooth distribution. Second, it can later be used as a way to estimate the goodness-of-fit of the real data prior to unblinding the physics parameter value, which can give an indication as to whether or not the hypotheses underlying the analysis template are reasonable. Such a procedure is called a *blind fit*, and is used to detect potential sources of mismodeling in the simulation prior to unblinding.

² Contrary to the pseudo-data generated in the sensitivity test, the pseudo-data of the ensemble test are only generated at a tau normalization of 1.00




Figure 75: Distribution of the tau neutrino appearance test statistic from the fit of 200 pseudo-experiments to the nominal template. The distribution is unimodal, which is the expected shape. Note that the absolute value of the likelihood is not relevant, as it is not compared to any null hypothesis.

The next set of interesting plots from an ensemble test is the set of distributions of every fitted parameter in the analysis, shown in Fig. 76a; fitted values can then be compared to the prior used in the fit to determine the average pull of a parameter in the analysis procedure. Figure 76b shows the pull distribution for one of the systematic parameter (the ratio of pions-to-kaons), along with the mean and width of the distribution for all parameters fitted on the pseudo-data template. The mean and width fitted to every systematic parameter is shown in Fig. 76c. These pulls average out to zero over multiple fits, which is the desired behaviour when the injected pseudo-data is derived from the nominal simulation template used in the fit.




Figure 76: Key information from individual parameter fits. a) Fitted value of a single systematic parameter (`pion_ratio`, the ratio of pions to kaons) across the 200 pseudo-experiment fits. The distribution is shown in yellow, while its mean and 1σ width are indicated by the vertical blue lines. Green lines show the initial prior for the parameter during the fit. b) Pull distribution for the same parameter. The yellow distribution is fit with a gaussian, and the fitted mean and 1σ width are shown as blue lines. c) Mean and 1σ width of the fitted gaussians, for all systematic parameters that have a gaussian prior imposed in the fit.

9.2 BLIND FIT AGREEMENT BETWEEN DATA AND SIMULATION

After passing the review by the Oscillation Working Group, approval is initially given to perform a fit of the Monte Carlo template to the real data, without looking at the physics results. This means that the fitted value of N_τ , θ_{23} and Δm_{31}^2 are kept hidden from view, along with the three-dimensional template showing the binwise event expectation that best matches the real data. This is done because it is often the case that unexpected disagreements in certain variables that feed into the analysis can be seen after a full fit; it can also happen that some systematic parameters fit beyond their expected prior. While these disagreements between data and Monte Carlo might be the result of a potentially interesting physics signal, they might also be caused by modeling errors whose impact on the analysis might have been overlooked (i.e. not included in the systematic uncertainties from the start), or underestimated (by attributing it a smaller prior than warranted). In some cases, this additional step can lead to the discovery of previously undiscovered bugs in either the data treatment or the simulation production.

For these reasons, it is often necessary to review certain parts of the analysis pipeline, and that review process could potentially be biased if one already knew where, in the analysis space, things should be tweaked in order to obtain a particular physics result with respect to the mixing angle, mass splitting, and the tau neutrino normalization. Blind fits thus constitute an additional phase of scrutiny

where one makes sure that:

- The data-MC agreements between the distributions of all analysis variables (energy, zenith, and PID) is good,
- The goodness-of-fit of the free fit is reasonable (based on the expected distribution of likelihood values of the ensemble test),
- The distribution of pulls among the fitted systematic parameters is statistically sound (i.e. that there is no excessive pulls compared to the number of parameters tested),
- In the event of a bad goodness-of-fit, distributions of auxiliary variables may reveal clues as to the origin of a potential mismodeling-induced disagreement.

Blind fits mark the point at which this thesis diverges from the official IceCube analysis.

9.2.1 Data-MC agreement in key variables

A fit to data was performed using the Monte Carlo template and analysis pipeline described in Chapter 8. As a first check, the 1D distributions of the analysis variables (energy, zenith, and PID) were plotted in data and Monte Carlo, to assess the level of agreement between the two. Figures 77, 78, and 79 show the outcome of this comparison in the case of the analysis variables.




Figure 77: Data-MC agreement, at the best-fit point, on the reconstructed total energy (in GeV) of the events (in logarithmic space).




Figure 78: Data-MC agreement, at the best-fit point, on the reconstructed zenith angle of the events (displayed as $\cos(\theta_{zenith})$).




Figure 79: Data-MC agreement, at the best-fit point, on the PID classifier output variable. Values close to one indicate that an event has a high probability of being a track.

Whereas energy and zenith distributions show very good agreements, as evidenced by their flat ratio plots, a slight disagreement is visible in the distribution of the PID classifier output. It can be seen in Fig. 79 that the fitted MC template tends to overestimate the real number of cascade events with an output score smaller than 0.3 in data. That region of the PID classifier roughly corresponds to neutrino cascade events (see Fig. 48 in Chapter 7).

Some disagreements were also visible in some of auxiliary variables that provide additional insight into the properties of the event sample. One of them is the output score of the Level 7 Muon classifier, which determines the probability of an event to be neutrino-like. That variable is plotted in Figure 80, where the mismatch between data and Monte Carlo appears to be localized to events that the Muonclassifier deems ambiguous (they are neither neutrino-looking nor muon-looking).




Figure 80: Data-MC agreement, at the best-fit point, on the Level 7 Muon classifier output of the events. A classifier output value close to 1.0 means that an event is likely a neutrino.

Another set of auxiliary variables showing signs of disagreements consist of the reconstructed vertex spatial location of events in the sample. Fig. 81b (and to a certain extent, Figures 81a and 81c) indicates a rather complex bias in the way events are reconstructed in MC compared to data. This bias was expected however, based on a study performed on the impact of DOM efficiency on the reconstructed event properties. Figure 82 shows, on the left, a mapping of the reconstructed vertices in the in the XY-plane for the nominal set, and on the right for a systematic set where the DOM efficiency was lowered to 0.90. The color scale on the plot indicates the binwise pull observed between the number of events in data and in MC: blue indicates a deficit in data with respect to MC, and red an excess. As can be seen on the left of Fig. 82, less events seemed to be reconstructing close to DeepCore strings, whereas on the right side of the figure, this bias went in the opposite direction. This suggested that the true DOM efficiency might be lower than the assumed nominal value used in the analysis. The impact of this disagreement in the analysis should thus, in theory, be mitigated by the use of hypersurfaces³.

³ Note that hypersurfaces only correct the DOM efficiency effect in the counts obtained in the binned analysis space. The disagreement seen in reconstructed vertices is expected to remain after the fit, since these variables are not directly corrected for.

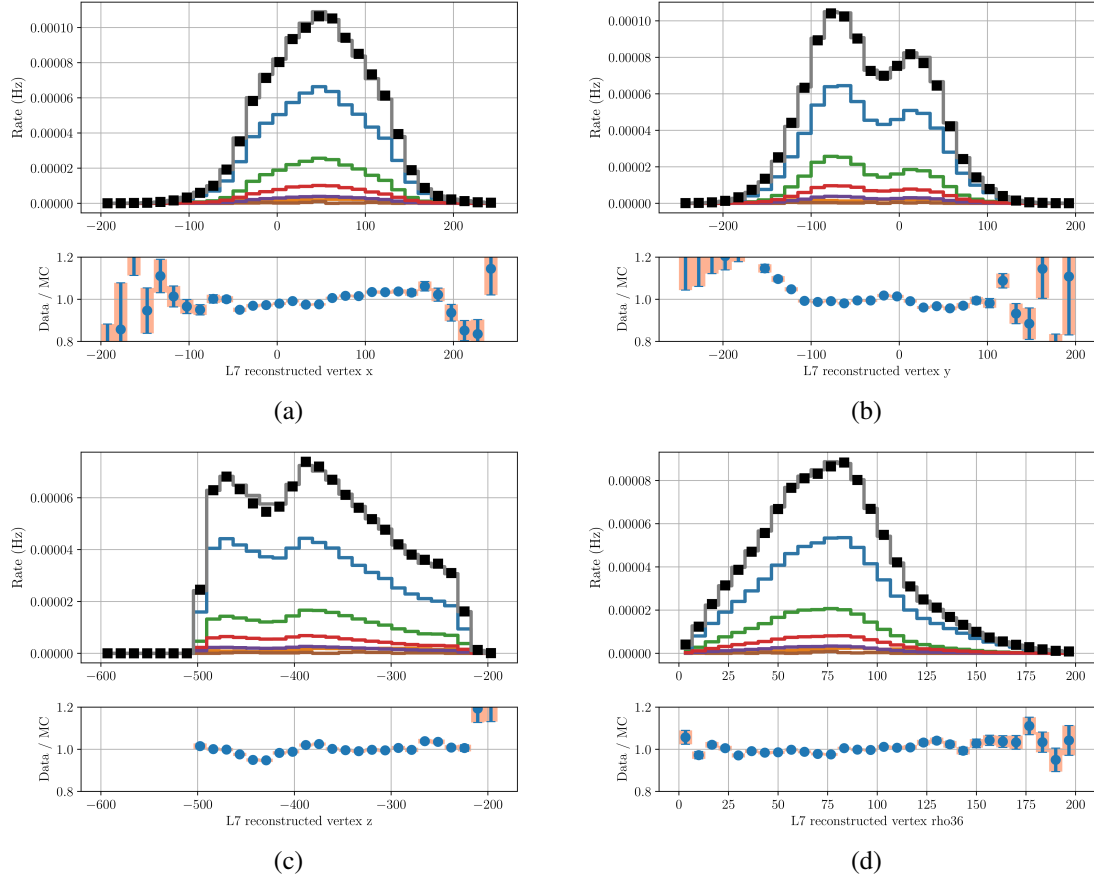


Figure 81: Data-MC agreement, at the best-fit point, on the position of the reconstructed vertex of the events. These are considered important auxiliary variables. a) X-position b) Y-position c) Z-position d) Radial distance from String 36. All distances are in meters.

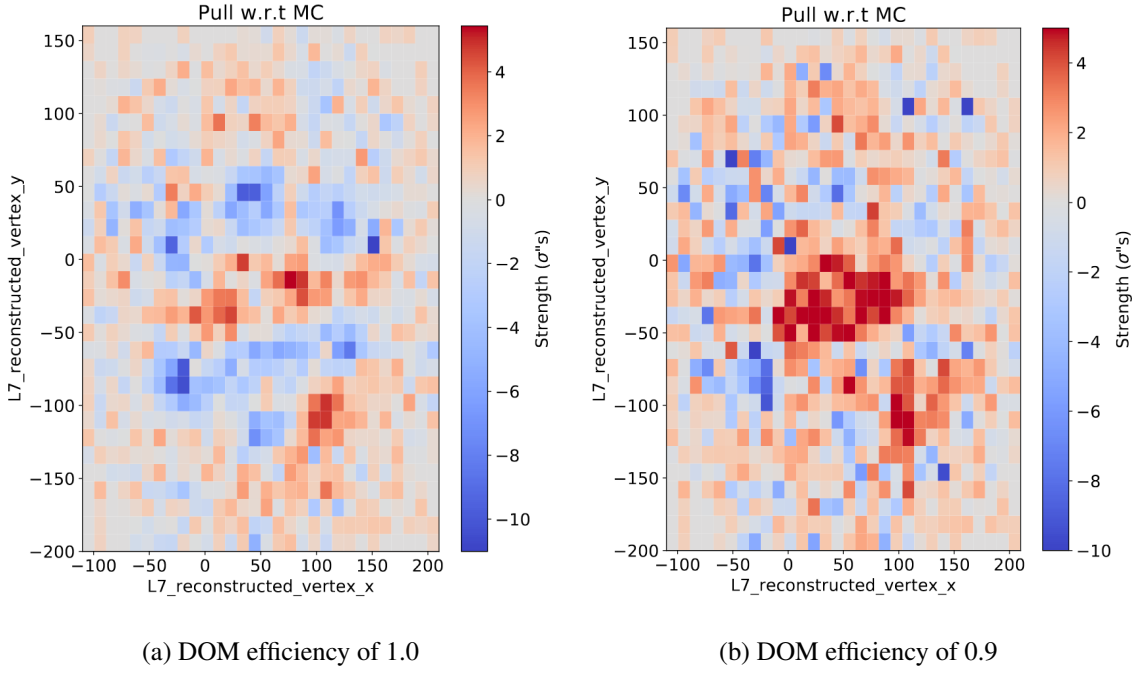


Figure 82: Two-dimensionnal XY-vertex agreement between data and two different DOM efficiency sets: the nominal set (a), and set 0001 (b). The disagreement is expressed in terms of pulls w.r.t. the Monte Carlo counts: negative pulls mean that there is a deficit in data.

9.2.2 Systematic parameter fits and pulls

Given that the disagreement seen in the PID and muon classifier distribution affected a rather small fraction of the event sample, and that the other analysis variables did not show any problems, the analysis result was deemed in sufficient agreement for unblinding the systematic uncertainty parameters of the fit to data, while still keeping the oscillation physics parameters blind. Figure 83 shows the total likelihood value obtained during the free fit to data (the black vertical line on the plot), which can be compared to the expected distribution of the same metric obtained in the ensemble test (the orange distribution on the same plot). As can be seen, the fitted value deviates significantly from the expectation, leading to a two-sided p-value of $\approx 0^4$.

Figure 84 shows the unblinded value of each systematic uncertainty, overlaid on top of the fitted distributions obtained in the ensemble test described in section 9.1.3. As for the goodness-of-fit plot from Fig. 83, several systematic uncertainties fit far away from the distributions obtained in the ensemble test. This portrait is however incomplete without looking at the priors accompanying some of the uncertainties. Those exist to allow certain parameters to vary beyond pure statistical fluctuations to consider systematic shifts allowed by the current external constraints on these quantities.

⁴ As a reminder, larger negative values of likelihood corresponds to a worse fit.




Figure 83: Goodness-of-fit of the fit to real data. The vertical black line indicates the value of likelihood obtained in the free fit of the MC template to data, and can be compared to the expected distribution of likelihood value obtained in the ensemble test. The fit to data has a poor goodness-of-fit.

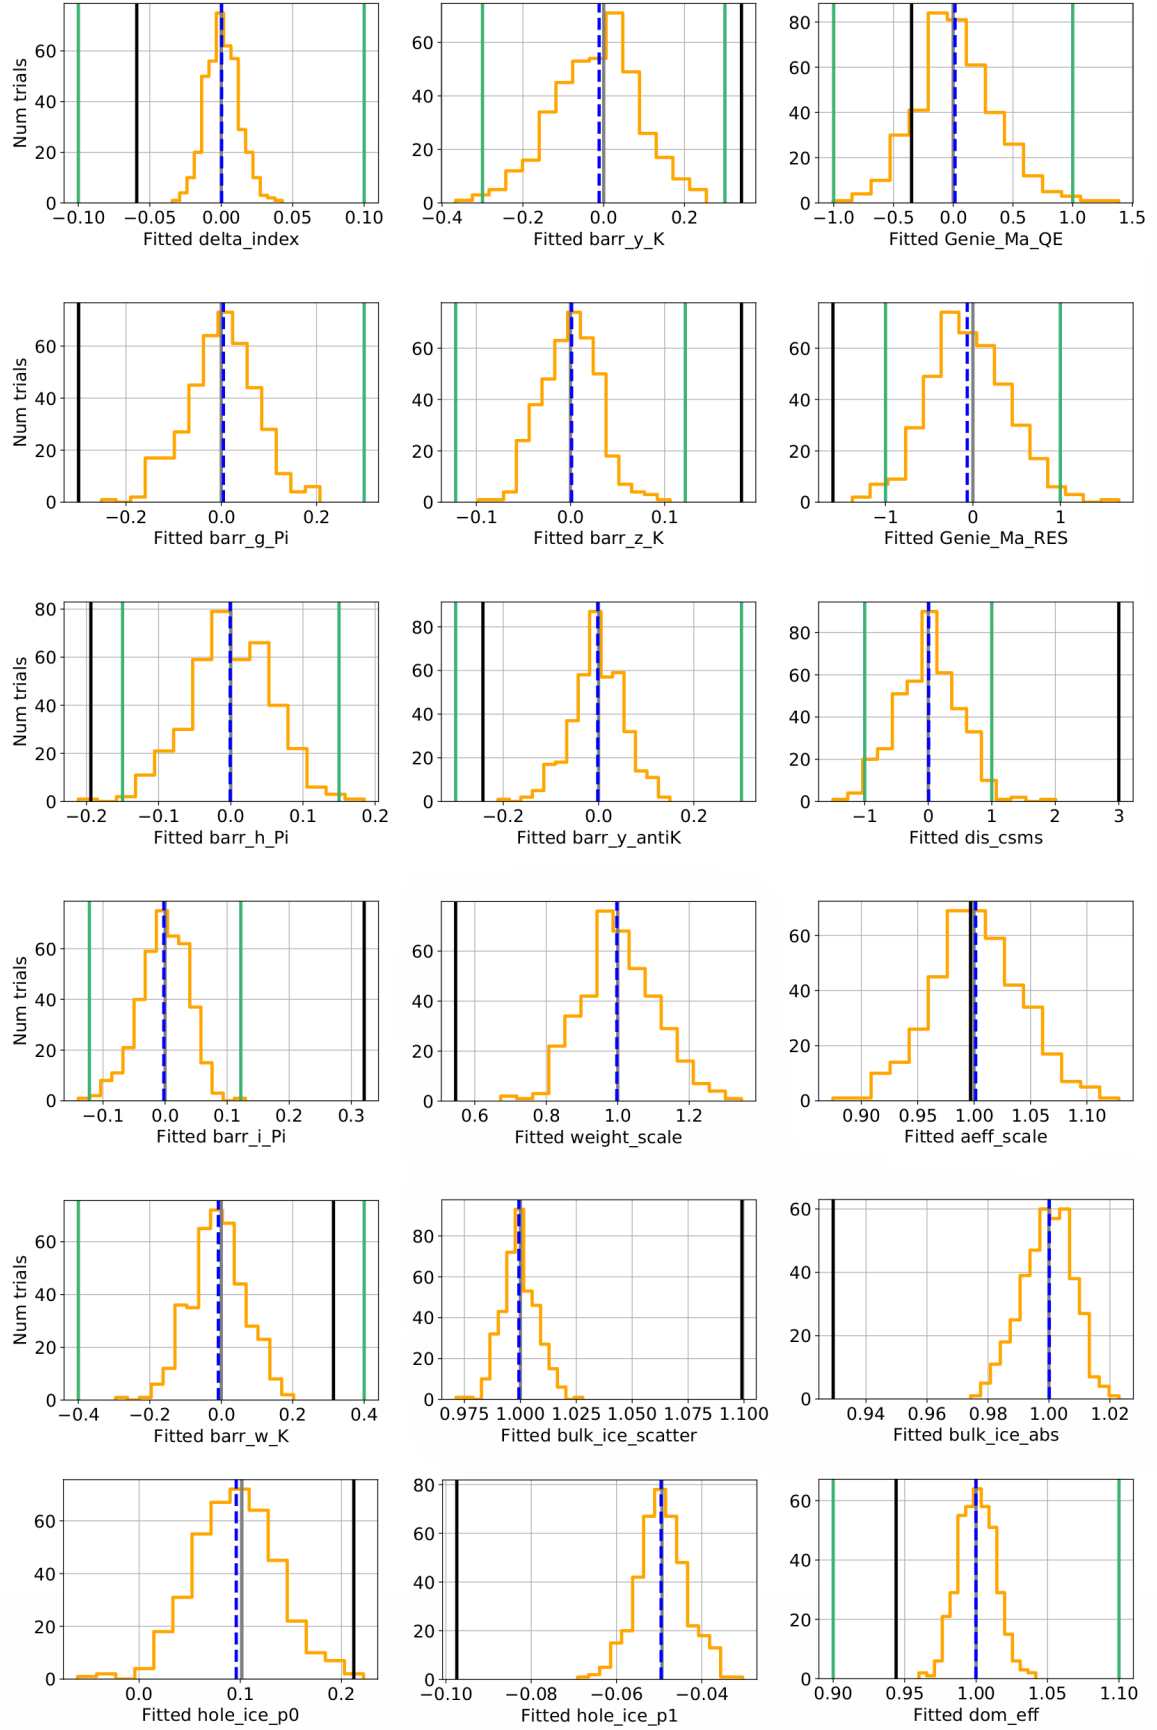


Figure 84: Blind fit result on the systematic uncertainty nuisance parameters included in the analysis.

The vertical solid black line indicates the value to which a parameter was fit, which can be compared to the distribution of fitted values obtained in the ensemble test (yellow histogram). Green vertical lines indicate the width of the gaussian prior, for the variables that had a gaussian prior defined. The oscillation parameter values (θ_{23} and Δm_{31}^2) and the tau normalization were kept blind at this stage.

Figure 85 shows the pulls recorded from the fit on all systematic uncertainties that had a prior assigned to their values. As is apparent from these plots, multiple systematics have large pulls beyond their expected prior: the probability of fitting only 5 out of 12 parameters within $\pm 1\sigma$ of their nominal values is relatively small ($\approx 5\%$). Moreover, the probability of fitting one parameter to extreme values of 3σ (as is the case for the DIS-CSMS parameter) is even smaller ($\approx 3\%$). This outcome strongly suggests that some form of mismodeling is affecting the outcome of the fit. Because of this, the tau neutrino normalization parameter, mixing angle, and mass splitting were therefore not unblinded.




Figure 85: Summary of the pulls obtained after the blind fit stage (green crosses), on each of the systematic parameter that had a defined gaussian prior. Several parameters pulled away from their nominal by more than 1σ , meaning that the physics result could not be unblinded. Blue dots indicate the pulls seen in the ensemble test.

In an attempt to better understand the origin of the discrepancy, additional ensemble tests and fits to data were performed using slightly different analysis configurations:

- First, a configuration in which the DIS systematic parameter is fixed to its nominal value, to determine if the parameter itself could explain the original p-value obtained.
- Second, a configuration in which the analysis is reduced to reconstructed energy bins below 54 GeV. This strategy had already been proven to be useful for improving the data-MC agreement of the previous tau appearance analysis.
- Third, individual fits to every detector season included in the sample, to identify any clear trend linking the fit p-values to known changes in the detector configuration.

The pre-trial p-values obtained for each of these extra fits are listed in Table 8. The table also includes the measured pulled on the DIS-CSMS systematic parameter, which still constituted the largest pull in most of the configurations tested⁵. As can be seen from this table, fixing the DIS systematic to its nominal value and reducing the energy range of the sample do not improve the goodness of fit significantly. Seasonal fits, meanwhile, do not conclusively indicate a connection with known changes in the detector configuration. While three later seasons appear to show better p-values than earlier seasons, no recorded change in the calibration of IceCube can explain a break

⁵ Note that the lower significance of the pulls in the yearly fits is less reflective of a better goodness-of-fit than the fact that these fits include less data than the full 8 year fit.

in 2016. Similarly, no coherent explanation can be given as to why the 2017 seasonal fit would be so much worse than its neighbouring seasons. Lastly, it can also be noted that good p-values are not associated with smaller DIS-CSMS pulls (and vice-versa), indicating that at least one other mismodeling source might be affecting the analysis.

Table 8: Outcome of additional fit configurations.

| Configuration Tested | Fit p-value | DIS pull |
|--------------------------------|---------------|--------------|
| Fixed DIS parameter | ≈ 0.0 | - |
| Reco Energies $< 54\text{GeV}$ | ≈ 0.0 | 3σ |
| Seasonal fit to 2012 | 0.002 | 2.5σ |
| Seasonal fit to 2013 | ≈ 0.0 | 0.07σ |
| Seasonal fit to 2014 | ≈ 0.0 | 2.2σ |
| Seasonal fit to 2015 | 0.002 | 0.91σ |
| Seasonal fit to 2016 | 0.222 | 1.1σ |
| Seasonal fit to 2017 | ≈ 0.0 | 1.3σ |
| Seasonal fit to 2018 | 0.088 | 0.45σ |
| Seasonal fit to 2019 | 0.004 | 1.1σ |
| Seasonal fit to 2020 | 0.384 | 0.63σ |

9.3 POTENTIAL ROADMAP TO UNBLINDING

Given the outcome of the blind fit, more investigation needs to be made on the analysis and on the event sample in order to resolve the tension seen in the many systematic parameters (particularly DIS-CSMS), and the disagreement in data/MC distributions following the fit. While this task goes beyond the scope of this thesis, two possible strategies could be attempted in order to resolve the problems mentioned in the last sections.

The first line of investigation that should clearly be pursued has to do with the current parametrization of the DIS CSMS uncertainty, which shows by far the largest pull away from its expected value. That pull appears to be independent of the season fitted, and is not improved by restricting the analysis to lower energies only. One reason that might explain this issue could be the choice of extrapolation that was used to link together the high-energy CSMS cross section prediction with the bulk of the high-energy cross section. The choice of using a constant extrapolation creates problems for scaling events with energies below 100 GeV, as the extrapolated ratio of CSMS to GENIE cross section is not assumed to be equal to 1.0. This treatment seems, in retrospect, to be physically unjustified, as there is no evidence that the disagreement between the two models remains constant at those energies. A better method for treating this uncertainty would be to introduce a smoothed step function at 100 GeV, which would transition from a full CSMS model to the GENIE one at low energy. The width of this step function could then be also included as a systematic uncertainty, which would thus encompass some of the uncertainty there is about the true nature of the transition between the two models.

A second strategy would be to change the simulation set that is used as the nominal template for one where the DOM efficiency is set to a value of 0.95. As was mentioned earlier, it was already suspected that DOM efficiency might have been overestimated in the nominal set, leading to biases in the location of the reconstructed vertices of the events in the selection. This might have had consequence in the selection process from section 7.2.5, as the muon classifier at Level 7 was trained on variables that mostly relied on event vertices. As a result, the final level events used in the analysis might not be representative of the true spatial distribution of neutrino events in data, which is some-

thing that can be difficult to fix using only event rate scaling such as hypersurfaces. A clear line of investigation here would be to re-train the Level 7 Muon classifier with the simulation set which is closest to the fitted dom efficiency (0.95), re-train subsequently the PID classifier on the new events and perform the fit again to data. This should in theory lead to a better agreements in the problematic plots mentioned in section 9.2.1.

Finally, a last observation can be made from having a closer look at how systematic uncertainties are correlated with each other (see Fig. 106 of Appendix B for more details). The mild anticorrelation of DIS-CSMS with the neutrino effective area might lead to pulls in other parameters, because the effective area itself is itself very correlated to some of the Barr parameters (See for instance Barr I and Barr H for pions). Hence it could be that the extreme pull in the DIS parameter is leading to compensating pulls from the other parameters in order to ease the tension with the data.

9.4 CONCLUDING REMARKS

- An event selection consisting of 8 years of neutrino data has been constructed. It covers reconstructed energy ranges of 5-300 GeV, and reconstructed zenith values smaller than 0.3.
- After a little optimization, events passing the selection have been binned across 20 energy bins, 10 zenith, and three PID bins.
- A multidimensional fit was designed, whereby the weights of a nominal Monte Carlo template are adjusted iteratively to best match the event count found in the detector data. Following a thorough investigation of the possible optimization metric, the Poisson likelihood was judged most appropriate to use in the analysis.
- The fit parameter of interest in this analysis is the tau normalization parameter N_τ , which scales up or down the contribution of tau neutrino events to the simulation template.
- A total of 21 systematic parameters are included in the fit, to account for uncertainties in the neutrino flux modeling, neutrino-nuclei interaction models, and the detector response.
- Blind fits of the Monte Carlo template to the data uncovered mild disagreement between the expected and observed event counts of the cascade PID bin. The fit to data returned a very poor goodness-of-fit, with values of several fitted systematic uncertainties pull much further away ($> 1\sigma$) from their nominal values than expected.
- The DIS-CSMS systematic uncertainty is likely not adequately designed based on its high pull value. It is however not the only culprit for the bad goodness-of-fit observed. The observed p-values from seasonal fits to data are inconsistent with known detector-configuration issues.
- As the result of the blind fits were not satisfactory, no meaningful measurements of the oscillation parameters were possible. However two potential solutions have been envisioned to resolve the issues observed at the blind fit stage.

Part III

2MRS CORRELATION SEARCH

NEUTRINO ASTRONOMY

The work presented so far in this thesis was aimed at understanding the behaviour and properties of neutrinos produced by cosmic rays. In a way, the origin of the physics that can be seen in DeepCore can be traced back to these charged particles that triggered the shower. But although we know a lot about the properties of this flux, the question of their origin is still unclear, particularly in regards to cosmic rays with energies greater than 1 EeV known as Ultra-High Energy Cosmic Rays (UHECRs).

Finding a plausible mechanism of production for UHECR is known to be a difficult undertaking [126]. Determining the source of these cosmic rays in the sky could help solve this issue by associating their site of production with known astrophysical phenomena seen by telescopes at various wavelengths. This chapter will describe how a putative relationship between UHECRs and neutrinos can provide the missing link between the sources of cosmic rays and the astrophysical reality perceived by traditional astronomy.

10.1 POTENTIAL SOURCES OF ULTRA-HIGH ENERGY COSMIC RAYS

Whereas no direct evidence exists to link UHECRs with a particular production process, the physics involved in accelerating a particle up to EeV energies do put some constraints on the type of astrophysical objects or structures that can be the source of cosmic rays. The main constraint was derived by Hillas in [127]: it connected the maximum energy at which a particle can be accelerated with the strength of the magnetic field used to trap the particle. Thus, potential UHECR sources are often proposed in the literature on the basis of the strength of the magnetic fields they can generate.

Supermassive Black Holes are one of these potential cosmic ray engines often cited [126]. These astrophysical objects are found at the core of galaxies, and are characterized by a matter density so high as to prevent light from escaping its vicinity. The environment surrounding black holes can provide plenty of interesting phenomenology to study, if not the properties of the black hole itself, the process by which matter that gets absorbed into it over time. This process known as *accretion* creates a high-energy environment in which streams of particles can collide with each other at energies far beyond anything that current particle accelerators are capable of generating.




Figure 86: Artist's depiction of an active galactic nuclei (AGN), with its associated relativistic jet. Both neutrino production and cosmic-ray acceleration is believed to occur in the shocks created within the jet by colliding blobs of particles. Image Credit: DESY, Science Communication Lab

Black holes with a particularly active accretion process are known as an *Active Galactic Nuclei*. The accretion flow of AGNs, which is depicted in Fig. 86, is often accompanied by *relativistic jets* at their center: it is within these jet structures that shock collisions between ultra-relativistic particles are believed to be capable of generating UHECRs[128].

It is important to note that all of the existing information about AGNs (and other potential cosmic ray engines) comes exclusively from light-based astronomy, which has its limitations (see the next section). Therefore, it cannot be excluded, at this point in time, that the true origin of UHECRs lies within unknown astrophysical objects beyond our electromagnetic sight.

10.2 LIMITATIONS OF PHOTONS IN ASTRONOMY

For centuries, the only accessible messengers for astronomers has been restricted to light. As a fast traveling, electrically neutral particle, the photon has proved to be quite useful for reconstructing the physical properties of objects in space. Indeed, most normal matter emits a certain amount of photons while interacting, which turns the vacuum of space into a large calorimeter where the energy from all physical processes is dumped. Moreover, photons conserve their original direction while they travel, which allows one to trace back the location of the engine that produced that light. Finally, given its high cross section, the photon is a particle that is easy to detect, collect, and characterize; astronomers have become quite good at collecting it over the entire range of energies at which photons are known to be produced.

Nevertheless, photon messengers do possess inconvenient traits. The main disadvantage comes from how photons interact with matter, particularly at high energies. Figure 87 shows the contribution of various processes to the interaction cross section of photons, both with light (C) and heavy (Pb) nuclei, as a function of energy. As photon energy increases beyond MeV energies, the domi-

nating contribution to the cross section becomes pair-production, a process by which the photon is converted into an electron-positron pair in the presence of matter. This process is destructive for the photon, and replaces the originally neutral messenger with two charged particles which can deviate from their original trajectory, due to the influence of galactic, interstellar, and local magnetic fields.




Figure 87: Interaction cross section between photons and two different atoms: carbon and lead. At the highest energies, interactions are dominated by pair production (the κ curves in both plots), which is a destructive process for the photon. Taken from Section 34 of [56].

The consequence of pair-production impacts the ability of seeing high-energy light in two ways: first, it means that most of the radiative information coming from high-energy astrophysical processes will have thermalized into lower energy products. Second, even the few gamma rays that manage to survive the initial environment of production will eventually pair-produce after propagating over large distances through the void of space, because that void actually contains a relatively large density of primordial, CMB photons. Given a sufficiently large propagation time, a gamma ray has a 100% probability of interacting with one of these CMB photons and triggering a pair-production event. Figure 88 shows the range of cosmological distances one can observe, with photons of a given energy. As can be seen, a large area of this parameter space, namely the one where UHECRs could be produced (energies above 10^{18} eV, located at extragalactic distances beyond $O(1 \text{ Mpc})$), are essentially inaccessible to photons.




Figure 88: Observable cosmological distances of the universe, as a function of energy. Particle interactions at ultra-high energies, located far away from Earth are not visible to light-based astronomy. Source: The IceCube Collaboration.

In light of all of this¹, the mystery of the origins of UHECR and its connection to AGNs can only be indirectly addressed with conventional astronomy, by correlating the arrival direction of cosmic rays with the visible low-energy component of the emission from those galaxies. Any direct information from the acceleration process requires the use of additional messengers, like neutrinos.

10.3 THE NEUTRINO / COSMIC RAY CONNECTION

As another lightweight, electrically neutral particle, the neutrino is a well-suited astrophysical messenger. Neutrinos also have smaller interaction cross section, which means that their mean free path is much longer than that of photons: this means that neutrinos probe deeper regions of space than light, thus circumventing the opacity problem mentioned in the previous section.

But for the purpose of this analysis, the most interesting feature about neutrinos is the direct connection they share with UHECRs: as the latter propagate across interstellar and intergalactic space, they interact with background photons, producing pions in the process:

$$p^+ + \gamma \rightarrow p^+ + \pi. \quad (114)$$

In this reaction, the pion produced can be either charged or neutral; the subsequent decay of charged pions will generate neutrinos, which can become tracers for the original direction of cosmic rays². Thus, neutrinos can serve as a clear evidence of hadronic interactions involving cosmic rays, bringing astronomers slightly closer to the physics responsible for UHECRs.

10.4 THE DIFFUSE ASTROPHYSICAL FLUX AND SEARCHES FOR NEUTRINO SOURCES

In 2013, IceCube discovered a flux of neutrinos that could not be explained by atmospheric production[129]. This discovery was made by studying a strict selection of 28 high-energy events satisfying

¹ ;)

² Neutral pions meanwhile will decay into gamma rays that will most likely pair-produce before reaching the Earth.

the High-Energy Starting Event (HESE) filter, characterized by a very bright imprint on the detector. Figure 89 shows the energy spectrum of the starting event neutrinos detected after 7.5 years of operations. As can be seen from the latter, a large fraction of the high-energy neutrino flux (above 6×10^4 GeV) cannot be accounted for by either muon background contamination or the conventional atmospheric neutrino flux, which is indicative of a flux of extra-terrestrial neutrinos.




Figure 89: Energy spectrum of the HESE event selection, for 7.5 years of IceCube data. The black points correspond to the number of observed events, while the coloured histograms represent the expected contribution of various physical process to the overall spectrum: atmospheric muons (purple), neutrinos from the atmospheric flux (orange), and the astrophysical component (yellow). That last component is fitted to the data, and assumes an isotropic flux of astrophysical neutrino shaped as a single power law. The shaded blue region is excluded from the fit. Source: [130].

Following the discovery of an astrophysical flux of neutrinos, the primary goal of IceCube naturally evolved toward identifying the origin of that flux, which amounts to locating *sources of neutrinos* in the sky. But how can a single experiment independently complete this task?

10.4.1 The Seven-Year Point Source Search

A first strategy to identify the source(s) of astrophysical neutrinos would be to look for statistically significant *clusters* of neutrinos in time and/or space, solely within the IceCube dataset. This involves comparing the strength of an observed cluster with the frequency at which clusters of the same strength are present in randomly scrambled datasets. If that frequency falls below a very strict threshold (i.e. if the observed cluster is an extremely rare occurrence in random scrambling), a discovery can be claimed.

One of the main attempts at finding spatial correlations between neutrino arrival directions was made using a sample of well-reconstructed, track-like events, spanning seven years of detector live-time[131]. The selection, which is referred throughout this set as the *Seven-Year Point Source (7yr-PS) Sample*, consists of track-like events only, which improves the precision of the reconstructed event direction at the expense of the sample's neutrino purity, by letting in more background muons than, for example, the HESE selection.

Since the background sources differ between the southern and northern sky, the selection criteria differ between events coming from different hemisphere. Events coming from the southern sky (with declination $\delta < -5^\circ$) are dominated by atmospheric muons; that background contribution is minimized at high energies (≈ 100 TeV), where the current models of the cosmic-ray flux is not expected to produce many muons³. In the northern sky, the background is dominated by mis-reconstructed atmospheric muons, which can be identified by poor goodness-of-fit values, and an irreducible background of atmospheric neutrinos. There, the selection criteria are less restrictive than in the southern sky.

In the original analysis performed on this dataset, a study searched for statistically significant clusters of neutrino events in the data. Figure 90 shows the result of the spatial correlation search performed on the sample. The map shows the significance obtained by calculating for every point in the sky a test statistic that was defined in [131]. The color scale represents the pre-trial significance obtained accross the sky, which give the probability at which a number n_s of signal neutrinos with a localized energy spectral index γ might can be found purely through random fluctuations of the neutrino arrival directions. The smaller the p-value, the more significant is the claim that a cluster has an astrophysical origin. Unfortunately, none of the clusters found in this search had significance corresponding to a discovery threshold, after correcting these p-values for the large trial factor involved in an all-sky search⁴.




Figure 90: Significance map of the 7-years Point Source Search, showing the pre-trial significance p-value obtained for modeling neutrino sources at every point in the sky. No significant neutrino cluster was found after applying trials factor corrections. From [131].

³ Note that this assumption is dependent on the cosmic-ray flux modeling.

⁴ One must correct the pre-trial significance to account for the fact that an hypothesis was tested at every point in the sky.

This example shows the difficulty of using such self-correlation methods, given the relatively low statistics of the astrophysical neutrinos detected and the large trial factor involved in performing a hypothesis test over the entire sky. Moreover, this technique is also susceptible to unknown biases related to how it is *assumed* that a purely random distribution of neutrino events is an accurate description of the null hypothesis: in reality, unmodeled anisotropies related to the cosmic-ray flux, or more fundamentally to the assumed distribution of matter in the universe could imply that our diffuse flux might actually be more diffuse in some regions of the sky than others...

This leads to an alternative strategy that can be employed to look for neutrino sources: exploit correlations between a neutrino dataset and other astronomical observations. The latter can be time-based, like periods where an AGN was seen in a flaring state by gamma-ray telescopes, or it can be spatial, like a catalog listing the position of nearby galaxies in the infrared bandwidth. Regardless, the additional information provided by multi-messenger partners can help boost the significance of a cluster observed in IceCube, and strengthens a discovery claim by providing corroborating evidence from more than one experiment.

The next chapter presents the outcome of a spatial correlation search, aimed at tying together neutrinos from the seven-year point source sample with galaxies mapped by the 2MRS catalog.

SEARCH FOR CORRELATIONS BETWEEN NEUTRINOS AND THE LOCAL UNIVERSE

The following section describes a search combining the two strategies discussed previously: the information from a spatial significance analysis of an IceCube dataset was used in a correlation analysis using a catalog of galaxies with redshift measurements. The result of this search was published in [2].

11.1 MOTIVATION

While both strategies have yet to provide clear evidence of a 5σ signal from a source in the accumulated IceCube dataset, they have nonetheless provided interesting results that limit what these sources might be, and where they can be located in the sky. We have evidence, for example, that the diffuse flux of astrophysical neutrinos cannot come entirely from gamma-ray bursts [132], nor from the known population of blazars seen by gamma-ray observatories such as Fermi [133]. Additionally, the absence of correlations between the diffuse flux and the galactic plane limits the contribution of galactic sources to the IceCube data [134, 135]. These results can lead us to look further into space, and wonder if neutrinos might be emitted in the *local universe*, roughly defined as a region of space of ≈ 500 Mpc in radius around Earth. This corresponds to the next logical place to look for neutrino sources, as Earth-based observatories would be more sensitive to emission from there than from places located further away (given that any source loses brightness by a factor of $1/r^2$).




Figure 91: Distribution of galaxies coloured by redshift. Source: [136].

A key fact about the local universe is that it is very *anisotropic*, as can be visible from Figure 91. The local universe corresponds to a size scale where large superclusters exist, leading to overdense regions that are visible by infrared observatories. This anisotropy is at the core of a *multiplet analysis*, which tests for the existence of low-luminosity neutrino sources that occur more frequently in dense regions of the local universe. Such sources will produce spatial clusters of neutrinos (referred to as *multiplets*) that directionally correlate with galaxy density [137].




Figure 92: Distribution of galaxy redshifts in the 2MRS catalog.

11.2 ANALYSIS METHOD

The multiplet analysis is sensitive to a directional excess of neutrino *multiplets* that correlate with the local galaxy density. It does so by comparing a pre-selected list of neutrino clusters from IceCube to a large catalog of galaxies, the 2MASS Redshift Survey [136] (2MRS). The latter is the result of decadal efforts made to obtain redshift measurements on nearly all galaxies previously identified by the Two-Micron All-Sky Survey (2MASS), a large, unbiased sky survey performed between 1997 and 2003 [138]. With the addition of redshift information, the 2MRS catalog adds a critical third dimension to our view of the galaxies surrounding us, opening the door to galactic superstructure studies [139], and matter density mapping. The survey contains the position, redshift, and K_s magnitude¹ of the galaxies of the local universe up to a redshift of $z \approx 0.10$. The distribution of redshift for all galaxies in the catalog is shown in Figure 92. Anisotropic features in this distribution have been used in other correlation searches, namely by the Auger collaboration [140, 141]. Above a redshift of $z = 0.03$, limitations in the telescopes' sensitivity limit the *completeness* of the catalog. Below that threshold however, the latter can be considered complete, meaning that it contains all luminous galaxies it is possible to see at those distances. For this reason, the multiplet analysis is thus restricting itself to galaxies with $z \leq 0.03$.

The data used to identify neutrino multiplets in IceCube comes from the seven-year sample of through-going events described in section 10.4.1. The starting point of the analysis was the significance map showed in Fig. 90. From that map, a list of multiplet locations had to be created, using the *hotspot* tool from HEALPix [142]. The algorithm located all local maxima within the significance map, but since the resolution of the latter was much higher than the true angular resolution of the events used, several maxima remained multiply-defined by groups of pixels. A cleaning algorithm is used to avoid double-counting a maximum: all pixels within a distance of 1.5° from the most significant local maximum were removed, and the procedure was recursively applied to the next most significant maximum remaining on the cleaned map. The distance cut of 1.5° was chosen to match the prescription that was followed in [131] to obtain a count of significance hotspots across both hemispheres; it relates the width of a 2D gaussian to the angular size of an equivalent circular bin.

From this initial selection process, local maxima are only kept as multiplets if their local pre-trial significance is greater than 2.0σ . Furthermore, the best-fit number of signal neutrinos at that location

¹ K_s refers to the brightness of an object as seen through a standardized filter bandwidth.

(the variable n_s in the 7yr-PS likelihood) is required to be greater or equal to 2. Finally, an additional cut on the local energy spectrum (the local spectral index $\gamma \geq -2.75$) is applied to filter out multiplets which are most likely coming from atmospheric neutrino background events (as astrophysical neutrinos are more likely to come from high-energy events). The choice of cut thresholds is based on an optimization study defined by:

$$m = \frac{D}{\sqrt{N}} \cdot \frac{\Gamma(n_s)}{\Gamma(n_s - 3/2)}, \quad (115)$$

where the optimization metric m is composed of three terms. The first one is the *discovery potential* D , defined by the number of signal multiplets required to achieve a discovery, while the second one is the statistical power of the selection, taken to be the square-root uncertainty associated with the total number of multiplets surviving the selection process N . The third term of the metric appears in Eq. 129 (see Appendix C) and describes how the sensitivity upper limit of the analysis scales with the multiplicity of the multiplet (which is approximated in this analysis by the variable n_s).

The combination of all three terms is chosen such that the optimization metric is optimal when it is minimized (one wants to minimize the discovery threshold, sensitivity upper limit, and statistical error). Figure 93 shows the value of this optimization metric for a set of 108 combinations of cuts tested. The minimal value obtained is identified by the red square. Values of 0.0 at the bottom of the graph corresponds to such harsh cuts on the spectral index that no multiplets are left in the selection.




Figure 93: Optimization metric value as a function of multiplet selection cuts.

Since the 2MRS catalog does not cover the Galactic Plane, the analysis excludes a region of the sky between $-10^\circ \leq b < 10^\circ$, where b is the galactic latitude. Furthermore, the regions around both celestial poles (with declinations $|\delta| < 85^\circ$) are also excluded, as they were not included in the 7yr-PS search [131]. Finally, a cutoff of $z \leq 0.03$ is applied to the redshift of the objects from the 2MRS catalog, since it becomes magnitude-limited beyond that distance.

The degree of correlation between IceCube multiplets and baryon density (as mapped by the 2MRS catalog) is evaluated on a pixel-by-pixel product of all N_m multiplets, using the following likelihood function:

$$L(n_a) = \prod_{i=1}^{N_m} \left(\frac{n_a}{N_m} S_i + \left(1 - \frac{n_a}{N_m}\right) B_i \right). \quad (116)$$

In this likelihood, the number of multiplets coming from the local universe n_a , normalized by the total number of selected multiplets N_m , is maximized for a given background (B_i) and signal (S_i) probability in all pixels i containing a multiplet. The signal term S_i of this likelihood is a normalized count of 2MRS objects contained in a HEALPix-defined pixel area of 3.36 square degrees. The full skymap of the signal PDF is shown in Fig. 94a. The background term, B_i , is assumed to be constant across the map, and is thus the fraction of the sky covered by a pixel on the map, i.e. $A_{\text{pixel}}/4\pi$.

The test statistic for the multiplet analysis is the likelihood ratio $\text{TS} = \text{Sign}(N_A) \cdot 2 \ln \left(\frac{L(n_a=0)}{L(\hat{n}_a)} \right)$, where the null hypothesis implies that $n_a = 0$ (i.e. that it cannot be excluded that the multiplet distribution is purely random). In order to properly account for statistical underfluctuations in the fitted number of multiplets, the usual likelihood ratio is modified by the sign of the fitted value of n_a , and the p-value used to quantify the significance of the signal is two-sided.

The significance of the correlation is estimated by comparing the value of the maximized likelihood ratio to the TS distribution calculated from a *scrambled* distribution of dataset: in this scrambling process, the original coordinates of each multiplet is randomized in right ascension (RA) to simulate the random distribution of neutrino clusters one would expect under the null hypothesis. Note that the declination coordinate of the multiplets are kept the same, to maintain the integrity of IceCube's zenith-dependent sensitivity.

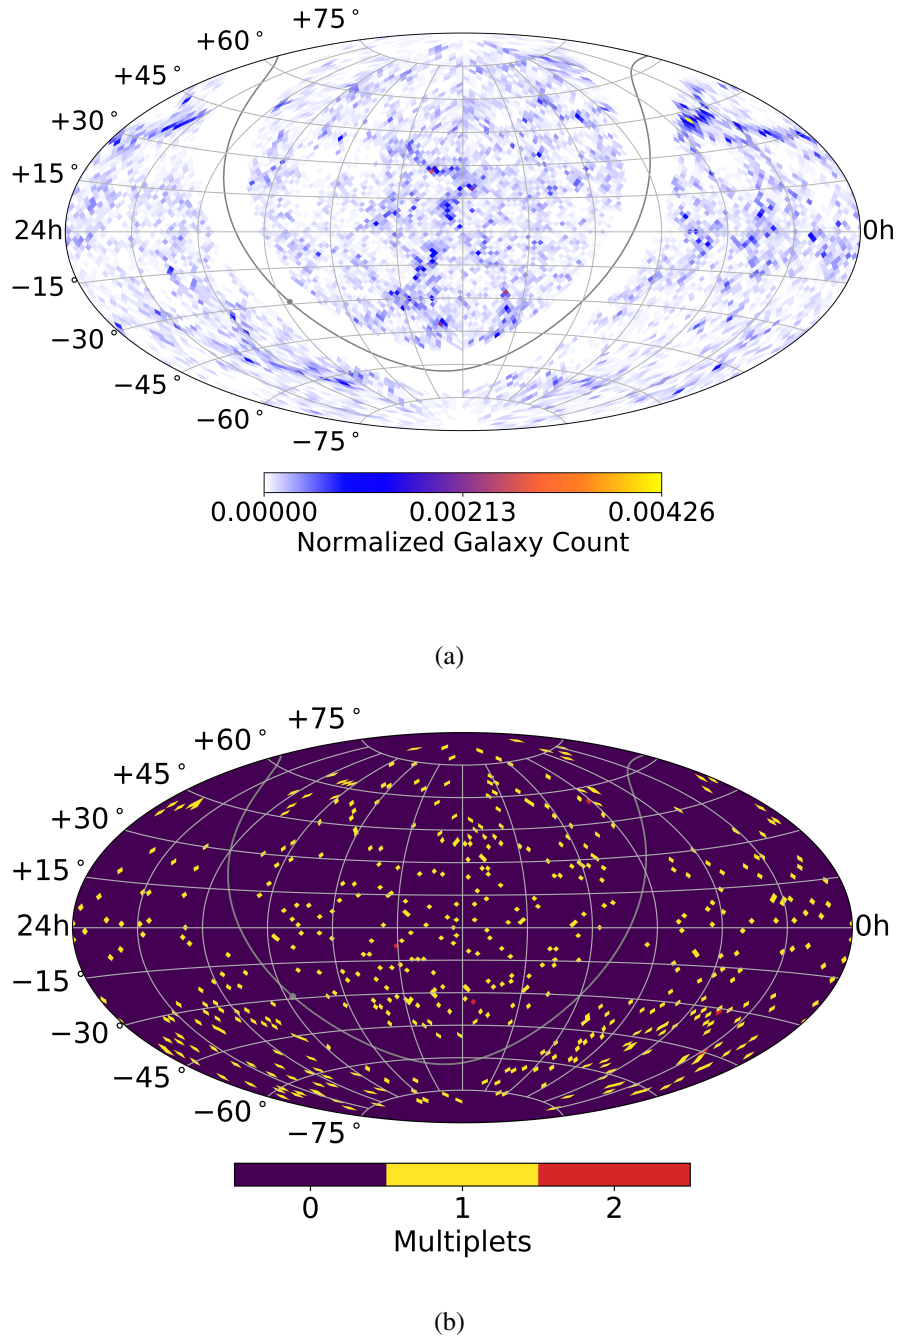


Figure 94: a) Normalized distribution of galaxies up to a redshift of 0.03, taken from the 2MASS Redshift Survey catalog [136]. b) Location of the selected subset of multiplets. Each yellow tile represents the location of a local maximum from the 7yr-PS significance map which satisfies the selection criteria of $n_s \geq 2.0$ and $\gamma \geq -2.75$. Red tiles denote the five instances in which two multiplets fall into the same pixel. These are therefore counted twice in the likelihood calculation.

11.3 RESULTS

Figure 95 shows the test statistic distribution for the null hypothesis constructed out of the scrambled maps, along with the unblinded value of the test-statistic obtained when running the analysis on data. The unblinded analysis returned a value of 0.238, with a corresponding p-value of 80%, which is

consistent with the null hypothesis of multiplets being uncorrelated with the 2MRS catalog.




Figure 95: Test-statistic distribution of the null hypothesis, where neutrino multiplets are scrambled randomly in right ascension. The vertical green line corresponds to the observed test statistic value obtained in the fit to data.

Given the absence of a detectable population of neutrino sources correlated with the selected multiplets, an upper limit is placed on the number of multiplet originating from a 2MRS galaxy N_A , using the Feldman-Cousins prescription from [143]. The interval is constructed by fitting N_A to multiple pseudo-experiments, built with a range of injected true signal.

This process is illustrated by the 2D histogram shown in Figure 96: each row of the histogram corresponds to a distribution showing the value of N_A fitted for multiple pseudo-experiment trials. These pseudo-experiments are constructed by substituting a number μ of multiplet locations from the original selection by locations drawn out of the 2MRS signal map². The background events are then randomly scrambled in RA, which allows one to produce a distribution of expected N_A fit values for various levels of true signal strength. Those distributions determine the color scale of each rows of Fig. 96; 10000 trials are performed for each signal event hypothesis (i.e. there are 10000 trials in every row of the plot).

Confidence intervals are then constructed to determine the range of true μ values that would be covered 90% of the time, given an observed value of n_a . This is done for each column of Fig. 96, where the particular values of μ to include into the interval are chosen based on the likelihood ratio:

$$R = \frac{P(x|\mu)}{P(x|\mu_{best})}, \quad (117)$$

² This ensures that the total number of multiplets considered remains the same in every trial

where, given a potentially observed value x , the rank R compares the probability of measuring x for a given true signal μ with the *most probable probability* of observing x , a number that is associated with a specific value of true signal μ_{best} . In this analysis, these probabilities are calculated as follow:

$$P(x|\mu) = \frac{N_x(\mu)}{\sum_{\mu=0}^{\infty} N_x(\mu)}, \quad (118)$$

where $N_x(\mu)$ is the number of trials for which value x is observed when injecting signal μ in the pseudo experiments. Once the total probability obtained by the selected interval of μ reaches the desired coverage of 90%, no more values are added. The result of this procedure is the set of black shaded boxes in the background of Fig. 96. These intervals can be compared with a traditional Neyman interval, shown as the white lines in the same figure: these are constructed by determining the 90% central containment limits on the distributions of fitted n_a (the rows of the figure). In cases where the observed/fitted values are small or negative, using the Neyman construct leads to overconservative limits. In the end, an upper limit of 21 on the number of signal multiplet was found by intersecting the unblinded n_a value from data (the vertical red line) with the black box confidence intervals appearing in the background of the figure.




Figure 96: Feldman-Cousins interval to determine the upper limit on the number of signal multiplets. The colored map shows the distribution of fitted n_a values obtained in pseudo-experiments generated with different injected true signal μ . Background black boxes show the resulting Feldman-Cousins 90% confidence interval for each possible observable value of n_a , while the red lines mark the same confidence interval if one were to use a Neyman construct. The unblinded value is shown by the black line.

Using the relationships defined in Appendix B, the limit on the number of signal multiplet can be converted into limits on the density n_0 and neutrino luminosity L_ν of the hypothesized population of sources. Figure 97 shows the 90% limit in this parameter space, derived using Eq. 2.17 of [137]. This equation relates the best-fit value of N_a obtained with the probability of seeing $n \geq 2$ neutrinos from

sources distributed within a comoving volume reaching out to the boundary redshift of $z = 0.03$. This probability can be analytically derived if one makes certain assumptions regarding the nature of the unseen signal from 2MRS galaxies:

- The hypothesized flux comes from a uniform population of sources (described by a single energy spectra and intrinsic luminosity),
- The detectable neutrino emission from these sources follows poisson statistics,
- The source density of the hypothesized population of sources is such that its neutrino flux falls just short of being detected by IceCube as localized multiplets.

In the construction of the limit on Fig. 97, the energy spectrum of the sources is assumed to be a power law with an index of -2.2, corresponding to the best-fit spectral index of the diffuse astrophysical neutrino flux [144]. The limit is also tuned to the sensitivity of the event selection used, which defines the median distance at which single neutrinos can be observed and thus the location of the turnover point in the figure. Beyond this turnover point, corresponding to neutrino luminosities of $L_\nu \approx 10^{42} \text{ erg} \cdot \text{s}^{-1}$, sources at distances beyond $z = 0.03$ are expected to contribute to the observed multiplet population. Taking the conservative assumption that the universe becomes isotropic beyond that point, the multiplet analysis is then expected to become less sensitive, as correlations with the 2MRS catalog are expected to wash out.

The multiplet search limit can be compared to the n_0 - L_ν limit one would obtain from the non-observation of statistically significant neutrino clusters, which has been studied by [137] on the 7yr-PS dataset, and by [145] in an earlier IceCube sample. Using a multiplet correlation technique instead of a full sky cluster search is slightly more advantageous in the case where the observed flux of neutrinos comes from a population of low-luminosity sources: this is illustrated in Fig. 97 by the fact that the multiplet limit is lower than that of the 7yr-PS search, for luminosities smaller than $10^{42} \text{ erg} \cdot \text{s}^{-1}$.

Both limits can also be compared to the lower (higher) density of source population expected, if 1% (100%) of the diffuse flux observed by IceCube in [144] came from 2MRS objects. This range is illustrated by the shaded bands in Fig. 97. The green band assumes that the density of neutrino sources does not evolve with redshift, while the red band incorporates a correction factor to account for a more plausible evolution, as derived in [146]. Given that both bands fall below the multiplet sensitivity line, it can be concluded that this search would not have been able to detect any sources, if it turned out that the origin of the diffuse flux was solely created by a homogeneous population of sources in the local universe.




Figure 97: Limits on the density of neutrino source populations in the local universe, as a function of that population’s neutrino luminosity L_ν . The analysis upper limit is shown by the green line, and can be compared to the 90% sensitivity (blue) and discovery potential (pink). The solid black line is a recalculation of the limit derived in [137], using a critical redshift matching the bound used in selecting 2MRS objects ($z=0.03$). The green band shows the parameter space covered if 2MRS objects were generating between 1% to 100% of the observed IceCube diffuse flux. The red band adds a correction factor, if one assumes a redshift evolution of neutrino sources that would match currently observed star formation rates (SFR).

11.4 CONCLUDING REMARKS

- A time-integrated correlation search was performed between a selection of neutrino multiplets detected in IceCube and galaxy density map derived from the 2MRS astronomical catalog.
- Neutrino multiplets were selected based on the significance map of a previous analysis, which used a sample of well-reconstructed track events with seven years of data livetime.
- No significant correlations were observed between the two datasets.
- Limits on the properties of a hypothetical population of neutrino sources located within the local universe were derived from the absence of observed signal

Neutrino astronomy could play a major role in explaining interesting phenomena such as the origin of ultra-high energy cosmic rays, but in order to do so, it must be able to identify the sources of the astrophysical neutrino flux. The correlation search presented here was an example of such search, but unfortunately it did not lead to significant evidence of neutrino sources hiding in the local universe.

The absence of signal multiplets may suggest that some of the assumptions regarding the uniformity of the neutrino source population are oversimplified. Alternatively, the result could also be seen as additional evidence that neutrino sources might be located further away than the local universe.

Since the first evidence for a neutrino source was the detection of neutrinos correlated with a flaring blazar located at much larger redshift of $z=0.33$ (≈ 1.75 Gpc)[147][148], it could very well be that IceCube observes mostly neutrinos coming from very distant objects.

Part IV

CONCLUSION AND FUTURE WORK

CONCLUSION

"Ma thèse? Ça été de l'excavation, du labourage, du défrichage..."

Adapted from François Pérusse

This thesis presented work done in three different areas of experimental neutrino physics: a characterization of the noise properties of the IceCube DOMs, a neutrino oscillation analysis aimed at measuring the appearance of tau neutrinos in the DeepCore sub-array, and a search for correlation between a dataset of high-energy astrophysical neutrinos and a catalog of galaxies representing baryonic matter distribution in the Local Universe.

An extensive laboratory characterization of the noise production mechanism inside the IceCube DOMs showed that the simulated modeling qualitatively matched the observed timing distribution of noise pulses in the low- Δt region. A more quantitative treatment of noise systematics within the oscNext event selection, meanwhile, showed that noise did not pose a significant problem at analysis-level.

The tau neutrino appearance analysis endeavoured to test the assumption that neutrino oscillation involves only three flavours of particles. A measurement of tau neutrino appearance was done by finding the ν_τ fraction contained in an event sample covering 8 years of data from the IceCube-Deepcore experiment. The analysis required the development of a new event selection making use of new reconstruction techniques and machine-learning classification tools to reduce the levels of background events in favour of high-quality signal neutrinos. The analysis itself consisted of a Poisson-likelihood minimization procedure involving a fit of a tri-dimensionnal, binned Monte Carlo template representing the expected event counts in reconstructed energy, zenith angle and PID. The fit included 21 free parameters in addition to the tau normalization, to account for various systematic uncertainties related to the neutrino flux, oscillation probabilities, interaction cross sections and detector response.

Unfortunately, results from the blind fit phase showed a very poor goodness-of-fit, and many systematic parameters fitted to values that significantly pulled away from their expected priors. Some disagreement was also observed between the PID distribution of the best-fit Monte Carlo and detector data. These issues prevented a trustworthy unblinding of N_τ and its associated muon disappearance parameters θ_{23} and Δm_{31}^2 . A strategy for solving these issues was proposed, involving a re-training of the Level 7 classifier to simulation produced with lower DOM efficiency, and a re-engineering of the systematic parameter used to model uncertainties in the DIS cross section.

Finally, a search for correlations between neutrinos of the Seven-Year Point Source Sample and galaxies of the 2MRS surveys was performed, using a multiplet technique that restricted the search to regions in the sky where the high-energy neutrinos of the sample were the most clustered. No significant correlations were found between the galaxy catalog and the selected list of multiplets. Limits on the density of low-luminosity neutrino sources in this region of space was derived for a

wide range of potential intrinsic luminosity.

As can be seen from part II, work in the immediate future need to be devoted to improving the oscNext event selection to reach better agreement between the nominal Monte Carlo expectation and the detector data. Additional work is also needed to better model cross section uncertainties, work that might need to be extended to other systematics if it turns out that DIS-CSMS is not the sole reason for the large pulls seen in the blind fits. Though it was certainly not the expected outcome, the result of the ν_τ appearance analysis certainly underline the challenges posed by high-statistics samples. With more than 8 years of cumulated livetime collected, one can no longer hide behind statistical uncertainties to avoid treating systematic uncertainties properly, something which is becoming increasingly true of other areas of physics probed by the IceCube Collaboration. Though greater is the pain for the analyzer, the greater is the reward in the end: overcoming the problems in this analysis could after all lead to the world's best constraints on N_τ which, combined with results from the future generation of neutrino detectors, could bring us closer to solving the question of neutrino unitarity.

Part V

APPENDICES

FULL RESULTS OF THE NOISE MEASUREMENT STUDY AT SNOLAB




Figure 98: Distributions of hit times between successive noise pulses in the four tested DOMs, at -20°C




Figure 99: Distributions of hit times between successive noise pulses in the four tested DOMs, at -30°C

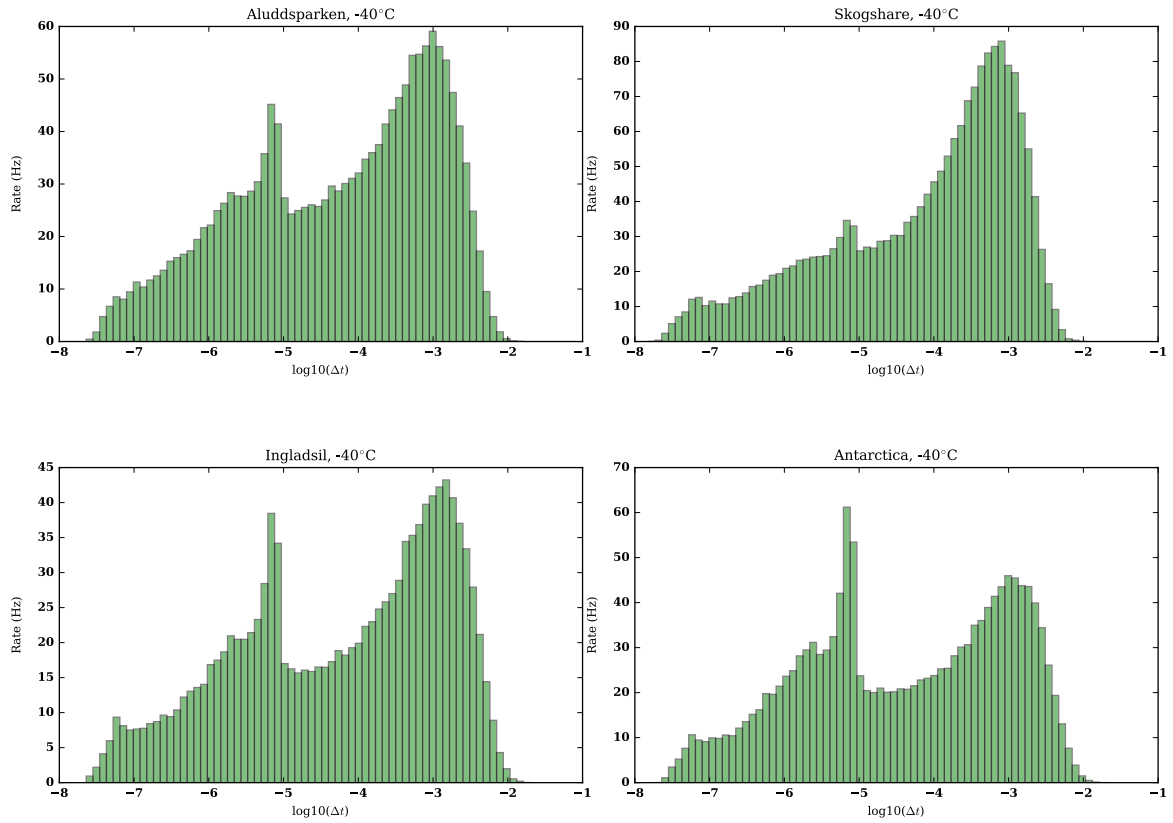


Figure 100: Distributions of hit times between successive noise pulses in the four tested DOMs, at -40°C

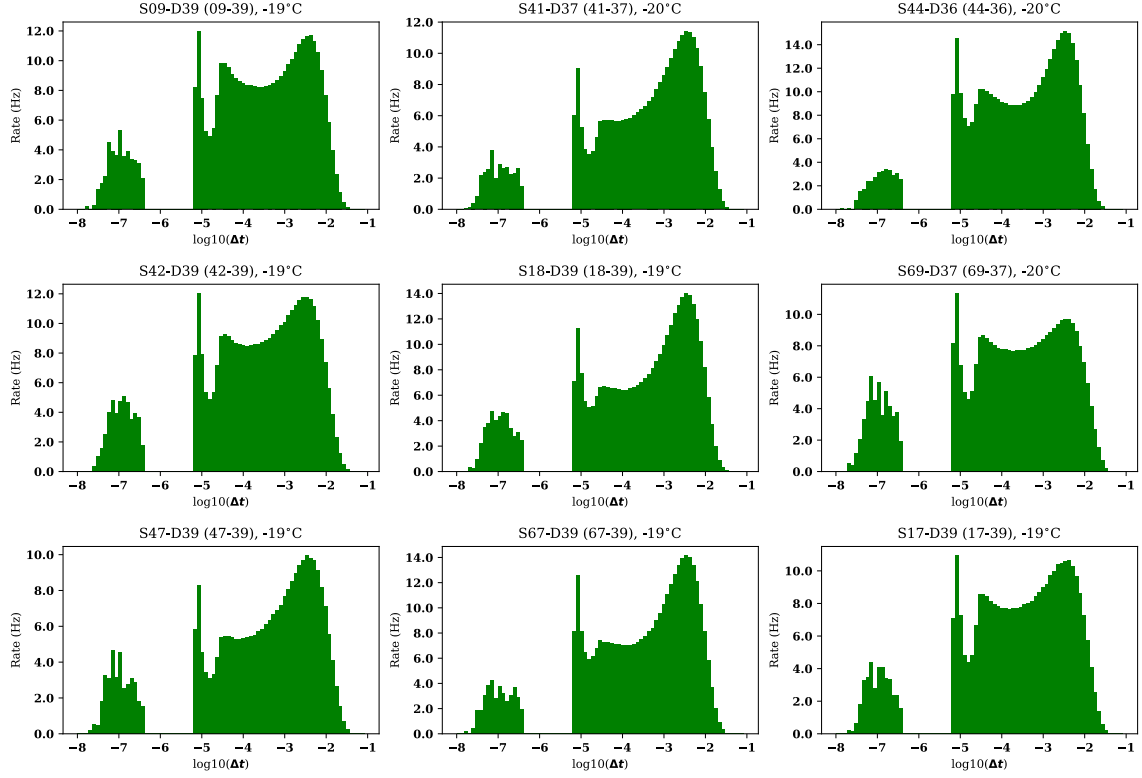


Figure 101: Distributions of ATWD hit times between successive noise pulses in simulated DOMs, at -20°C




Figure 102: Distributions of FADC hit times between successive noise pulses in simulated DOMs, at -20°C

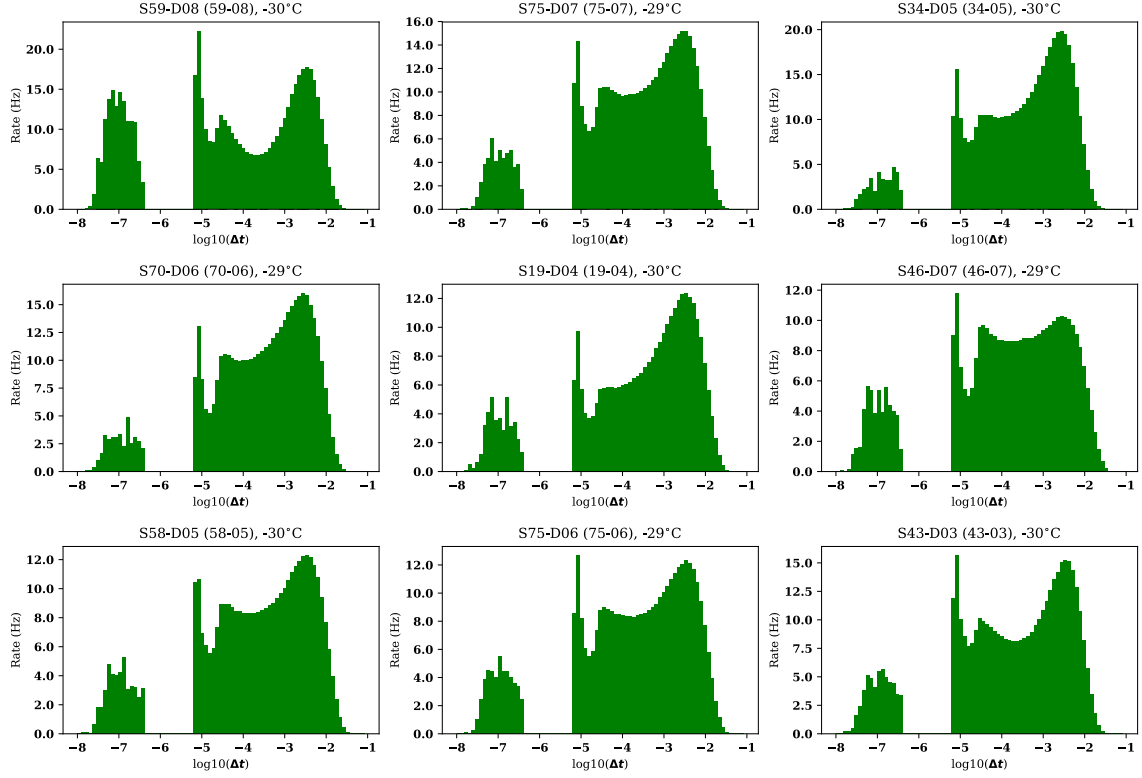


Figure 103: Distributions of ATWD hit times between successive noise pulses in simulated DOMs, at -30°C




Figure 104: Distributions of FADC hit times between successive noise pulses in simulated DOMs, at -30°C

ADDITIONAL INFORMATION ABOUT THE TAU APPEARANCE ANALYSIS

B.1 DEFINITION OF THE VETO-IDENTIFIED CAUSAL HIT REGION

Figure 105 illustrates the causality region used by the VICH variable in oscNext. Given a reference pulse time and location (which is taken to be the pulse that is closest to the trigger time), every other pulse of the event can be placed in a distance/time plot such as this one. If a DOM has at least one pulse located in the red causal region of the figure, this DOM counts as a veto-identified hit. This means that the DOM triggered at least once in a space-time region that would be causal with a physics event arriving into DeepCore.




Figure 105: Region of causality defined by the Veto-Identified Causal Hit (VICH) variable. The reference time is the event trigger time (t_{trig}), while the reference distance is calculated with respect to the DOM that triggered closest to the trigger time (r_0). Regions that are not considered valid include acausal regions, causal regions corresponding to events going out of DeepCore, and noise events occurring far away from the trigger. Source: [80].

B.2 FULL LIST OF SYSTEMATIC UNCERTAINTIES

Table 9 presents the full list of systematic parameters considered in the analysis, including their nominal value and priors used in the fit to data. Meanwhile, Figure 106 shows the correlation matrix of all parameters that were allowed to float in the analysis.




Figure 106: Correlation matrix showing the observed correlation coefficients between each pair of systematic uncertainties, in the ensemble test.

Table 9: Systematic uncertainties considered in the analysis

| | Parameter Name | Nominal Value | Prior | Full range | Fixed? |
|-------------|-------------------------|--------------------------|------------------|----------------|--------|
| | Neutrino Effective Area | 1.0 | uniform | [0., 3] | No |
| Flux | Neutrino Spectral Index | 0.0 | gaussian (0.1) | [-0.5, 0.5] | No |
| | Barr Y (K) | 0.0 | gaussian (0.3) | [-1.5, 1.5] | No |
| | Barr H (pi) | 0.0 | gaussian (0.15) | [-0.75, 0.75] | No |
| | Barr I (pi) | 0.0 | gaussian (0.122) | [-0.61, 0.61] | No |
| | Barr W (K) | 0.0 | gaussian (0.4) | [-2, 2] | No |
| | Barr Y (antiK) | 0.0 | gaussian (0.3) | [-1.5, 1.5] | No |
| | Barr G (pi) | 0.0 | gaussian (0.3) | [-1.5, 1.5] | No |
| | Barr Z (K) | 0.0 | gaussian (0.122) | [-0.61, 0.61] | No |
| | Barr D (pi) | 0.0 | gaussian (0.3) | [-1.5, 1.5] | Yes |
| | Barr W (antiK) | 0.0 | gaussian (0.4) | [-2, 2] | Yes |
| | Barr E (pi) | 0.0 | gaussian (0.05) | [-0.25, 0.25] | Yes |
| | Barr Z (antiK) | 0.0 | gaussian (0.122) | [-0.61, 0.61] | Yes |
| | Pion +/- pion - ratio | 0.0 | gaussian (0.05) | [-0.25, 0.25] | Yes |
| | Pion +/- pion - ratio | 0.0 | gaussian (0.05) | [-0.25, 0.25] | Yes |
| Propagation | Deltam31 | 2.38e-3 GeV ² | Uniform | [0.001, 0.007] | No |
| | Theta13 | 8.61 | gaussian (0.13) | [8.22, 8.98] | Yes |
| | Theta23 | 45.6 | uniform | [0, 90] | No |
| | DeltaCP | 217 degrees | gaussian (40) | [0., 360] | Yes |
| Interaction | MA (QE) | 0.99 GeV | gaussian (0.198) | [-40%, 40%] | No |
| | MA (RES) | 1.12 GeV | gaussian (0.224) | [-40%, +40%] | No |
| | DIS CSMS | 0.0 | gaussian (0.1) | [-3, 3] | No |
| Detector | Hole ice p0 | 0.101569 | uniform | [-2, 1] | No |
| | Hole ice p1 | -0.049344 | uniform | [-0.2, 0.2] | No |
| | Dom Efficiency | 1.0 | gaussian (0.1) | [0.8, 1.2] | No |
| | Bulk ice sca | 1.0 | uniform | [0.95, 1.15] | No |
| | Bulk ice abs | 1.0 | uniform | [0.9, 1.1] | No |
| Background | Muon spectral index | 0.0 | gaussian (1.0) | [-3, 3] | Yes |
| | Muon Effective Area | 1.0 | uniform | [0, 3] | No |
| Physics | Tau normalization | 1.00 | uniform | [0, 3] | No |

B.3 DEFINITION OF THE EXTENDED POISSONIAN LIKELIHOOD CONSIDERED

This section describes the difference in the mathematical definitions of the Poissonian likelihoods mentioned in Section 8.4.1 of Chapter 8. As it was mentioned there, each likelihood has the form:

$$L_{bin\ i} = \int_0^\infty Poi(k_{data,\ i}; \lambda_i) \cdot P(\lambda_i) d\lambda_i, \quad (119)$$

Where $P(\lambda_i)$ is the prior encoding the likelihood of the Monte Carlo template to give, in bin i , an expectation count λ_i . The *effective* (mcllh_eff) and *mean* likelihoods [122] use the prior:

$$P(\lambda_i) = \Gamma(\lambda_i; \alpha, \beta), \quad (120)$$

Where Γ denotes the generalized gamma function, which depends on two parameters α and β . In the effective likelihood we have that:

$$\alpha = \frac{(\sum_n w_n)^2}{\sum_n w_n^2} + 1, \quad (121)$$

$$\beta = \frac{\sum_n w_n}{\sum_n w_n^2}, \quad (122)$$

Where $\sum_n w_n$ is the sum of individual event weights n in a bin of the analysis. For the mean likelihood, β is the same as Eq. 122 and $\alpha_{mean} = \alpha_{eff} - 1$. The third formulation is referred to as the *generalized Poisson likelihood*, and defines its prior as:

$$P(\lambda_i) = [\Gamma(\lambda'_{i,1}; \alpha_1, \beta_1) * \dots * \Gamma(\lambda'_{i,N_s}; \alpha_{N_s}, \beta_{N_s})] (\lambda_i) d\lambda_i. \quad (123)$$

In this formulation, $P(\lambda_i)$ is defined as a convolution of N_s generalized gamma functions, where N_s corresponds to the number of simulated particle sets included in the simulation. The origin of this formulation comes from work in [123] and [124]. Each gamma function has associated α_j and β_j parameters, which are defined as:

$$\begin{aligned} \alpha_j &= \frac{(\mu_{k_{mc}})^2}{Var[Z]}, \\ &= \frac{(\mu_{k_{mc}} \cdot E[W])^2}{\mu_{k_{mc}} \cdot (Var[W] + (E[W])^2)}, \end{aligned} \quad (124)$$

$$\begin{aligned} &= \frac{\mu_{k_{mc}} \cdot (\sum_i w_i)^2}{k_{mc} \cdot \sum_i w_i^2}, \\ \beta &= \frac{\mu_{k_{mc}} \cdot E[W]}{\mu_{k_{mc}} \cdot (Var[W] + (E[W])^2)}, \\ &= \frac{\sum_i w_i}{\sum_i w_i^2}. \end{aligned} \quad (125)$$

Where $E[W]$ and $Var[W]$ denote respectively the expectation value and variance of the weight distribution associated with particle weights of type j . The other variable $\mu_{k_{mc}}$ is a bias-correcting variable described in the next section.

B.3.1 Corrections specific to the generalized likelihood

The generalized likelihood still neglects additional uncertainty associated with the finite MC statistics, which is related to the number of Monte Carlo events k_{mc} obtained in each bin. Particularly, in the definition of the parameters α_j and β_j used, the expectation value of the number of MC events ($\mu_{k_{mc}}$) is simply approximated as $\mu_{k_{mc}} \approx k_{mc}$, the observed number of MC events in this particular realization of the simulated weight distribution. This approximation is biased when the simulation statistics are low (which tends to happen in several bins on particular simulation sets), and to correct for this bias, an adjustment is made to the mean such that:

$$\mu_{k_{mc}} = \begin{cases} k_{mc} & \text{if } \langle k_{mc} \rangle_{all \text{ bins}} \geq 1 \\ k_{mc} + (1 - \langle k_{mc} \rangle_{all \text{ bins}}) + 0.01 & \text{if } \langle k_{mc} \rangle_{all \text{ bins}} < 1 \end{cases} \quad (126)$$

Where the value of the mean adjustment depends on the average number events in all bins, $\langle k_{mc} \rangle_{all \text{ bins}}$, for a given dataset. Since different datasets can have very different statistics in any particular bin, the adjustment to the mean is performed on each dataset individually. Note that the particular choice of threshold to switch between case has been taken ad hoc from [124], and is rather arbitrary.

Finally, cases involving empty bins can often be problematic, especially if a bin has many events from one simulated dataset, but none from the others. This situation can lead to biases and computational issues with the gamma functions, as the parameters α and β are required to be strictly positive and finite. To solve these problems, [124] advocated for the use of a pseudo-noise weight to fill the empty bins that have at least one Monte Carlo event in them from another dataset. The value of that weight was suggested to be defined, for each dataset individually, as the maximal value of all individual event weights. However, this created large bias issues in the tau appearance analysis, because the muon dataset simulated contains several instances of weights with values that were much larger than any other weights in the simulation. Therefore, instead of using the maximum weight value, the pseudo-noise term has been defined as the mean weight of the first 90 percentiles of the weight distribution.

B.3.2 Toy Monte Carlo study of the optimization metrics

An implementation of the generalized likelihood defined in the previous section was integrated into the PISA analysis framework. To validate the implementation, a simple one-dimensional case similar to the one performed in [124] was tested, and compared with the other metrics available for minimization (that is to say, the modified χ^2 , the effective and mean log-likelihood and the conventional Poisson likelihood).

In this toy Monte Carlo used as a test, pseudo-data is generated as a sum of a uniform background and a gaussian-shaped signal with true parameters $\mu_{true} = 20$ and $\sigma_{true} = 3.1$ (black histograms in Figures 107a to 107d). That pseudo-data is then fit with simulation templates constructed out of two simple PISA stages: one that populates the analysis histogram with randomly distributed background and signal weights, and another which re-weights the signal events according to a gaussian PDF with parameters $\mu_{MC} = 20$ and $\sigma_{MC} = 3.1$. Throughout these tests, the mean of the MC template μ_{MC} can be varied either manually or in the course of a fit to pseudo-data. The Monte Carlo events already created are thus re-weighted in a process that is identical to the one used in the real analysis.

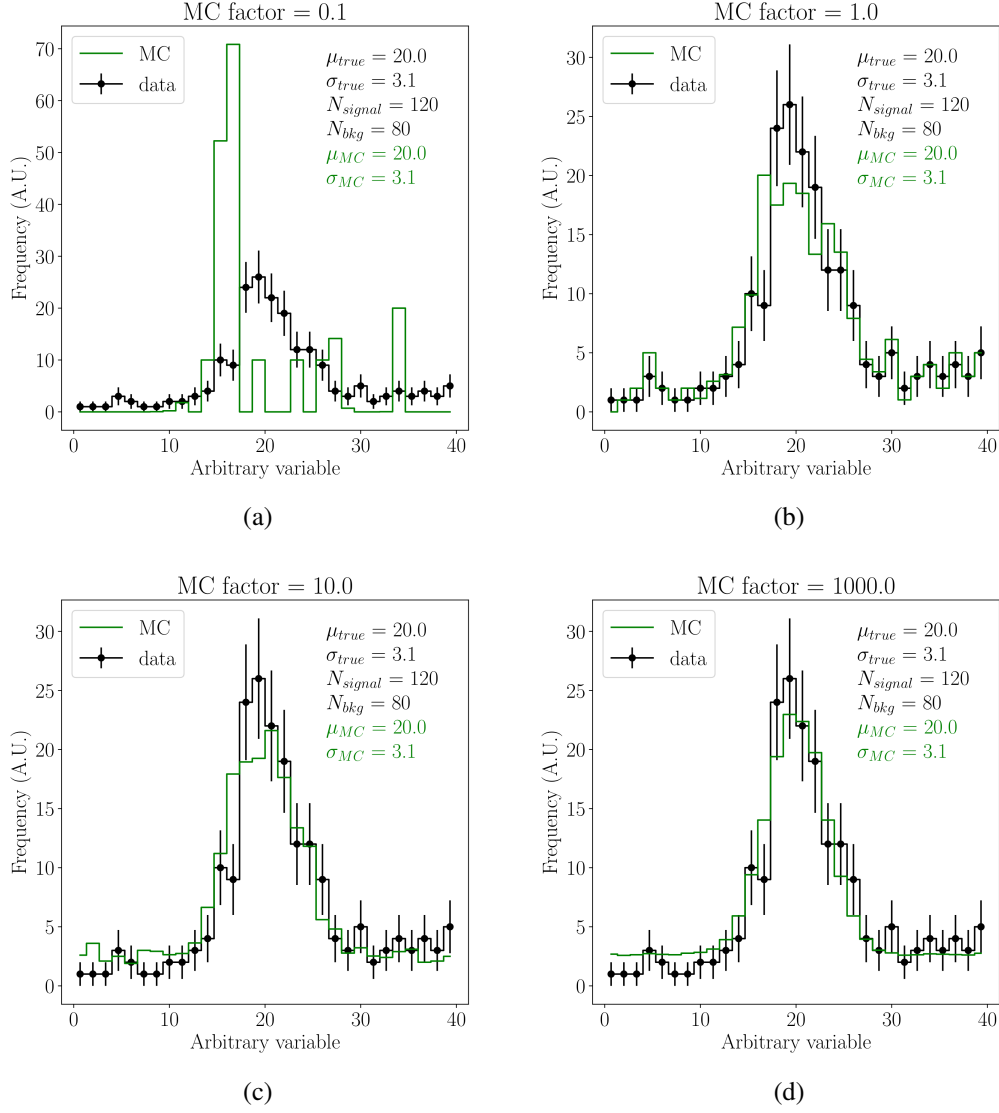


Figure 107: Histograms showing the distribution of events in the 1D Toy Monte Carlo test case. A total of 500 pseudo-data events are binned in a histogram (in black in all the figures), and compared to simulated Monte Carlo templates produced with 0.1 (a), 1.0 (b), 10 (c) and 1000 (d) times the number of events in the pseudo-data. The resulting weighted MC template is shown as a green histogram in each of these plots.

Figure 107 shows the distribution of 500 pseudo-data events binned in a one-dimensional analysis binning. That distribution is then compared to MC templates built out of varying levels of MC statistics: Fig. 107a to 107d shows how these templates (plotted in green) look like, for cases where the MC statistics reaches 10%, 100%, 1000% and 10000% of the pseudo-data's statistics level.

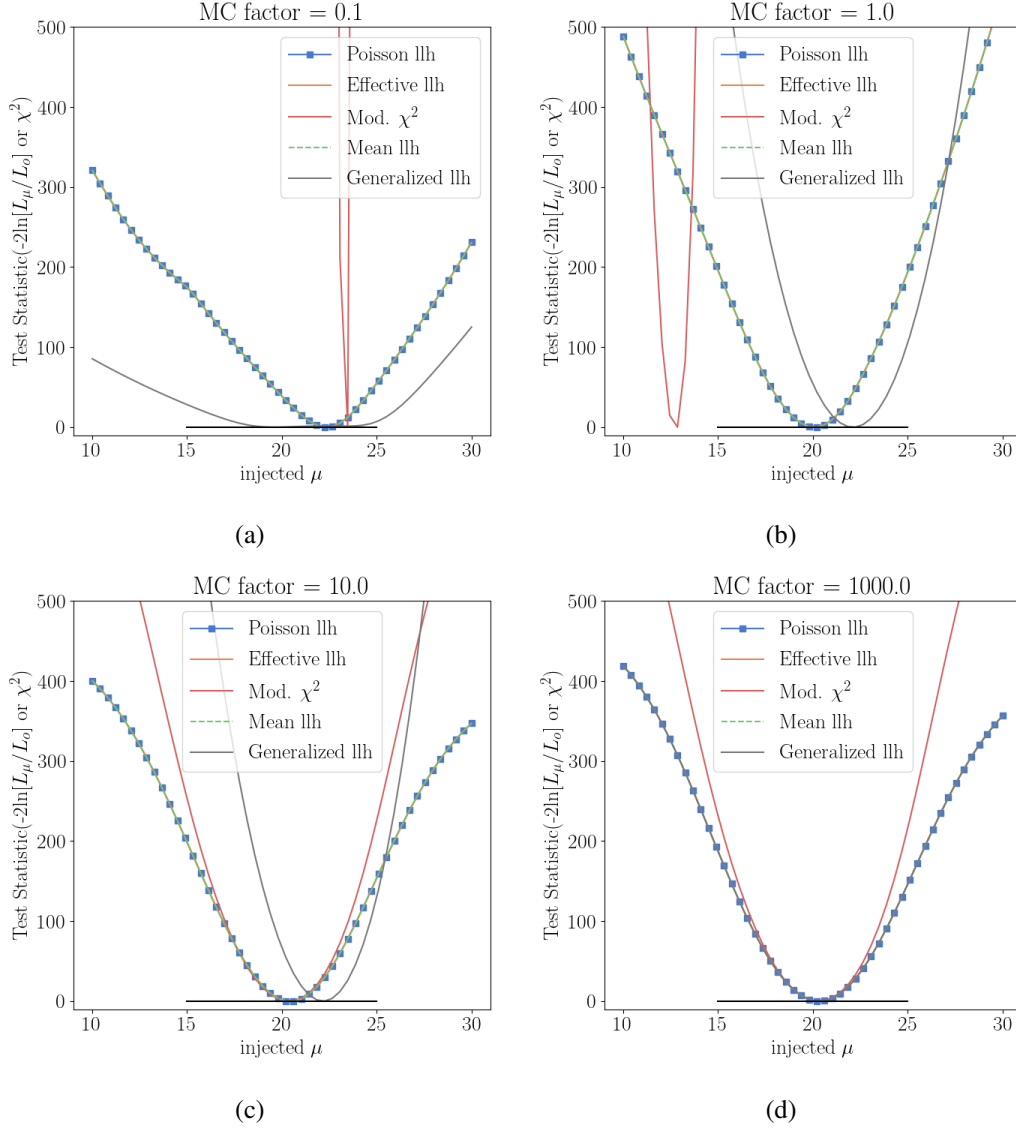


Figure 108: Likelihood scans performed on the 1D toy Monte Carlo test. Each available metric in PISA is evaluated on the MC templates shown in Fig. 107. These templates are reweighted as the mean signal parameter μ_{MC} is changed across the scan, and the absolute value of the metric is compared to the minimal value of the metric calculated across the scan. The black horizontal line is meant as a guide to locate a test statistic value of 0.

The first test performed is a likelihood profile scan, whereby the likelihood value (or modified χ^2) is calculated for simulated templates built with μ_{MC} values ranging from 10 to 30; that likelihood value is normalized by the minimum value obtained in the curve, to facilitate comparison between the different metrics. Figures 108a to 108d show the profiles obtained for the different levels of MC statistics mentioned above. Several observations can be made from these profile scans, the first of which being that every likelihood metric tends to converge to a Poisson-like profile once the MC statistics is much higher than the data's. This serves as a good control test for making sure that the software implementations of each metric is correct, in the limit of infinite MC. Note that in the case of the modified χ^2 , Poisson convergence is obtained only once the statistics of the data itself is sufficiently small, as can be seen in Fig. 109.




Figure 109: Special Toy MC case with very large data statistics (1000 points), showing the convergence of the modified χ^2 to the Poisson likelihood profile. The black horizontal line is a visual guide indicating a likelihood ratio value of 0.0.

The initial profile scan also shows anecdotal evidence that the generalized likelihood and modified χ^2 can lead to biases in cases of very low MC statistic as a function of the MC statistics (see for example the location of the minimum of their profile in Fig. 108a and Fig. 108a). To properly evaluate bias against statistically fluctuated pseudo-data, a second test was performed where 500 pseudo-data trials are fit with the same MC template once for each metric being tested. The same procedure was repeated for the same levels of MC statistics as in the profile scan.

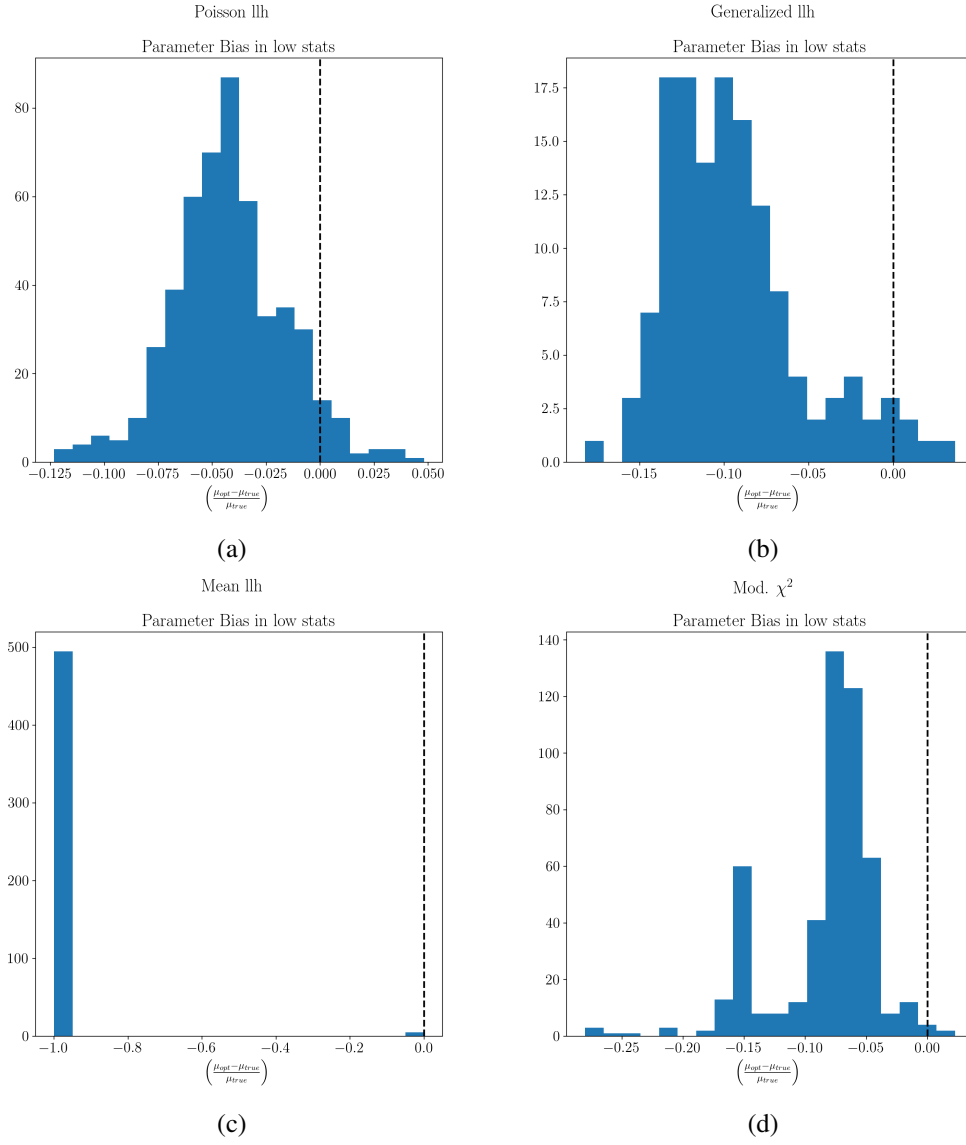


Figure 110: Bias of each likelihood metric in the determination of the signal's mean, for a level of MC statistics corresponding to 10% of the data's. Note that both the effective and mean likelihood biases are very similar, which is why only one of them is shown.




Figure 111: Bias of each likelihood metric in the determination of the signal's mean, for a level of MC statistics corresponding to 1000% of the data's. Note that both the effective and mean likelihood biases are very similar, which is why only one of them is shown.

Figures 110 and 111 show the distribution of biases obtained when fitting for the mean μ of the gaussian toy Monte Carlo signal, in different cases of MC statistics. This test showed that in low MC statistics regime, the effective/mean llh (and to a lesser extent, the modified χ^2) seemed to diverge significantly from the correct value, which is probably due to their lack of mechanism for handling empty bins, something that is taken care of in the generalized poisson llh. The latter still had overall larger biases than the pure Poisson construct. In cases of larger MC stats, the mean/effective likelihood seemed to fit back values with the least bias. Given its poor performances in the first cases however, that likelihood formulation was deemed unsafe to use in the main analysis.

From the same set of Poisson-fluctuated data, a coverage test was performed, similar to the one described in Section 8.4.1. Figure 112 shows an example of TS distribution obtained for one of the metric, in the case of a 1:1 statistics between data and MC. The blue histogram shows the distribution obtained from using the metric, whereas the black line shows the ideal distribution expected from a χ^2 probability density function (itself draw in green). In this particular case, it is possible to observe larger frequencies of larger TS values in the blue histogram compared to the black one, which is a

sign of *undercoverage*: if one were to assume Wilk's theorem applied to this blue distribution, the predicted uncertainty contours (derived from the black histogram) would underestimate the probability of measuring large TS values, which would lead to overconfident uncertainty contours.






Figure 112: Test statistic distribution obtained after fitting 500 pseudo-data trials to a normal Poisson likelihood fit. The Monte Carlo template used in this fit was built using 10% of the pseudo data statistics. The resulting TS distribution (blue), has a noticeably higher number of large TS values than what would be expected from a purely χ^2 -distributed test statistic (the black histogram). This leads to an overestimation of a result's significance, which is a case of test statistics undercoverage.


In order to better visualize the coverage properties of a TS distribution, plots like the ones in Figure 113 are used: percentile intervals on the x-axis correspond to TS values for which 0%, 1%, 10%... of the trials from an ideal χ^2 -distributed experiment are included. These percentiles are matched to the actual cumulative fraction of observed pseudo-experiments falling into that range, in the test statistics obtained from the various metric optimization. In this type of plot, a line falling perfectly on the diagonal dashed line would be deemed to have a perfect coverage, meaning that its test statistic distribution can be well approximated by simple χ^2 statistics.




(a) 10% MC statistics



(b) 100% MC statistics



(c) 1000% MC statistics






(d) 100000% MC statistics

Figure 113: Coverage test of the 1D toy montecarlo pseudo-data using various levels of MC statistics. For each of the MC statistics scenario previously plotted, the experimental cumulative distribution function (CDF) of a metric's fitted TS values is compared to the expected CDF of a χ^2 function with one degree of freedom. Coverage is determined by looking at how close (or how far) the line of a particular metric is from the 1:1 diagonal.

From the outcome of the coverage test, it seemed that this time the generalized likelihood offered slightly better performances, whereas the modified χ^2 appeared to be worse choice in terms of coverage. These discrepancies in coverage did not translate however, when performing the full tau neutrino appearance fit.

B.4 FULL SET OF SYSTEMATIC PULL PLOTS

The following plots show the expected pull on the MC event counts in every bin of the chosen template, when a particular systematic parameter is pull away from its nominal value by 1σ .

B.4.1 *Neutrino Flux Uncertainties*


barr_i_Pi, $+1\sigma$, $0 \rightarrow 0.122$ 

Figure 117: Barr parameter I


barr_w_K, $+1\sigma$, $0 \rightarrow 0.4$ 

Figure 118: Barr parameter K


barr_y_antiK, $+1\sigma$, $0 \rightarrow 0.3$ 

Figure 119: Barr parameter Y (applied to anti-kaons)




Figure 120: Barr parameter G








Figure 121: Barr parameter Z (applied to kaons)

B.4.2 Neutrino Oscillation Uncertainties

Figure 122: Neutrino Mass splitting, Δm_{31}^2

Figure 123: Neutrino mixing angle θ_{23}

B.4.3 Cross Section Uncertainties

Figure 124: Quasielastic axial mass, $M_{A, \text{QE}}$ Figure 125: Resonant axial mass, $M_{A, \text{RES}}$






Figure 126: CSMS Cross-section scaling (1: CSMS-like, 0:genie-like)

B.4.4 Discrete Detector Uncertainties

Figure 127: Hole ice acceptance parameter p_0 Figure 128: Hole ice acceptance parameter p_1


dom_eff, +1 σ , 1 \rightarrow 1.1

Figure 129: DOM efficiency


bulk_ice_scatter, +10%, 1 \rightarrow 1.1

Figure 130: Bulk ice scattering coefficient


bulk_ice_abs, +10%, 1 \rightarrow 1.1

Figure 131: Bulk ice absorption coefficient




Figure 132: Neutrino effective area

B.4.5 Muon Uncertainty










Figure 133: Overall Muon rate (scaled effective area)



B.5 FULL RESULTS OF THE SYSTEMATIC IMPACT TEST



This section presents the detailed outcome of the systematic impact test described in Section 8.5. Each plot shows the impact scale value I_s obtained (and the corresponding measured offset of the fitted N_τ parameter) for different signs of pull (positive vs negative), and for different true values of N_τ . Horizontal bars colored in red mean that the sign of the impact scale was negative, which can happen when the systematically pulled set is so far from the nominal template that even the free fit leads to large likelihood ratios.



B.5.1 Test where N_τ is set to 0.75Figure 134: Positive pull, $N_\tau=0.75$.Figure 135: Tau normalization offset for a positive pull, $N_\tau=0.75$.

Figure 136: Negative pull, $N_\tau=0.75$.Figure 137: Tau normalization offset for a negative pull, $N_\tau=0.75$.

B.5.2 Test where N_τ is set to 1.00Figure 138: Positive pull, $N_\tau=1.00$.Figure 139: Tau normalization offset for a positive pull, $N_\tau=1.00$.

Figure 140: Negative pull, $N_\tau = 1.00$.Figure 141: Tau normalization offset for a positive pull, $N_\tau = 1.00$.

B.5.3 Test where N_τ is set to 1.25Figure 142: Positive pull, $N_\tau=1.25$.Figure 143: Tau normalization offset for a positive pull, $N_\tau=1.25$.

Figure 144: Negative pull, $N_\tau = 1.25$.Figure 145: Tau normalization offset for a positive pull, $N_\tau = 1.25$.

B.6 ADDITIONAL INFORMATION ABOUT THE LEVEL 7 MUON CLASSIFIER

Table 10 lists the hyperparameters used in the Level 7 muon classifier training algorithm.

Table 10: Hyperparameters used in the training of the Level 7 Muon Classifier

| | |
|-------------------------|---------------|
| Algorithm | LightGBM |
| Objective | regression_l2 |
| Learning rate | 0.1 |
| Number of leaves | 10 |
| Max depth | 4 |
| Min data in leaf | 20 |

ADDITIONAL MATERIAL FOR THE 2MRS CORRELATION SEARCH

C.1 COMPUTING A LIMIT ON THE DENSITY OF NEUTRINO SOURCES

Given a number of signal multiplets that correlate with the 2MRS catalog, a relationship between the luminosity and density of the underlying source population can be derived. This relationship used in this work has been previously derived by [137]. It depends on three main assumptions regarding the nature of the neutrino emission suspected to originate from the local universe:

- The hypothesized flux comes from a uniform population of sources (described by a single power spectra and intrinsic luminosity)
- The detectable neutrino emission from these sources is assumed to follow poisson statistics
- The source density of the hypothesized population of sources is such that its neutrino flux falls just short of being detected by IceCube as localized multiplets.

At the base of the limit calculation lies the probability for a source to emit a multiplet of multiplicity m greater than a threshold μ^1 , which is given² by:

$$P_{m \geq \mu}(\lambda) = 1.0 - \frac{\Gamma(\mu, \lambda)}{\Gamma(\mu)}. \quad (127)$$

That probability depends on λ , the expectation value for the number of detectable neutrinos emitted by the source, which is set by the experiment's sensitivity, the intrinsic luminosity L assumed and its distance with respect to Earth $d(z)$. The latter quantity is constrained by the third assumption made about the source population: since [131] has found no significant neutrino clusters, the bulk of the population must lie at least as far as the distance at which single neutrinos become visible in IceCube, $d_{m=1}$, which in the case of the 7yr-PS sample is equal to:

$$d_{m=1} = 180 \text{ Mpc} \cdot \sqrt{\frac{L_v}{10^{42} \text{ erg} \cdot \text{s}^{-1}}}. \quad (128)$$

The number of neutrinos correlating with sources from the 2MRS catalog can be estimated by integrating the source density multiplied by $P_{m \geq \mu}(\lambda)$, over a comoving volume of space up to the desired critical redshift $z_c = 0.03$. Since the integral occurs over a very small cosmological distance (ie $d_{m=1} \ll (c/H_0)\sqrt{(\mu)}$), the number of contributing source can be simplified to:

$$N_{limit} = n_0 \Delta \Omega d_{m=1}^3 \frac{\Gamma\mu + \lambda_c^{3/2} \Gamma(\mu - 3/2, \lambda_c) - \Gamma(\mu, \lambda_c)}{3\lambda_c^{3/2} \Gamma(\mu)} \quad (129)$$

¹ In the correlation analysis, the threshold is set to $\mu = 2$.

² Note that this representation allows for non-integer number of photons to be observed, which suits the case when n_s is derived from a fit to data.

The relationship above introduces n_0 , which corresponds to the source density at redshift of $z = 0$. From a derivation in the appendix of [137], the average test statistic expected to be obtained for finding N_{limit} signal multiplet, given N_{ws} total multiplets is given by:

$$TS = \frac{N_{limit}^2}{N_{ws}} \cdot 16\pi^2\sigma^2 \quad (130)$$

where σ is the variance of the 2MRS density map. merging Eq.129 into Eq.130 and isolating for n_0 finally yields:

$$n_0 < \frac{TS}{2(4\pi\sigma)^2} \cdot \frac{1}{\Delta\Omega} \cdot \left(\frac{c}{H_0} z_c \right)^{-3} \cdot f(\mu, \lambda_c) \quad (131)$$

where $f(\mu, \lambda_c)$ is a function that can be divided into two regimes. For high-luminosity populations ($\lambda_c \gg \mu$), f becomes:

$$f(\mu, \lambda_c) \approx 9\lambda_c^{3/2} \frac{3\Gamma(n_s - 3/2)}{\Gamma(n_s)} \cdot 2 \quad (132)$$

while for low-luminosity populations ($\lambda_c \ll \mu$), f becomes:

$$f(\mu, \lambda_c) \approx \lambda_c^{-3/2} \frac{3\Gamma(\mu)}{\Gamma(\mu - 3/2)} \cdot \left(1 + \sqrt{1 + \frac{4N_{bkg}(4\pi\sigma)^2}{TS(p)}} \right) \quad (133)$$

These two equations determine the "V" shape of the limit seen in Fig.97.

BIBLIOGRAPHY

- [1] C. L. Cowan et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124.3212 (1956), pp. 103–104. ISSN: 0036-8075. DOI: 10.1126/science.124.3212.103. eprint: <https://science.sciencemag.org/content/124/3212/103.full.pdf>. URL: <https://science.sciencemag.org/content/124/3212/103>.
- [2] IceCube Collaboration. “Constraints on neutrino emission from nearby galaxies using the 2MASS redshift survey and IceCube”. In: *Journal of Cosmology and Astroparticle Physics* 2020.07 (July 2020), pp. 042–042. URL: <http://dx.doi.org/10.1088/1475-7516/2020/07/042>.
- [3] *Wikipedia page on the Standard Model*. https://en.wikipedia.org/wiki/Standard_Model. Accessed: 2020-08-11.
- [4] M. Srednicki. *Quantum Field Theory*. Cambridge University Press, 2016.
- [5] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.
- [6] ALEPH Collaboration et al. “Precision Electroweak Measurements on the Z Resonance”. In: *Physics Reports* 427.5-6 (May 2006), pp. 257–454. URL: <http://dx.doi.org/10.1016/j.physrep.2005.12.006>.
- [7] AMANDA Collaboration. “The AMANDA neutrino telescope: principle of operation and first results”. In: *Astroparticle Physics* 13.1 (2000), pp. 1–20. ISSN: 0927-6505. DOI: [https://doi.org/10.1016/S0927-6505\(99\)00092-4](https://doi.org/10.1016/S0927-6505(99)00092-4).
- [8] ANTARES Collaboration. “ANTARES: The first undersea neutrino telescope”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 656.1 (Nov. 2011), pp. 11–38. ISSN: 0168-9002. DOI: 10.1016/j.nima.2011.06.103. URL: <http://dx.doi.org/10.1016/j.nima.2011.06.103>.
- [9] IceCube Collaboration. “The IceCube Neutrino Observatory: Instrumentation and Online Systems”. In: *JINST* 03 (12 2017), P03012. DOI: 10.1088/1748-0221/12/03/P03012. eprint: 1612.05093.
- [10] P. A. Čerenkov. “Visible Radiation Produced by Electrons Moving in a Medium with Velocities Exceeding that of Light”. In: *Phys. Rev.* 52 (4 Aug. 1937), pp. 378–379. DOI: 10.1103/PhysRev.52.378. URL: <https://link.aps.org/doi/10.1103/PhysRev.52.378>.
- [11] Ig. Tamm. “Radiation Emitted by Uniformly Moving Electrons”. In: *Selected Papers*. Ed. by Boris M. Bolotovskii, Victor Ya. Frenkel, and Rudolf Peierls. Berlin, Heidelberg: Springer Berlin Heidelberg, 1991, pp. 37–53. URL: https://doi.org/10.1007/978-3-642-74626-0_3.
- [12] Adam Bernstein et al. *Report on the Depth Requirements for a Massive Detector at Hometown*. 2009. arXiv: 0907.4183 [hep-ex].
- [13] Hamamatsu Corporation. *Photomultiplier Tubes: Basics and Applications*. 2007. URL: https://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf.
- [14] IceCube Collaboration. “In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes”. In: *Journal of Instrumentation* 15.06 (June 2020), P06032–P06032. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/06/p06032. URL: <http://dx.doi.org/10.1088/1748-0221/15/06/P06032>.

- [15] IceCube Collaboration. “The design and performance of IceCube DeepCore”. In: *Astroparticle Physics* 35.10 (May 2012), pp. 615–624. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2012.01.004. URL: <http://dx.doi.org/10.1016/j.astropartphys.2012.01.004>.
- [16] M. A. Unland Elorrieta. “Studies on dark rates induced by radioactive decays of the multi-PMT digital optical module for future IceCube extensions”. MA thesis. WWU Münster, 2017.
- [17] D. F. H. von Zuydtwyck. “HitSpooling: An Improvement for the Supernova Neutrino Detection System in IceCube”. PhD thesis. ULB Bruxelles, 2015.
- [18] M. Larson. “Simulation and Identification of non-Poissonian Noise Triggers in the IceCube Neutrino Detector”. MA thesis. University of Alabama, 2013.
- [19] D.-M. Mei and A. Hime. “Muon-induced background study for underground laboratories”. In: *Physical Review D* 73.5 (Mar. 2006). ISSN: 1550-2368. DOI: 10.1103/physrevd.73.053004. URL: <http://dx.doi.org/10.1103/PhysRevD.73.053004>.
- [20] D. Hanna et al. “An LED-based flasher system for VERITAS”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 612.2 (Jan. 2010), pp. 278–287. ISSN: 0168-9002. DOI: 10.1016/j.nima.2009.10.107. URL: <http://dx.doi.org/10.1016/j.nima.2009.10.107>.
- [21] R. Abbasi et al. “Calibration and characterization of the IceCube photomultiplier tube”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 618.1-3 (June 2010), pp. 139–152. ISSN: 0168-9002. DOI: 10.1016/j.nima.2010.03.102. URL: <http://dx.doi.org/10.1016/j.nima.2010.03.102>.
- [22] R. Saldanha et al. *Cosmogenic activation of silicon*. 2020. arXiv: 2007.10584 [physics.ins-det].
- [23] E. Armengaud et al. “Measurement of the cosmogenic activation of germanium detectors in EDELWEISS-III”. In: *Astroparticle Physics* 91 (2017), pp. 51–64. ISSN: 0927-6505. DOI: <https://doi.org/10.1016/j.astropartphys.2017.03.006>. URL: <http://www.sciencedirect.com/science/article/pii/S0927650517301044>.
- [24] Susana Cebrián. “Cosmogenic activation of materials”. In: *International Journal of Modern Physics A* 32.30 (Oct. 2017), p. 1743006. ISSN: 1793-656X. DOI: 10.1142/s0217751x17430060. URL: <http://dx.doi.org/10.1142/S0217751X17430060>.
- [25] Aya Ishihara. *The IceCube Upgrade – Design and Science Goals*. 2019. arXiv: 1908.09441 [astro-ph.HE].
- [26] H. O. Meyer. *Dark Rate of a Photomultiplier at Cryogenic Temperatures*. 2008. arXiv: 0805.0771 [nucl-ex].
- [27] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.9.36>.
- [28] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (Apr. 2001), pp. 218–224. ISSN: 0370-2693. DOI: 10.1016/S0370-2693(01)00307-0. URL: [http://dx.doi.org/10.1016/S0370-2693\(01\)00307-0](http://dx.doi.org/10.1016/S0370-2693(01)00307-0).
- [29] *The 2015 Nobel Prize in Physics*. Oct. 2015. URL: <https://www.nobelprize.org/prizes/physics/2015/press-release/>.
- [30] B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. In: *Sov. Phys. JETP* 7 (1958), pp. 172–173.

- [31] Maki Z., Nakagawa M., and Sakata S. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. ISSN: 0033-068X. DOI: 10.1143/PTP.28.870. eprint: <https://academic.oup.com/ptp/article-pdf/28/5/870/5258750/28-5-870.pdf>. URL: <https://doi.org/10.1143/PTP.28.870>.
- [32] E. Kh. Akhmedov and A. Yu. Smirnov. “Paradoxes of neutrino oscillations”. In: *Physics of Atomic Nuclei* 72.8 (Aug. 2009), pp. 1363–1381. ISSN: 1562-692X. DOI: 10.1134/s1063778809080122. URL: <http://dx.doi.org/10.1134/S1063778809080122>.
- [33] KATRIN Collaboration. “Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN”. In: *Physical Review Letters* 123.22 (Nov. 2019). ISSN: 1079-7114. DOI: 10.1103/physrevlett.123.221802. URL: <http://dx.doi.org/10.1103/PhysRevLett.123.221802>.
- [34] N. Aghanim et al. “Planck 2018 results”. In: *Astronomy and Astrophysics* 641 (Sept. 2020), A6. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201833910. URL: <http://dx.doi.org/10.1051/0004-6361/201833910>.
- [35] IceCube Collaboration. “Development of an analysis to probe the neutrino mass ordering with atmospheric neutrinos using three years of IceCube DeepCore data”. In: *The European Physical Journal C* 80.1 (Jan. 2020). DOI: 10.1140/epjc/s10052-019-7555-0. URL: <http://dx.doi.org/10.1140/epjc/s10052-019-7555-0>.
- [36] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson”. In: *Phys. Rev. Lett.* 13 (4 July 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.138>.
- [37] The T2K Collaboration. “Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations”. In: *Nature* 580.7803 (2020), pp. 339–344. DOI: 10.1038/s41586-020-2177-0. URL: <https://doi.org/10.1038/s41586-020-2177-0>.
- [38] M. A. Acero et al. “First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA”. In: *Physical Review Letters* 123.15 (Oct. 2019). ISSN: 1079-7114. DOI: 10.1103/physrevlett.123.151803. URL: <http://dx.doi.org/10.1103/PhysRevLett.123.151803>.
- [39] B. Aharmim et al. “Electron energy spectra, fluxes, and day-night asymmetries of 8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory”. In: *Physical Review C* 72.5 (Nov. 2005). DOI: 10.1103/physrevc.72.055502. arXiv: nucl-ex/0502021. URL: <http://dx.doi.org/10.1103/PhysRevC.72.055502>.
- [40] G. Alimonti et al. “The Borexino detector at the Laboratori Nazionali del Gran Sasso”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 600.3 (Mar. 2009), pp. 568–593. DOI: 10.1016/j.nima.2008.11.076. URL: <http://dx.doi.org/10.1016/j.nima.2008.11.076>.
- [41] S. Abe et al. “Precision Measurement of Neutrino Oscillation Parameters with KamLAND”. In: *Physical Review Letters* 100.22 (June 2008). ISSN: 1079-7114. DOI: 10.1103/physrevlett.100.221803. URL: <http://dx.doi.org/10.1103/PhysRevLett.100.221803>.
- [42] F.P. An et al. “A side-by-side comparison of Daya Bay antineutrino detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 685 (Sept. 2012), pp. 78–97. ISSN: 0168-9002. DOI: 10.1016/j.nima.2012.05.030. URL: <http://dx.doi.org/10.1016/j.nima.2012.05.030>.

- [43] RENO Collaboration and J. K. Ahn. *RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang*. 2010. arXiv: 1003.1391 [hep-ex].
- [44] The Double Chooz Collaboration. “Double Chooz θ_{13} measurement via total neutron capture detection”. In: *Nature Physics* 16.5 (2020), pp. 558–564. DOI: 10.1038/s41567-020-0831-y. URL: <https://doi.org/10.1038/s41567-020-0831-y>.
- [45] K. Abe et al. “The T2K experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (Dec. 2011), pp. 106–135. ISSN: 0168-9002. DOI: 10.1016/j.nima.2011.06.067. URL: <http://dx.doi.org/10.1016/j.nima.2011.06.067>.
- [46] D.S. Ayres et al. “The NOvA Technical Design Report”. In: (Oct. 2007). DOI: 10.2172/935497.
- [47] Yoichiro Suzuki. “The Super-Kamiokande experiment”. In: *The European Physical Journal C* 79.4 (2019), p. 298. URL: <https://doi.org/10.1140/epjc/s10052-019-6796-2>.
- [48] Ivan Esteban et al. *The fate of hints: updated global analysis of three-flavor neutrino oscillations*. 2020. arXiv: 2007.14792 [hep-ph].
- [49] Sebastian A. R. Ellis, Kevin J. Kelly, and Shirley Weishi Li. *Current and Future Neutrino Oscillation Constraints on Leptonic Unitarity*. 2020. arXiv: 2008.01088 [hep-ph].
- [50] K Abe et al. “Physics potentials with the second Hyper-Kamiokande detector in Korea”. In: *Progress of Theoretical and Experimental Physics* 2018.6 (June 2018). ISSN: 2050-3911. DOI: 10.1093/ptep/pty044. URL: <http://dx.doi.org/10.1093/ptep/pty044>.
- [51] DUNE Collaboration. *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 1: The LBNF and DUNE Projects*. 2016. arXiv: 1601.05471 [physics.ins-det].
- [52] JUNO Collaboration. “Neutrino physics with JUNO”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.3 (Feb. 2016), p. 030401. ISSN: 1361-6471. DOI: 10.1088/0954-3899/43/3/030401. URL: <http://dx.doi.org/10.1088/0954-3899/43/3/030401>.
- [53] Sébastien Descotes-Genon and Patrick Koppenburg. “The CKM Parameters”. In: *Annual Review of Nuclear and Particle Science* 67.1 (Oct. 2017), pp. 97–127. ISSN: 1545-4134. DOI: 10.1146/annurev-nucl-101916-123109. URL: <http://dx.doi.org/10.1146/annurev-nucl-101916-123109>.
- [54] Parke S. and Ross-Lonergan M. “Unitarity and the three flavour neutrino mixing matrix”. In: *Phys. Rev. D* 113009 (93 2016).
- [55] Sebastian A. R. Ellis, Kevin J. Kelly, and Shirley Weishi Li. *Leptonic Unitarity Triangles*. 2020. arXiv: 2004.13719 [hep-ph].
- [56] Particle Data Group. “Review of Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 2020). 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104. eprint: <https://academic.oup.com/ptep/article-pdf/2020/8/083C01/33653179/ptaa104.pdf>. URL: <https://doi.org/10.1093/ptep/ptaa104>.
- [57] Jr. H. R. Hart and R. L. Fleischer. “Apollo 16 Cosmic Ray Experiment – Solar Flare Energy Spectrum and Heavy Element Enrichment at Low Energy”. In: *Lunar and Planetary Science Conference*. Vol. 4. Lunar and Planetary Science Conference. Mar. 1973, p. 336.

- [58] A. C. Cummings et al. “GALACTIC COSMIC RAYS IN THE LOCAL INTERSTELLAR MEDIUM: VOYAGER 1 OBSERVATIONS AND MODEL RESULTS”. In: *The Astrophysical Journal* 831.1 (Oct. 2016), p. 18. URL: <https://doi.org/10.3847/0004-637x/831/1/18>.
- [59] Rikard Enberg, Mary Hall Reno, and Ina Sarcevic. “Prompt neutrino fluxes from atmospheric charm”. In: *Physical Review D* 78.4 (Aug. 2008). ISSN: 1550-2368. DOI: 10.1103/PhysRevD.78.043005. URL: <http://dx.doi.org/10.1103/PhysRevD.78.043005>.
- [60] I. Storehaug. “Improving the Atmospheric Neutrino Flux Estimation in IceCube”. MA thesis. Niels Bohr Institute, 2019.
- [61] A Yu Smirnov. “The MSW Effect and Matter Effects in Neutrino Oscillations”. In: *Physica Scripta* T121 (Jan. 2005), pp. 57–64. ISSN: 1402-4896. DOI: 10.1088/0031-8949/2005/t121/008. URL: <http://dx.doi.org/10.1088/0031-8949/2005/T121/008>.
- [62] M. Chizhov, M. Maris, and S. T. Petcov. “On the oscillation length resonance in the transitions of solar and atmospheric neutrinos crossing the earth core”. In: (July 1998). arXiv: hep-ph/9810501.
- [63] E.Kh. Akhmedov. “Parametric resonance of neutrino oscillations and passage of solar and atmospheric neutrinos through the earth”. In: *Nuclear Physics B* 538.1-2 (Jan. 1999), pp. 25–51. ISSN: 0550-3213. DOI: 10.1016/S0550-3213(98)00723-8. URL: [http://dx.doi.org/10.1016/S0550-3213\(98\)00723-8](http://dx.doi.org/10.1016/S0550-3213(98)00723-8).
- [64] J. A. Formaggio and G. P. Zeller. “From eV to EeV: Neutrino cross sections across energy scales”. In: *Reviews of Modern Physics* 84.3 (Sept. 2012), pp. 1307–1341. ISSN: 1539-0756. DOI: 10.1103/revmodphys.84.1307. URL: <http://dx.doi.org/10.1103/RevModPhys.84.1307>.
- [65] *Neutrino Interactions*. <http://danielscully.co.uk/thesis/interactions.html>. Accessed: 26-12-2020.
- [66] Omar Benhar et al. “Neutrino–nucleus interactions and the determination of oscillation parameters”. In: *Physics Reports* 700 (July 2017), pp. 1–47. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2017.07.004. URL: <http://dx.doi.org/10.1016/j.physrep.2017.07.004>.
- [67] Artur M. Ankowski and Alexander Friedland. “Assessing the accuracy of the GENIE event generator with electron-scattering data”. In: *Physical Review D* 102.5 (Sept. 2020). ISSN: 2470-0029. DOI: 10.1103/PhysRevD.102.053001. URL: <http://dx.doi.org/10.1103/PhysRevD.102.053001>.
- [68] Konstantin S. Kuzmin, Vladimir V. Lyubushkin, and Vadim A. Naumov. *Axial masses in quasielastic neutrino scattering and single-pion neutrino production on nucleons and nuclei*. 2006. arXiv: hep-ph/0606184 [hep-ph].
- [69] Dieter Rein and Lalit M. Sehgal. “Neutrino Excitation of Baryon Resonances and Single Pion Production”. In: *Annals Phys.* 133 (1981), pp. 79–153. DOI: 10.1016/0003-4916(81)90242-6.
- [70] J. D. Bjorken. “Asymptotic Sum Rules at Infinite Momentum”. In: *Phys. Rev.* 179 (5 Mar. 1969), pp. 1547–1553. DOI: 10.1103/PhysRev.179.1547. URL: <https://link.aps.org/doi/10.1103/PhysRev.179.1547>.
- [71] Amanda Cooper-Sarkar, Philipp Mertsch, and Subir Sarkar. “The high energy neutrino cross-section in the Standard Model and its uncertainty”. In: *Journal of High Energy Physics* 2011.8 (Aug. 2011). ISSN: 1029-8479. DOI: 10.1007/jhep08(2011)042. URL: [http://dx.doi.org/10.1007/JHEP08\(2011\)042](http://dx.doi.org/10.1007/JHEP08(2011)042).

- [72] S. Kretzer and M. H. Reno. “Tau neutrino deep inelastic charged current interactions”. In: *Physical Review D* 66.11 (Dec. 2002). DOI: 10.1103/physrevd.66.113007. URL: <http://dx.doi.org/10.1103/PhysRevD.66.113007>.
- [73] N. Agafonova et al. “Discovery of Tau Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment”. In: *Physical Review Letters* 115.12 (Sept. 2015). ISSN: 1079-7114. DOI: 10.1103/physrevlett.115.121802. URL: <http://dx.doi.org/10.1103/PhysRevLett.115.121802>.
- [74] K. Abe et al. “Evidence for the Appearance of Atmospheric Tau Neutrinos in Super-Kamiokande”. In: *Physical Review Letters* 110.18 (May 2013). ISSN: 1079-7114. DOI: 10.1103/physrevlett.110.181802. URL: <http://dx.doi.org/10.1103/PhysRevLett.110.181802>.
- [75] IceCube Collaboration. “Measurement of Atmospheric Tau Neutrino Appearance with IceCube DeepCore”. In: *Phys. Rev. D* 99 (2019).
- [76] Costas Andreopoulos et al. *The GENIE Neutrino Monte Carlo Generator: Physics and User Manual*. 2015. arXiv: 1510.05494 [hep-ph].
- [77] Alfonso Garcia and Aart Heijboer. *High-energy neutrino event simulation at NLO in Genie for KM3NeT and other observatories*. 2019. arXiv: 1908.10077 [hep-ex].
- [78] J. Van Santen. “Neutrino Interactions in IceCube above 1 TeV”. PhD thesis. University of Wisconsin-Madison, 2014.
- [79] Thomas K. Gaisser, Todor Stanev, and Serap Tilav. *Cosmic Ray Energy Spectrum from Measurements of Air Showers*. 2013. arXiv: 1303.3565 [astro-ph.HE].
- [80] M. Larson. “A Search for Tau Neutrino Appearance with IceCube-DeepCore”. PhD thesis. University of Copenhagen, 2018.
- [81] J.H. Koehne et al. “PROPOSAL: A tool for propagation of charged leptons”. In: *Comput. Phys. Commun.* 184 (2013), pp. 2070–2090. DOI: 10.1016/j.cpc.2013.04.001.
- [82] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [83] C. Kopper. *clsim on GitHub*. <https://github.com/claudiok/clsim>. Accessed: 2020-09-25.
- [84] M. Rongen, R. C. Bay, and S. Blot. “Observation of an optical anisotropy in the deep glacial ice at the geographic South Pole using a laser dust logger”. In: *The Cryosphere* 14.8 (2020), pp. 2537–2543. DOI: 10.5194/tc-14-2537-2020. URL: <https://tc.copernicus.org/articles/14/2537/2020/>.
- [85] Dmitry Chirkin. *Likelihood description for comparing data with simulation of limited statistics*. 2013. arXiv: 1304.0735 [astro-ph.IM].
- [86] M. Ackermann et al. “Optical properties of deep glacial ice at the South Pole”. In: *J. Geophys. Res.* 111.D13 (2006), p. D13203. DOI: 10.1029/2005JD006687.
- [87] Jerome H. Friedman. “Greedy Function Approximation: A Gradient Boosting Machine”. In: *The Annals of Statistics* 29.5 (2001), pp. 1189–1232. URL: <http://www.jstor.org/stable/2699986>.
- [88] Tianqi Chen and Carlos Guestrin. “XGBoost”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Aug. 2016). DOI: 10.1145/2939672.2939785. URL: <http://dx.doi.org/10.1145/2939672.2939785>.

- [89] Guolin Ke et al. “LightGBM: A Highly Efficient Gradient Boosting Decision Tree”. In: *Advances in Neural Information Processing Systems 30*. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 3146–3154. URL: <http://papers.nips.cc/paper/6907-lightgbm-a-highly-efficient-gradient-boosting-decision-tree.pdf>.
- [90] M. Aaboud et al. “Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV”. In: *The European Physical Journal C* 79.8 (Aug. 2019). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-7140-6. URL: <http://dx.doi.org/10.1140/epjc/s10052-019-7140-6>.
- [91] CMS Collaboration. “Measurement of properties of $B_s^0 \rightarrow \mu + \mu^-$ decays and search for $B^0 \rightarrow \mu + \mu^-$ with the CMS experiment”. In: *Journal of High Energy Physics* 2020.4 (Apr. 2020). URL: [http://dx.doi.org/10.1007/JHEP04\(2020\)188](http://dx.doi.org/10.1007/JHEP04(2020)188).
- [92] IceCube Collaboration. “Improvement in fast particle track reconstruction with robust statistics”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 736 (Feb. 2014), pp. 143–149. URL: <http://dx.doi.org/10.1016/j.nima.2013.10.074>.
- [93] J. P. Yáñez Garza. “Measurement of neutrino oscillations in atmospheric neutrinos with the IceCube DeepCore detector”. PhD thesis. Humboldt-Universität zu Berlin, 2014. DOI: <http://dx.doi.org/10.18452/17016>.
- [94] P. Eller and J. Lanfranchi. *Retro Reco*. <https://github.com/IceCubeOpenSource/retro>. Accessed: 2020-08-11.
- [95] M. Leuermann. “Testing the Neutrino Mass Ordering with IceCube DeepCore”. PhD thesis. RWTH Aachen University, 2018.
- [96] D.E. Groom, N. V. Mokhov, and S. I. Striganov. “MUON STOPPING POWER AND RANGE TABLES 10 MeV–100 TeV”. In: *Atomic Data and Nuclear Data Tables* 78.2 (2001), pp. 183–356. ISSN: 0092-640X. DOI: <https://doi.org/10.1006/adnd.2001.0861>. URL: <http://www.sciencedirect.com/science/article/pii/S0092640X01908617>.
- [97] M. Kowalski. “Search for Neutrino-Induced Cascades with the AMANDA-II Detector”. PhD thesis. Humboldt-Universität zu Berlin, 2004.
- [98] Leif Rädcl and Christopher Wiebusch. “Calculation of the Cherenkov light yield from electromagnetic cascades in ice with Geant4”. In: *Astroparticle Physics* 44 (Apr. 2013), pp. 102–113. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2013.01.015. URL: <http://dx.doi.org/10.1016/j.astropartphys.2013.01.015>.
- [99] T.A. Gabriel et al. “Energy dependence of hadronic activity”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 338.2 (1994), pp. 336–347. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(94\)91317-X](https://doi.org/10.1016/0168-9002(94)91317-X). URL: <http://www.sciencedirect.com/science/article/pii/016890029491317X>.
- [100] M. G. Dunkman. “Measurement of Atmospheric Muon Disappearance with IceCube-DeepCore”. PhD thesis. Penn State, 2015.
- [101] IceCube Collaboration. “Measurements using the inelasticity distribution of multi-TeV neutrino interactions in IceCube”. In: *Physical Review D* 99.3 (Feb. 2019). ISSN: 2470-0029. DOI: 10.1103/physrevd.99.032004. URL: <http://dx.doi.org/10.1103/PhysRevD.99.032004>.

- [102] Scott M Lundberg and Su-In Lee. “A Unified Approach to Interpreting Model Predictions”. In: *Advances in Neural Information Processing Systems 30*. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 4765–4774. URL: <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- [103] IceCube Collaboration. *Computational Techniques for the Analysis of Small Signals in High-Statistics Neutrino Oscillation Experiments*. 2018. arXiv: 1803.05390 [physics.data-an].
- [104] M. Honda et al. “Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model”. In: *Physical Review D* 92.2 (July 2015). ISSN: 1550-2368. DOI: 10.1103/PhysRevD.92.023004. URL: <http://dx.doi.org/10.1103/PhysRevD.92.023004>.
- [105] G. D. Barr et al. “Uncertainties in atmospheric neutrino fluxes”. In: *Physical Review D* 74.9 (Nov. 2006). ISSN: 1550-2368. DOI: 10.1103/PhysRevD.74.094009. URL: <http://dx.doi.org/10.1103/PhysRevD.74.094009>.
- [106] Anatoli Fedynitch et al. *Calculation of conventional and prompt lepton fluxes at very high energy*. 2015. arXiv: 1503.00544 [hep-ph].
- [107] Hans Dembinski et al. *Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10¹¹ GeV*. 2017. arXiv: 1711.11432 [astro-ph.HE].
- [108] Felix Riehn et al. *The hadronic interaction model SIBYLL 2.3c and Feynman scaling*. 2017. arXiv: 1709.07227 [hep-ph].
- [109] Roger Wendell. *Prob3++ software for computing three flavor neutrino oscillation probabilities*. <http://www.phy.duke.edu/~raw22/public/Prob3++/>. 2012.
- [110] V. Barger et al. “Matter effects on three-neutrino oscillations”. In: *Phys. Rev. D* 22 (11 Dec. 1980), pp. 2718–2726. DOI: 10.1103/PhysRevD.22.2718. URL: <https://link.aps.org/doi/10.1103/PhysRevD.22.2718>.
- [111] Adam M. Dziewonski and Don L. Anderson. “Preliminary reference Earth model”. In: *Physics of the Earth and Planetary Interiors* 25.4 (1981), pp. 297–356. DOI: [https://doi.org/10.1016/0031-9201\(81\)90046-7](https://doi.org/10.1016/0031-9201(81)90046-7). URL: <http://www.sciencedirect.com/science/article/pii/0031920181900467>.
- [112] Ivan Esteban et al. “Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ_{CP} , and the mass ordering”. In: *Journal of High Energy Physics* 2019.1 (Jan. 2019). DOI: 10.1007/jhep01(2019)106. URL: [http://dx.doi.org/10.1007/JHEP01\(2019\)106](http://dx.doi.org/10.1007/JHEP01(2019)106).
- [113] IceCube Collaboration. “Measurement of Atmospheric Neutrino Oscillations at 6–56 GeV with IceCube DeepCore”. In: *Phys. Rev. Lett.* 120.7 (2018), p. 071801. DOI: 10.1103/PhysRevLett.120.071801. arXiv: 1707.07081 [hep-ex].
- [114] Arie Bodek and Un-ki Yang. *Axial and Vector Structure Functions for Electron- and Neutrino-Nucleon Scattering Cross Sections at all Q^2 using Effective Leading order Parton Distribution Functions*. 2013. arXiv: 1011.6592 [hep-ph].
- [115] M. Glück, E. Reya, and A. Vogt. “Dynamical parton distributions revisited”. In: *Eur. Phys. J. C* 5 (1998), pp. 461–470. DOI: 10.1007/s100520050289. arXiv: hep-ph/9806404.
- [116] V. Radescu and H1/ZEUS collaborations. *Combination and QCD Analysis of the HERA Inclusive Cross Sections*. 2013. arXiv: 1308.0374 [hep-ex].
- [117] M. Liubarska, J. P. Yanez, and J. Hignight. *Systematic uncertainty in neutrino-nucleon deep inelastic scattering cross sections for low energy analyses in IceCube*. Tech. rep. U. Alberta, 2020.

- [118] Jakob Feintzeig. “Searches for Point-like Sources of Astrophysical Neutrinos with the Ice-Cube Neutrino Observatory”. PhD thesis. University of Wisconsin-Madison, 2014.
- [119] N. Kulacz. “In Situ Measurement of the IceCube DOM Efficiency Factor Using Atmospheric Minimum Ionizing Muons”. MA thesis. University of Alberta, 2019.
- [120] Roger Barlow and Christine Beeston. “Fitting using finite Monte Carlo samples”. In: *Computer Physics Communications* 77.2 (1993), pp. 219–228. ISSN: 0010-4655. DOI: [https://doi.org/10.1016/0010-4655\(93\)90005-W](https://doi.org/10.1016/0010-4655(93)90005-W). URL: <http://www.sciencedirect.com/science/article/pii/001046559390005W>.
- [121] G. Bohm and G. Zech. “Statistics of weighted Poisson events and its applications”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 748 (June 2014), pp. 1–6. ISSN: 0168-9002. DOI: 10.1016/j.nima.2014.02.021. URL: <http://dx.doi.org/10.1016/j.nima.2014.02.021>.
- [122] C. A. Argüelles, A. Schneider, and T. Yuan. “A binned likelihood for stochastic models”. In: *Journal of High Energy Physics* 2019.6 (June 2019). ISSN: 1029-8479. DOI: 10.1007/jhep06(2019)030. URL: [http://dx.doi.org/10.1007/JHEP06\(2019\)030](http://dx.doi.org/10.1007/JHEP06(2019)030).
- [123] Thorsten Glüsenskamp. “Probabilistic treatment of the uncertainty from the finite size of weighted Monte Carlo data”. In: *The European Physical Journal Plus* 133.6 (June 2018). ISSN: 2190-5444. DOI: 10.1140/epjp/i2018-12042-x. URL: <http://dx.doi.org/10.1140/epjp/i2018-12042-x>.
- [124] T. Glüsenskamp. “A unified perspective on modified Poisson likelihoods for limited Monte Carlo data”. In: *Journal of Instrumentation* 15.01 (Jan. 2020), P01035–P01035. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/01/p01035. URL: <http://dx.doi.org/10.1088/1748-0221/15/01/P01035>.
- [125] I. Asimov and D. Shannon. *Franchise*. Creative Short Stories Series. Creative Education, 1989. ISBN: 9780886822323.
- [126] M. Kachelrieß and D.V. Semikoz. “Cosmic ray models”. In: *Progress in Particle and Nuclear Physics* 109 (Nov. 2019), p. 103710. ISSN: 0146-6410. DOI: 10.1016/j.pnpnp.2019.07.002. URL: <http://dx.doi.org/10.1016/j.pnpnp.2019.07.002>.
- [127] A. M. Hillas. “The Origin of Ultra-High-Energy Cosmic Rays”. In: *Annual Review of Astronomy and Astrophysics* 22.1 (1984), pp. 425–444. DOI: 10.1146/annurev.aa.22.090184.002233. eprint: <https://doi.org/10.1146/annurev.aa.22.090184.002233>. URL: <https://doi.org/10.1146/annurev.aa.22.090184.002233>.
- [128] Frank M. Rieger. *UHE Cosmic Rays and AGN Jets*. 2019. arXiv: 1911.04171 [astro-ph.HE].
- [129] “Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector”. In: *Science* 342.6161 (Nov. 2013), pp. 1242856–1242856. ISSN: 1095-9203. DOI: 10.1126/science.1242856. URL: <http://dx.doi.org/10.1126/science.1242856>.
- [130] IceCube Collaboration. *The IceCube high-energy starting event sample: Description and flux characterization with 7.5 years of data*. 2020. arXiv: 2011.03545 [astro-ph.HE].
- [131] IceCube Collaboration. “All-sky search for time-integrated neutrino emission from astrophysical sources with 7 years of IceCube data”. In: *Astrophys.J.* 2 (835 2017), p. 151. DOI: 10.3847/1538-4357/835/2/151. eprint: 1609.04981.
- [132] IceCube Collaboration. “Extending the Search for Muon Neutrinos Coincident with Gamma-Ray Bursts in IceCube Data”. In: *The Astrophysical Journal* 843.2 (July 2017), p. 112. ISSN: 1538-4357. DOI: 10.3847/1538-4357/aa7569. URL: <http://dx.doi.org/10.3847/1538-4357/aa7569>.

- [133] IceCube Collaboration. “The contribution of Fermi-2LAC blazars to the diffuse TeV-PeV neutrino flux”. In: *ApJ* 1 (835 2017), p. 45. DOI: 10.3847/1538-4357/835/1/45. eprint: 1611.03874.
- [134] IceCube Collaboration. “Constraints on Galactic Neutrino Emission with Seven Years of IceCube Data”. In: *Astrophys. J.* 67 (849 2017). DOI: 10.3847/1538-4357/aa8dfb. eprint: 1707.03416.
- [135] ANTARES and IceCube Collaborations. “Joint Constraints on Galactic Diffuse Neutrino Emission from the ANTARES and IceCube Neutrino Telescopes”. In: *Astrophys. J.* 2 (868 2018), p. L20. DOI: 10.3847/2041-8213/aaeecf.
- [136] J. P. Huchra et al. “The 2MASS Redshift Survey - Description and Data Release”. In: *Astrophys. J. Suppl.* 26 (199 2012). DOI: 10.1088/0067-0049/199/2/26. eprint: 1108.0669.
- [137] M. Rameez P. Mertsch and I. Tamborra. “Detection prospects for high energy neutrino sources from the anisotropic matter distribution in the local universe”. In: *JCAP* 03 (1703 2017). DOI: 10.1088/1475-7516/2017/03/011. eprint: 1612.07311.
- [138] M.F. Skrutskie et al. “The Two Micron All Sky Survey”. In: *The Astronomical Journal* 131 (2), pp. 1163–1183.
- [139] Steffen Heß, Francisco-Shu Kitaura, and Stefan Gottlöber. “Simulating structure formation of the Local Universe”. In: *Monthly Notices of the Royal Astronomical Society* 435.3 (Sept. 2013), pp. 2065–2076. ISSN: 0035-8711. DOI: 10.1093/mnras/stt1428. URL: <http://dx.doi.org/10.1093/mnras/stt1428>.
- [140] Pierre Auger Collaboration. “Observation of a large-scale anisotropy in the arrival directions of cosmic rays above 8×10^8 eV”. In: *Science* (357 2017), pp. 1266–1270. DOI: 10.1126/science.aan4338. eprint: 1709.07321.
- [141] Pierre Auger Collaboration. “Indication of anisotropy in arrival directions of ultra-high-energy cosmic rays through comparison to the flux pattern of extragalactic gamma-ray sources”. In: *ApJL* (853 2018), p. L29. DOI: 10.3847/2041-8213/aaa66d. eprint: 1801.06160.
- [142] K. M. Gorski et al. “HEALPix – a Framework for High Resolution Discretization, and Fast Analysis of Data Distributed on the Sphere”. In: *Astrophys. J.* (622 2005), pp. 759–771. DOI: 10.1086/427976. eprint: 0409513.
- [143] Gary J. Feldman and Robert D. Cousins. “Unified approach to the classical statistical analysis of small signals”. In: *Physical Review D* 57.7 (Apr. 1998), pp. 3873–3889. ISSN: 1089-4918. DOI: 10.1103/physrevd.57.3873. URL: <http://dx.doi.org/10.1103/PhysRevD.57.3873>.
- [144] IceCube Collaboration. “Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using Six Years of IceCube Data”. In: *The Astrophysical Journal* 833.1 (Dec. 2016), p. 3. ISSN: 1538-4357. DOI: 10.3847/0004-637x/833/1/3. URL: <http://dx.doi.org/10.3847/0004-637X/833/1/3>.
- [145] E. Waxman K. Murase. “Constraining high-energy cosmic neutrino sources: Implications and prospects”. In: *Phys. Rev. D* 103006 (94 2016). DOI: 10.1103/PhysRevD.94.103006.
- [146] M.D. Kistler H. Yuksel and A.M. Hopkins. “Revealing the High-Redshift Star Formation Rate with Gamma-Ray Bursts”. In: *Astrophys. J.* (683 2008), pp. L5–L8. DOI: 10.1086/591449.
- [147] IceCube et al. “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. In: *Science* 6398 (361 2018), eaat1378. DOI: 10.1126/science.aat1378. eprint: 1807.08816.

- [148] IceCube Collaboration. “Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert”. In: *Science* 6398 (361 2018), pp. 147–151. DOI: 10 . 1126/science.aat2890. eprint: 1807.08794.